

Janet Site E-mail Requirements

Janet Site E-mail Requirements

What's this all about?

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 any case there are some requirements which will ensure seamless interworking with other Janet organisations and with the wider Internet community, and they are outlined here.

This note is primarily intended for managers of e-mail services within Janet organisations, particularly those who are considering major changes to their design or implementation and may be procuring related products or services, or are setting up e-mail within the organisation for the first time. It may also be of interest to suppliers and others who provide or support e-mail services for a Janet organisation.

The note assumes a little familiarity with Internet ways of working. In particular common terms and abbreviations such as IP, DNS, TCP or RFC are used in the text without explanation. A very brief [glossary](#) is included at the end. The terms "e-mail" and "mail" are used without distinction to refer to electronic mail.

- [Nature of the requirements](#)
- [External view](#)
 - [E-mail addresses](#)
 - [Required E-mail addresses](#)
 - [DNS](#) - use of the Domain Name Service
 - [Message format](#)
 - [Relaying](#)
 - [Gateways](#)
 - [Dialup accounts](#)
 - [Web-based e-mail](#)
- [Internal view](#)
 - [Addresses](#)
 - [Audit trail](#)
 - [Distribution lists](#)
 - [Privacy](#)
 - [User support](#)
 - [Service scaling](#)
- [Service agreement](#)
- [Glossary](#)

[Contents](#) |

Nature of the requirements

[RFC 2119](#) ^[1] sets out 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.

Where keywords in this note are presented in **BOLD UPPER CASE**, the conventional meanings of RFC 2119 apply in assessing whether a Janet site's e-mail services conform to Janet's requirements. Certain technical or management requirements may seem rather severe in this interpretation, but the "MUST" and similar keywords are only used where there is some specific and compelling need.

Where the same words are used without the above emphasis, they have their ordinary informal meanings.

[Up](#) | [Contents](#)

External view

E-mail addresses

Internet e-mail addresses are of the form `<local-part>@<domain>` in which the *domain* name broadly distinguishes your 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 your 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), but the correspondence between Janet and *ac.uk* is not complete. Janet(UK) publishes [rules](#) ^[2] on the choice of domain names immediately below *ac.uk*.

Published destination addresses **SHOULD** be of forms such as

- *Karl.Marx@college.ac.uk* ^[3] - simple and explicit, but may be hard to keep unambiguous;
- *k.marx@dept.college.ac.uk* ^[4] - less personal identity revealed, more organisational structure;
- *fr03200@college.ac.uk* ^[5] - no personal name included

Addresses **SHOULD NOT** be of forms such as

- *karlm-chemsrv2@ntserver2.chem.college.ac.uk* ^[6]

- 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; administrative or technical changes can make these addresses misleading.
- the information they appear to contain is of more value to someone considering breaching your security than to any legitimate user.

E-mail addresses **SHOULD NOT** be case-sensitive. If *K.Marx* is the local part of the

published form of one of your e-mail addresses then it is normal to recognise forms such as *k.marx*, *K.MARX* and even odd mixtures of case, and to regard them all as referring to the same address. Some people prefer to publish addresses which use capitals in the conventional way; some prefer all lower case. Whatever you publish on paper will sometimes be entered wrongly, and if the error is only to forget an initial capital you probably want the e-mail to be correctly delivered.

Most e-mail software will default to this case-insensitive behaviour.

[Up](#) | [Previous](#) | [Contents](#)

Addresses used in the mail protocol (the "SMTP envelope") enable e-mail systems to route messages and to report failures. Addresses in the message header which precedes the contents of the message are for the use of end users' mail programs, enabling them to do such things as showing users who a message (ostensibly) came from and devising reply addresses.

Any address in the protocol "envelope" or the header of a message sent from your organisation's e-mail service which appears to be in one of the organisation's domains **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.

[Up](#) | [Contents](#)

Required e-mail addresses

You **MUST** implement the *postmaster* and *abuse* addresses for all domains within your management. [RFC 2142](#) ^[7] describes other role addresses which you should provide under certain circumstances; in particular if you support the facilities which any of those addresses provide you **MUST** use the names prescribed for them.

You **MUST** arrange for a timely response to messages from Janet(UK) or its contractors sent to the *postmaster* and *abuse* addresses for your organisation.

You may, of course, wish to use alternative names as well; you then need to ensure that your e-mail systems support such aliasing and that the originator set in e-mail sent from role accounts is valid and appropriate. Equally importantly, you need to ensure that e-mail to these role addresses is routed to one or more individuals who have the skills and resources to deal with it, and that response to such mail does not depend on the attendance of a single individual whose duties may sometimes make her or him unavailable. It may be appropriate to direct role mail to a help desk or similar team or function able to reliably receive messages and to ensure that they are attended to.

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

[Up](#) | [Previous](#) | [Contents](#)

DNS

The IP address of a sending mailer **SHOULD** have a *PTR* (pointer - address to name) record in the *IN-ADDR.ARPA.* zone of the DNS. If you are not sure who has the delegated authority to make entries for your network, [Janet Service Desk](#) ^[8] will be able to find out for you.

The sending mailer HELO (or EHLO in [ESMTP](#) ^[9]) **SHOULD** be fully qualified and **SHOULD** correspond to an *A* (Address) record in the DNS which matches the mailer's IP address.

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

[Up](#) | [Previous](#) | [Contents](#)

Message format

The header part of every message sent from your organisation **MUST** meet the requirements of [RFC 822](#) ^[10] in which the following lines are mandatory:

Date:

From:

To:

(see above on [E-mail addresses](#)).

The header **SHOULD** also include a Message-ID: line.

Whether or not the headers of messages sent from your organisation have

Received: lines

depends on the software you use and the structure of your e-mail service. You may be able 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 1 second or better, with the correct timezone indicated.

The Date: line is usually supplied by the user program which generates the mail messages, so that you may need to keep accurate the clocks of numerous desktop or public computers. Various network technologies have proprietary ways to synchronize clocks within your own network; the [Janet Network Time service](#) ^[11] enables you to keep your network's time in step with those of other organisations in Janet and throughout the Internet.

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

Messages **MUST** conform to the MIME specifications in [RFC 2045](#) ^[12] and related RFCs (possibly by having no MIME features at all).

[Up](#) | [Previous](#) | [Contents](#)

Relaying

In many 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 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 messages from outside through the organisation's mailers and back to the outside, except where such messages are explicitly authorised (see [Dialup accounts](#) below).

Similar considerations apply to TCP port 587, which [RFC 2476](#) ^[13] assigns for message submission.

[Up](#) | [Previous](#) | [Contents](#)

Gateways

Some e-mail systems use open Internet standards ([SMTP](#) ^[14], [POP3](#) ^[15], [IMAP](#) ^[16]) throughout and will have no basic difficulty meeting the criteria listed here.

Others 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; but to exchange Internet mail with other organisations they need a gateway system which presents to the Internet behaviour exactly like that of a native Internet mail system as described above. It then accepts in some way messages from the proprietary e-mail system which 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 your 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 are poor implementations of their Internet side and each detail mentioned in this note needs checking carefully.

[Up](#) | [Previous](#) | [Contents](#)

Dialup accounts

You may wish to allow your e-mail users to work in a limited way as if they were sitting at a computer on your network when in fact they have connected to the Internet elsewhere. They may be at home connected through a dialup ISP, away at a conference, working with colleagues at another organisation or on holiday using an Internet café.

If you have a proprietary e-mail system this is likely to be impossible. If your e-mail uses open standards there are a number of security issues, including the danger of operating an open e-mail relay as mentioned above. The main difficulty is authentication of your individual users.

There is nothing in the most common open e-mail standards ([POP3](#) ^[15], [SMTP](#) ^[14]) that will securely allow a site mailer to tell the difference between a genuine user working from home, and a spammer or other abuser willing to forge addresses. Both will attempt to connect to your site mailer from another network. Your user will be acting legitimately in preparing e-mail which looks as if it comes from your organisation domain, whereas identical address details from the spammer will be a forgery.

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

Web-based e-mail (see [below](#)) overcomes most of the problems. Other approaches using [SSH](#) ^[17], [SSL](#) ^[18] or proprietary secure connections, possibly in conjunction with an [IMAP](#) ^[16] 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.

[Up](#) | [Previous](#) | [Contents](#)

Web-based e-mail

The approach of [Hotmail](#) ^[19] and since then of 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 and sending, and storing in folders. The e-mail system then comprises

- a Web server supplying a variety of forms for the tasks which 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 which certain of the forms will manipulate to support an "In box" from which each user can read their own incoming e-mail, and usually folders named by each user to allow them to organize their mail as they wish;
- an external mailer which 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.

If you consider obtaining a Web-based e-mail service from an external supplier, you should ensure that they can meet your requirements for the appearance of outgoing mail, for the management of user accounts, and for usage reports. A public (free) service is unlikely to do so.

[Up](#) | [Previous](#) | [Contents](#)

Internal view

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.

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. If your internal e-mail is of a proprietary nature it will be difficult or impossible to arrange that, and you will need to clearly document good practice.

Even where internal e-mail uses open Internet technology, there may be operational or historical reasons for the use of addresses which are different from the standard published ones. For instance, you 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 which include the names of those servers. The danger then arises that one of these internal addresses will escape to the outside world, and you should check that this is acceptable. Your outgoing e-mail system or systems may be able to rewrite internal addresses to a suitable external form; or you may have to accept that such addresses will start to be used for incoming mail.

[Up](#) | [Previous](#) | [Contents](#)

Audit trail

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

- 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);
- or 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 some particular time is or was authorised to do so; so you **MUST** also:

- publish clear instructions to all e-mail users about security of accounts, passwords, use of shared or unattended computers and other related matters.

You **MUST** adopt software and management procedures to result in a log of all messages sent out from the organisation which is retained for a suitable and agreed period (*eg* 3 months).

You **MUST** ensure that such logs are kept secure against unauthorised examination, alteration or accidental loss.

You **MUST** ensure that Janet(UK) or their contractors can contact your e-mail administrators if any difficulty arises. Janet Service Desk ^[8] maintain a list of Technical Contacts, with telephone numbers. Depending on the nature of the enquiry, we may use the *postmaster* role address or the person identified in the RIPE database ^[20], and it is desirable that all these contact details are kept up to date.

[Up](#) | [Previous](#) | [Contents](#)

Distribution lists

Your e-mail system **SHOULD** support the expansion and management of internal distribution lists. You are likely to want to maintain such lists so as to easily write to all your staff or students, or all those in certain departments. Most e-mail products will support lists, possibly by purchasing and installing additional software or components.

Check that:

- you are satisfied with the 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 you may wish to extend this to certain managers or others, to anyone inside your organisation, or even to certain individuals or roles outside your organisation.

[Up](#) | [Previous](#) | [Contents](#)

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;
- and the way in which originators and recipients will then be advised of the disclosure of this information.

The content of this statement will be influenced by the [Data Protection Act](#) ^[21] and the [Regulation of Investigatory Powers Act](#) ^[22].

[Up](#) | [Previous](#) | [Contents](#)

User support

Users can expect support of various sorts:

- routine requests for changes to their account details;
- information on the status of your e-mail service;
- Advice on the e-mail software they use;
- advice on reports (often failure reports) from your own e-mail system or from one somewhere else;
- action on what they perceive as abuse through the e-mail service, including both unwanted bulk e-mail (UBE, spam etc) and abuse which appears to be personal;
- advice on good practice in using e-mail;
- advice on the use of e-mail to access external services (eg mailing lists);
- advice on the interaction between your e-mail service and dialup or other connection services they may wish to use;
- directions on security and acceptable use (with reference to the Janet [Acceptable Use Policy](#) ^[23] as appropriate).

Much of the advice and information will best be provided through internal Web pages or other documentation; effort put in to the maintenance of a Frequently Asked Questions list is likely to be effective.

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

[Up](#) | [Previous](#) | [Contents](#)

Service scaling

The system or systems providing 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;
- the central computer accepting incoming e-mail connections;
- the central computer sending e-mail outside your organisation;
- the central computer storing delivered e-mail for your users to read;
- the computers and software with which your users read and compose their e-mail.

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

For a great variety of reasons, very few Janet 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. If your organisation needs specific comment on a proposal, contact [Janet Service Desk](#) [8].

[Up](#) | [Previous](#) | [Contents](#)

Service agreement

Whether your e-mail service is resourced internally or from outside, both providers and consumers of the service need to be clear what they can expect.

For an outside contractor you will normally have at least 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. Some of the above headings may be covered elsewhere or may be thought unnecessary.

Where facilities are provided away from your own site or networks, each of the headings may have impact on your internal e-mail as well as external traffic, and you may assess some of the risks differently.

[Up](#) | [Previous](#) | [Contents](#)

Glossary

DNS

The Domain Name Service. IP addresses are impracticable for people to type, recognise or remember and the DNS makes it possible to use names instead. Domain names have several components separated with dots forming a hierarchy, with the last components representing the top of the tree; for example, ns0.ja.net. The DNS is the Internet's distributed lookup service associating IP addresses with names. It is specified in [RFC 1035](#) [24] and related RFCs.

Both management and operation of the DNS are delegated and distributed. Janet(UK) allocates all names ending in ac.uk and assigns responsibility for each domain immediately below there to some organisation. The organisation operates (or arranges for a contractor to operate on its behalf) two or more nameservers - computers with DNS software - from which all Internet users can obtain the IP address corresponding to a name in that domain.

Other DNS information includes PTR (pointer) records through which IP addresses can be converted to domain names. The DNS also contains data supporting its own

management and operational integrity.

ESMTP

SMTP extensions allow two mail systems to negotiate facilities beyond those of SMTP, both for management and for kinds of message content not provided for in SMTP. The first item in an ESMTP exchange is EHLO, which takes the place of the HELO used otherwise.

Header

An Internet mail message is transferred between mail systems in the form of a text file. The first lines of the file are the header; next comes a single blank line, followed by the lines which make up the body of the message (some of which may also be blank). The header part is created by the sending user's mail program and subsequently added to by each mail system through which the message passes; it is closely defined by [RFC 822](#) ^[10] and it is the responsibility of all those mail programs and systems to follow the rules. The body of the message can obviously include whatever the sending user wants, although for anything except the simplest typed text messages the MIME standards require that the mail program encodes it in specific ways and adds header lines which will enable the recipient's mail program to recover what was intended.

IMAP

Internet Message Access Protocol. An alternative to POP in which messages can remain in a store on a mail server and users can read and manipulate them there. Usually IMAP supports folders and possibly address books in the message store so that users have the same view of their mail, no matter where they are when they connect to the server. IMAP connections support password or similar authentication. This arrangement needs more server capacity than the combination of POP for fetching mail and SMTP for submitting it.

The version in most common use at the time of writing is IMAP4, described in [RFC 2060](#) ^[16]

IP

Internet Protocol. [RFC 791](#) ^[25] and related RFCs define the interpretation of information packets. The use of IP addresses to indicate the source and destination of each packet is part of the specification. The RFCs also specify how a computer should react to packets with particular kinds of information, although this is at a very crude level and, for instance, the transfer of an e-mail message involves the exchange of many IP packets.

IP address

In a reasonable but rather simplified view, each computer connected to the Internet has a unique number allocated to it while it is connected; this is its address and much of the machinery of the Internet is concerned with finding a route so as to pass packets of information from one address to another.

Present (version 4) IP addresses are 32 bits long, roughly corresponding to numbers up to about 4 thousand million. They are usually written using four 8-bit numbers (between 0 and 255 in ordinary decimal notation) separated with dots; for example,

128.80.230.187.

ISP

Internet Service Provider. Janet provides Internet service to customer organisations, but most Janet users obtain Internet connections from other providers at certain times. For instance, many individuals will have dialup accounts which they use from home and possibly from elsewhere, and the extent to which your organisation's Janet facilities are then available to them may be limited by security considerations.

MIME

Multipurpose Internet Message Extensions. Internet mail relying only on RFC 821 (SMTP) and RFC 822 supports simple text messages with no formatting. Although this is adequate for many purposes, people now expect to use e-mail for transferring electronic information in other forms such as spreadsheets or word processor documents.

The MIME standards in RFC 2045 and related RFCs set out how an e-mail program can encode such things in a text form, not for recipients to read directly but for their e-mail programs to decode and deal with in some appropriate way.

The standards specify how e-mail programs can cooperate to deliver what their users want, while the messages which pass across the Internet are of the kind required by the basic standards and present no difficulties to the e-mail systems in between.

POP

Post Office Protocol. In most cases a message is delivered to its recipient by writing it into a file or database in a mail server where it is stored. It is not necessary for the recipient to be using her or his computer at the time; they can collect the message when it is convenient for them. POP is a protocol with which user e-mail programs can retrieve delivered messages from the mail server.

Since user programs can easily retrieve messages at regular intervals without attention, users do not always appreciate the distinction.

Commonly, a POP server program responds to user requests by manipulating the same store of messages as the main mail program may be delivering message into so a degree of cooperation is necessary and some POP products are bound in with mailer products. At present POP3, described in [RFC 1939](#) ^[15] and related RFCs, is the most widely deployed version.

RFC

Request For Comments. Documents making public the Internet's protocols and many of its procedures. The authoritative repository for RFCs is <http://www.rfc-editor.org/> ^[26]. Some RFCs have the status of standards; others are for information or are experimental. *Internet Drafts* are documents under development, some of which later become RFCs.

SMTP

Simple Mail Transfer Protocol. [RFC 821](#) ^[14] describes the negotiation between two mail systems which is the basis for Internet e-mail.

The sending mail system first identifies itself (with the HELO command) and provides

the e-mail addresses of the originator and recipients of the message. The receiving system acknowledges each item. The sender then sends the complete message without interruption, with a special code to mark the end of it.

SSH

Secure SHell. SSH is a collection of client programs and corresponding servers which have similar functions to well-established and convenient network applications but which use encrypted TCP connections. The IETF *secsh* Working Group has produced several Internet Drafts and is preparing RFCs.

SSL

Secure Socket Layer. A protocol set which encrypts TCP traffic from a low level, requiring system software configuration at both ends of the secure path. SSL version 3 is described in an Internet Draft dating from 1996, but in no published RFC.

Submission

A message may be transferred from one mail system to another several times in its journey. Most of these transfers are between pairs of mailers which are essentially peers to each other, and neither has any special interest either in the message itself or in its originator and recipient. However, the first transfer is special and is called submission; at that stage the message moves from a program under the control of the originating user into the general Internet transfer scheme.

The submission stage is normally the only moment at which it is possible to know which individual or account sent the message, so it is an opportunity to authenticate the originator; it is usual also for user e-mail programs to leave messages incomplete (perhaps without Date: or Message-ID: header lines) and the mail system accepting submission will need to supply such information.

RFC 2476 ^[13] formalises arrangements for submission and in particular assigns TCP port 587 for the purpose. At the time of writing there is little use of the submission port or facilities; most user e-mail programs use the SMTP port (port 25) and the mail systems to which they submit ("*SMTP servers*") apply *ad hoc* authentication and completion.

TCP

Transmission Control Protocol. For many applications including electronic mail, it is convenient to ignore the details of the exchange of individual IP packets, and to work as if a stream of information was flowing from one computer to another.

TCP keeps track of packets so as to provide this service. TCP packets include information to which the destination is expected to respond so that information arrives complete and in the right order, and if that is not possible TCP will report an error.

The proper name is TCP/IP - TCP over IP - since the packets forming the stream of information are IP packets. Other arrangements are possible in principle.

TCP ports

A computer with a single IP address may be able to provide several services using TCP connections. As well as the IP address, TCP connections carry with them a number between 1 and about 30 000, A different number is assigned to each service;

the computer can identify the expected behaviour from the port number and respond appropriately.

For e-mail, a system able to receive mail normally accepts connections on TCP port 25 and will make SMTP responses to information directed to that port. If the system is also a Web server it will accept connections on TCP port 80 and provide there responses according to the Web protocol HTTP.

UBE

Unsolicited Bulk E-mail. Also called spam, junk mail, UCE. Certain individuals and businesses send mail indiscriminately to e-mail addresses for which they have not in any sense obtained permission; associated with UBE is a practice of concealing the identity of the individual or organisation responsible. All this causes problems of privacy, security, and unauthorised use of resources.

[Up](#) | [Previous](#) | [Contents](#)

Source URL: <https://community.jisc.ac.uk/library/janet-services-documentation/janet-site-e-mail-requirements>

Links

- [1] <http://www.ietf.org/rfc/rfc2119.txt>
- [2] <https://community.ja.net/library/janet-services-documentation/dns-faqs>
- [3] <mailto:Karl.Marx@college.ac.uk>
- [4] <mailto:k.marx@dept.college.ac.uk>
- [5] <mailto:fr03200@college.ac.uk>
- [6] <mailto:karlm-chemsrv2@ntserver2.chem.college.ac.uk>
- [7] <http://www.ietf.org/rfc/rfc2142.txt>
- [8] <https://community.ja.net/library/janet-service-desk>
- [9] <http://www.ietf.org/rfc/rfc1869.txt>
- [10] <http://www.ietf.org/rfc/rfc822.txt>
- [11] <http://community.ja.net/library/janet-services-documentation/network-time-service>
- [12] <http://www.ietf.org/rfc/rfc2045.txt>
- [13] <http://www.ietf.org/rfc/rfc2476.txt>
- [14] <http://www.ietf.org/rfc/rfc821.txt>
- [15] <http://www.ietf.org/rfc/rfc1939.txt>
- [16] <http://www.ietf.org/rfc/rfc2060.txt>
- [17] <http://www.ssh.org/>
- [18] <http://www.openssl.org/>
- [19] <http://www.hotmail.com/>
- [20] <https://apps.db.ripe.net/search/query.html>
- [21] <http://www.legislation.hmso.gov.uk/acts/acts1998/19980029.htm>
- [22] <http://www.legislation.hmso.gov.uk/acts/acts2000/20000023.htm>
- [23] <https://community.ja.net/library/acceptable-use-policy>
- [24] <http://www.ietf.org/rfc/rfc1035.txt>
- [25] <http://www.ietf.org/rfc/rfc791.txt>
- [26] <http://www.rfc-editor.org/>