Janet Support Manual

Janet Support Manual

There are many sources of information about Janet, but it can take time to access specific details. This manual has been produced to place all essential material in one document. It is aimed at systems administrators and technical contacts at Janet sites. It is designed to work on a number of different levels and contains both non-technical and technical details.

Most sections of this manual provide an introduction to a topic and in some cases more detailed information, plus pointers to other reference sources. The appendices contain either specific examples of documents such as a completed licence application form or further technical information.

Janet manages the operation and development of the Janet network. This is the United Kingdom's education and research network, managed by Janet on behalf of the combined UK Further and Higher Education Funding Councils represented by JISC (the Joint Information Systems Committee).

Fault Reporting

Fault Reporting

Janet operates a fault reporting process to deal with all problems at both a network and a site level, as quickly and efficiently as possible. However, this process can only work if all Janet-connected organisations are familiar with the correct reporting routes.

Reporting Problems

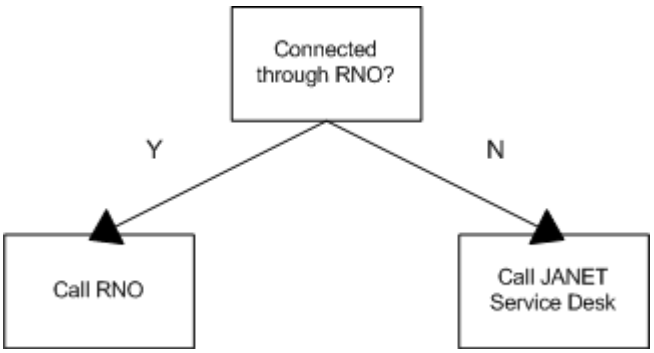

All customers who wish to report a problem with their connection to Janet should follow the route set out in their fault reporting letter. Contact information is sent to the management and technical contacts by e-mail on an annual basis and includes the current telephone, fax and e-mail address of the appropriate fault reporting contacts. For organisations connected to Janet through a Regional Network, this will be the appropriate RNO. For ease of reference, site contacts may wish to print off a Record of Network Information and Contacts form from Appendix 3 and add the fault reporting information. If, for any reason, a site is unable to contact their RNO helpdesk, please contact the Janet Service Desk.

Sites not connected to Janet through an RNO should report all faults to the Janet Service Desk. Examples of the fault reporting information supplied may be found in Appendix 4.

If site contacts require confirmation of the appropriate fault reporting route for their organisation, they should contact the Janet Service Desk.

The Janet Service Desk

The Janet Service Desk is the point of contact for reporting faults with the Janet Operational Services provided on the Janet backbone under the SLA with JISC [1]. This includes problems with:

- external circuits and all the associated routing and switching equipment
- services such as Mailer Shield, Web Mail and Web Hosting.

If the fault to be reported is known to concern the Janet backbone, it should be reported to the Janet Service Desk.

Who Should Report Problems

Problems may only be reported by staff in the computer department of an organisation or one of the designated Janet contacts, not by individual end-users. In order for the fault reporting mechanism to work efficiently, it is essential for each organisation to establish a clear reporting structure, and make users aware that the correct route for reporting problems is via the technical contact.

Please note that the contact information provided by the Janet Service Desk is only to be used by designated contacts. It must not be passed to end-users at the site.

Response Times

Under the provisions of the current Janet SLA [1], the Janet Service Desk will respond to fault reports within one hour during working hours.

Emergency Cover

Emergency cover is provided outside normal working hours. Contacts requiring assistance will be asked to specify whether the call is urgent or non-urgent. Urgent calls are relayed to an on-call engineer for immediate action; non-urgent calls are dealt with at the beginning of the next working day. The fault reporting information provided by the Janet Service Desk includes the telephone number for the emergency service. If the problem reported requires attention outside normal working hours then the site contact must be available to discuss the matter with the engineer when the call is returned.

Please note that the out-of-hours emergency telephone number is not advertised on the web and should only be given to staff who are Janet contacts.

Escalation Procedure

Janet provides an escalation mechanism for customers who are unable to contact their nominated fault reporting point or are dissatisfied with the general performance of the fault reporting process. In either instance, sites should provide full details of their problem to the Janet Service Desk [2]. The request will be logged and steps will be taken to address the problem. Janet will also investigate the reasons for the unsatisfactory service and provide a report of the action taken.

Trouble Tickets

The trouble ticket system was originally set up to share information about network problems and planned work between the networking experts who maintained JANET. Although access to trouble tickets is now freely available to support staff at sites with a Primary Connection to Janet, the terminology used in the tickets still reflects their very specialised origin.

All Janet-connected organisations are asked to set up a generic e-mail address using the format **operations@sitename.ac.uk** [3] and then use this as their internal mailing list for all technical staff at their site who should receive this type of information. Please contact the Janet Service Desk [4] once this address has been set up.

Janet Scheduled Maintenance Period

Scheduled maintenance work within Janet is normally timed to take place in the period between 07:00 and 09:00 on Tuesdays. All such work is advertised via the trouble ticket system. The Janet Service Desk will give prior warning to organisations directly affected by scheduled maintenance work.

Scheduled Work on Janet Site Networks

Please note that all Janet-connected organisations should advise their fault reporting point of any planned work on their local network that may make the site temporarily inaccessible.

Arrangements may then be made to circulate the details to other Janet sites as necessary.

Disaster Recovery Plans

All organisations should consider producing a network disaster recovery document as part of their overall site disaster recovery plans.

Points to consider include:

- identify staff responsible for each action required by the Recovery Plan
- keep machine rooms tidy to minimise fire risk
- maintain offsite secondary servers
- agree to exchange systems with other sites
- keep offsite copies of:
    - backups
    - the router configuration
    - critical mailing lists
    - configuration tables for vital services.

Further advice may be found at UCISA [5] and ACU [6] (Janet takes no responsibility for the content or accuracy of external web sites).

Connecting to Janet

Connecting to Janet

Primary Connections

All centrally funded HE organisations in the UK and centraly funded FE organisation, who take up the Jisc FE subscrition, are entitled to a connection to Janet

- Centrally funded HE organisations are entitled to a Primary Resilient Janet IP Connection, with resilinec up to 10GBits/s
- Centraly funded FE organisations, who take up the Jisc FE subscrition, are entitled to a Primary Janet IP Connection
- Organisations assiciated with the education and reserach sector also benefit from a preferential tarriff for Jisc services

A Janet IP Connection provides access to the full range of network services and support at levels specified in the SLC [7].

Other organisations may connect to Janet, provided they fulfil the eligibility requirements detailed in the Janet Network Connection Policy [8].

Overview of Connection Process for Additional Connectuions

The connection process is complex and involves many different parties, both internally and externally. The main point of contact during this process is the Connections Management Group (CMG). You can contact us on connect@ja.net [9]. The following paragraphs provide a brief overview of the connection process, which is more fully described at:

/support/connections/types-of-connection.html [10]

Once the formal quotation process has been completed and agreement to proceed is received a contract in the form of a declaration letter will be sent to the organisation wishing to connect, this also contains login details for completing the web-based JCUR. Guidance is provided throughout the online process.

The completed JCUR is submitted online and the signed declaration returned to the CMG. All organisations should contact the CMG for advice and assistance with completing the online form.

In signing up to Janet, an organisation is also signing up to the Terms and Conditions for provision of a Janet service. These can be found at https://community.ja.net/library/janet-policies/terms[illegible]on-janet-se... [11]

T[illegible]ess is illustrated by this flowchart:

[12]

The Connection Process

As part of the connection process each newly connected organisation must apply for IP addresses from Janet and if required a domain name under one of the .uk domains, at no charge subject to availability and eligibility. The Connections Management Section provides customers with information about the progress of their connection and can also deal with any queries that arise during this process.

Making a Connection Ready for Use

A connection is considered to be ready for use after:

- the circuit has been delivered
- IP addresses have been assigned
- all the cabling from the router to the NTU at both the connecting organisation and RNO are in place
- the organisation's router has been configured
- the backbone routing is in place
- a new domain name, if required, has been approved and delegated, and the customer has supplied the nameserver names and IP addresses.

However, it is recognised that the customer may not be using the circuit at this point if there are still unresolved issues or testing is required. The Connections Management Section will therefore contact the customer and ask for confirmation that the connection is in use. Once this is received, a date is agreed for invoicing purposes (for those organisations that pay for a connection to Janet).

Once the circuit has been delivered a Fault Reporting letter is sent in PDF format to the technical and management contacts at the organisation, which provides information about the fault reporting process. Fault reporting is covered in more detail in Section 1 of this manual and traffic monitoring is described in Appendix 9..

If the organisation is required to pay for the connection, an invoice for the charges agreed at the beginning of the process is sent after the Connections Management Section has checked that the customer is satisfied with the connection.

Further information about DNS, IP addresses, routers and e-mail is available in more detail elsewhere in this manual.

Multihoming

All organisations connecting to Janet must discuss any plans to retain or install an additional Internet connection through a commercial ISP with Janet. There are a number of implications and technical issues associated with multi-homing arrangements of which an organisation should be fully aware before an agreement is made with another ISP. Janet may make a charge for the increased amount of central administration and technical effort required (e.g. effective filters, aggregation rules) if a multihoming facility is provided.

Please contact the CMG for further information.

If an organisation connecting to Janet does not wish to set up a web server, it may be permitted to maintain an existing domain name and web site through a commercial ISP, subject to the preceding provisos about multihoming. Customers should contact the CMG.

Responsibilities of Janet-Connected Organisations

Once a site is connected to Janet it is expected to take responsibility for the following:

- ensuring that users accessing Janet through the organisation's connection are aware of the Janet AUP [13] and Security Policy [14]
- providing a secure environment for the Janet equipment installed on site
- providing technical staff with access to the Janet equipment on site
- monitoring the performance of its connection and reporting suspected faults promptly through the agreed channels
- providing assistance to Janet and law enforcement agencies in tracking down criminal activity
- appointing listed contacts (and where possible, deputies) for liaison functions and advising the Janet Service Desk when these change.

Acceptable Use of the Network

All organisations connected to Janet are subject to the Janet AUP [13]. This permits Janet to be used for any purpose that is legal, that is socially acceptable to the community Janet serves, and that does not cause degradation of the performance of the network beyond that which might reasonably be expected. All use of the network does, of course, have an impact on performance; the intent is to prevent reckless or inconsiderate activities by users at one organisation causing inconvenience to others. It is important that all users of the network are aware of the AUP. Connected organisations are therefore encouraged to set up a local AUP and require users to sign up to it. Some pointers on how to set up an AUP are provided in Section 9.

Further information about the Janet AUP is provided in Section 9 and at https://community.ja.net/library/acceptable-use-policy [13]

Recording Essential Information

It is recommended that all organisations maintain a record of information that is essential to running their Janet connection. The Record of Network Information and Contacts form at Appendix 3 can be used to record information such as IP addresses and domain names. This record will not only be a quick reference source for the member of staff who is responsible for the day-to-day running of the site network, but will also provide vital background information for new or temporary staff.

Contacts at Janet-Connected Organisations

The Janet Service Desk should be advised of any changes to site contact names or addresses. Janet cannot ensure that important information about the network reaches each organisation if contact details are not kept up to date. For example, changes to the fault

reporting information may not reach management and operational/technical contacts, or information about security threats might not be passed to the security contact.

The Janet Service Desk sends Janet customers a Contacts form to amend each year but ideally sites should provide details of changes as they occur. Customers may wish to consider using generic e-mail addresses for the roles, e.g. technical@customer.ac.uk [15].

Responsibilities of Janet Contacts

At least one named contact must be provided for the technical, management and security roles, even if the same person is covering each one. The following notes explain what is generally expected of each contact.

Management Contact

The management contact must be able to make decisions relating to, and have the authority to contract for, the Janet connection. This contact is also required to sign the declaration of compliance with the Janet AUP [13] and Security Policy [14] (the Connection Agreement). Janet(UK) usually passes this contact's details to RIPE [1], the Internet registry for Europe, as the 'owner' of the IP addresses that are being requested, because this is normally the person in overall charge of IP address assignment for the organisation. However, you can specify an alternative RIPE contact on the JCUR if this is not appropriate. The management contact will normally be added to the mailing list for *Janet News* and similar publications, and to e-mail lists for news and support items.

Technical Contact

The technical contact is usually responsible for day-to-day decision making on technical issues concerning the organisation's network and the Janet connection. The name of this contact will normally be passed to RIPE, unless you provide an alternative contact on the JCUR.

Computer Security Contact

The Janet Security Policy [14] requires organisations with Primary Connections to nominate one or more people to be the first point of contact in dealing with any security incident that affects the organisation. The nominated contact is expected to be someone with technical knowledge and also with management authority since they may be asked to disconnect a computer from the network as a matter of urgency. Since it may be necessary to speak to a security contact at short notice, organisations may wish to nominate more than one person. However, in these circumstances they are expected to share information about current incidents. Ideally, an out-of-hours telephone number should be provided for these contacts.

When a message relating to an incident is sent, it is important that the security contact acknowledges that it has been received and that the problem is being investigated. If Janet CSIRT [16] does not hear from the security contact that progress is being made to contain and resolve the incident then they are likely to escalate the problem through management to protect the operation of the network. The security contact will also receive general security information from Janet CSIRT relating to the prevention of security incidents. The contact is expected to pass these to appropriate people within their organisation and ensure that

appropriate preventive action is taken.

Security contacts are also responsible for the activities of any Sponsored and Proxy Connections hosted by their organisation, and need to ensure that they are able to communicate quickly with those sites in case of problems. Security contacts are added to the UK Security JISCmail list and the mailing list for Janet News. Names and phone numbers of security contacts may also be given to law enforcement agencies if their investigations involve a Janet site.

Out-of-hours Contact

Technical, operational or security problems may arise outside normal working hours. It is therefore essential that the out-of-hours contact is available and has sufficient authority to take the appropriate action when such problems occur, such as Janet CSIRT needing to make an organisation aware of a security threat. If it is not possible to contact an organisation when necessary then Janet(UK) will take whatever action is needed to protect the Janet network, even if this involves disrupting the service to an individual organisation.

Upgrading a Janet Connection

All enquiries about upgrading a connections (irrespective of the method of connection) must be directed to the CMG. They will be able to advise on the procedure and funding arrangements for the upgrade. Note that the cost of the upgrade might be the responsibility of the site requesting the increase in bandwidth. In these instances the CMG can also provide details of the tariff that applies.

Overview of the Process

The upgrade procedure has many similarities with the standard connection process. In each case a formal request for an upgrade should be sent by e-mail to CMG, connect@ja.net [9] who will acknowledge the enquiry, and confirm the arrangements and, if appropriate, the charges. Some upgrades require the installation of a new circuit whereas others can be achieved by the circuit supplier making a few adjustments to the existing access circuit. The organisation will be asked to confirm acceptance of the offer by e-mail and will also be required to complete the appropriate paperwork.

The CMG advises the nominated installation contact of the progress of the order at each stage of the upgrade process and deals with any queries. When the upgrade is complete, steps will be taken to cancel the redundant circuit (if appropriate). The customer is then sent an invoice for the cost of the upgrade, if a charge is to be made.

Moving a Janet Connection

If it becomes necessary for a customer to relocate the equipment connected to its access circuit to another building, or to another part of the building in which it is housed, the Connections Management Section must be advised. In many cases the access circuit is leased by Janet and the supplier will not take any action unless the arrangements for the relocation are made through Janet. If there is no record of who leases the access circuit, the customer should contact the CMG for clarification.

The customer will be required to meet the charges levied by the circuit supplier, if the move is made at the customer's request. In these circumstances the Connections Management Section will obtain a quote for the cost of the work and ask the customer to accept responsibility for the charges in writing, or by e-mail, before the work is carried out.

Disconnection from Janet

An organisation may cancel its connection to Janet at any stage. However, if the notice of cancellation is received during the first year of service, a charge may be made equal to the residual rental (from the circuit supplier) of the access circuit for the remainder of the contract, plus a charge to cover administrative costs. After the end of the agreed contract term, a connection may be cancelled upon 30 days notice being given. Notice of cancellation must be made in writing by the nominated Management contact to connect@ja.net [9].

Janet will make arrangements for the return of the IP addresses originally assigned from the Janet allocation, at the time of disconnection. All queries regarding the return of IP addresses should be sent to: ipaddress@ja.net [17]

Interconnect Connections

These are available to organisations responsible for the operation of a network, to connect that network to Janet subject to Janet and JISC's agreement. The network will normally be supporting the broader education and/or research community or be delivering educational services to communities that are not directly connected to Janet, such as schools. These are not subject to the Janet SLA [18]. Only an IP service is normally available — other Janet services are subject to further agreement.

Further information about Interconnect Connections can be obtained from the Connections Management Section.

Display as Single Column?:

Domain Name System (DNS)

Domain Name System (DNS)

The Domain Name System (DNS) allows a computer presented with a textual name to convert or map it to the numeric IP address of another computer with which it needs to

communicate, say to fetch a web page or deliver an e-mail. The process is called DNS resolution.

There are also occasions where the reverse is required and a known IP address needs to be resolved to the corresponding domain name. Such reverse lookups are often performed as part of an automated security check. Mail exchangers are a common example. For further information on Reverse Delegations, see Section 4: IP Addresses.

The DNS was originally conceived as a worldwide database capable of storing many types of data. There is no single authority responsible for the entire database. To enable manageability and distributed administration, domains are broken into separately managed units known as zones. A domain encompasses both the parent zone (e.g. ac.uk) and all child zones (e.g. site1.ac.uk, site2.ac.uk). The maintainer of the parent domain can delegate authority for a child zone to an individual or organisation, which then becomes responsible for the child zone's data.

The DNS is made up of a collection of Resource Records, containing all the Internet addresses and names in the world, together with two types of computer program that process these records and convert between them: nameservers and resolvers. There are various types of Resource Record:

- Address (name to number)
- PoinTeR (number to name)
- Mail-eXchanger (identifies a mail server)
- NameServer (identifies a nameserver)
- Start-Of-Authority (SOA; contains information about a set of resource records).

The Janet Technical Guide *The Domain Name System* is a good starting point if you want to know more about the DNS, or are considering setting up your own.

Nameservers

Nameservers are server programs, often running on dedicated computers, which hold the primary copies of information about the names and addresses within a particular Internet domain (for example **yoursite.ac.uk**). Their main purpose is to let other people look up the names of computers within your domain (e.g. the name of your web server, mail server, etc.) and convert them into the numeric IP addresses that let their computers communicate with yours. Programs such as web browsers running on computers outside your network will find where your nameservers are from the Janet nameservers for ac.uk, and will send simple requests to your primary or secondary nameservers for the DNS records they hold (but for no other records).

Primary and secondary nameservers

The nameserver that holds the primary copy of a zone file in which changes can be made to records is called the primary nameserver for that zone. The zone file contains the most accurate information about a specific domain over which this server has authority. Copies of the zone file will usually also be held on one or more other nameservers, known as secondary nameservers, which automatically update their information from the primary server when the zone file is changed. Consequently, both primary and secondary nameservers can answer

queries about the domain with authority, so they are referred to as authoritative nameservers.

To apply for the JANET primary and secondary nameserver services, please go to:
https://community.ja.net/library/janet-services-documentation/primary-na... [19]
and
https://community.ja.net/library/janet-services-documentation/secondary-... [20]

Before an organisation sets up its nameservers, it needs to choose a domain name and agree it with the administrators of the parent zone. This can be found by looking up the SOA record for the parent domain. If the domain is immediately under ac.uk — for example, **yoursite.ac.uk** — then the parent zone is JANET(UK)'s responsibility. If the domain is under another site — for example, **physicsdepartment.yoursite.ac.uk** — then the organisation itself is responsible for the relevant parent zone.

A domain only becomes visible to the Internet when its name has been registered and the parent domain contains the delegation (pointers) to its nameservers.

For information on obtaining a domain name, see the section Obtaining Domain Names.

Off-site Resolver

Resolvers are programs that handle the other end of the DNS resolution process. A client program, such as a web browser, will contact a resolver with a request for a lookup, for example to find the numeric IP address equivalent to a given Internet name. The role of the resolver is to formulate a DNS query that will answer the client's request and send that query to the appropriate nameserver to find the required information. When the resolver receives the answer to the query, it returns the information to the original client computer. Every computer on a local network must be able to contact a resolver before it can look up information in the DNS; the IP address of the resolver (and possibly also a backup resolver) must be entered into the computer as part of its initial configuration. Resolvers must be able to contact nameservers elsewhere on the Internet so they can follow any referrals and work through the tree of Internet names to find the nameserver able to answer each individual query.

Resolver activity is therefore quite different from authoritative nameserver activity, though the two functions can often be provided by a single computer. If a primary or secondary nameserver is within the local network then it may be possible to have it act as a resolver for local clients. It is not recommended to allow a local nameserver to act as a general resolver for external clients as this may conflict with its most important function, and subject the server to possible spoofing attacks as described in the next section.

To apply for the JANET Off-site Resolver Service please go to
https://community.ja.net/library/janet-services-documentation/site-resol... [21]

Security Matters and the DNS

A malicious third party that compromises an organisation's nameserver could modify DNS resource records, causing traffic to the organisation's other servers (e.g. web and mail) to be redirected elsewhere. This redirection would probably be to hosts under the control of the attacker. All network managers should ensure that they receive security advisories from Janet CSIRT [22] and from their operating system manufacturers, and that operating systems

are patched in accordance with the manufacturer's guidelines. Apart from these general precautions, there are several actions that may be taken to improve the security of your nameservers.

- **Restrict Zone Transfers**. A nameserver should never accede to a request for a zone transfer from just any device on the Internet. Generally speaking, a primary server should only perform zone transfers with its secondary. A secondary nameserver should not be configured to respond to any zone transfers requests at all.
- **Restrict Dynamic Updates**. A nameserver that is exposed to the Internet should not generally accept dynamic updates. If this is unavoidable for some reason, then the server should never accept updates from an unknown source.
- **Restrict Recursive Queries**. An Internet-visible nameserver is vulnerable to spoofing attacks if it answers recursive queries from any source. In this type of attack, the cracker directs a query about a zone under his control to the nameserver he wishes to compromise. The target nameserver is then forced to query the cracker's server and receives bogus data, which it stores in its cache. Sites may also wish to protect their network resources by prohibiting their nameserver from acting as a general resolver for anybody on the Internet.

Further, more detailed information is available in the JANET Technical Guide *The Domain Name System* in the section 'Securing a Public DNS Server'.

Obtaining Domain Names

Each Janet customer is entitled to one free name registration under a .uk domain as part of the connection package. The majority of organisations connected to Janet have at least one name registered in the ac.uk domain, if they are eligible, and may also have names registered in other domains, e.g. org.uk. The Janet Service Desk is responsible for administering this service for Janet.

All organisations connecting to Janet are required to indicate whether they wish to register a new domain name on the JCUR, which is then submitted to the Janet Service Desk for processing.

Eligibility for an ac.uk Domain Name

The Policy is available here. [23]

Choosing a Domain Name

An eligible organisation may register as many names within the ac.uk domain as it wishes, provided payment is received for all but the first name registered and the following rules about the format of the name are met:

- a request will not be allowed if it is for a name that is either one or two characters in length
- a request will not be allowed if it is for a name that is currently a second level domain name under the .uk domain or a top level domain name in the DNS: e.g. a name such as org.ac.uk is not allowed because 'org' is both a second level domain name within the .uk country code [org.uk], as well as being a generic top level domain name [.org].

Similarly, com.ac.uk is not allowed because 'com' is also a generic top level domain name [.com]
- the domain name must, in JANET(UK)'s opinion, be representative of the requesting organisation's name; if not, a detailed explanation is required
- the name requested must also be unlikely to present a substantial risk of confusion with other similarly named organisations or activities already registered under ac.uk
- organisations requesting generic domain names that could be applicable to a number of eligible sites must provide evidence that they have the backing and approval from the majority of relevant members of the UK academic and/or research community, in order to be permitted to have that generic domain name
- a project or service must be centrally funded and of wide relevance to the ac.uk community; it must be of at least two years duration and be UK-based
- internationalised domain names that start with the characters 'xn- -' (ie. 'xn' followed by two hyphens) may not be registered
- domain names must not coincide with internet protocols such as ‘www’, ‘ftp’, ‘dns’ or ‘whois’.

Subject to these constraints, names will be approved on a 'first come, first served' basis.

Registering Additional Domain Names

Once an organisation has connected to JANET, it may need to register additional domain names. All such requests must be channelled through the computing services department at the organisation to avoid confusion. The standard procedure is outlined on the JANET web site at https://community.ja.net/library/janet-services-documentation/domain-nam... [24]

Note that a fee will be charged for each successful request for registration.

Additional Names under ac.uk

Each name request should be made on the standard template that may be found on the JANET web site at https://community.ja.net/library/janet-services-documentation/register-acuk [25].

The template should be returned to naming@ja.net [26]

JANET customers are charged a one-off standard fee for additional domain names and are not required to register for the biennial maintenance charge applied to commercial hosts of ac.uk domains, for as long as they remain connected to JANET. They will also not be charged for any modifications provided that the changes made keep the domain name within the JANET network. These special arrangements exist because the JANET Service Desk handles the day-to-day administration of the ac.uk domain. The Janet Service Desk will accept payments by cheque (made payable to JANET(UK)), credit card or BACS, and will also be able to answer any queries. Customers are advised of the outcome of their request within five working days. Information about the current fees for ac.uk name registrations may be found at https://community.ja.net/library/janet-services-documentation/payments-a... [27].

An example of a completed domain name request template may be found in Appendix 7 .

Additional Names under .uk

Janet can arrange additional name registrations for Janet customers under other .uk domains, e.g. org.uk. All requests should be submitted on the standard template that is available at https://community.ja.net/library/janet-services-documentation/register-acuk [25]

Completed templates should be sent by e-mail to naming@ja.net [26].

Please note that JANET(UK) will add a handling charge to the standard fee for these domain names and it would therefore be cheaper for JANET sites to apply directly to Nominet [28]. Details of the current handling charge may be obtained from the Janet Service Desk.

Amending Domain Name Details

It may sometimes become necessary for an organisation to change the names or IP addresses of the nameservers that it uses. In these circumstances a member of the computing services department should complete the modification template at https://community.ja.net/library/janet-services-documentation/modify-dns-entries [29]

The template should be returned to naming@ja.net [26].

The JANET Service Desk will notify the individual who requested the change once the domain records have been updated.

If a third party has been running an organisation's nameservers and that arrangement is to be terminated, details of the new nameservers should be sent by e-mail using the same modification template to naming@ja.net [26].

In addition, a fax or scanned email with an original signature must also be sent to the Janet Service Desk on the organisation's headed paper to authorise the move to the new nameservers.

A domain name does not become active until it is matched to an IP address and that cannot happen until the Janet Service Desk are provided with full details of the nameservers, as specified on the template, and those nameservers are correctly set up. FE and specialist colleges may contact their JISC RSC if they require assistance with this process. All other organisations should contact the Janet Service Desk for advice.

Display as Single Column?:

IP Addresses

IP Addresses

Every organisation that wishes to send and receive e-mail, or gain access to the Internet, needs a globally unique address, known as an IP address. These addresses are numeric and uniquely identify one network interface on a computer. Each address is written as four fields, separated by dots, and each field can be a number ranging from 0 to 255, e.g.

193.63.117.225.

The address is divided into two sections: the network number and the host (computer) number. The network number will be the same for all computers on the same network. Each network interface on each computer on the same network must, however, have a unique host number.

Classless IP Addresses

The method of allocating address space according to need is known as CIDR. Further information about this strategy may be found in RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy* [30].

All information sent across the Internet is encapsulated into IP datagrams. At present IPv4 is used: see RIPE-405 [31], the RIPE document for IPv4 Address Allocation and Assignment Policies.

Since RIPE NCC documents are frequently updated, the status of this document can be checked at http://www.ripe.net/ripe/docs/titletoc.html [32].

IPv6 (Internet Protocol version 6)

IPv6 is the new version of IP, the common protocol underpinning all Internet communications. It is expected ultimately to supersede the current version, IPv4, in order to accommodate the rapid growth of the Internet. The transition from IPv4 to IPv6 will take some years, but in the meantime the two protocols can and will coexist and operate together in various ways.

Janet has been experimenting with IPv6 services for a number of years and has deployed dual-stack services on the JANET core which have been stable since their introduction in 2003. Steps are now being taken to introduce IPv6 into the JANET SLA.

The full Janet IPv6 Policy statement can be found at http://www.ja.net/products-services/janet-futures/ipv6 [33]

Applying for IP Addresses

IP Addresses are provided for use with a Primary or Sponsored Connection, but not for use with a Proxy Connection (see Section 2. Connecting to Janet [34] for information about the different types of connection).

- If you are connecting to JANET, to obtain an initial assignment of JANET addresses simply complete the relevant section within the JCUR or the Sponsored Licence Application Form.
- If you have an existing Primary JANET connection and need further address space to facilitate additional network services, send an initial email to ipaddress@ja.net [17] – you will then be provided with login credentials for the web-based request form.

The JANET LIR Team administers all IP address allocation requests, assessing the requirements of the customer against the appropriate policies before assigning any address space.

Internet Registeries

ICANN [35], based in the USA, controls the global allocation of IP address space. This body has allocated blocks of IP addresses to five Regional Internet Registries, one of which is RIPE NCC [36]. JANET and other networks in Europe may apply to RIPE NCC for IP address space.

The Janet Registry

Janet is a Local Internet Registry and has authority from RIPE NCC to assign IP address space to its customers. These assignments are made under the classless system. Under the current guidelines, each organisation needs to demonstrate that it will use at least 25% of the address space applied for immediately and 50% in the first year. Additional address space can be obtained if the requirements of an organisation change and they can demonstrate a physical need for more addresses. There is no charge for this service.

Assignment of IP Address Space

Janet customers in the process of connecting to the network apply for IP addresses using the web-based JCUR. FE and specialist colleges and PCDL centres may enlist the help of their JISC RSC [37] in completing the form.

Existing customers who wish to request additional address space should contact the JANET Service Desk for advice.

FE and specialist colleges and PCDL centres should consult their JISC RSC before applying for additional IP addresses.

Organisations wishing to use IPv6 address space for either experimental or production traffic can also contact the Janet Service Desk for login details for the web-based request form. The current policy allows a /48 prefix to be assigned per organisation without requiring any justification. Organisations that request a prefix greater than a /48 will need to provide suitable justification to illustrate their intended use of the address space. All IPv6 address space for the Janet network is assigned from within the 2001:0630::/32 prefix.

Returning IP Addresses

IP addresses assigned by JANET belong to the network, not the customer and cannot be routed across a network run by another ISP. The addresses must therefore be returned to JANET if an organisation decides to leave the network or has redundant address space. Organisations with a Primary Connection that enter into sponsoring arrangements with third party organisations must not use part of their address space for the sponsored site. Address space for this purpose will be assigned separately. All enquiries about returning address space should be addressed to Janet Service Desk.

Reverse Delegations

Janet are responsible for reverse delegations within the JANET IP address space, as allocated by the RIPE NCC. Therefore, Janet will undertake delegations within the following zones:

60.193.in-addr.arpa	80.194.in-addr.arpa
61.193.in-addr.arpa	81.194.in-addr.arpa
62.193.in-addr.arpa	82.194.in-addr.arpa
63.193.in-addr.arpa	83.194.in-addr.arpa
66.194.in-addr.arpa	87.81.in-addr.arpa
194.195.in-addr.arpa	219.212.in-addr.arpa
195.195.in-addr.arpa	0.3.6.0.1.0.0.2.ip6.arpa

Reverse zones for newly assigned IP ranges will be automatically delegated in the DNS as part of the IP assignment process. Alternatively, requests for the delegation or modification of a reverse domain should be submitted to naming@ja.net [26] using the form available at http://www.ja.net/forms/reverse-delegation-service-application-form/28 [38]
Completed examples may be found at https://community.ja.net/library/janet-services-documentation/example-re... [39]

All sites may contact the Janet Service Desk for advice and assistance when dealing with reverse delegation problems. FE and specialist colleges may also receive assistance from their JISC RSC with this process.

Hardware Addresses

Computers have a unique hardware address that is allocated to the network interface when the machine is manufactured. When a computer is connected to a LAN, the Address Resolution Protocol provides a mapping of IP addresses to the hardware or MAC address. This allows each computer on that network to recall the address of every other machine connected to the same LAN. The benefit of this system is that it reduces traffic on the LAN, because there is no need to query the address of another machine every time correspondence is exchanged. A computer will usually retain the same MAC address if it is moved to another network, although it may change if the hardware is upgraded or altered.

IP Subnet Addressing

The subnetting of IP addresses can help to make better use of IP address space, for example on expansion of a network. It also allows IP addresses on networks to be divided into multiple smaller networks or subnets. The addresses for the subnets are derived from the main network address by applying a subnet mask, and as such belong to that organisation. Effective use of subnets may remove the need to obtain additional address space. Subnets also provide some advantages over one large network:

- smaller networks are easier to manage and troubleshoot

- network traffic overall is reduced
- network security can be applied more easily at the interconnections between the subnets.

Further information about subnetting is provided in Appendix 8.

Private IP Addresses and Network Address Translation

The use of private addressing and NAT permits even fairly large organisations to make the best use of a small block of addresses allocated from the globally unique IP address space, and assists in conserving this limited Internet resource.

RFC 1918, *Address Allocation for Private Internets* [40], sets out the rules for using a set of reserved numbers (private addresses) for hosts on a local network. IANA has guaranteed that these addresses will never be used publicly on the Internet.

Since privately addressed nodes on a network do not have a presence on the Internet, there has to be a method by which these addresses are translated to globally routable numbers. This translation service is provided by NAT, which allows packets to be sent and received from outside the local network. Queries about private addressing and NAT may be sent to the Janet Service Desk. FE and specialist colleges may contact their JISC RSC for assistance on the use of private address space.

Routers

Routers

Routers are network devices that forward packets of data between different networks. A router between an organisation's LAN and JANET will not have a direct connection to every other router on the Internet. It is however possible to set up that router to forward packets to their destinations by the most efficient path. The router does this by referring to its routing tables, which list all the possible paths that data can take to get from source to destination IP address. Each router on the path repeats this process until the data reaches its final destination.

Routers permit each connected network to maintain its independent identity and IP address and can also facilitate the implementation of security procedures. It is possible for example to control access to a network from the outside world by using packet filtering (see Appendix 9).

Access Routers at JANET Sites

Customers are expected to have an on-site IP router to connect their LAN to JANET. The choice of router will be determined by the capacity of the access link and the organisation's specific requirements. Organisations requiring informal advice on router products that have

been successfully used by other sites on JANET should contact the JANET Service Desk for assistance.

The majority of FE colleges connecting to JANET under recent Government initiatives were supplied with a router as part of their new connection package. The JISC RSCs configured most of their routers and assisted in bringing their connections into operation.

As part of the process of setting up the router, customers need to select a suitable link-level protocol to carry IP traffic over the access link between their router and the router on the JANET core. The process of routing across the access link is sometimes referred to as encapsulation. Two link-level protocols are supported for an IP connection across JANET:

- PPP (Point to Point Protocol)
- HDLC (High Level Data Link Control).
- All organisations connecting to JANET are required to confirm which method of encapsulation will be used, by completing the appropriate section of the JCUR form.

Additional information about setting up a router may be found in Appendix 10.

Interfa?ces

The customer's access router needs a suitable WAN interface to connect to the JANET access link. If the bandwidth of an organisation's connection is 2Mbit/s or less, the PTO will normally present the leased line at the site's NTU with an X.21 DCE interface. Most routers are compatible with this type of interface and it should be possible to connect the circuit interface to the access router with an X.21 cable.

If it is not possible for the PTO to deliver a 2Mbit/s connection with an X.21 interface, the leased line will be presented at the NTU with a G.703 interface. This type of interface may also be used for 34Mbit/s connections. In these circumstances however, it will be necessary to install a DSU converter between the G.703 interface and the interface on the access router (for further information, see Appendix 11).

Organisations requiring a DSU will be supplied with a recommended box and a pair of three metre coax cables to connect it to the telecommunications termination point. If longer cables are required, the installation contact should give the JANET Service Desk prior notice of the length needed. The connecting organisation is responsible for the provision of an X.21 cable to connect its routing equipment to the DSU. The cost of the DSU and cables may be included in the connection package or charged separately, depending on the funding arrangements for the connection.

Once the connection is operational, the DSU becomes the property of the connecting organisation, which then assumes responsibility for arranging suitable maintenance cover for the DSU and cable. If there are any faults on the DSU or cable after it has been installed, the organisation is responsible for any repairs.

However, for FE colleges the DSU maintenance arrangements may vary. The organisation responsible for maintaining the site router is usually also responsible for maintaining the DSU.

It is acceptable for a connecting organisation to purchase and install a DSU on-site. However, the DSU must match the unit installed at the Janet Point of Presence and the organisation must also purchase all of the cables required. The Janet Service Desk can provide details of recommended DSU equipment for use on Janet.

Additional information about interfaces may be obtained via the Janet Service Desk. FE and specialist colleges may consult their JISC RSC if they require further advice.

Router Setup and the Janet Netsight Service

Organisations that have received assistance from contractors in setting up their access router should be aware of the requirements of the Janet Netsight network monitoring system. In order to monitor access links, the Janet NOC and the RNOs need to ping the IP address of the router interface that supports the customer's access link. A small number of sites have in the past set up routers that do not respond to pings. Please note that if a site router is set up in this way, there will be problems in providing the link status and traffic statistics via Netsight. There may also be time delays in identifying faults on these access links.

Janet and the RNOs will be happy to discuss which IP addresses should be allowed to ping a site access router. Please contact the Janet Service Desk for further information. FE and specialist colleges may initially seek advice from their JISC RSC.

Ownership and Maintenance of Routers

Some Janet sites own the router that connects their LAN to the Janet access link. In this case they are also responsible for the management and configuration of these routers and for maintenance arrangements.

There are, however, a large number of FE colleges whose router was supplied and is owned by Janet. In these circumstances the router is covered by a Janet maintenance contract. Most of these sites also receive assistance in supporting their access router from their JISC RSC or RNO. Janet informs the RSC of the fault reporting number for router problems; the RSC decides whether or not to pass the number on, depending on the level of support they provide to colleges.

Router maintenance includes both hardware and software support. The manufacturer Cisco® periodically releases newer IOS images with bug fixes and/or new features. Upgrading a router's software should result in improved system performance and reliability.

Reconfiguring a Site Access Router

Please note that reconfiguration of a router should only be undertaken by experienced technical staff. The site network may become inaccessible to the outside world or the

equipment may be badly damaged if a mistake is made. It is also possible that the maintenance contract may be invalidated. Please contact the Janet Service Desk if it is not clear who should be responsible for making changes to a configuration of a site access router after the site's JANET connection has been brought into service.

Janet Managed Router Service

Janet provides a Managed Router Service for organisations connected to the network, under which the site access router is monitored and managed as part of the Janet service. Janet's contractor undertakes all fault diagnosis and resolution work on these routers, either remotely or by an engineer visiting the site. Maintenance of the hardware is also included as part of the service. Maintenance of routers provided to FE colleges by Janet through the RSCs is available as a chargeable, opt-in service. Further information is available at http://www.ja.net/products-services/janet-connect/managed-router-service [41].

Router Security

Even if routers are only being used to transfer IP traffic, it is imperative that their security is not compromised. If intruders manage to obtain control of a router or firewall then they will be able to remove (temporarily or permanently) any traffic management rules used to protect the network. They may also be able 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, such as those made by Cisco®, have two levels of privilege, each protected by a separate password. The lower level gives 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 for privileged accounts on other IT facilities.

If a router can be managed over a network connection, it is possible for hostile attackers to try to access the management function, just as legitimate administrators would. Such attacks may come from inside or outside the organisation. Wherever possible, the router should be set to refuse connection requests that 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 in the following sections) 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 and HTTP.

Simple Network Management Protocol

This has an authentication mechanism that uses 'community strings'. In effect these are passwords, so they 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 the 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 are configured to ensure that SNMP can only be used to read information and not to update it.

Web-based Interfaces

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 the SSL protocol to provide encrypted communications, but unless the network device also supports this, 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 protocol in place of Telnet, SNMP version 3, SSL for web interfaces - and these should be chosen whenever possible.

Enquiries about the security aspects of setting up an access router may be directed to JANET CSIRT or the JANET Service Desk. FE colleges may also contact their JISC RSC for advice.

Further information is available in Section 7: Security.

Open Mail Relays in Janet

Open Mail Relays in Janet

Open relays allow any combination of origin and destination address, and are frequently abused by advertisers and others to distribute UBE. This will usually overload an organisation's mail 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 blocklists used by many network operators and ISPs to reject all e-mail and other traffic. Advice on preventing relaying is available from the MAPS website:

http://www.mail-abuse.com [42].

Another term used in this context is UCE, but for most purposes the 'bulk' aspect of the problem is more important than the 'commercial'.
There are now a number of bodies that seek to make UBE unworkable and one of their most effective activities is the maintenance of public blocklists of systems or networks that are in some way sources of UBE, particularly open relays. Of the well-known services, Janet uses certain services from the MAPS RBL™, a subscription system for creating intentional network outages ('blackholes') for the purpose of limiting the transport of known-to-be-unwanted mass e-mail. The policies and methods of some other lists are more controversial.

Janet has subscribed to some MAPSSM services on behalf of all Janet customer organisations and they are encouraged to use the RBLSM. Guidance notes on how to use these services are available to assist managers or administrators of mail services within Janet-connected organisations from https://community.jisc.ac.uk/library/janet-services-documentation/dns-al... [43]

E-mail and Security Issues

Janet-connected orgnaisations MUST NOT allow any comoputers on thier networks to be used to relay UBE.
If a Janet-connected organisation is relaying messages without authorisation, the manager of the e-mail system should take appropriate action to repair the open relay.

E-mail is one of the most widely used services on the Internet 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 an e-mail system is therefore to disable all these unnecessary services and install extra protection at the network level, to help avoid problems that will undoubtedly spring up in future. For those hosts that do need to provide a mail service there are less powerful alternatives to sendmail available, which may be sufficient for many

situations while also being easier, and therefore less error-prone, to set up. Further information about security issues may be found on the Janet CSIRT web pages.

All the usual network level threats and consequent countermeasures apply to computer and networking equipment associated with e-mail provision. As with other application services, there are also issues specific to the nature of mail and the way it is used and abused on the Internet.

General Countermeasures

Ensure that an intruder cannot take control of mail systems by:

- limiting connections with packet filters at firewalls and routers
- disabling unnecessary services on servers and workstations
- configuring servers to accept connections only as authorised
- installing patches and updates promptly for operating systems, mail and other applications and anti-virus software.

Specific Threats and Countermeasures

Monitor the service to establish what is a normal level of activity, and to recognise signs of overload before they cause difficulty. Document actions that might need to be taken if a problem occurs which requires disconnecting the mail server.

Unsolicited Bulk E-mail for their Own Users

Consider configuring the mail server to consult one or more DNS Block Lists about the source IP address before accepting each connection. The Janet mirror of the MAPS RBL+™ is one such list conveniently available for organisations connected to Janet.

Consider central filtering of incoming messages by pattern matching, or support so that the users can filter their own mail.

Relaying Through the Mail or Proxy Servers by Bulk Mailers

This can lead to overload, damage to reputation and blocking of mail to other places. Consider restricting access to TCP port 25 (SMTP) so that e-mail traffic can only travel by the intended route through the network. Similar considerations apply to TCP port 587, which RFC 2476 assigns for message submission.

For other issues surrounding e-mail relays, see the separate Janet documentation at https://community.ja.net/library/janet-services-documentation/janet-csirt

[22]

Open proxies are systems not primarily for mail use that accept some sort of inward connection and allow it to set up an ongoing connection that may be a mail transfer. Typically the incoming connection is web or HTTP on TCP port 80, and it is intended that client computers within the network can send all their web requests through it. SOCKS on port 1080 is another proxy protocol. Other ports are sometimes used, and an incorrectly configured web server can show the same behaviour.

Elimination of open proxies follows much the same pattern as elimination of open relays: examination of possible paths through one or more systems in the network and careful configuration and checking of firewalls, servers and client computers.

Introduction of Viruses and Other Malicious Software Through E-mail

If possible, use anti-virus software to scan incoming messages both at the mail server and on client computers; keep it up to date and regularly scan all the computers as viruses may arrive by routes other than e-mail.

Trojan Software Performing Bulk Mail Abuse

Software introduced into servers or client computers as a worm or virus by user indiscretion or by some intrusion may act as a proxy or may originate bulk mail on its own. Such rogue software installed through a system compromise can be very hard to detect on the machine affected, but routine monitoring of patterns of network traffic can alert the administrator to an incident and the headers of any mail sent will normally give some pointers to the source. It will often be necessary to rebuild a machine after such damage, and then try to find how the intrusion occurred to reduce the likelihood that it will happen again.

Further information on the use of e-mail, for both organisations and individuals can be found in Appendix 12.

Display as Single Column?:

Security

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 [14].

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.hmso.gov.uk/acts/acts1998/19980029.htm [44]

The Regulation of Investigatory Powers Act 2000 can be found at
http://www.hmso.gov.uk/acts/acts2000/20000023.htm [45]
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_... [46]
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.

Janet Services

Janet Services

Please visit the Products and Services section [47] of the Janet website.

Janet Policies and Legal Requirements

Janet Policies and Legal Requirements

Janet manages Janet in accordance with policies set by JISC.

Janet Network Connection Policy

The Janet Network Connection Policy [8] defines explicitly who can connect to Janet.

Primary Connections

All FE and HE organisations and Research Councils are entitled to a Primary Connection to Janet. This type of connection provides access to the full range of network services and support at levels specified in an SLA between JISC and Janet. JISC can also allow other bodies to have a Primary Connection if they are primarily engaged in education or research, or will only use the connection for collaborative research.

Further information about connecting to Janet as a primary site may be found in Section 2 of this manual.

Interconnect Connections

These are available to organisations responsible for the operation of a network, to connect that network to Janet subject to Janet and JISC's agreement. The network will normally be supporting the broader education and/or research community or be delivering educational services to communities that are not directly connected to Janet, such as schools. Interconnect Connections are not subject to the Janet SLA. Only an IP service is normally available - other Janet services are only provided by arrangement.

Further information about Interconnect Connections can be obtained from the Janet Service Desk [48].

Janet Acceptable Use Policy

All organisations connected to Janet are subject to the Janet AUP. This permits Janet to be used for any purpose that is legal, that is socially acceptable to the community Janet serves, and that does not cause degradation of the performance of the network beyond that which might reasonably be expected. All use of the network does, of course, have an impact on performance; the intent is to prevent reckless or inconsiderate activities by members of one organisation causing inconvenience to others.

A key concept of the AUP is that activities, other than those that are detailed below, are permissible when they are in accordance with the aims and policies of the organisation concerned. An organisation may therefore make its own policy regarding recreational use of Janet by its staff and students, or regarding the use of Janet by departments engaged on external contracts, or by staff engaged upon authorised private consultancy work. Most categories of unacceptable activity are designed to take account of relevant legislation.

Janet may not be used for any of the following:

- the infringement of copyright
- hacking, or other deliberately disruptive activity
- the transmission or creation of obscene or offensive material

- the transmission or creation of material of a threatening nature, or intended to harass, frighten etc
- the transmission or creation of material of a libellous nature
- the transmission of unsolicited commercial or advertising material or similar activities (spamming).

Where advertising material is embedded within, or is otherwise part of a service to which the user has chosen to subscribe, it is deemed to be permissible. Similarly, Janet would be prepared to sanction the use of Janet to transmit obscene or offensive material where this was necessary in pursuance of a properly supervised research project. In these circumstances the organisation must obtain permission from Janet prior to commencing the project, accept responsibility for the legal restrictions that exist and ensure that they are adhered to.

Note that the organisation may be legally liable for serious breaches of these restrictions. Where an organisation is shown to have breached the AUP, JISC may decide to either temporarily suspend the Janet service or withdraw the service for an indefinite period.

A copy of the AUP is sent to each organisation when an initial enquiry is made about the connection process. Additional copies may be obtained from the Janet Service Desk, or the Janet web site at https://community.ja.net/library/acceptable-use-policy [13]

Producing a Local Acceptable Use Policy

All Janet-connected organisations should formulate a local AUP and ask staff and students to sign a declaration to confirm that they will abide by its rules. Students under the age of 18 should have their forms countersigned by their parent or legal guardian.

Suggested Headings for Formulating an AUP

- Introduction (on the need for users to conform to standards of use and behaviour)
- Whom the policy affects
- List of equipment, networks, and data to which the policy refers (i.e. all, but specify items such as 'all servers, PCs, laptops, systems')
- Access to the network(s) (mention use and changing of passwords)
- Viruses and virus checking
- Access to and use of the Internet
- Use of discussion groups and chat (if allowed) (protocols and etiquette should be included)
- Use of software (cover licensing aspects and specify the screensaver software to which users should be limited)
- Copying software (the illegality should be stressed and consequences for students/employees made clear)
- Use of e-mail
- The Data Protection Act 1998
- The Computer Misuse Act 1990
- The Copyright, Designs and Patents Act 1988
- Physical security of computing equipment
- Backing-up strategies and rules.

- UCISA publishes model regulations for use of organisational IT facilities and systems at http://www.ucisa.ac.uk/publications/modelregs/modelregs.aspx [49].

Advertising on Janet

There is some limited advertising on the network at present. The formal policy on advertising is set by JISC in the AUP, and Janet is responsible for applying that policy. The Janet Factsheet Advertising describes how this policy is interpreted.

Janet Security Policy

Copies of the Janet Security Policy are available at https://community.ja.net/library/janet-policies/security-policy [14] or from the Janet Service Desk.

Janet is an open network that can be accessed worldwide via the Internet. As such it is subject to security threats from both external and internal sources. Security problems on the Internet or at specific Janet sites can easily spread their impact around the network. Many organisations now rely on the network and connected systems to do their teaching, research and administration. It is therefore increasingly important to protect against security incidents, for example:

- breaches of confidentiality, ranging from intrusion of privacy to theft of intellectual property
- loss of integrity - computers and the information they contain may be modified, casting doubt on the accuracy of the results produced, whether they relate to scientific research or course assessments
- failures of availability, if vital information is lost or destroyed by accident or malice, or if the networks or systems are unavailable due to failure or inappropriate use
- damage to reputation, if intruders boast of their success in attacking the organisation, or use its systems to disseminate unsolicited, unwanted and possibly offensive and/or illegal material
- legal liability, if the organisation is unable to meet its legal obligations as a result of computer failure or misuse.

Each organisation and user connected to Janet is required to comply with the Janet Security Policy. Organisations with a Primary Connection and others who connect third parties to the network have particular responsibility for the security of both their connection to Janet and that of any Sponsored or Proxy Connections made via their site. They must also ensure that information about security problems can be communicated both within the organisations they provide connections for, and between those organisations and Janet.

All organisations with a Janet connection have a duty to:

- enable users, through training, procedures and systems, to use the network safely
- manage and be accountable for access to Janet by individual users
- manage the risk of insecure network devices and take recommended security measures

- investigate, contain and resolve breaches of security.

All users are required to abide by both Janet-wide policies and those local to their organisation and location, and must cooperate with their organisation and the network operator. In particular, they must follow good security practice and not act in a way that puts the network or connected systems at risk.

The Janet Security Policy recognises that different approaches to security will suit different organisations, and leaves it up to each organisation to choose an appropriate way to meet its obligations under the Policy.

Security Contacts

It is essential that organisations with a Primary Connection have at least one nominated security contact available to provide and receive information on behalf of their organisation. It is accepted that the level of cover will vary depending on the size of the organisation concerned. However, all organisations must accept that, in an emergency, it may become necessary to temporarily disconnect the site if the security contact cannot be reached by Janet CSIRT.

Further information about the Janet Security Policy and advice on setting up a suitable security system for an organisation may be obtained from Janet CSIRT. Section 7 also covers security issues in more detail.

Legal Requirements

The operation and use of networks is subject to various legal requirements. Current information about requirements that are particularly relevant to network and system managers can be found at https://community.ja.net/blogs/regulatory-developments [50]

More general information is available from the JISC Legal Information Service, including a brief guide to IT Law for FE and HE Senior Management. Network and system managers should be familiar with at least the relevant provisions of the following Acts:

Computer Misuse Act 1990

http://www.hmso.gov.uk/acts/acts1990/Ukpga_19900018_en_1.htm [51]

Data Protection Act 1998

http://www.hmso.gov.uk/acts/acts1998/19980029.htm [52]

Regulation of Investigatory Powers Act 2000

http://www.hmso.gov.uk/acts/acts2000/20000023.htm [53]

Malicious Communications Act 1988

http://www.hmso.gov.uk/acts/acts1988/Ukpga_19880027_en_1.htm [54]

The Law Relating to Third Party Access to Data

The London InterNet eXchange (LINX®) has published a Best Current Practice document on privacy, which contains useful guidelines on dealing with statutory notices.

Further information about logfiles and third party access to data, including the new RIPA powers, may be found in the Janet Guidance Note on Logfiles.

See also the Data Protection Act 1998 and the Regulation of Investigatory Powers Act 2000.

JISC Legal Information

http://www.jisclegal.ac.uk/ [55]
Display as Single Column?:

Appendix

Appendix

1. The Structure of Janet [56]

2. Managing and Supporting Janet

3. Record of Network Information and Contacts

PDF available

4. Examples of Fault Reporting Letters

Example of a Fault Reporting Letter for Janet-connected Organisations

Fault Reporting Letter for Sites Subscribing to the Managed Router Service

5. No longer used

6. No longer used

7. Example of a Completed Request for a New Domain Name under ac.uk

PDF available

A blank template is available from:

/support/domain-name-registration/register.ac.uk/

A blank template for requesting a Domain Name under gov.uk is available from:

/services/domain-name-registration/register.gov.uk/

A blank template for requesting a modification to a Domain Name under ac.uk or gov.uk is available from:

/support/domain-name-registration/modify-dns-entries/modify-dns-entries.html

8. IP Subnet Addressing

9. No longer available

10. No longer available

11. No longer used

12. E-mail for Users and Organisations

Managing and Supporting Janet

Managing and Supporting Janet

Janet is responsible for the overall management of Janet. However, there are also other organisations that are engaged in running all or part of the network, or providing support. The following paragraphs describe how these bodies are related.

Janet

Janet is responsible for managing and developing Janet, and provides a comprehensive support system dedicated to the needs of the education and research community.

The Janet Service Desk

The Janet Service Desk is the first point of contact for queries about Janet. They can be contacted by e-mail, telephone, fax or post. All enquiries are handled in accordance with Janet's SLA and will be acknowledged within two hours of receipt during working hours.

Contact details for the Janet Service Desk [48]

Janet CSIRT (Computer Security and Incident Response Team)

Janet operates Janet CSIRT, one of the services provided for Janet customers. The team

provides emergency assistance when Janet sites are involved in security incidents, and publishes advice to raise awareness and improve the security of computers and networks. Core office hours are 08:00 to 18:00 Monday to Friday. A reduced service is available from 18:00 to 00:00 Monday to Friday and 09:00 to 17:00 at weekends, via a message answering service and a call-out process. There is no cover on UK-wide public holidays.

Contact details for Janet CSIRT [57]

Further information about the services provided by Janet CSIRT can be found in Section 7 and at https://community.ja.net/library/janet-services-documentation/about-csirt [58]

Regional Networks

The majority of sites are connected to Janet through a Regional Network. These networks are contracted by Janet to provide certain telecommunication services in their region and access to the Janet backbone.

A PDF diagram showing the Regional Networks is provided here.

The Regional Networks are run by local partnerships of educational organisations that often work with local businesses and regional development agencies. They provide a very cost-effective means of providing network connections into colleges and universities and can support connections into local businesses, learning centres, schools and other organisations within their geographical area.

Contracts have been negotiated with the Regional Networks under which they are expected to perform in accordance with a standard Service Level Definition agreed with Janet.

JISC (the Joint Information Systems Committee)

Janet is publicly funded and has a responsibility to implement the policies established by JISC. This is a strategic advisory committee working on behalf of the funding bodies for Further and Higher Education in England, Scotland, Wales and Northern Ireland. It also works in partnership with the Research Councils and the Learning and Skills Council. Further information about JISC may be found at http://www.jisc.ac.uk/ [59]

JISC RSCs (Regional Support Centres)

The RSCs were set up and are managed by JISC. There are 13 JISC RSCs: nine in England, one in Wales, one in Northern Ireland and two in Scotland. All are run by HE/FE partnerships or consortia. Their initial role was to assist in the connection of FE colleges to Janet and to provide ongoing technical support and training on JISC services. From August 2003 their remit has included supporting FE organisations to realise their ambitions in deployment of ICT in order to achieve their organisational mission. They:

- support FE organisations in the development of e-learning in each region
- act as a two-way, prime communication link between JISC and FE organisations
- work in partnership with regional and national agencies to gain maximum value from support activity.

Further information about the services provided by the JISC RSCs may be found at http://www.jisc.ac.uk/whatwedo/services/as_rsc.aspx [60]

Details of individual RSCs are available from http://www.jisc.ac.uk/whatwedo/services/as_rsc/rsc_home/rscs_contact.aspx [37]

The Structure of Janet

The Structure of Janet

The Janet backbone

The Janet network is based around a high-speed, high bandwidth backbone. The latest version of the backbone, which was designed to provide a high level of reliability as well as substantially increased capacity, came into full service in early 2007. It currently runs at 100Gbit/s.

A PDF map of the present Janet backbone is provided here and a schematic is available here.

Regional Networks

There are currently 12 Regional Networks in operation in England, one in Wales, five in Scotland and one in Northern Ireland. These networks provide the principal connection points for Janet in their areas.

Each of the Regional Networks has dual, diversely-routed connections to the core network, either via two point-to-point connections between each Regional Network entry point and separate locations on the core, or via a collecter arc where the network sweeps through a number of Regional Network entry points, connecting to two different core locations.

A PDF map showing the location of the Regional Networks is provided here.

Core Points of Presence (C-PoPs)

The traffic flow around Janet is controlled by routers situated at the C-PoPs, which are sited at critical 'crossroads' on the network. The C-PoPs house the network's switching and content equipment.

The Technology

The Janet backbone uses a technology called SDH that provides a high degree of resilience by re-routing traffic around a break in the network within 50 milliseconds of the break being detected. This resilience guarantees an availability of 99.95% for each link, together with a target time to repair of five hours.

IP Subnetting

IP Subnetting

This appendix discusses what subnetting is and gives some examples of how IP addresses can be separated into a network part and a host part. It describes a simple college subnet and shows how subnet masks are used to decode IP addresses. The final section covers binary numbers and converting to and from decimal.

Subnetting provides a mechanism for dividing a large unwieldy network into smaller sections that can be managed more effectively. It also offers better network security as measures can be applied at network boundaries and may reduce overall network traffic.

Subnetting does not provide additional IP addresses - only private IP addressing using NAT does that.

What is Subnetting

Subnetting is a way of partitioning a network at the IP level by dividing a block of addresses into a number of smaller sets.

To understand what is happening, first look at some small numbers rather than a full IP address.
Take the set of 'addresses' 100 to 129.
If computers used decimal representation, this network of 30 addresses might naturally be split into:

100 - 109

110 - 119

120 - 129

The first two digits represent the subnet number (10,11,12), while the final digit gives the address within the subnet.

Since computers use binary representation (base 2) rather than decimal (base 10) an equivalent pattern only emerges if the addresses are put into their binary form and the sets are broken into powers of 2 rather than powers of 10. (See below for an explanation of binary numbers.)

Looking at the first few addresses in the set, if the numbers are broken after the fifth digit then the two subnets, 01100 and 01101, are clearly visible. (Note that the network does not split into the same subnets as above.)

Decimal	Binary	Subr
100	01100100	0110
101	01100101	0110
102	01100110	0110
103	01100111	0110
104	01101000	0110
105	01101001	0110
106	01101010	0110
107	01101011	0110

etc.

Since an address could be split into network address and host number at any point, the boundary for a particular network must be known. This is achieved using a subnet mask that is the same length as the address, and has ones in the subnet address position and zeroes afterwards. For the above example the subnet mask would be 11111000 - five ones for the five digits of the subnet and zeroes for the remainder of the address - or equivalent to 248 in decimal.

Note that since each subnet has an associated subnet mask, the sizes of the subnets, and consequently the numbers of addresses in them, can vary.
This form of mask enables the address of the subnet to be extracted by performing a logical AND on an address and the subnet mask, which in turn makes it easy for the routing software to decide whether a destination address is in the local network or not and route the packets accordingly.

Full Length Examples

If the principle is extended to full length IP addresses, it produces longer addresses and masks. An IP address is four bytes long and these four bytes hold both the network address and the host number within the network. The associated mask defines the extent of each part of the IP address.

Considering the IP addresses 193.62.83.10 and 193.62.83.108 and an associated mask of 255.255.255.224:

The first becomes:	11000001 00111100 01010011 00001010
With a mask of	11111111 11111111 11111111 11100000
ANDing gives	
Subnet address of	11000001 00111100 01010011 00000000

So IP address 193.62.83.10 lies within the subnet that starts at 193.62.83.0.

The second becomes:	11000001 00111100 01010011 01101100
With a mask of	11111111 11111111 11111111 11100000
ANDing gives	
Subnet address of	11000001 00111100 01010011 01100000

So IP address 193.62.83.108 lies within the subnet that starts at 193.62.83.96

Note that the mask used above contains 27 ones. An alternative way of specifying the mask that is associated with an address is to append that number, which is known as the 'prefix length', to the IP address: e.g. 193.62.83.10/27

Defining Subnets

Anycollege has decided to subnet their facilities, to separate the administrative functions from the faculties and to separate the computer department from the rest of the network.
They have been allocated Janet IP addresses starting at 193.62.83.0. They want to allow for a maximum of 128 hosts for the faculties and 32 each for administration and computing. To conserve address space, the start positions of the subnets have been selected to minimise padding in the address space by starting with the largest subnets and working down in size.
They have therefore defined the following structure.

Name	The relative position* in this address space at which this subnet starts	The subnet mask of this sub
Faculties	0.0.0.0	255.255.255.128
Administration	0.0.0.128	255.255.255.224
Computing	0.0.0.160	255.255.255.224

* The relative position of the first subnet is always given as 0.0.0.0, i.e., the start of the address space allocated to the college (but see note below). For each subsequent subnet the start position is selected to allow for the maximum possible number of hosts in the subnets that precede it.

NOTE: Historically the use of 'all ones' or 'all zeros' in either the subnet or the host part of the address was not allowed. The restriction on using 'all ones' in the host part remains as this represents the subnet broadcast address (i.e. 193.62.83.127 is a broadcast on 193.62.83.0/25). The restriction on the use of the 'all zeros' address may still be in place for older equipment.

The subnet mask defines the number of hosts available on a subnet. The mask 255.255.255.224 leaves five bits for the host address, which could theoretically hold 32 addresses. However, in practice there are only 30 generally available addresses on the two small subnets and 126 on the large one.

This reduction in the number of available addresses is the price that is paid for subnetting.

Decoding IP Addresses

The following table gives examples of IP addresses in Anycollege's small networks.

		Network
Mask	255.255.255.224	11111111.11111111.111111
Address (1)	193.62.83.135	11000001.00111100.010100
Address (2)	193.62.83.164	11000001.00111100.010100
Address (3)	193.62.83.144	11000001.00111100.010100

ANDing the mask and addresses shows clearly which subnet the IP-address belongs to: addresses 1 and 3 belong to the subnet starting at (1000 0000) or .128; address 2 belongs to the subnet starting at (1010 0000) or .160.

The following table shows an IP address in the larger subnet. Note that the mask used is shorter, thereby allowing for more hosts.

		Network
Mask	255.255.255.128	11111111.11111111.111111
Address (4)	193.62.83.96	11000001.00111100.010100

ANDing the mask and address shows clearly that IP address 4 belongs to the subnet starting at (0000 0000) or .0, i.e. the 'start' of the network (but see note above for restrictions on older equipment).

RFC 1219 On the assignment of subnet numbers [61] describes a way of flexibly allocating subnets that defers having to decide where to draw the subnet boundary.

Binary Numbers

Binary numbers use base 2 rather than base 10, which is used for decimal numbers. So rather than needing 10 characters to represent a number (various combinations of 0-9), the binary system needs only two (0 and 1) and similarly the position of a character indicates multiplication by powers of 2 rather than powers of 10.

So 101 in decimal means 1*100 + 0*10 + 1*1, making a hundred and one.

And 101 in binary means 1*4 + 0*2 + 1*1, making five.

The rightmost character of a binary number represents the number of ones (20), the next digit to the left the number of twos (21), the next digit the number of fours (22), and so on.

To convert binary to decimal

As an example - the binary number 1110 0101 is to be converted to decimal.

Identify what each position represents (work from right to left starting at 1 and multiplying by two on each left shift):

1	1	1	0	0
128	64,	32,	16,	8,,

Multiply and sum:

1*128 + 1*64 + 1*32 + 0*16 + 0*8 + 1*4 + 0*2 + 1*1 = 128+64+32+4+1 = 229 base 10 (decimal)

For speed, ignore the multiplication and just add the value of columns with a one in them, but the full method works for any base and is worth knowing.

To convert decimal to binary

For those familiar with it, this conversion is most quickly done using the hexadecimal system (see below). Otherwise it can be simply achieved by repeated division by two, recording the

remainder each time.

Starting with 229 base 10 (decimal) so that we can check the result with the previous example:

229 / 2 = 114 remainder 1

114 / 2 = 57 remainder 0

57 / 2 = 28 remainder 1

28 / 2 = 14 remainder 0

14 / 2 = 7 remainder 0

7 / 2 = 3 remainder 1

3 / 2 = 1 remainder 1

1 / 2 = 0 remainder 1

The binary equivalent reads from the bottom up, so 229 base 10 (decimal) = 11100101 base 2 (binary).

Hexadecimal Numbers

Hexadecimal numbers use base 16 rather than base 10, which is used for decimal numbers. So rather than needing 10 characters to represent a number (various combinations of 0-9), the hexadecimal system needs 16 and similarly the position of a character indicates multiplication by powers of 16 rather than powers of 10.

So	101 in decimal means	1*100 + 0*10 + 1*1 making

And 101 in hexadecimal means 1*256 + 0*16 + 1*1 making

The rightmost character of a hexadecimal number represents the number of ones (16^0), the next digit to the left the number of 16s (16^1), the next digit the number of 256s (22), and so on.

Since hexadecimal requires 16 characters to represent its digits rather than the 10 used for decimal notation, the first six letters of the alphabet are also used.

In addition, because 16 is a power of 2 (2^4) there is a direct relationship between hexadecimal and binary numbers: each hexadecimal digit representing a group of four binary digits as shown in the following table.

Decimal	Binary	Hexa
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9

10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

Any binary number can therefore be represented as a (much shorter and more memorable) string of hexadecimal digits, e.g. using the previous example, the binary number 11100101 appears in hexadecimal as E5 (using the conversion table above, 1110 = E, 0101 = 5).

To Convert Hexadecimal to Decimal

As an example - the hexadecimal number E5 is to be converted to decimal. As for the conversion from binary to decimal, identify what each position represents by working from right to left, starting at 1 and multiplying by the base value (in this case 16) on each left shift:

16 1
E 5
Multiply and sum:

E (or 14)*16 + 5*1 = 224 + 5 = 229 base 10 (decimal)

To Convert Hexadecimal to Binary

Using decimal or hexadecimal rather than the underlying binary is easier to understand and much less prone to transcription errors. Hexadecimal, however, has the added advantage of giving easy access to the underlying binary pattern, which can be particularly useful when working with masks.
That the mask 1111 1111 .1111 1111 .1111 1111 .1000 0000 is rendered in decimal notation as 255.255.255.128 is not obvious, but the mapping to FF.FF.FF.80 is trivial using the table above.

1111	1111	1111	1111	1111
F	F	F	F	F

Equally, the reverse conversion from the address C1.3E.53.96 to binary is a simple matter of substituting the binary equivalents from the same table - much more straightforward than converting from the decimal, 193.62.83.96.

C	1	3	E	5
1100	0001	0011	1110	0101

E-mail for Users

E-mail for Users

The following notes provide some hints and tips about the use of e-mail. They cover the content and format of e-mails, sending and receiving attachments, courtesy and general housekeeping tips.

Common Courtesy

Always consider the feelings of recipients when composing e-mails. Good manners are also important when dealing with incoming e-mail.

E-mail should be read at least once a day if possible. If users are away for long periods, try to let people know.

Think whether an acknowledgement of receipt would be appropriate (and be aware that automatic acknowledgements are not always reliable).

Say 'Please' and 'Thank you'.

If an annoying e-mail is received, do not react with a 'knee-jerk' response that is likely to make the situation worse. It is possible that the sender did not realise the effect their message would have. Alternative actions could be pursued. These include:

- go and talk to the sender
- wait until the next day before replying
- delete the e-mail and ignore it.

Requests for action should be acknowledged and it may be appropriate to also confirm completion of the task. Reply to questions as soon as possible, even if it is only to confirm that help cannot be provided.

Do not forward chain letters - they clog up the network.

Content

There are many ways of composing e-mails and some are better than others. It is worth remembering that it is easy to offend recipients, particularly when a message has been composed in a hurry. Many users of e-mail have discovered the hard way that messages meant for the eyes of one individual have been broadcast to a wider audience because the 'reply to all button' was used in error.

It is good practice to:

- use a descriptive title in the subject field - but do not rely on people having read the title. If it is important, repeat the subject in the text (cut and paste is quick and easy)
- keep the message short and to the point but remember that the reader may need some background information. You may need to explain the following, which are also useful if mail goes to the wrong person, as the context lets them guess who it was really for and pass it on:
- why the e-mail is necessary
- why you are sending it
- why they are receiving it.

Do not use:

- abbreviations unless you are sure your reader will understand them
- all capitals - it makes the message harder to read and equates to 'shouting' in the e-mail world
- all lower case - it looks sloppy.

Format

Ideally an e-mail should not exceed a screen of information and should be simple text only, as formatted mail will probably not appear the same on all machines:

- keep the line length to about 60 characters
- use plain text - no special characters (even pound signs can cause problems)
- do not use formatting: no colours, bold, italic, underline, etc.
- do not use special fonts: rely on the one provided by your mail software
- do not use tabs: if you want a table effect, use a fixed-width (non-proportional) font that uses the same space for all characters including blank spaces, such as Courier or Mishawaka, and insert spaces.

In most cases it is best to send a minimally-formatted message with line ends manually inserted, as the damage done to such messages when received in a richer environment is less than the effect of a complex message that cannot be displayed properly.

Signature File

Most packages allow users to set up signature files. These files should be short (not more than four lines) and professional.

Sending Attachments

Do not attach something that can easily be sent in the text of the e-mail (meeting agendas etc.):

- only send attachments if it is clear that the intended recipient can read them (make sure their word processor/spreadsheet software can read the file)
- if the attachment is large, it is well worth doing a trial with a small one first
- mailing lists generally object to attachments - consider whether the item can be published somewhere else and let readers choose when (and whether) to fetch it
- in general, do not attach html files - give the web address
- give attachments meaningful names (and the correct extensions).

Unlike the text of an e-mail, all attachments are sent encoded. Although with PC mailers this happens automatically, there are several coding formats in use. If the mailer at the sending site is incompatible with the mailer at the recipient's site, it will not be possible to read any attachments. Attachments can also contain viruses and some people will therefore not accept them.

It is generally acceptable to send attachments within an organisation, but remember that people who do not use the same e-mail program may have problems. When sending attachments to other people, a MIME (Multipurpose Internet Mail Extensions) compliant encoding format is most likely to be readable. In Internet mail the MIME standards RFC 2045 Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies, and others specify how to assemble a message.

Text may be encoded if it contains special characters or features. The parts of a message (e.g. some typed text and an attachment) are linked together into a single object that will pass through all Internet mail systems without damage, and reversing the MIME processes allows the recipient's mail program to recover the various parts, but different mail products will present the delivered message with different appearances and behaviours. MIME ensures that all the parts of a message are delivered, but not that they will be displayed identically on

the receiving machine. Try to configure the sending program to avoid:

- proprietary or obsolete encodings such as 'binhex' or 'uuencode'
- sending '8-bit' characters without encoding ('quoted-printable' (QP) and 'base64' encodings cover most needs)
- sending both plain and encoded versions
- encoding simple text ('US-ASCII') for which no encoding is necessary.

The process of encoding a document or file for attachment, and of assembling the message that includes it, is quite complicated and can easily result in large messages even when sending very little information.

Before sending any document or file, check that it does not contain a virus. The mail program used may be able to do this automatically, provided anti-virus software is kept up to date.

Privacy

E-mail is not a private or secure medium. Remember that messages may be read by people other than those intended to receive them, by accident or malice; and that a received message may not be what it seems either in its origin or its contents. The sort of messages received from regular correspondents will be familiar, and this is a good guide to their authenticity.

If it is important to send a message that remains private then it will have to be encrypted, subject to regulations at the sending and receiving organisations. Readily available products include PGP® (Pretty Good Privacy®), GNU Privacy Guard and S/MIME (Secure/Multipurpose Internet Mail Extensions). The same technology and products ensure that the contents of a message have not been changed and to some extent confirm the identity of the sender. There are, of course, other media for communication such as post, fax and telephone that can be used instead or can provide verification of a received e-mail message.

When sending a message to multiple recipients, consider what personal or other information is being revealed. If the addresses are entered in the usual and simplest way, each recipient will see the addresses and possibly the names of all the others, and they may not have agreed to that. Most mail programs have a Blind Copy facility (Bcc:) that allows the inclusion of an address to which the message will be delivered but which will not appear in any delivered copies. This facility does not work in exactly the same way in all mail systems and programs, and local support staff should be asked for detailed advice (or test with a few friends who have agreed in advance). Note too that the Bcc: recipients will not normally see their addresses in the copy they get, and they may not understand why they have got the message.

Forged E-mails

Any part of an e-mail message, including the From: address, can be forged, for example to commit fraud or distribute viruses. Any message that is unexpected or has unusual content should be treated with suspicion, even if it appears to come from a known source. In particular beware of opening attachments to such messages, as these may contain viruses, backdoors or other hostile code. If there is any doubt about the authenticity of a message, consult your

local support staff. It may be appropriate to delete suspect messages unread, even though this runs the risk of losing a legitimate communication. When writing messages yourself, include some obviously 'human' text: do not just send blank messages with an attachment as these may well be deleted unread by recipients who are concerned about viruses.

When replying to an e-mail, remember that it may not have been written by the person from whom it claims to come.

Receiving Attachments

Attachments in almost any format can contain viruses. All machines belonging to JANET-connected organisations should therefore be running a virus checker and any suspect items received should be checked before they are opened. If there is any doubt whether an attachment is genuine, do not open it without seeking advice from your local computer support staff.

Replying

When replying to an e-mail, ensure that text taken from the original message is clearly distinguished from the text forming your reply. Most mail clients can automatically insert a > character before each line of the original.

There is currently no general agreement over whether replies should come above, below or interleaved with the original message. There is no right answer here and in general users should comply with the norm of the group with which they are exchanging e-mails. It is however advisable to type the points in the reply close to the points in the original message for clarity. The guiding principle is to consider the convenience of the intended reader(s) rather than that of the sender. One widely recognised style is to show each point (or part of it) from the original separately, followed by your comments in response to that point. Blank lines are usually enough to separate points.

When replying to e-mails:

- check the subject field is still relevant
- include ticket numbers in correspondence with help desks including the JANET Service Desk and JANET CSIRT;
- check that the recipients are those for the message is intended - beware when using 'reply to all'
- delete unnecessary quoted text but leave in enough so that the reply makes sense even to a recipient who has not got the original - remember when sending a message to a help desk (e.g. the Janet Service Desk or Janet CSIRT) that the person who reads the reply may not be the person who sent the original message.

Housekeeping

Users can easily become overwhelmed by the volume of e-mail received. To keep it down to sensible proportions:

- make sure that the mailboxes for incoming and sent mail are cleared out regularly

- file what needs to be kept using a folder or similar filing system provided by the e-mail system and delete the rest, checking that they are not hiding in a 'Trash' or Deleted Items' folder
- consider unsubscribing from high volume e-mail lists if away for an extended period.

E-mail for Organisations

Existing and new JANET customer organisations may choose to operate their own e-mail systems and services, or to procure all or parts of them externally in various ways. In all cases there are some requirements that will ensure seamless interworking with other JANET-connected organisations and with the wider Internet community.

These notes are primarily intended for managers of e-mail services within JANET-connected organisations where either e-mail is being set up for the first time or major changes in the design or implementation of e-mail services are being considered. They may also be of interest to suppliers and others who provide or support e-mail services for a JANET-connected organisation.

Conventions

The RFC document 2119 Key Words for Use in RFCs to Indicate Requirement Levels provides a convention for the use of keywords such as 'MUST' and 'SHOULD' to indicate what degree of discretion is allowed (if any) in interpreting directions and requirements in this area.

Where keywords are presented in BOLD UPPER CASE, the conventional meanings of RFC 2119 apply in assessing whether e-mail services conform to those required for an organisation connected to JANET. Certain technical or management requirements may seem rather severe in using this interpretation, but MUST and similar keywords are only used where there is some specific and compelling need.

External View

E-mail Addresses

Required E-mail Addresses

Domain Nameserver

Message Format

Relaying

Gateways

Dial-up Accounts

Web-based E-mail

E-mail Addresses

Internet e-mail addresses are of the form:

`<local-part>@<domain>`

The domain name broadly distinguishes an organisation's network from all others in the Internet and possibly specifies some department or division within the organisation, and the local-part identifies a particular individual or role within the organisation. There are normally no spaces in either part, nor on either side of the '@' separator.

Most organisations connected to Janet are entitled to a domain in ac.uk (the imaginary domain college.ac.uk is used for illustration here). JANET(UK) publishes rules on the choice of domain names immediately below ac.uk which are explained in Section 3 of this manual.

Published destination addresses SHOULD be of forms such as:

`Fred.Bloggs@college.ac.uk`

simple and explicit, but may be hard to keep unambiguous

`f.bloggs@dept.college.ac.uk`

less personal identity revealed, more organisational structure

`fr03200@college.ac.uk`
no personal name included.

Addresses SHOULD NOT be of forms such as:

fredb-chemsrv2@ntserver2.chem.college.ac.uk [62]

in which local private information about usernames, machine names and possibly the operating system in use are visible. There are two kinds of difficulty with such addresses:

- they are not stable since administrative or technical changes can make these addresses misleading
- the information they appear to contain is of more value to someone considering breaching an organisation's security than to any legitimate user.

E-mail addresses SHOULD NOT be case sensitive. If F.Bloggs is the local part of the published form of one of an organisation's e-mail addresses then it is normal to recognise forms such as f.bloggs, F.BLOGGS and even odd mixtures of case, and to regard them all as referring to the same address. Some people prefer to publish addresses that use capitals in the conventional way; some prefer all lower case. Whatever is published on paper will sometimes be entered incorrectly, and if the error is only to forget an initial capital then it is reasonable to expect the e-mail to be correctly delivered. Most e-mail software will default to this case-insensitive behaviour.

Addresses used in SMTP (Simple Mail Transfer Protocol) enable e-mail systems to route messages and to report failures. Addresses in the message header preceding the contents of the message are for the use of mail programs at the recipient's site, and enable them to do such things as show users who a message (ostensibly) came from and devise reply addresses.

Any address in the protocol envelope or the header of a message sent from an organisation's e-mail service MUST have a fully-qualified domain name (all components included up to top level .uk or similar), and MUST be valid for delivery.

MAIL FROM: (envelope) SHOULD be as published (but see the exceptions below)

From: header line MUST be as published

Sender: header line SHOULD NOT normally be used (but see below)

Reply-To: header line SHOULD NOT normally be used (but see below).

Exceptions: the requirements may be different for messages sent by an automatic process such as a web form or a mailing list.

Required E-mail Addresses

Organisations MUST implement the postmaster and abuse addresses for all domains within their management. RFC 2142, Mailbox Names for Common Services, Roles and Functions, describes other role addresses that should be provided under certain circumstances. If a site supports the facilities that any of those addresses provide, it MUST use the names prescribed for them.

Each organisation MUST arrange for a timely response to messages from JANET(UK) or its contractors sent to the postmaster and abuse addresses.

An organisation may, of course, wish to use alternative names as well. In these circumstances, e-mail systems must support such aliases and the originator that is specified in e-mail sent from role accounts must be valid and appropriate. Equally important, e-mail to these role addresses must be routed to one or more individuals who have the skills and resources to deal with it.

It is quite acceptable for the same individual or team of individuals to be responsible for messages addressed to more than one role.

Domain Nameserver

The IP address of a sending mailer SHOULD have a PTR (PointTeR - address to name) record in the IN-ADDR.ARPA. zone of the Domain Nameserver. The JANET Technical Administration Group can advise who has the delegated authority to make entries for a site network if this is not clear.

The sending mailer HELO (EHLO in Extended Simple Mail Transfer Protocol (ESMTP)) SHOULD be fully qualified and SHOULD correspond to an A (Address) record in the Domain Nameserver that matches the mailer's IP address.

Domains and subdomains for which the domain name of the appropriate inbound mailer differs from the domain name in e-mail addresses MUST have MX (Mail eXchanger) records in the Domain Nameserver indicating the mailer.

Message Format

The header part of every message sent from a JANET-connected organisation MUST meet the requirements of RFC 2822 Internet Message Format.

This states that the following lines are mandatory:

Date:

From:

To:

(see E-mail Addresses above).

The header SHOULD also include a Message-ID: line. The headers of messages sent from an organisation may or may not have Received: lines, depending on the software used and the structure of the e-mail service. It may be possible to use Received: lines to identify the person originating each message (see Audit Trail below).

Timestamps in Date: and Received: header lines MUST be accurate to one second or better, with the correct time zone indicated. Timestamps SHOULD use the format +0100.

The Date: line is usually supplied by the user program that generates the mail messages, and it may be necessary to maintain the accuracy of clocks on numerous desktop or public computers. Various network technologies have proprietary ways to synchronise clocks within a network. The JANET Network Time Service enables an organisation to keep its network's time in step with those of other organisations on JANET and throughout the Internet.

Message-ID: and Received: are not usually shown to recipients, so deficiencies in these header lines may not easily be spotted.

Messages MUST conform to the MIME specifications in RFC 2045 Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies and related RFCs (possibly by having no MIME features at all).

Relaying

It is notionally possible for an e-mail message to travel by a single Internet connection from the computer on the sender's desk to the recipient's computer, with no other computer or service directly involved. In practice this is very rare and most users of e-mail are familiar with the idea that their messages travel in distinct hops, first within the sending organisation that has one or more mail servers, perhaps between service providers outside the sending or receiving organisation and then through servers close to the recipient.

A function of most Internet mail servers is to accept messages from some places and pass them to other places. In cases where the server makes no important change to a message, this is called relaying and all major e-mail server products provide such a facility. It is essential to configure this relaying in a secure way, allowing some combinations of source and destination but denying others.

In most cases only a very small number of systems will be expected to send e-mail out from the organisation, and very few will be expected to listen for incoming mail connections (SMTP, TCP port 25). The sending and listening systems need not be the same but they will all be known to and managed by staff responsible for the organisation's e-mail as a whole. Systems with this external access in either direction are exposed to open relaying attempts, and such attempts will only be defeated by a combination of technical and administrative arrangements.

The organisation's router or firewall SHOULD reject outbound packets to port 25 (the port used to send e-mail) at external addresses and inbound packets to port 25 at internal addresses, except for the above sending and listening managed mailers.

All e-mail systems with port 25 (SMTP) accessible from outside the organisation MUST be configured so that they will reject attempts to relay incoming messages through the organisation's mailers and back to the outside, except where such messages are explicitly authorised (see Dial-up Accounts below).

Similar considerations apply to TCP port 587, which RFC 2476 RSVP Operation over IP Tunnels assigns for message submission.

For further information, see Section 6 of this manual.

Gateways

Some e-mail systems use open Internet standards such as SMTP, POP3 (Post Office Protocol3), or IMAP (Internet Message Access Protocol) throughout and will have no basic difficulty meeting the criteria listed here.

Other systems are primarily designed for use within an organisation. They can offer features not available in the Internet at large by a variety of proprietary techniques. However, in order to exchange e-mail with other organisations they need a gateway system that behaves exactly

like that of a native Internet e-mail system as described above. The gateway system accepts messages from the proprietary e-mail system that are intended for outside Internet addresses and vice versa, and in each case makes any changes necessary to the messages concerned.

In such an environment the external behaviour of an organisation's e-mail (including message formatting and control of relaying) is almost entirely determined by the gateway. While this should in principle make management easy, many gateway products have poor implementations on their Internet side and each detail mentioned needs checking carefully.

Dial-up Accounts

An organisation may wish to allow its e-mail users limited access when they are away from their normal place of work. They may be at home, connected through a dial-up ISP, away at a conference, working with colleagues at another organisation or on holiday using an Internet café.

Organisations that use a proprietary e-mail system may find it impossible to offer this type of access. If the e-mail system uses open standards then there are a number of security issues, including the danger of operating an open e-mail relay as mentioned above. However, the main difficulty is authentication of individual users.

There is nothing in the most common open e-mail standards (POP3, SMTP) that will allow a site mailer to differentiate securely between a genuine user working from home, and a spammer or other abuser able to forge addresses. Both will attempt to connect to the site mailer from another network. An authorised user will be acting legitimately in preparing e-mail that looks as if it comes from an organisation domain, whereas identical address details from the spammer will be a forgery.

The normal advice to users is to send e-mail through their ISP's outbound mailer instead. The ISP normally has additional knowledge, such as the telephone number from which the dial-up call originated, and can in most cases justify the relaying required even if the user chooses to use their organisation address from home. If dial-up access is important for an organisation and its users, any claims by a supplier that a standard on-site product will get it right without leaving an open e-mail relay in the site network or the supplier's should be checked very carefully.

Web-based e-mail (see below) overcomes most of the problems. Other approaches using the SSH (Secure Shell) and SSL (Secure Sockets Layer) protocols or proprietary secure connections, possibly in conjunction with an IMAP message store, are technically satisfactory but it is difficult to ensure that the end-user client software on which they depend is available in arbitrary external places.

Web-based E-mail

The approach of Hotmail® and many other service providers, particularly where there is no payment for the e-mail service, is to make a web browser the interface to all user e-mail facilities - authentication, reading, composition, sending and storing in folders. The e-mail system then comprises:

- a web server supplying a variety of forms for the tasks that a user can perform (web

forms are pages with provision for user input)
- an authentication database against which one of the forms will validate users
- a message store that certain of the forms will manipulate to support an 'Inbox' from which each user can read their own incoming e-mail, and usually folders named by each user to allow them to organise their mail as they wish
- an external mailer that delivers incoming and internal e-mail to message stores, and formats and transmits outgoing messages.

For resilience and ease of management in all except the smallest services, the functions may be spread across two or more computers. Benefits include:

- the complete absence of e-mail software and data (messages) from all end-user computers
- concentration of management and technical resources for e-mail at a central system or collection of systems
- access to organisational e-mail (both reading and sending) from other locations.

Against this:

- the software is relatively complex and needs significant management
- it may be difficult to integrate existing user databases
- fewer features and facilities may be available to end-users than through dedicated e-mail software
- the service may interact more slowly with users than dedicated e-mail software
- some service providers are particularly vulnerable to malicious activity, Hotmail® for example having been broken into on several occasions.

An organisation considering obtaining a web-based e-mail service from an external supplier should ensure that it can meet all requirements for the appearance of outgoing mail, for the management of user accounts, and for usage reports. A free public service is unlikely to do so.

Internal View

E-mail Addresses

Audit Trail

Distribution Lists

Privacy

User Support

Service Scaling

Service Agreement

E-mail Addresses

All addresses used in outgoing messages MUST be valid and MUST have fully-qualified domain names. User e-mail programs SHOULD provide some address book or similar facility

so that users need only supply short or easily remembered versions of addresses, or can select from a list. This is not the same as allowing mail programs to supply a default domain, which is less satisfactory and should not be used.

It is highly desirable that all the addresses used internally are the same as those published for external use, so that users do not need to choose which address to give to their correspondents. Organisations using proprietary systems for internal e-mail may find it difficult or impossible to arrange this.

Even where internal e-mail uses open Internet technology, there may be operational or historical reasons for the use of addresses that are different from the standard ones published. For instance, an organisation may have departmental or location e-mail servers and it may appear more efficient in network and machine resources to route e-mail directly between them by using addresses that include the names of those servers. The danger then arises that one of these internal addresses will escape to the outside world; organisations using this system should check that this is acceptable. The outgoing e-mail system or systems may be able to rewrite internal addresses to a suitable external form; or sites may have to accept that such addresses will start to be used for incoming mail. In this case special arrangements of MX records are likely to be needed to ensure that messages are delivered.

Audit Trail

Each organisation MUST adopt software and management procedures that make it possible to identify the person responsible for sending each message, independently of any information in the message itself that an end-user might supply falsely. This might be achieved in various ways:

- record an IP address in the message header along with the timestamp (possibly as a Received: line) and refer to access logs
- record an authenticated login in the message (only a partial solution)
- ensure that e-mail system logs are adequate and are not lost or damaged.

It is recognised that such technological procedures alone offer no assurance that the person using an account name and password at any particular time is or was authorised to do so. Organisations MUST therefore also publish clear instructions to all e-mail users about security of accounts, passwords, use of shared or unattended computers and other related matters.

Software and management procedures that result in a log of all messages sent out from the organisation MUST be adopted. The log must be retained for a suitable and agreed period (e.g. three months). The logs MUST be kept secure against unauthorised examination, alteration or accidental loss.

Janet or their contractors MUST be able to contact organisation e-mail administrators if any difficulty arises. The Janet Service Desk maintains a list of technical contacts, with telephone numbers. Depending on the nature of the enquiry, Janet may use the postmaster role address or the person identified in the RIPE database, and it is desirable that all these contact details are kept up to date.

Distribution Lists

An e-mail system SHOULD support the expansion and management of internal distribution lists. This allows individuals within an organisation to correspond easily with all staff or students, or all individuals in particular departments. Internal lists SHOULD NOT otherwise be accessible to mail sent from outside the organisation.

Most e-mail products will support lists, possibly by purchasing and installing additional software or components.

Check that:

- there are satisfactory arrangements for managing membership of internal lists
- access to internal lists is controlled.

Normally only members of the list will be able to send messages to it, but for some lists it may be desirable to extend this to certain managers or others inside an organisation, or even to certain individuals or roles outside the organisation.

Privacy

The organisation MUST publish internally a privacy statement setting out the circumstances in which e-mail and access logs and stored messages will be made available to persons or agencies other than the originator and recipients of the messages concerned.

The requirements of the Data Protection Act and the Regulation of Investigatory Powers Act will influence the content of this statement.

User Support

Users can expect support of various sorts:

- routine requests for changes to their account details
- information on the status of the e-mail service
- advice on the e-mail software they use
- advice on reports (often failure reports) from the e-mail system or from somewhere else
- action on what they perceive as abuse through the e-mail service, including both UBE and abuse that appears to be personal
- advice on good practice in using e-mail
- advice on the use of e-mail to access external services (e.g. mailing lists)
- advice on the interaction between the organisation's e-mail service and dial-up or other connection services they may wish to use
- information on an organisation's policies and procedures with regard to relevant current legislation, such as data protection, retention and destruction of data, or handling of requests from law enforcement agencies
- directions on security and acceptable use with appropriate reference to the JANET Acceptable Use Policy.

Much of the advice and information will be best provided through internal web pages or other documentation. Maintaining a list of FAQs (Frequently Asked Questions) is likely to be effective.

Where the user support function is separate from the operation of the e-mail service, communication between the two activities must be adequate. This is likely to be particularly important if an organisation is spread across two or more physical locations, or if all or part of the e-mail service is outsourced.

Service Scaling

The system or systems providing an e-mail service for a small organisation can be very simple, whether provided in-house or outsourced. Elements should include:

- a firewall barring access to unused ports on e-mail systems
- a central computer accepting incoming e-mail connections
- a central computer sending e-mail outside the organisation
- a central computer storing delivered e-mail for users to read
- computers and software with which users read and compose their e-mail.

For more information see the JANET Technical Guide Designing Reliable Mail Systems.

The central functions can be combined in a single system. Indeed, one computer may be able to manage parts of the e-mail system for several organisations, and this arrangement is quite normal for an outsourced mail service.

For a great variety of reasons, very few Janet-connected organisations have e-mail systems as simple as this. The need for gateways (for resilience in case of certain failures, for operation across multiple locations, for management within separate departments) and the size of an organisation all increase the complexity of the service. It is not practicable to give general advice on scaling for these conditions. The Janet Service Desk can provide assistance to organisations that require specific comment on a proposal. FE organisations and specialist colleges may also consult their JISC RSC for advice.

Service Agreement

Whether an e-mail service is resourced internally or from outside, both providers and users of the service need to be clear about what to expect.

Where the service is provided by an outside contractor, an organisation will normally have the most critical items closely linked to the agreement under which the contractor does the work. Headings might include:

- capacity of the service
- performance of the service (time taken for messages to pass through the system and network)
- availability of the service (and of certain specific parts of it)
- assurance of security
- response to requests for changes
- response to fault reports
- escalation and penalties.

For internal support, this may be regarded as documentation of the service or may be

included in Quality or other documentation.

Where facilities are provided away from the site or networks, each of the headings may have an impact on internal e-mail as well as external traffic. This may affect the way in which some of the risks are assessed.

Links
[1] http://webarchive.dev.ja.net/services/publications/supportmanual/glossary.html
[2] http://www.ja.net/forms/fault-escalation-form
[3] mailto:operations@sitename.ac.uk
[4] http://www.ja.net/forms/trouble-tickets-mailing-list
[5] http://www.ucisa.ac.uk/%7E/media/Files/publications/toolkits/ist/ISTEd3_Section_B%20pdf.ashx
[6] http://www.acu.edu/technology/is/recovery.html
[7] http://repository.jisc.ac.uk/7326/1/janet-backbone-and-janet-ip-connections-availability-and-service-level-commitments.pdf
[8] http://repository.jisc.ac.uk/7562/1/janet-network-connection-policy-november-2019.pdf
[9] mailto:connect@ja.net
[10] http://webarchive.dev.ja.net/services/connections/types-of-connection.html
[11] https://community.ja.net/library/janet-policies/terms-provision-janet-service
[12] https://community.jisc.ac.uk/system/files/public_images/connectflowchart-new3.png
[13] https://community.ja.net/library/acceptable-use-policy
[14] https://community.ja.net/library/janet-policies/security-policy
[15] mailto:technical@customer.ac.uk
[16] http://www.ja.net/products-services/janet-connect/csirt
[17] mailto:ipaddress@ja.net
[18] https://community.ja.net/library/janet-policies/service-level-agreement
[19] https://community.ja.net/library/janet-services-documentation/primary-nameserver-service
[20] https://community.ja.net/library/janet-services-documentation/secondary-nameserver-service
[21] https://community.ja.net/library/janet-services-documentation/site-resolver-service
[22] https://community.ja.net/library/janet-services-documentation/janet-csirt
[23] https://community.ja.net/library/janet-services-documentation/eligibility-policy
[24] https://community.ja.net/library/janet-services-documentation/domain-name-registration
[25] https://community.ja.net/library/janet-services-documentation/register-acuk
[26] mailto:naming@ja.net
[27] https://community.ja.net/library/janet-services-documentation/payments-and-charges
[28] http://www.nominet.org.uk
[29] https://community.ja.net/library/janet-services-documentation/modify-dns-entries
[30] http://www.ietf.org/rfc/rfc1519.txt
[31] http://www.ripe.net/ripe/docs/ipv4-policies.html
[32] http://www.ripe.net/ripe/docs/titletoc.html
[33] http://www.ja.net/products-services/janet-futures/ipv6
[34] https://community.ja.net/library/janet-services-documentation/connecting-janet
[35] http://www.icann.org/
[36] http://www.ripe.net/
[37] http://www.jisc.ac.uk/whatwedo/services/as_rsc/rsc_home/rscs_contact.aspx
[38] http://www.ja.net/forms/reverse-delegation-service-application-form/28
[39] https://community.ja.net/library/janet-services-documentation/example-reverse-delegation-forms
[40] http://www.ietf.org/rfc/rfc1918.txt
[41] http://www.ja.net/products-services/janet-connect/managed-router-service
[42] http://www.mail-abuse.com
[43] https://community.jisc.ac.uk/library/janet-services-documentation/dns-allow-and-deny-lists
[44] http://www.hmso.gov.uk/acts/acts1998/19980029.htm
[45] http://www.hmso.gov.uk/acts/acts2000/20000023.htm

[46] http://www.ico.gov.uk/Home/what_we_cover/data_protection/guidance/codes_of_practice.aspx
[47] http://www.ja.net/products-services
[48] http://www.ja.net/contact-us
[49] http://www.ucisa.ac.uk/publications/modelregs/modelregs.aspx
[50] https://community.ja.net/blogs/regulatory-developments
[51] http://www.hmso.gov.uk/acts/acts1990/Ukpga_19900018_en_1.htm
[52] http://www.hmso.gov.uk/acts/acts1998/19980029.htm
[53] http://www.hmso.gov.uk/acts/acts2000/20000023.htm
[54] http://www.hmso.gov.uk/acts/acts1988/Ukpga_19880027_en_1.htm
[55] http://www.jisclegal.ac.uk/
[56] https://community.ja.net/library/janet-services-documentation/structure-janet
[57] https://community.ja.net/library/janet-services-documentation/contact-csirt
[58] https://community.ja.net/library/janet-services-documentation/about-csirt
[59] http://www.jisc.ac.uk/
[60] http://www.jisc.ac.uk/whatwedo/services/as_rsc.aspx
[61] http://www.faqs.org/rfcs/rfc1219.html
[62] mailto:fredb-chemsrv2@ntserver2.chem.college.ac.uk