

BEST PRACTICE GUIDE

Exploiting and protecting the network

Edition 4.0



Universities and Colleges
Information Systems Association

Contents

| | | |
|----|--|----|
| 1 | Summary and recommendations | 3 |
| 2 | Introduction and aim of document | 7 |
| 3 | Institutional policy and practice | 8 |
| 4 | Legal and responsible use | 12 |
| 5 | Registration, authentication and authorisation | 16 |
| 6 | Network configuration | 18 |
| 7 | Network vulnerability assessment | 22 |
| 8 | Encryption | 23 |
| 9 | Mobile computing | 25 |
| 10 | Cloud computing | 28 |
| 11 | Wireless access and security | 30 |
| 12 | Email | 32 |
| 13 | The CERT, security policy and practice | 38 |
| 14 | Business and community engagement | 40 |
| 15 | References in order | 41 |
| 16 | References by publisher | 44 |
| | Acknowledgements | 46 |



Universities and Colleges
Information Systems Association

University of Oxford
13 Banbury Road
Oxford OX2 6NN

Tel: +44 (0)1865 283425
Fax: +44 (0)1865 283426
Email: admin@ucisa.ac.uk
www.ucisa.ac.uk

1 Summary and recommendations

An institution's network and its link to JANET and thence to the global internet comprise a valuable resource to help it achieve its aims. Networking is mission critical, carrying data and increasingly voice/video traffic. However, the network also exposes the institution to various dangers. This document draws attention to these dangers, makes recommendation on the management of the network infrastructure and associated facilities, and provides pointers to sources of further information to help institutions to protect themselves against those dangers whilst, at the same time, making full and appropriate use of the resource.

Recommendations

Institutional policy and practice

- 3.1 Institutions should ensure that a set of policies and regulations regarding use of the network are established, that these are well communicated and are regularly reviewed in order to ensure that they are adequate in the current environment.
- 3.2 Staff contracts and the regulations signed by students on enrolment must include sections relating to computer and network misuse; senior management must be prepared to take action if any breach is detected.
- 3.3 Institutions should do their best to ensure that file store areas are kept free of offensive material and that they have procedures in place to respond promptly when the presence of inappropriate material is brought to their attention.
- 3.4 The institution should review their policy with regard to monitoring of network traffic, reserving the right to monitor traffic, subject to the constraints of national legislation, and the policy should be fully communicated to users.
- 3.5 Institutions should ensure that their IT staff have sufficient authority, in writing, to carry out any necessary investigation of misuse, that they have at their disposal the means to carry out appropriate monitoring and that they are adequately trained. Training must include warnings against misuse of these facilities and that they are not to monitor content unless explicitly instructed to do so by the appropriate authority.

Legal and responsible use

- 4.1 Institutions should ensure there is a policy on the use of the internet, including email, and should consider their policy on the provision of information on the network.
- 4.2 Institutions should have procedures in place to respond promptly when the presence of material that might infringe intellectual property rights is brought to their attention.
- 4.3 Institutions should consider organising training programmes incorporating guidelines on responsible use for students and staff.
- 4.4 Institutions should maintain a watching brief with regards to the implications of the Digital Economy Act. This can be facilitated through reviewing the updates being circulated by JANET(UK).

Registration, authentication and authorisation

- 5.1 Institutions should ensure that all users are allocated a unique user ID and password, and that they are registered and de-registered in a timely manner.
- 5.2 The registration process should ensure that users acknowledge that they understand the conditions of access to facilities for which they are authorised.
- 5.3 Institutions should consider investigating providing additional levels of client security at layer 2, such as DHCP snooping, port-security and 802.1X technology for identity based authentication systems.
- 5.4 Institutions should ensure that all use of IT facilities, whether centrally managed or departmental facilities, is restricted to authorised users and that appropriate authentication mechanisms are in place.

Network configuration

- 6.1 The institutional network should consider dividing the network into separate logical or physical domains and controls introduced to manage the traffic between the domains.
- 6.2 The physical security of the network should be reviewed to ensure that all components of the network are adequately protected.
- 6.3 Implementing a form of Network Access Control is desirable, particularly for open access areas and Halls of residence networks.
- 6.4 Consider implementing a segregated VLAN for the VOIP network.
- 6.5 Institutions should consider familiarising themselves with IPv6 and consider implementing it.
- 6.6 These days it is not possible to only allow access to the institutional network from authorised equipment, however a suitable form of authentication should be implemented to only allow access from authorised users.
- 6.7 It is essential to implement firewalling at external gateways, organisations should consider establishing a DMZ for internet facing services, and to consider implementing a default deny policy, particularly for incoming connections. Firewall rules should be regularly reviewed.
- 6.8 Institutions should review the logs being kept of network and user activity and ensure that they are adequate for the needs of investigating misuse and for the tracking of problems.
- 6.9 Institutions should consider installing a time server system, which should be synchronised with the JANET Network Time Service, and synchronise the clocks on all systems.
- 6.10 Institutions should consider implementing IDS/IPS for their campus connection to the Internet.
- 6.11 Institutions should consider implementing DNSSEC.

Network Vulnerability Assessment

- 7.1 Institutions should consider implementing a Vulnerability Assessment policy that specifies which servers should be monitored and how often they will be scanned.
- 7.2 Institutions should consider investigating routine external accredited vulnerability assessment (penetration testing) for key servers.
- 7.3 Institutions should consider scanning devices highlighted by PCI DSS, at least quarterly, with a recognised vulnerability scanner that complies with the PCI DSS standards.

Encryption

- 8.1 Encryption should be used to store, transmit or transport critical data.
- 8.2 Commercial encryption products chosen for implementation should comply with the latest standards i.e. is FIPS compliant.
- 8.3 Institutions should have a policy in place for remote workers that require them to ensure they use appropriate encryption software, preferably provided by the institution, to store, transmit and transport information.
- 8.4 A defence-in-depth approach while choosing an encryption solution will enable organisations to provide stronger information security controls for their data and the flexibility for it to be applied at various levels e.g. client, network/infrastructure etc.

Mobile computing

- 9.1 Facilities for remote or mobile access should be carefully designed and implemented. All remote access should be properly authorised and all users of remote or mobile equipment should be given guidance in mandatory good practice.
- 9.2 Access to any sensitive information or service should only be permitted over secure, encrypted connections.
- 9.3 Institutions should consider implementing IMAP (or IMAPS) and authenticated SMTP services (with TLS encryption enabled) to facilitate secure remote access to email. This is important for users with devices such as smartphones over mobile 3G networks secure access their email. Institutions using Microsoft Exchange should also consider supporting ActiveSync and RPC over http/https for remote access.
- 9.4 Institutions implementing VPN facilities should consider L2TP over IPsec, or a SSL VPN solution.

Cloud Computing

- 10.1 An appropriate evaluation needs to be carried out to have a clear understanding of the key business and technical drivers such as scalability, cost, platform independence etc. before any application or services can be moved to a cloud based environment.
- 10.2 Security evaluation should form a key criterion to choosing a suitable platform from a cloud service provider similar to the one suggested in this document.
- 10.3 Institutions should consider a risk assessment of their connection to JANET. Is a resilient connection available with diverse fibre routes?
- 10.4 Institutions should consider evaluating the spare capacity required on JANET connections to engage in Cloud Computing activities, this may need to be in partnership with JANET(UK) and/or Regional Network Operators (RNOs).
- 10.5. Institutions should consider monitoring and graphing their Internet connections to JANET to establish the trends and requirements associated with cloud computing.

Wireless access

- 11.1 Any deployment of wireless networking must be carefully planned and the expectations of users carefully managed. Currently, wireless networks should be promoted as complementing wired networking rather than an alternative.
- 11.2 Access points that support multiple SSIDs and VLANs should be used, with SSIDs being used as community identifiers and VLANs employed to permit differentiation of services.
- 11.3 Removal of 802.11b support from access points should be considered.
- 11.4 Encryption should be mandated for all sensitive information, including user IDs and password, and institutions should employ WPA2 for secure connections.
- 11.5 For secure access, institutions should deploy WPA2 with 802.1X and a RADIUS authentication server – WPA2 Enterprise.
- 11.6 Institutions should consider joining eduroam.
- 11.7 Procedures should be in place for the identification of any rogue wireless devices. Organisations should consider making a decision on the use of active mitigation technology, including a risk assessment of which access points may be mitigated.

Email

- 12.1 Institutions should ensure that appropriate policies and procedures are in place to meet their legal obligations and covering issues such as avoiding defamation, disclaimers, avoiding creation of unauthorised contracts, and bulk emailing.
- 12.2 For every mail domain supported by an institution, mail to postmaster@domain and abuse@domain should be checked frequently by a mail system administrator.
- 12.3 Mail systems must be set up so as to prevent relaying from outside the domain to outside the domain except when the incoming connection has been properly authenticated as coming from an authorised user.
- 12.4 Institutions should make use of reputable block listing sites for configuring their mail systems to minimise the amount of spam delivered to the institution's users. Consideration should also be given to implementing other techniques such as greylisting or content filtering, including the possibility of outsourcing these functions.
- 12.5 Institutions should take account of relevant legislation, including the *Data Protection Act* and the *Regulation of Investigatory Powers Act*, particularly regarding scanning the content of emails, whether for virus-protection or for other reasons, and should make their users aware of the conditions when their incoming and outgoing emails might be monitored.
- 12.6 Institutions should be aware that email is likely to be mission critical and should take appropriate measures to protect the facility from being completely, or partially, disabled through malicious or accidental action. Institutions should consider installing an emergency back-up facility (e.g. an additional connection to the Regional Network or an ADSL line) to protect against external network failure.

The CERT, security policy and practice

- 13.1 Institutions should ensure that they appoint a Computer Emergency Response Team (CERT) contact who has adequate internal powers to take actions recommended by JANET CSIRT.
- 13.2 An institution's CERT staff should inform themselves of the issues relating to computer and network security and should keep themselves updated.
- 13.3 An institution should take what preventative action is possible to protect their institution against those with malicious intent.
- 13.4 Institutions should develop a suitable Incident response plan with clear definitions of scope of security incidents that will be handled. The plan can be based on the recommendations made in this document.
- 13.5 Institutions should keep JANET CSIRT informed of security incidents involving their JANET traffic.
- 13.6 Institutions should ensure that their network operations provide adequate local support for CERT activities and out of hours cover sufficient to support their main operations.

Business and Community Engagement

- 14.1 Institutions should carefully consider the legal implications and publicity issues of possible abuse of a network connection by an external organisation.
- 14.2 Institutions should keep a watching brief of Business and Community Engagement activities within JANET(UK) and the sector.
- 14.3 Institutions should consider making realistic charges, covering both direct costs (including any network charges) and indirect costs (including support), to the connected organisation.
- 14.4 Institutions should ensure that their own bandwidth requirements have first call on the access bandwidth allocation so that external organisations are not subsidised at the expense of the host institution.
- 14.5 Institutions should consider what network-related services (such as domain name allocation, nameservers, electronic mail, etc) that they are prepared to provide to the connected organisation.

2 Introduction and aim of document

This is the fourth edition of this document, which has been produced by the *University and Colleges Information Systems Association* Networking Group, and it supersedes earlier versions. The document takes into account the further growth in network use since the original versions, changes in legislation, and changes in the allowable uses of the JANET and regional networks. Much of the structure and text of the current document has been taken from the previous versions and the work of the original authors and editors is hereby acknowledged.

The term the network should be understood to mean the collection of cables, wireless access points, network switches, IP enabled telecommunications equipment and routers that link the various computers and other IP devices together within an institution and onward to the Joint Academic Network (JANET) and thence the global internet. Also included within the term, but depending on the context, are various related network service hosts such as web and email servers. A single network is usually considered to be delimited by the logical boundaries of the organisation that has management responsibility for the network. Thus the JANET network, which is managed by JANET(UK), is considered a separate network from an institution's network, which is generally managed by the institution's IT Services, even though JANET and the institution's network may be interconnected.

This document attempts to strike several balances:

- between protection of both the network and the institutions connected to it and the benefits of widespread use of networking;
- between being prescriptive and *laissez faire*;
- between being all encompassing and easy to read.

The reader is cautioned against assuming that this document provides all the answers – it does not. Nor does it even raise all the questions. Even if it did, it would gradually become dated. The material in this document has not been vetted by lawyers and therefore, where there is any doubt, individuals and institutions must obtain their own legal advice on whether they should follow the guidance given, especially in areas where the civil and criminal justice systems are involved. Where advice and recommendations are given, this is believed by the UCISA Networking Group to represent good practice at the time of writing.

The UCISA Networking Group would appreciate being told about any errors in this document and any areas that the reader believes should be covered, but are not. Should you have any comments about this document, please forward them to the Executive Secretary via execsec@ucisa.ac.uk.

3 Institutional policy and practice

The provider of communications facilities has a number of responsibilities established by law and there are issues, such as confidentiality, that are of concern to all institutions. There are also a number of responsibilities that are, by convention, expected to be assumed by those providing services over the internet, for example supporting email servers. Reference should be made to *RFC 1173: Responsibilities of Host and Network Managers*¹.

It is clear that a set of policies and regulations regarding use of the network are needed, as are disciplinary procedures to cover cases where persuasion and education have failed. The policies should include acceptable use policies (AUPs), connection policies, and security policies for the various network services. An institution's policies and practice should be ratified by the institution's governing body, or the development and promulgation of those policies and practice should be explicitly delegated by the governing body to an appropriate part of the institution. The policies and practice must be applicable to all staff, whether academic, academic-related or administrative, to students, to conference attendees and to other visitors and partners, although some tailoring may be necessary for specific groups. It is pointless for IT Services to develop and promulgate policies if those policies do not have the necessary authority and support from the institution. The policies must have sufficient authority to ensure that they are followed and for disciplinary action to be initiated if they are contravened.

The institution's policies on network use might well be part of its policies for use of IT in general and these policies could best be framed in general terms. UCISA has published an *Information Security Toolkit*², which could be used in drawing up or revising your own regulations.

Being too specific is likely to lead to the need to revise the policies frequently as technology develops. The general policies should be backed up by a guidance document, which could well be produced and promulgated by IT Services, to explain the implications of the policies. Thus, if the policy stated that the institution's IT facilities may be used only by properly authorised persons, there might be guidance to indicate that use of another person's identification and password would be treated as a serious offence. It is helpful to phrase policies in terms of maintaining good order for the benefit of the institution and its members as a whole, rather than as a set of *diktats* apparently written to take the fun out of IT.

It is important that the institution's policies and practice are both relevant and respected, and that they are consistently enforced. It is essential that regulations are clear and that they are regularly communicated. It is also recommended that they are incorporated into employment contracts, into conditions when registering students, and into contracts when appointing contractors.

Many institutions publish their IT policies and guidance openly on the web and most will have no objection to other institutions benefiting from them, although it would be prudent to ask before engaging in outright plagiarism.

The following is taken from the University of Cambridge's *Use and Misuse of Computing Facilities*³:

"It must also be understood that computer systems and networks are not designed to prevent every form of misbehaviour and it is therefore naive to think that just because something is possible it is necessarily permitted.

Users should understand that both the Rules and these notes have been drafted with a view to maintaining good order, which means not only preventing illegal or undesirable behaviour, but also ensuring that the use of shared facilities such as a computer or a network for bona fide academic work is neither jeopardised nor disrupted.

The maintenance of good order for the sake of the majority requires constant vigilance by those responsible for the operation of shared facilities. Small, even trivial, misdemeanours repeated on a large scale can result in the waste of large amounts of valuable staff time. For this reason, actions that result in significant waste of effort can be just as unacceptable as more flagrant breaches of the Rules. Being inebriated at the time of the misdemeanour is not an excuse.

Finally, users are also expected to be guided by common sense. Over-pedantic interpretation of the Rules or these guidelines is no substitute for common sense; a failure to act sensibly may in itself be regarded as a breach of the Rules."

Regulation

There are two basic approaches to ensuring that facilities of any kind are used in a responsible manner. One is to rely primarily on a set of regulations that say precisely what is and is not allowed, and to set out procedures to be invoked when the regulations are broken. The other is to educate members of society to behave well, using formal regulations as a backstop when things go wrong. These two approaches are not mutually exclusive.

1. <http://www.ietf.org/rfc/rfc1173.txt>

2. <http://www.ucisa.ac.uk/ist>

3. <http://www.cam.ac.uk/cs/iss/rules/guidelines.html>

Since our networks are interconnected, there is considerable advantage in making the regulatory environment as uniform as possible. The UCISA *Model regulations for use of institutional IT facilities and systems*⁴ were written to provide advice to institutions in this area. These are deliberately both harsh and clear, because they are likely to be invoked only in clear cases of misuse. Most institutions will already have a similar set of regulations, but it is recommended that these are regularly and formally reviewed by an appropriate body within the institution, in order to ensure that the latest applicable laws, as well as newer uses of computers, are included. Sites are strongly recommended also to read the UCISA *Information Security Toolkit*⁵ when reviewing their regulations and might also wish to refer to RFC2196, *The Site Security Handbook*⁶.

As well as any regulations that an institution may have put in place, the provisions of the law will also apply. See section 4, *Legal and responsible use*, for more details. Additionally, an institution's connection and use of JANET are subject to *Terms for the Provision of the JANET Service*⁷ and various policy documents approved by the Joint Information Systems Committee (JISC):

JANET Acceptable Use Policy (<http://www.ja.net/services/publications/policy/aup.html>)

JANET Security Policy (<http://www.ja.net/documents/publications/policy/security.pdf>)

Some institutions are connected to the JANET network via a Regional Network Operator and there are likely to be similar contracts and policies relating to use of the Regional Network.

Disciplinary procedures

Inevitably, such regulations tell users basically what they may and may not do. They typically also describe the judicial and disciplinary mechanisms that can be invoked in cases of an apparent breach. They are often issued together with a more informal summary that tries to outline the basic principles and may give some idea of what is allowed; the usual guidance is that if an individual's usage comes to the attention of the network service providers, e.g. from an institution's network monitoring activities or from a complaint, there is probably a *prima facie* case of misuse. The institution will investigate and, if there is actual misuse, will invoke the rules. However, explicitly looking for trivial breaches of the rules is unlikely to be a sensible use of staff time.

When material or network use which is probably illegal is identified, it must be referred to the police. Their advice on evidential requirements must be taken before any approach to the suspected person. In a case where the police have been involved because of suspected or actual criminal activity, an institution must ensure that its disciplinary proceedings follow, and do not precede, criminal proceedings.

Acceptable and unacceptable use: censorship or selectivity

This is a difficult area, one which some might wish to avoid but this is not possible. There is a duty to ensure that the reputation of the institution is maintained and that the passage of material that is in breach of the JANET AUP is not knowingly permitted. Whilst it is contrary to the ethos of higher education and scholarship to engage in true censorship, some institutions may choose not to facilitate access to some materials. Selectivity can take several guises:

- blocking material that would be illegal to transmit;
- blocking material that would contravene relevant acceptable use policies;
- blocking material that would damage the good name of the institution;
- blocking unauthorised material that purports to be official;
- blocking material for reasons other than its content, e.g. because of its volume.

The classic area where selectivity might be invoked is the transmission and storage of adult material, but there are others. It is regularly argued that, since this material is reasonably freely available on the network, we should not impede its transmission. The argument that the user might be using it in their research is regularly raised. However, well publicised incidents have shown the dangers to the reputation of the institution. Whether an institution implements selectivity or not, there is a need for a clear statement of what is acceptable and what might give rise to disciplinary action. Such a policy must be communicated widely.

Most material in this category is accessed whilst browsing the web but there are other means of access; these include groups, such as *Google Groups*, and discussion forums. Those institutions that, for example, support discussion forums should remain alert to the possibility of inappropriate material being held on their servers and ensure that material is adequately filtered. Needless to say, the filter mechanism must be subject to ongoing review.

4. <http://www.ucisa.ac.uk/publications/modelregs.aspx>

5. <http://www.ucisa.ac.uk/ist>

6. <http://www.ietf.org/rfc/rfc2196.txt>

7. <http://www.ja.net/documents/publications/policy/service-terms.pdf>

Intervention is necessary in two situations. First, it cannot be acceptable to use institutional equipment (or personal equipment on an institution's premises) to store defamatory, obscene or other proscribed material whether or not this is for onward transmission. This is also the most likely activity to give rise to unfavourable publicity. Second, onward transmission of material to a recipient who has not requested it explicitly can be harassment and, as such, must also be unacceptable. Unfortunately, it is impossible to police this activity (in the sense of preventive action), except in repeated cases.

The policing of traffic is a point that cannot be conceded. It is not feasible to expect the network service provider to take responsibility for the contents of every packet that is transmitted across the network. It has to be accepted that the determined searcher for obscene, racist or other offensive material, who knows where it is stored, might go undetected; this is akin to receiving such material at home in an unmarked envelope, in which case the postal service is not held responsible. This principle, of considering whether network misuse would also be misuse if another transmission medium were being used, is a sound one.

However, an institution must take prompt action when the presence of inappropriate material on their servers is brought to their attention. Any protection afforded to the institution by the *E-Commerce Directive*⁹ is lost once there is actual knowledge of the material. Procedures must be established for assessing material and for removing it if it is deemed inappropriate. Such procedures might lead to the initiation of action under disciplinary regulations.

There is a need for balance. An institution might choose not to implement selectivity on web accesses, relying on a clearly stated policy on use, but might find it advantageous to filter material being delivered in email. Offensive material arriving in emails is usually unsolicited and implementing good spam filters will protect the users. Each institution is likely to adopt slightly different policies in this area; however, each should bear in mind that major differences in policy will lead to situations where material that is banned in one institution could easily be obtained from a neighbouring institution – this could make both institutions look somewhat foolish.

Investigation and monitoring

In addition to the institution's regulations making clear what is acceptable, they must also clearly state that any alleged misuse will be investigated and under what circumstances the contents of messages and files might be monitored. The *Regulation of Investigatory Procedures Act* (and its associated *Lawful Business Practice Regulations*), the *Data Protection Act* and the *Human Rights Act* are all relevant to this area (see section 4, *Legal and responsible use*). The *JISC Senior Management Briefing Paper 14: The Regulation of Investigatory Powers (RIP) Act 2000: Email and Telephone Monitoring*⁹, although now a few years old, contains information that remains relevant to this topic. The Information Commissioner's *Employment Practices Data Protection Code*¹⁰, in part 3, also contains recommendations in relation to the monitoring of employees' use of IT facilities.

Users must be warned that the institution reserves the right to monitor suspected misuse, and to supply properly audited evidence for use in disciplinary or legal proceedings, subject to applicable legislation. This warning should be widely published and be present on forms that the end users sign to request network access; the warning should be worded carefully to give the institution adequate authority to monitor, but without purporting to take away any of the end users rights.

The Higher Education Information Directors Scotland (HEIDS), in partnership with the JISC Legal Information Service, have compiled a document¹¹ that describes an investigation process and incorporates general guidance notes that set the background against which any process must operate.

Since the *Regulation of Investigatory Procedures Act* (RIPA) requires that staff undertaking certain kinds of monitoring be properly authorised, institutions should ensure that their staff have received adequate training and been given written authorisation. This should be made a matter of institutional policy. In this context, institutions should note that if they are asked by the police to undertake monitoring or provide information from traffic logs etc, they should only do so if given a properly completed RIPA notice.

Guidelines like the following might be envisaged:

- Monitoring data will only be collected for proper purposes – including but not limited to: capacity planning purposes; protecting against the risk of harm from viruses and other known threats; security checking, or investigating a suspected breach of security, a suspected breach of regulations or other suspected misuse.
- Data might be obtained that identifies individual users and which may be passed to the appropriate institutional authorities for disciplinary action. Such data will otherwise not be made available, except to law enforcement or other government agencies who properly request it, or if ordered by a court.

8. <http://www.opsi.gov.uk/si/si2002/20022013.htm>

9. <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/tabid/243/ID/1235/RIPA-Part-2-New-Codes-of-Practice-22012010.aspx>

10. http://www.ico.gov.uk/upload/documents/library/data_protection/detailed_specialist_guides/employment_practices_code.pdf

11. <http://www.jisclegal.ac.uk/publications/Inappropriateuse.htm>

Recommendations

- 3.1 Institutions should ensure that a set of policies and regulations regarding use of the network are established, that these are well communicated and are regularly reviewed in order to ensure that they are adequate in the current environment.
- 3.2 Staff contracts and the regulations signed by students on enrolment must include sections relating to computer and network misuse; senior management must be prepared to take action if any breach is detected.
- 3.3 Institutions should do their best to ensure that file store areas are kept free of offensive material and that they have procedures in place to respond promptly when the presence of inappropriate material is brought to their attention.
- 3.4 The institution should review their policy with regard to monitoring of network traffic, reserving the right to monitor traffic, subject to the constraints of national legislation, and the policy should be fully communicated to users.
- 3.5 Institutions should ensure that their IT staff have sufficient authority, in writing, to carry out any necessary investigation of misuse, that they have at their disposal the means to carry out appropriate monitoring and that they are adequately trained. Training must include warnings against misuse of these facilities and that they are not to monitor content unless explicitly instructed to do so by the appropriate authority.

4 Legal and responsible use

Legal issues

Legal issues, civil and criminal, are relevant to both the individual and the institution: the individual because it is the individual, or a group of individuals, whose action may attract civil liability or directly contravene some legislation; the institution because it may similarly become liable because of the actions or inaction of its members; and also because the institution may be held responsible for not adequately overseeing the individual actions of its members. An institution should also be concerned about the effects on its reputation if one of its members, whether student or staff, is the subject of civil or criminal action with this being reported unfavourably by the Press.

Although this document is too short to explain these matters fully, the following are important areas of legislation likely to be of relevance:

Law relating to use and misuse of computers

Computers may be the targets of crime; they may also be instruments in committing crime. Institutions may be the victims of attacks against (or using) their computers and, regrettably, their users may also commit crimes against others. Most of the laws that apply to actions in the physical world also apply to those actions when carried out through the medium of computers, so online theft, fraud and harassment, for example, are all crimes. Computers can also be the means for committing new crimes involving access to data or systems; these may not cause physical damage, or deprive the rightful owner of their property, so laws have had to be created to express society's disapproval of such activities. The JISC Legal Information Service has a paper on cybercrime¹² that discusses the various types of such crime and the laws that apply.

Law relating to content of files and communications

In many situations, the law regards computer files and communications as being equivalent to objects or activities in the real world. In these cases, the same law will often be held to apply to the online versions; it is safest to assume that if a file or communication would be unlawful on paper or telephone, it will also be unlawful in electronic form. There are also a smaller number of laws that apply specifically to online content or activities.

There are, therefore, a large number of situations where institutions may be legally liable for information stored on or transmitted across their IT systems. The JISC Legal Information Service has an article summarising these areas of liability¹³ and the measures that organisations should take to protect themselves. Email and other electronic communication systems are a particular problem, as they are often perceived by users as informal even though they can easily create documents with long lasting legal effect. A binding contract can now be created through email alone and a typed name at the end of an email can function as a signature. Considerable harm to the organisation and individuals can be caused by thoughtless use of communications. To protect themselves and their users, all organisations should have a policy on use of email and the internet. The Advisory, Conciliation and Arbitration Service have a useful paper on how to develop an internet and email policy¹⁴.

Material in electronic form will normally be protected by intellectual property law, whether it was originally produced in that form or converted from some other medium. Copyright and related legislation is a fruitful area for both legislators and lawyers. Computer programs and databases, text files, graphics, sound and video recordings among others have associated rights and breaching those rights, even unintentionally, may be a serious offence in the eyes of the courts. As both creators and users of intellectual property, institutions need to be well informed about developments in this area: the JISC Legal Information Service has an introduction to intellectual property¹⁵ with links to more detailed documents.

Law relating to personal and other data held by the institution

Any data held by the institution that relates to an identifiable living individual will be subject to Data Protection legislation if it is stored on a computer or in a structured manual filing system. The JISC Legal Information Service has an introduction to Data Protection law¹⁶ and JISC has also developed a *Data Protection Code of Practice*¹⁷.

As public authorities, academic institutions also come within the scope of Freedom of Information legislation. The JISC Legal Information Service has a collection of links on freedom of information, and the JISC Records Management project is also very relevant. Both Data Protection and Freedom of Information are covered by the Information

12. <http://www.jisclegal.ac.uk/cybercrime/cybercrime.htm>

13. <http://www.jisclegal.ac.uk/publications/legalRisks.htm>

14. <http://www.acas.org.uk/index.aspx?articleid=808>

15. <http://www.jisclegal.ac.uk/ipr/IntellectualProperty.htm>

16. <http://www.jisclegal.ac.uk/dataprotection/dataprotection.htm>

17. http://www.jisc.ac.uk/publications/generalpublications/2001/pub_dpacop_0101.aspx

Commissioner's website¹⁸, which includes guidance on compliance with both Acts. Note that Scotland has its own separate legislation in the area of Freedom of Information and guidance may be found on the Scottish Information Commissioner's website¹⁹.

Law relating to the operation of communications networks

As discussed in JANET(UK)'s factsheet on User Authentication²⁰, society expects those who provide communications and publishing facilities to behave responsibly. As operators of networks and computer systems, it is often asserted that universities and colleges may be held liable for the actions of their users if they have not taken appropriate steps to prevent misuse of the facilities causing harm. The JISC Legal Information Service has a paper on the developing law of Internet Service Provider liability²¹, explaining the potential problems and what precautions need to be taken.

However, computer and network operators are also required to respect the privacy of users. The *Regulation of Investigatory Powers Act* attempts to strike a balance between the needs of individual privacy and protection of the community. The Act is explained in a JISC Senior Management Briefing Paper²², which includes a number of example situations that may arise. More generally, the Information Commissioner's *Codes of Practice on Monitoring at Work*²³ sets out the obligations of organisations that wish to monitor the activities of their users.

UK law requiring network and computer operators to keep records of the use of their systems only applies to public networks at present. However, the keeping of appropriate logfiles is good practice and recommended for all networks; logfiles are one of the most useful tools in detecting and investigating problems with computer systems, and can provide information about system faults and misuse as well as early warnings of problems. The JANET(UK) Guidance Note *Logfiles*²⁴ contains recommendations on what logs should be kept, the laws that apply to log files, and the circumstances in which they may be disclosed to law enforcement and other authorities.

Copyright

There are now various software systems that enable users to publish selected material from their computer over the Internet, either via particular servers or via distributed systems. Peer-to-Peer still remain the most common method using software such as BitTorrent clients.

In the UK, copyright holders and their industry organisations have taken a twin-track approach to copyright breach using peer-to-peer networks: sending infringement reports to the networks whose IP Addresses are used to exchange material in breach of copyright and seeking court orders (known as *Norwich Pharmacal Orders*) to force ISPs to disclose the identities of those considered serious infringers. Institutions connected to JANET have always been expected, under the JANET Acceptable Use Policy, to deal effectively with complaints of infringements: guidance on good practice in handling such reports is available as a JANET Factsheet²⁵. On occasion, reporters may request the details of the person committing the alleged breach but it is usually unwise to disclose user details without a court order. We are not aware that a *Norwich Pharmacal* order has ever been sought to disclose the identity of JANET user.

The Digital Economy Act, which came into force from April 2010, will require Qualifying ISPs to handle infringement reports in a similar way to that already practised by most institutions on JANET. If passing those reports on to subscribers proves insufficient to reduce the level of infringement, the Act permits the Government to require ISPs to impose technical measures – from site or protocol filtering up to suspension of Internet access – on those repeatedly accused of breaching copyright. JISC Legal has published a document that looks into the possible implications this might have on UK colleges and universities²⁶. Unfortunately it is still not clear (at the end of 2010) whether or not JANET or its connected universities and colleges will be classed as ISPs or Qualifying ISPs or which of the Act's provisions will apply to them. Further information is provided in a timely manner by the Chief Regulatory Officer of JANET(UK), Andrew Cormack in his blog: <http://webmedia.company.ja.net/edlabblogs/regulatory-developments/>.

If infringing material is hosted on an institution's own servers, then a report that a particular item infringes copyright (or any other law) is likely to constitute notice under the *Electronic Commerce (EC Directive) Regulations 2002*, and the institution needs to act promptly (defined as within two working days for material infringing the *Terrorism Act 2006*) if it wishes to avoid the possibility of being found liable for publication.

The JISC Legal Information Service provides an overview of Intellectual Property Rights²⁷, which briefly explains the landscape of copyright law and its application to further and higher education.

18. <http://www.ico.gov.uk>

19. <http://www.itpublicknowledge.info>

20. <http://www.ja.net/documents/publications/factsheets/041-user-authentication.pdf>

21. <http://www.jisclegal.ac.uk/ispliability/ispliability.htm>

22. <http://www.jisc.ac.uk/media/documents/publications/smbp11improvingnetwork.rtf>

23. http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/coi_html/english/employment_practices_code/part_3-monitoring_at_work_2.html

24. <http://www.ja.net/documents/publications/technical-guides/logfiles.pdf>

25. <http://www.ja.net/documents/publications/factsheets/077-investigating-copyright-complaints.pdf>

26. <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/tabid/243/ID/1603/The-Digital-Economy-Act-2010-Implications-for-UK-Colleges-and-Universities-23072010.aspx>

27. <http://www.jisclegal.ac.uk/ipr/IntellectualProperty.htm>

Responsible use

Responsible use covers wider ground than *legitimate use* – that is, something may be irresponsible without being illegal, for example:

- Use which contravenes the acceptable use policy of any of the networks involved;
- Use which may bring the name of the institution into disrepute;
- Use which causes unreasonably high traffic levels which
 - i. deprive other legitimate users of the facility,
 - ii. otherwise calls the institution's integrity into question;
- Use, for commercial activities or private gain (e.g. consultancy work) of facilities that were provided for academic purposes, without the agreement of the institution;
- Playing games, if this is against the institution's policy.

There is much guidance available on the internet itself. In particular, *RFC1855, Netiquette Guidelines*²⁸, provides a minimum set of guidelines for Network Etiquette, which organisations may take and adapt for their own use.

In addition simply to having guidelines, all users (students in particular) need to be told about them and, preferably, have their purpose explained, as this will increase the likelihood of the guidelines being followed. This requires a programme of education, supported by all relevant staff, particularly teaching and administrative staff. These latter two groups should be involved because the academic staff have more contact with students than do computer or network staff, but also because administrative staff, who have real concerns about the integrity of the institution's corporate information, must have confidence in the approaches being taken. They need to understand and sympathise with the adopted approach; IT Services cannot afford to have good practice undermined by staff in the institution sniping at it.

A multi-layer approach is desirable. First, an introductory course in IT skills should be made available to all students who need what it offers. One unit of this course should cover network navigational skills, with the opportunity taken to publicise good practice. The second layer should provide training for the increasing number of students who will already know enough not to need an introductory skills course. This could be billed along the lines of *How to make good use of your knowledge here*, but would be a further opportunity to publicise good practice. In parallel, there should be courses aimed at staff, and oriented to tell staff how they can exploit the level of knowledge students will have in their courses – another route for selling the policy. Another approach attracting interest in the USA is the requirement to pass an elementary ethics test as part of the registration procedures, giving students the opportunity to consider and discuss the issues rather than simply handing out regulations.

There is a good chance of the overwhelming majority of the university using the internet responsibly, if there is wide consultation and the whole community is persuaded to buy into the policy, possibly also by offering as a carrot, the advantages of being able to provide reasonably open access.

Guidelines for the provision of information itself are also desirable. Many institutions have provided their own guidelines and many of these can be found by using a search engine, searching on keywords *Web, Guidelines, ac.uk*.

Legal risks relating to information kiosks

A number of institutions use information kiosks in order to provide quick access to some types of information resources. These are typically PCs running a browser (and no other software) in order to access a limited set of services and do not require the users to authenticate themselves. A number of legal questions have been raised in relation to these kiosks.

There has been some discussion as to whether any use of JANET is required to be limited to authenticated users. If this were the case, it would effectively prohibit any use of the kiosks to access resources not physically located on the institution's network, ruling out access to most online information resources. The position is clarified in a JANET(UK) document²⁹, the key point is that access to JANET must be controlled responsibly and given only to those entitled to use it. In the context of kiosks, this would be satisfied by locating them in areas that are only accessible to members of the institution.

In many situations, kiosks are accessible to members of the general public and are primarily used for accessing information about the institution itself. For these kiosks, it might be prudent to restrict access to local information resources only. There are, however, two potential issues. The first is that this may block access to more services than expected, for example, online payments are often handled by a third party site and so would not be possible. The

28. <http://www.ietf.org/rfc/rfc1855.txt>

29. <http://www.ja.net/documents/publications/factsheets/041-user-authentication.pdf>

second issue is that some on-site resources will not be licensed for use by walk in users and so a further level of access control will also be required. The *whitelist* approach of allowing access only to permitted resources is preferred rather than simply blocking access to those resources to which access is not permitted, but this can require considerable effort.

Where unauthenticated users are able to use kiosks to access off-site resources, some care must be taken to reduce the risk of harm to others. Institutions that behave recklessly in allowing access to the internet are at risk of negligence suits and resulting awards of damages. Reasonable steps should therefore be taken to make it difficult for the kiosks to be used for improper purposes. These could include:

- Locating kiosks in staffed, supervised areas;
- Having a basic audit process to identify those using the area (e.g. controlled access to the area);
- Displaying acceptable use notices;
- Regularly checking that security settings are still effective.

In conclusion, there seems to be no reason why information kiosks should not be used by unauthenticated users provided that reasonable steps are taken to:

- Limit access to those information resources and services for which licence terms permit and
- Reduce the risk of kiosks being used for purposes that are likely to cause harm to others.

Recommendations

- 4.1 Institutions should ensure there is a policy on the use of the internet, including email, and should consider their policy on the provision of information on the network.
- 4.2 Institutions should have procedures in place to respond promptly when the presence of material that might infringe intellectual property rights is brought to their attention.
- 4.3 Institutions should consider organising training programmes incorporating guidelines on responsible use for students and staff.
- 4.4 Institutions should maintain a watching brief with regards to the implications of the Digital Economy Act. This can be facilitated through reviewing the updates being circulated by JANET(UK).

5 Registration, authentication and authorisation

In practice, an institution's users will be accessing facilities other than those that are made freely available. Accesses to most facilities need to be properly authorised and, in order to implement an institution's regulations and procedures or to investigate faults or breaches of security, it is necessary to know who is accessing a particular facility on the network at any given time. This implies that a registration procedure that allocates a unique identifier (user ID) for each individual user is required. Shared user IDs should only be permitted where there is good reason to do so and their passwords must be managed with great care.

Registration

Registration is the process by which details of a user are recorded and it is frequently coupled with *authorisation* for the use of a number of IT facilities. Often the user will be given one or more tokens, e.g. a user ID and password, to be used for *authenticating* access to computer systems in the institution. A user may have several different user IDs on different controlling computers each of which may have different token(s) associated with it. An example of another form of token that could be allocated is a smart card which can be used for two factor authentication, thin client access, physical access or printing.

As passwords remain the principal means of validating a user's authority to access a computer system, their allocation and communication should be controlled. Users must be made aware of the need for quality passwords and for keeping them confidential, and techniques encouraging the use of best practice guidelines should be implemented. Some guidance can be found in the JANET(UK) Factsheet on *Using Passwords*³⁰. The registration process should also ensure that the user acknowledges that they understand the conditions of access (unless this was done during an appointment or enrolment process), and that access is disabled in a timely manner when a user's rights cease (for example, by leaving the organisation or by changing job).

One of the more difficult areas with registration is that of de-registration, e.g. when staff or students leave. It is essential that they be de-registered from the system as soon as possible, particularly if they are able to access the institution remotely, e.g. using a VPN connection. Appropriate liaison and coordination with the Student Registry and HR departments is needed if this is to be effected efficiently.

Authentication

Authentication is the process by which a user proves to a relevant controlling system that they are the owner of a particular user ID or token. Typically, this requires them to supply a user ID (supplied by the institution at registration) and a password (initially supplied, but afterwards private) before allowing access.

Authorisation

Authorisation is the granting of permission to use a particular facility. Normally this involves checking some attributes associated with an authenticated user and allowing or denying permission to use a facility based on the values of those attributes.

Control of access

Access control standards are the rules that an organisation applies in order to control access to its systems and information. A clearly defined access policy statement defining the access rights of each user (or group of users) should be set out for each system on the network. Such standards should always be appropriate to the organisation's business and security needs. Inappropriate restrictions could result in individual users being unable to do their job effectively, and excessive privileges could allow a user to damage information systems and files. Records should be maintained of all users registered to use a service and of any privileges allocated.

The use of special privileges that enable a user to override system or application controls should be restricted and managed and a formal authorisation process should be defined. The unnecessary allocation and use of system privileges increases the vulnerability of a system and could lead to its being compromised inadvertently. Any system administrative access must be via a secure connection or logon process, and consideration should be given to limiting the locations, from which such sessions can be initiated and implementing an automatic, robust terminal identification system.

30. <http://www.ja.net/documents/publications/factsheets/026-using-passwords.pdf>

Network controls such as DHCP snooping, port security (a feature of the Cisco IOS) etc can also be explored to introduce an additional level of Layer 2 security. Rogue DHCP servers on the network could pose a serious security risk as clients on a network could be sent spoofed replies to DHCP queries and intercept the data being transmitted. DHCP snooping overcomes this issue through trusts, which can be setup on edge switches with a white list of IP's from which DHCP query replies would be allowed.

Port security is another feature, which allows administrators to configure a maximum number of MAC-addresses that can be accessed, i.e. number of machines that can be connected to a switchport on an edge switch. This enables key areas where users are not allowed to plug-in unauthorised machines, and when triggered would shutdown the switchport thereby disabling use of such unauthorised machines on the network.

Appropriate authentication measures should be employed for all network connected computers. Whilst in many cases this might be primarily an issue for systems design, this can be greatly assisted by appropriate planning and configurations of the network itself. Although there is a growing use of alternative authentication procedures, password security remains the most common. Institutions should consider the use of a centralised directory, such as *Active Directory* or *Novell eDirectory*, for holding users' details. Having a centralised directory, accessed by standard protocols such as LDAP or RADIUS as well as native protocols, will facilitate the maintenance of a common user ID and password across multiple systems. Further advancements in secure identity based authentication methodology such as 802.1X and client support (a.k.a. supplicant) are now available on almost all Operating Systems, this method is quickly being adopted by educational establishments at least on strategic wired network locations. This would enable more flexibility and centralised control as successful authentication can be used to automatically allocate appropriate network access privileges to a client according to their profile.

If it is not appropriate to duplicate, in a centralised directory, attributes held in databases managed by other departments: then provided protocols, such as Shibboleth, can access all appropriate attribute sources so their precise location is unimportant. As Shibboleth defines a set of protocols for the secure passing of attributes between institutions, users can access multiple resources with a user ID, provided by their home institution, while roaming. Similarly, this user ID may be used to get access to network resources when roaming between institutions that are participating in eduroam.

Remote users, either teleworkers or staff on business trips etc, may need to communicate directly with their organisations' systems. Such users are physically remote and often connecting through public (insecure) networks. This increases the threat of unauthorised access. Remote access was traditionally provided by means of dial-up or leased telephone lines but now Virtual Private Networks (VPNs) are frequently used to provide access across public networks, e.g. the internet. Connections via these networks should be authenticated by a robust secure user authentication mechanism and possibly also by a means of authenticating the remote computer. Encryption should be a requirement for any VPN connection that might be used to access sensitive information. Also see section 9, *Mobile computing*.

Recommendations

- 5.1 Institutions should ensure that all users are allocated a unique user ID and password, and that they are registered and de-registered in a timely manner.
- 5.2 The registration process should ensure that users acknowledge that they understand the conditions of access to facilities for which they are authorised.
- 5.3 Institutions should consider investigating providing additional levels of client security at layer 2, such as DHCP snooping, port-security and 802.1X technology for identity based authentication systems.
- 5.4 Institutions should ensure that all use of IT facilities, whether centrally managed or departmental facilities, is restricted to authorised users and that appropriate authentication mechanisms are in place.

6 Network configuration

Networks are generally designed to allow maximum scope for sharing of resources and flexibility of routing and, increasingly, are being extended beyond traditional institutional boundaries. However, these features might also provide opportunities for unauthorised access to IT facilities.

To reduce these risks, it is recommended that networks are divided into separate logical or physical domains and controls introduced within the network to segregate groups of users and computers that have different access requirements. Judicious use of routers, switches and VLANs can reduce the risk of serious compromise and limit the extent of any problems that do occur. Networks may require the incorporation of routing controls to ensure that computing connections and information flows do not breach the access policies. Such controls might include:

- Separation of networks for employees, third parties and students, or
- Separation of secure areas, wireless networks, open access rooms and remote access links, so that each network or subnet is only used by one target group;
- Control of the route between the user's equipment and the services that the user is wishing to access;
- Introduction of firewall zones to segregate groups of users or computers, to secure network areas and to increase the protection for both users and services;
- Segregating networks so that any fault or instability on one subnet will have minimal impact on other subnets.

Physical security

Where a network is built using fixed wiring, the prevention of unauthorised network access might be fairly straightforward. Networking and communications facilities, including wiring closets, data centres and computer rooms, must have good physical security and be protected from environmental hazards. Institutions should bear in mind the danger of unauthorised persons patching into the wiring where it is accessible or making connections to network equipment that is not located in a secure place. Institutions might wish to consider the use of equipment with the appropriate management facility to block network traffic from equipment whose connection has not been authorised, especially in areas where there is easy physical access to network sockets. This type of Network Access Control (NAC) system is commonplace these days for student halls of residence so that all users have to register their device and can be individually controlled, suspended, etc. and is becoming more and more commonplace for all areas of a campus network so that members of staff can move around the campus and still gain access to restricted resources they usually have access to. Implementing 802.1X on all network ports with a radius server is one of the most commonly used tools for controlling and auditing network access. This is good for institutions that have the resource to run and manage such a system but can be a headache for institutions, which don't have the resource, as such there are other systems available to purchase which can carry out all of this functionality with a much lower management overhead. Unused sockets should not be left in a *live* state.

The physical security of rooms containing network connected computers should be of concern, especially where it is common practice for the computers to be left running unattended. Computers deployed in open access classrooms and laboratories should require authenticated login and open access facilities operated in kiosk mode should be restricted to running safe applications. Data centres/Machine rooms should also implement a form of access control so as to restrict access to only authorised people. This can take the form of alarm systems and door access controllers, which can be connected together with appropriate software and hardware to monitor the whole environment so that doors can only be accessed after entering via another door, etc. There are systems which can be implemented to control this and also include CCTV cameras and environmental monitoring probes so as to detect high temperatures, water leaks, power failures, etc.

Cables carrying data requiring protection from interception or damage should maintain appropriate separation from power cables, and safety measures, such as fire stops, should be employed. Cabling within buildings should be protected, by using conduit or by avoiding routes through public areas. Cables between buildings should be underground where possible (or subject to adequate alternative protection) and there should be adequate security for the covers on inspection chambers. Where cables form part of a loop or are providing a resilient link, consideration should be given to using separate physical routes in order to reduce loss in the event of damage.

Wireless networks also need protection from unauthorised access. The following measures can be applied to reduce these risks:

- Wireless access points should, if possible, be positioned so that they are out of physical reach, and can be maintained only by authorised personnel;
- Power and network connections to the wireless units should be adequately protected;
- Wireless aerial power levels should be reduced to prevent the signal being broadcast beyond designated areas, while still providing coverage;

External access

The points of access to an institution's network from the outside world should be given special attention from the point of view of security. Institutions that do not already do so should have a strategic aim to control these routes. Security risks can be significantly reduced if firewalling techniques are employed at the institution's external gateway(s). It is recommended that a *default deny* policy is adopted for incoming traffic. This policy prohibits all incoming traffic except that which has been specifically authorised; for example, incoming email connections may only be made to specifically nominated institutional email servers. Institutions should also consider prohibiting incoming connections, of any type, to standard desktop PCs or other systems that should not be providing external services.

There are advantages to implementing a *default deny* policy for outbound traffic as well. Most user desktops only need to use a fairly limited set of the commonly used ports and activity on unusual ones, e.g. outbound SMTP, may indicate an infected PC being used as a mail relay or in a *denial of service* attack. If a stateful, application-aware firewall is in use, a default outbound deny policy should not affect legitimate use at all.

To assist in managing firewall rules, institutions should consider establishing a network *de-militarised zone* (DMZ), a separate subnet that is used to accommodate servers that provide external services, such as the website, DNS servers, email relays and VPN concentrators. In this way, the number of routes into the institution's internal network can be minimised, along with the risk that these might be compromised.

A firewall can also be used to separate wireless networks with the main campus wired network in order to reduce the risk of the wireless network being used for unauthorised purposes.

When a firewall is in used to protect the local network some computers, typically servers, will need special rules to protect them, or allow them privileged access outside. Over time, the set of *holes* through the firewall can rise significantly. It is essential that an institution retains track of firewall rules to ensure that they are still needed. It is, for instance, not unknown for servers to be replaced and the old one to be reassigned to another role, in which case firewall rules should be checked. Quality control techniques for firewall rules include closing holes that have not been used for some time and assigning an owner to every rule and periodically asking them to confirm that it is still necessary.

Voice over IP networks should also be kept separate to main campus networks and should be on their own segregated network/VLAN. This is to stop people being able to join the voice network and carry out malicious activity. It is common place for institutions to implement a voice VLAN and have sockets on switches available for use by either a client or a VOIP phone however a suitable risk assessment should be carried out with this type of port as there are pieces of software which can emulate a VOIP phone and give users and interface onto the voice network. There are emerging technologies like multi-domain, which will allow telephony handsets to authenticate to a 802.1X network as well as attached computers; however this has not been seen in production within the sector as yet.

It is also worth investigating the use of IPv6, as with the exhaustion of the IPv4 address space continuing, more and more networks are being IPv6 enabled. It is also worth adding rules to institution firewalls to drop such traffic if you haven't yet started enabling it. It is worth noting that some products these days are coming IPv6 enabled and disabling the IPv6 stack on these servers may invalidate your warranty/support contract, one particular example of this is Microsoft Exchange 2007 which requires IPv6 to be able to run on Microsoft server 2008.

Care should be taken with VPN concentrators to ensure that only authorised users may gain access to the institutional network and access to any sensitive systems should only be permitted over connections that have been encrypted. Logs should be kept of both successful and unsuccessful access attempts.

Occasionally, departments or individuals within the institution install their own external access arrangements, e.g. *GoToMyPC*, *LogMeIn*, etc. Users will normally install these utilities using a default configuration and, through a lack of knowledge, be unaware of the security implications. This can be a particular (and often unknown) threat to the network. Steps should be taken to locate and identify any unauthorised external access facilities, and these should then either be disconnected or brought under proper institutional management.

Unprotected diagnostic ports, out-of-band management ports and in-band access to diagnostic and management services (e.g. SNMP or web interfaces on routers) could provide a means of unauthorised access. These ports should be

protected by appropriate security mechanisms (e.g. a key lock, secure logon, dial-back or limiting the locations that can gain access) and procedures should be established to ensure that they are only accessible by agreed arrangements. In-band diagnostic or management services that are not used should be disabled on the device.

Conventional DNS servers can be vulnerable to cache poisoning attacks, which allow remote attackers to reply with their own answer to queries before the official DNS server can send its answer back to the requesting client. One particular example of this was discovered by Dan Kaminsky and was made public in July 2008, before this time the vulnerability was more theoretical and had not actually been exploited widely. DNSSEC (DNS Security Extensions) addresses this issue by signing the zone a server is authoritative for with a private key and then publishing the public key within the zones DNSSEC records so that remote sites can verify the reply they get back is from the official DNS server for that zone.

Problem tracing and logging

When network problems occur or when there is a need to investigate alleged misuse of systems, it is important to be able to identify where activity has occurred; it needs to be possible to track who, when, and where use of, and access to, the network has been made. It is strongly recommended that appropriate logs of user and network activity are maintained; logfiles are one of the most useful tools in detecting and investigating problems with computer systems, and can provide information about system faults and misuse as well as early warnings of problems. The JANET(UK) Guidance Note *Logfiles*³¹ contains recommendations on what logs should be kept.

Institutions might also wish to consider the use of automatic software protection, such as network anomaly detection and intrusion detection/prevention systems/software. IDS/IPS installations typically connect on a span port for key network connections (eg the main internet connection) or inline. These devices typically undertake pattern matching on the traffic scanner, to identify known bad traffic, exploit attempts command and control virus infections, etc. An IPS would then block these types of traffic whereas an IDS would only warn about the traffic. Network anomaly detection systems can take a span port feed or a type of network flow information (netflow, sflow, etc) and create a baseline of usual traffic. Activity outside of this baseline will be detected and can alert users of the problem or carry out remedial action (tell a Network Access Control system to disable that users access, etc). These types of system can often catch rogue machines on the network, which may get missed in typical day to day network administration and as such is a common auditor requirement. Next generation firewalls have the IPS/IDS functionality as an integral component so that separate systems do not need to be used, although depending on one system alone may also have security implications for defence in depth.

During the investigation of problems, it is good practise to correlate together the records logged by different systems and, for this relationship to be reliable, the date and time being recorded by the various systems must be synchronised. It is preferable for the time to be synchronised by reference to a consistent time server system and, if comparisons are being made with external log files, such as those maintained by JANET(UK), a common time reference is needed. The JANET Network Time Service³² delivers a stable time reference to institutions using the Network Time Protocol (NTP).

Undesirable protocols

File sharing systems are now common and their use is particularly prevalent among students. Such systems can be, and are, used to distribute copyright material and rights holders are taking legal action to prevent such abuse. To reduce the threat of legal action, institutions might choose to block ports specifically employed by file sharing software. However, this has the danger that it does not block the traffic but the software begins to operate via an external port mapping proxy, thereby making it effectively untraceable by the institution's network operators. Such action can also be countered by the file sharing service providers using a well known port assigned to some other, acceptable, service.

Other systems are continually being developed that might be seen as attractive or valuable to many users. Whilst offering value, some systems also bring risks with them.

The general advice, stated earlier, of implementing a *default deny* policy on an institutional firewall, particularly for incoming connections, and prohibiting incoming connections to desktop PCs, can help to minimise any adverse impact of these protocols.

31. <http://www.ja.net/documents/publications/technical-guides/logfiles.pdf>

32. <http://www.ja.net/services/ntp/>

Recommendations

- 6.1 The institutional network should consider dividing the network into separate logical or physical domains and controls introduced to manage the traffic between the domains.
- 6.2 The physical security of the network should be reviewed to ensure that all components of the network are adequately protected.
- 6.3 Implementing a form of Network Access Control is desirable, particularly for open access areas and Halls of residence networks.
- 6.4 Consider implementing a segregated VLAN for the VOIP network.
- 6.5 Institutions should consider familiarising themselves with IPv6 and consider implementing it.
- 6.6 These days it is not possible to only allow access to the institutional network from authorised equipment, however a suitable form of authentication should be implemented to only allow access from authorised users.
- 6.7 It is essential to implement firewalling at external gateways, organisations should consider establishing a DMZ for internet facing services, and to consider implementing a default deny policy, particularly for incoming connections. Firewall rules should be regularly reviewed.
- 6.8 Institutions should review the logs being kept of network and user activity and ensure that they are adequate for the needs of investigating misuse and for the tracking of problems.
- 6.9 Institutions should consider installing a time server system, which should be synchronised with the JANET Network Time Service, and synchronise the clocks on all systems.
- 6.10 Institutions should consider implementing IDS/IPS for their campus connection to the Internet.
- 6.11 Institutions should consider implementing DNSSEC.

7 Network Vulnerability Assessment

Vulnerability assessment should be common practise for institutions. It involves the scanning of the network, or a specific server, by a security specialist in order to detect any errors present on a system connected to the network that could be used by attackers to gain extra privileges or gain access to restricted content.

There are three types of vulnerability assessment:

- **Black box** – The penetration tester carries out the attack with no prior knowledge of the setup of the target organisation. This method reflects the typical script kiddie attacks an organisation may be subject to.
- **Grey box** – The penetration tester performs the attack with limited knowledge of the setup of the target organisation. Grey box testing reflects an attack by well prepared attackers or insiders with limited access and privileges.
- **White box** – The penetration tester performs the attack with full knowledge of the setup of the target organisation. White box testing reflects an attack by well prepared attackers or insiders with largely unlimited access.

Vulnerability assessments can be carried out by non accredited specialists, but most auditor requirements ask for CREST or CHECK accreditation to be the standard against certain systems. These accreditations guarantee that the assessment is carried in such a way that it meets the high standards specified by these companies.

Vulnerability assessments do not, however, need to be accredited and many institutions find that routine scans of existing servers or new servers yield interesting results. Many institutions will carry out automated scans against their servers using tools such as Nessus to identify misconfigured or unpatched servers. This is particularly of use when the server in question is new or run by a department outside of the central IT Services department.

Institutions should consider implementing a policy dictating that any new servers added to the network would be routinely scanned to check their security compliance. This would be carried out by a member of the institutions Information Security (or network) team, using automated software such as Nessus.

A vulnerability assessment is a key requirement for some levels of PCI DSS compliance and all servers/tills, and network equipment which can carry out payment card transactions may need to be scanned at least quarterly in order to comply.

Recommendations

- 7.1 Institutions should consider implementing a Vulnerability Assessment policy that specifies which servers should be monitored and how often they will be scanned.
- 7.2 Institutions should consider investigating routine external accredited vulnerability assessment (penetration testing) for key servers.
- 7.3 Institutions should consider scanning devices highlighted by PCI DSS, at least quarterly, with a recognised vulnerability scanner that complies with the PCI DSS standards.

8 Encryption

Encryption is a method of securing data either being stored or transmitted that can be achieved by using a suitable encryption algorithm, which transforms the original data into an encrypted (unreadable) format and a reverse process, called decryption to obtain the original data. Original information is transformed into its encrypted form using a cipher, also referred to as a key and the strength and methodology depends on the type of algorithm it uses. The algorithms used are broadly classified as symmetric and asymmetric.

The symmetric algorithm uses the same key for both encryption and decryption process. There are mainly two standards for this type i.e. Data Encryption Standard (DES) and Advanced Encryption Standard (AES). These standards are implemented in encryption algorithms, which generate ciphers that are mathematically computed with the original information resulting in encrypted data. The key size used for the cipher and the process of generating and storage are also determined by the standards used. The DES standard is known to use weak ciphers (key strength) i.e. of 56-bit, which can now easily be cracked by brute force attacks and hence has been phased out from being used. The latest encryption software's use AES standard that use up to 256-bit of key size making it strong and highly improbable to break. The asymmetric algorithm uses a pair of keys i.e. private and public keys to encrypt and decrypt information. The private key is stored locally and used to encrypt the information and requires a public key at the destination in order to decrypt the information. Key management is fairly simpler in asymmetric algorithms compared to symmetric algorithms that require a key management infrastructure to be present to ensure keys are exchange securely.

There are various reasons to use encryption, the main one being to ensure confidentiality of critical information along with which added functions that can help maintain integrity of such information through, e.g., digital signatures. It also helps ensure data is inaccessible in the event of loss or theft.

The need for encryption can be varied depending on business needs and the information security policy employed by the institution. In order to have an effective encryption solution, it is crucial to identify which type of data being handled within different areas of an institution that needs to be secured and also requires participation by the owners of the data who will be able to do this. This would then enable an appropriate encryption solution or product to be identified based around other factors, such as OS platform used. Another key aspect in considering an encryption solution requires an effective key management process depending on the scale of deployment. This would include a process of creating encryption/decryption keys its storage and a suitable recovery process.

Due to this varied requirements across businesses and industries it sometimes could be challenging at the product selection stage to choose the right solution that can achieve the desired level of security. This process is slightly eased by an internationally recognised standard for encryption products known as FIPS (Federal Information Processing Standards). This is a US government body for computer security standards that sets out stringent guidelines for cryptographic based solutions and, depending on the level achieved by a product, are awarded the relevant level of certification.

The most recent of the FIPS standard is the FIPS 140-2. This standard describes four different levels of security requirements for cryptographic modules, both on hardware and software components. Level 1 dictates the bare minimum requirements and provides the lowest level of security. As the levels progress there are more stringent guidelines for stronger physical security mechanisms with the highest level, i.e. Level 4 requiring high standards of operational environment security etc., which is normally required for defence related operations. The higher standards also have high hardware and computational power requirements which could make them very expensive. Products and solutions applicable to desktop machines, and relevant to the educational sector, generally fall under Level 2 categories, and are able to provide desired level of security and integrity, depending on specific requirements.

Encryption can be achieved at different levels:

- a. Client level: This would cover encrypting various aspects that directly relate to the end users device. This could be a Full Disk Encryption, which encrypts all information stored on the device, i.e. operating system, drives and files and folders stored on it. An example is PGP desktop, which enables a user to encrypt the whole system, which complied with the FIPS 140 standard and is available for multiple OS platforms. Apart from this scenario, encryption can also be applied only to certain areas within a client machine, e.g. only encrypt required files and folders if the requirement is limited to only storing or transmitting the data safely.
- b. Application level: This corresponds to the application being accessed e.g. web server and encryption can be achieved in the session that a client machine establishes with the server at this level. This is done through secure HTTP (Hyper Text Transmission protocol) which implements SSL (Secure sockets layer) that encrypts the information transmitted over the wire using a certificate (private key). This is useful to protect critical data, such as login information etc., from being intercepted by eavesdroppers. This function can be easily implemented on web servers and requires programmers to provide support for this security protocol within their application. SSL version 2 is the most conventional method used and administrators should look at

migrating to version 3, which addresses some of the inherent weaknesses of its predecessor, such as being vulnerable to man-in-the-middle type attacks etc. This technology is now rapidly being replaced by TLS (Transport Layer Security) which offers more resilient encryption to the network connection.

- c. Network/Infrastructure level: The underlying network infrastructure is a key component for data communication, and institutions that require high levels of data confidentiality must look into implementing appropriate encryption technologies at this level. Merely securing communication between client and a web server (e.g.) may not be sufficient in this case and would require an end-to-end solution. This can be achieved by tunnelling traffic between the two trusted locations, so all information between these points are encrypted and appear as private traffic to other sections of the network. This can be achieved using a dedicated VPN device between the two locations, which can encrypt and decrypt information between them. There are a few other factors that would need to be considered before implementing such a solution. E.g., if the solution employed encrypts the whole packet, instead of the data, and there is a packet inspection device on the encrypted link, it would take longer to process introducing delays.

Following are a few key areas to look for while choosing a suitable encryption product or solution:

- Have a clear scope when choosing endpoint encryption solutions e.g. whether you would like Whole Disk Encryption (WDE) or only file encryption etc.
- When choosing WDE products look for features such as pre-boot authentication.
- Centralised deployment, administration for better control and simple management of devices across the network. The solution must support and easy integration with existing user directory infrastructure e.g. Microsoft Active directory, Novell e-Directory etc.
- A suitable encryption password recovery mechanism to handle cases where a user may have lost or forgotten a password.
- Certain advanced features such as support for non domain users, remote recovery of encryption password, repair of corrupted encrypted disks could be good to have features as well.

Implementing and enforcing a suitable encryption policy for mobile users is always a challenge. A few factors that can be considered in this area are:

- A suitable mechanism in place which will ensure critical data made available to remote/mobile users is stored securely in encrypted form;
- All data that needs to be accessed from the organisations network should be through an encrypted channel e.g. a VPN to the institution with appropriate access privileges to required areas only. This will provide two levels of security to ensure even if the first layer were to be compromised data would still be inaccessible;
- All essential data communication such as email or file transfers must be made over a secure encrypted channel.

Recommendations

- 8.1 Encryption should be used to store, transmit or transport critical data.
- 8.2 Commercial encryption products chosen for implementation should comply with the latest standards, i.e. is FIPS compliant.
- 8.3 Institutions should have a policy in place for remote workers that require them to ensure they use appropriate encryption software, preferably provided by the institution, to store, transmit and transport information.
- 8.4 A defence-in-depth approach while choosing an encryption solution will enable organisations to provide stronger information security controls for their data and the flexibility for it to be applied at various levels, e.g. client, network/infrastructure etc.

9 Mobile computing

Modern computing and telecommunications devices make it increasingly easy to work when away from the office. Portable computing devices such as laptops and smart mobile phones, along with the greatly increased availability of networked computers, from 3G phone connections to visitor facilities in other organisations (eduroam, for example), encourage staff and students to access the institution when away from the office. This can be of significant benefit to an institution, but facilities to support remote working need to be designed, implemented and promoted to maximise the potential benefits while minimising the risks.

Since the computers and networks used for mobile computing may not be owned by the organisation and may be shared with other users, those devices cannot be assumed to implement any security so must be treated with some suspicion. When they are used on other network connections, particularly for non-work purposes, they may acquire viruses or other malware and subsequent connection to the institution's network, either directly or over a VPN, might introduce a virus or compromise the network's security.

Technical controls, therefore, need to be implemented on the systems within the organisation supporting mobile working. These technical controls must be complemented by mandatory good practice by users of mobile computing systems, for example personal computers should not be used at home for business activities if virus controls are not in place. It is essential that guidance be given to all users of remote or mobile equipment. Some VPN servers and wireless access systems are available with network admission control, where the connecting device is checked to ensure that it has appropriate antivirus measures and that its patches are up to date before allowing it to connect. This, of course, affords little protection against a laptop being brought into the institution and being connected directly to a network point in an office unless similar controls are implemented on those network points.

For some systems it will be impractical to provide adequate protection for access by mobile computing. It is, therefore, likely and reasonable that the organisation will need to select which systems and services may be used through mobile computing systems.

All mobile access needs to be authorised, although some uses, such as web access to email, might automatically be granted to all users. However, specific authorisation might be needed for remote access to more sensitive systems and these systems must implement robust authentication methods to control access.

Institutions should also consider joining eduroam³³, as a home institution, so that mobile users with wireless access facilities who are visiting other academic institutions might benefit from a consistent and easy method of wireless remote access.

Remote access to email

The most common need for remote connections is the desire to gain access to electronic mail. It is common for institutions to provide access to email by supporting a web-based interface. With more and more remote users' access emails from devices, such as tablet PCs, smartphones etc., along with conventional interfaces (e.g. Outlook) from their home PC, institutions should consider supporting more direct methods for users to gain access to their email. Most smartphones these days provide suitable methods to connect built in mail clients to most email systems, a lot of them also provide groupware connectivity.

Whilst offered by many ISPs, the protocol POP3 is not recommended for the support of email services, as messages are normally downloaded to the client computer, making it difficult to access email from multiple devices. Systems that support a central mail store are to be preferred. The Internet protocol IMAP should be considered for users to gain access to their email; this protocol is supported by most mail servers and most client devices and is normally quite easy to configure. The institution should also consider using the secured or encrypted version of this protocol, IMAPS, otherwise passwords will be passed over the internet in clear text. In conjunction with IMAP, there needs to be an SMTP service available, so that users are able to send emails. For a computer at home that is connected via an ISP, the SMTP service of that ISP may be used, but for truly mobile devices, identifying an SMTP service might prove difficult and might result in frequent reconfiguration of the client software. This difficulty may be overcome if the institution supports an authenticated SMTP service to which users may connect from any location. Client devices can then be configured, once, to use this service and, by requiring authentication, the issues around open mail relays are overcome. For SMTP authentication, Exchange servers can be configured to encrypt the communication channel before the authentication request and credentials are transmitted between server and client. Configuring the SMTP server to *Require TLS encryption* under the Authentication section does this. However, this does not encrypt all other communications between the client and SMTP server. To achieve this functionality, administrators should look into configuring their SMTP servers to use SSL for all communications. In Exchange 2003 and above, selecting the *Require secure channel* under the communications tab can configure it.

33. <http://www.ja.net/services/authentication-and-authorisation/janet-roaming.html>

Institutions that use Microsoft Exchange for their mail services can configure IMAP and authenticated SMTP but also have an additional option. Users of Exchange should consider implementing *ActiveSync* on their service. This will enable handheld devices to access email and will also enable the synchronisation of calendars and contacts. Google also provide *Activesync* access to their mail system to provide access to gmail and its calendars and contacts. Furthermore, remote users should consider using the *RPC over http (now known as Outlook Anywhere)* feature, configured with SSL as their default method to connect to Exchange, as this a more secure method to connect to an enterprise when using MAPI clients, such as Microsoft Outlook. These recommendations require that the version of Exchange being used is Exchange Server 2003 or newer.

Remote desktop services

The provision of remote desktop services (sometimes known as terminal services), using services such as *Windows Terminal Server* or *Citrix*, have been found to be valuable by many institutions; such services can provide the user with an interface remarkably similar to that which they experience on their office desktop, including access to networked filestore. If the service becomes popular, the servers will need to be configured with sufficient resources to be able to support the load generated by several concurrent users. These services are particularly valuable when there is a need to provide access to specific applications and there are difficulties in supporting those applications on remote PCs. Difficulties that might be addressed include restrictions imposed by licensing conditions, technical complexity of installation or support on remote equipment, or the need for protecting the security of data accessed by the application. These services can also overcome problems with supporting applications on PCs not owned by the institution, or those with incompatible hardware or operating systems, e.g. Apple Macintosh.

Virtual Private Networks

Two forms of VPN can be established. One form is where a secure connection is made, across a public network, between two LANs so that they may appear to have direct interconnectivity; a *LAN-to-LAN VPN*. With the second, a *remote VPN*, a mobile user establishes a remote connection from their PC to an organisational LAN in a way that makes their PC appear to be directly connected to that LAN. This section deals with the second of these.

To establish a remote VPN, the remote PC needs to have suitable client software installed and there needs to be a compatible server at the institution to receive the connection. The institution's Network Manager may configure the routing to control which IT facilities the remote VPN user has access to, and may provide that access with the same functionality as a PC directly connected to the LAN. Mobile users with particularly demanding requirements might benefit from using a remote VPN, but it should also be clear that security is an issue, and that any VPN server must be configured with great care.

An early VPN protocol that was widely used is *Point-to-Point Tunnelling Protocol (PPTP)*. Although this protocol is not an official standard, it is very widely available and supported. It is the only VPN protocol built into early versions of Windows (9x and NT), there is built in support for PPTP in Mac OS X 10.2 and later, and clients are readily available for Linux. Unfortunately, with this protocol, the data being carried across the connecting network (the internet) is not encrypted unless it is used in conjunction with an encryption protocol, such as *Microsoft Point-to-Point Encryption (MPPE)*; even then it still continues to be criticised for security flaws. Despite these concerns, its wide availability and relative ease of deployment might mean it is acceptable for a number of services. For these reasons, institutions implementing VPNs should consider offering support for this protocol but should limit the IT facilities that can be accessed when using this protocol to non-sensitive systems or data.

The *Layer Two Tunnelling Protocol (L2TP)* includes the features of PPTP but has the advantage that it is an internet standard³⁴ and also can be used on non-IP networks. Like PPTP, it does not, itself, include data encryption and would appear not to provide any benefit over PPTP when used on IP networks. However, for IP networks L2TP is now normally found being used in conjunction with IP Security (IPSec) to provide the necessary levels of encryption and security.

Before the agreement of the Internet standards, many vendors implemented secure VPNs based on IPSec but there was a need for proprietary extensions. For example, the Cisco VPN system is based on IPSec and has been adopted by a number of institutions, but it has the requirement that PCs need to have the Cisco VPN client installed. This solution is, however, quite common in industry and it is not clear that any benefit is to be gained by changing to a different protocol for those organisations with a large community of remote users. Institutions with only a small group of VPN users, and expecting that community to grow in size, might wish to consider adopting a non-proprietary solution.

The use of L2TP in conjunction with IPSec is an Internet standard³⁵ and implementations are built into later versions of Windows (2000, XP, 7) and in Mac OS X 10.3 and later. Using this combination of protocols provides encryption and data integrity and should be considered by those wishing to implement a secure VPN service. Having software built into the operating system means implementing and supporting the client is simplified, but there is a cost in that

34. <http://www.ietf.org/rfc/rfc3931.txt>

35. <http://www.ietf.org/rfc/rfc3193.txt>

configuring the institution's server is more complex. There is also a higher performance overhead for this protocol (than PPTP, for example), but this is a price paid for adopting a standards based solution and for the security that it provides.

Using IPSec, whether with proprietary extensions or with L2TP, through network equipment that is implementing Network Address Translation (NAT), requires additional configuration, as IPSec can fail in these circumstances. The NAT device might be an institutional firewall but it may also be a router at the remote location. The protocol to overcome this issue, NAT-Traversal (NAT-T), is an Internet standard³⁶ and is normally implemented in both client and server software, although it might need to be enabled. In addition, use of NAT-T will require any institutional firewall to have the relevant ports opened for the VPN server.

In recent years, SSL VPN systems have become more prevalent. One example of this is CISCO Anyconnect. Once more these systems require the installation of a vendor specific client on the remote PC in order to do the creation of the tunnel to the server. This type of VPN works by first creating an SSL Tunnel to the VPN server. When the SSL Tunnel is established, the client will then try to negotiate a UDP DTLS (Datagram Transport Layer Security) Tunnel. Once the DTLS Tunnel is established the traffic from the client is sent via this and the SSL Tunnel is used for control traffic. However if the DTLS Tunnel fails all traffic passes over the SSL Tunnel. An immediate benefit of this kind of VPN is that no special firewall rules need to be configured, as the only connection from the client machine is on port TCP/443 that is used by SSL encrypted web servers and will be open across the Internet from any ISP. A number of *clientless* VPN systems for remote users are now on the market; essentially these systems use a web browser, with secure socket layer (SSL) as the client. The protocols that can be supported are limited, as are the applications. For example, they might be limited to Windows Terminal Services, web based enterprise applications, applications like Microsoft Exchange, and terminal based systems. Smart tunnel applications can also be provided, via this method, to allow local applications access across the SSL VPN to remote services that are not available via a web browser interface. Whilst there is little or no configuration of the client machine required, configuring the server can be complex and the load imposed on the server can be very high. Currently, such options are quite costly.

Recommendations

- 9.1 Facilities for remote or mobile access should be carefully designed and implemented. All remote access should be properly authorised and all users of remote or mobile equipment should be given guidance in mandatory good practice.
- 9.2 Access to any sensitive information or service should only be permitted over secure, encrypted connections.
- 9.3 Institutions should consider implementing IMAP (or IMAPS) and authenticated SMTP services (with TLS encryption enabled) to facilitate secure remote access to email. This is important for users with devices such as smartphones over mobile 3G networks secure access their email. Institutions using Microsoft Exchange should also consider supporting ActiveSync and RPC over http/https for remote access.
- 9.4 Institutions implementing VPN facilities should consider L2TP over IPSec, or a SSL VPN solution.

36. <http://www.ietf.org/rfc/rfc3947.txt>

10 Cloud Computing

Cloud computing is quickly being realised as prospective infrastructure solution for major corporations and is rapidly being adopted in the educational sector. Virtualisation, lower costs and managed secure end-to-end services have been some of the key driving factors for this. Cloud Computing has mostly been synonymous to the Internet referring to its characteristics of large interconnected networks spread across various locations.

NIST define cloud computing as, “... a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”

Some of the key characteristics and features that this technology offers, are dynamism in the amount of resources available, high reliability and availability, devolved ownership hence lower costs, platform independence etc. So the technical aspects and the economics certainly make sense, but is this enough? This is a key question being asked mainly in the education sector due to varied requirement by each institution.

Cloud computing is quickly evolving with technology integration firms offering varied solutions under different umbrellas such as Software as a Service (SaaS) that hosts a complete application on equipment managed by the cloud service provider, Infrastructure as a Service (IaaS) that comprises of a significant section of the corporate network such as server infrastructure, mostly on virtualised platforms that provide a cost effective infrastructure among many another services out of which these are more relevant to the educational sector. These mainly relate to the architecture that is used to deliver the service. The most commonly known models used in cloud computing are:

- a. Public clouds: These are subscription based, provided by service providers, e.g. Amazon, Google, that have large infrastructures of integrated systems across various locations that can be hired based on requirements, and resources can be added or removed dynamically making them a useful option for small scale requirements, such as off-site storage etc.
- b. Community clouds: These can be deployed by multiple organisations that can share the resources, making it easier to implement policies and achieve other requirement, such as security compliance etc., easily. On the other hand, it could prove to be expensive, as the infrastructure would have to be owned and operated by these institutions.

One of the key aspects that need to be evaluated before procuring such a solution is security. This is primarily because the concept of a centralised infrastructure owned by a third party, which could potentially be based in any location across the world, may raise concerns with IT security professionals, regarding aspects of information security, data confidentiality etc.

The following are some recommendations on criteria that can be investigated when evaluating a cloud computing service:

- Does the cloud service provider implement data security controls, such as a perimeter firewalls, network communication security, i.e. between client and the cloud service provider?
- Does the cloud service provider implement an appropriate vulnerability assessment schedule against its systems and any applications it uses to ensure data protection and integrity?
- Does the service provider implement complementary security functions, such as Intrusion detection, maintain logs etc?
- For applications hosted on the cloud that accept credit card details, ensuring the service provider complies with the required standards, such as PCI-DSS etc?
- Clear contractual agreements to cover legal aspects that apply mainly to the education sector, such as Data protection Act, intellectual property etc?

In the education arena, some of the key services that can be seen being migrated to a cloud-based environment could be certain web based applications, such as HR systems etc. The above criteria can be used to ensure the platform being outsourced is secure and certain security controls within the application itself, such as ensuring SSL communication using trusted certificates etc., to ensure data integrity, which would provide an additional layer of security.

Educational institutions within the United Kingdom are connected to the Internet through a fast, reliable and low cost connection provided by JANET(UK) and/or a RNO (Regional Network Operator). It is over this connection that many organisations will want to procure cloud services. Whilst a high bandwidth connection is already provided to sustain the data transfers that are typically associated with cloud computing, it is the business continuity aspect of this model which needs to be considered. Many organisations will be procuring cloud computing for the very purpose of business

continuity. Therefore, what are the implications if a site loses its main JANET connection? Is there a resilient secondary connection? Is there a diverse fibre route if there are problems with unforeseen underground work?

It is important not to underestimate the bandwidth implications of adopting cloud based services, from simple migration of student email to Google or Live@edu, through to the full replication of filestores in the cloud. Student email which would have previously been handled in-house, will transverse the JANET connection. This is noticeable when students are emailing attachments to many colleagues and the traffic is coming in and out of the organisation over the JANET connection. Simple networking monitoring and graphic of these links is recommended to provide a trend analysis of the traffic. Replication of filestores will typically utilise a higher bandwidth, even if a first stage replication is started out of hours, at a weekend, or in a physically adjacent location. Some replication technology may require dark fibre infrastructure for fibre channel connectivity, others can operate over traditional IP technology.

A facet of cloud service procurement, in the event of a latency sensitive or high bandwidth services, should be the investigation of a JANET peering for the service provider to avoid *last mile* networking problems.

Recommendations

- 10.1 A clear evaluation needs to be carried out to have a clear understanding of the key business and technical drivers, such as scalability, cost, platform independence etc., before any application or services can be moved to a cloud based environment.
- 10.2 A clear security evaluation a criterion needs to be in place similar to the one suggested in this document to evaluate the platform an application is being outsourced to.
- 10.3 Institutions should consider a risk assessment of their connection to JANET. Is a resilient connection available with diverse fibre routes?
- 10.4 Institutions should consider evaluating the spare capacity required on JANET connections to engage in Cloud Computing activities, this may need to be in partnership with JANET(UK) and/or Regional Network Operators (RNOs).
- 10.5. Institutions should consider monitoring and graphing their Internet connections to JANET to establish the trends and requirements associated with cloud computing.

11 Wireless access and Security

Wireless networking is increasingly being used to provide access to network resources and onward connections to the JANET network. Although new protocols and standards for wireless communications continue to be developed, the current standards are now well established and form a sound basis on which to build a reliable service. Well structured wireless networks can be valuable assets, particularly when used to complement and extend the wired network, but as there remains much hyperbole about what can be achieved, institutions need to be aware of the true functionality and limitations of the technology. In order to aid the JANET community in wireless networking, JANET(UK) has established a Wireless Technology Development Area, which is steered by their Wireless Advisory Group (WAG), and have they have produced a number of *Factsheets*³⁷.

Good planning is important before a wireless network is deployed. It is important to know how many users will need support, where they will be given access, how they might move around, and their anticipated bandwidth requirement. The network will also need to be designed around the physical properties of the building as the fabric of the building, such as steel columns or foil-backed plasterboard, will affect the range and coverage of the network. Other equipment, such as security devices, might be using the same radio frequency, which is unlicensed, and wireless networks in adjacent properties might also interfere. The plan needs to recognise that *access points* (APs) in close proximity should not use the same channel in the allocated spectrum and that there are only four non-overlapping channels. In all but the simplest of installations, this will require a site survey to confirm signal strength, signal quality and potential data rate around the proposed location of an AP. As signal quality deteriorates, equipment is forced to select a lower data rate. Multiple APs might be required to get the necessary coverage.

Most current networks support both 802.11b (up to 11Mbit/s) and 802.11g (up to 54Mbit/s), which operate with a transmission frequency of 2.4GHz, and this is also standard on most current wireless enabled laptops. Implementations that use 802.11a, which operates up to 54Mbit/s in the 5GHz band are much less common. It is recommended to remove 802.11b access, as any 802.11b clients will cause an access point to service all of its clients at the slower data rate. As of October 2009 the Wifi Alliance approved the 802.11n amendment to the 802.11 Wireless LAN standards. The 802.11n amendments updated the previous wireless standards by adding multiple-input multiple-output (MIMO) and 40 MHz channels to the PHY (physical layer), and frame aggregation to the MAC layer meaning that you can get raw data rates of up to 300Mbps which means approximate useable data rates of up to 180Mbps.

When a number of concurrent users are connected via an access point, the maximum realistically achievable throughput is around half of the nominal bandwidth, shared amongst them. For example, an office of ten people connected via a shared 802.11g access point can expect, on average, no more than 2.5Mb each. A bandwidth-hungry application, such as downloading a large file, can significantly impact the network performance of other users of an access point. Streaming data, such as voice, music or video, might not perform well when a wireless network becomes moderately loaded.

In many situations, wireless networking cannot be seen as an alternative to a fully wired network; it would be installed to complement the wired network and extend its functionality. Wireless networks have different security characteristics from wired ones: each has its own threats and benefits. Controls should, therefore, be implemented where network traffic passes between wired and wireless media to prevent threats from one side undermining benefits on the other.

By using a shared transmission medium, the radio waves, in addition to the problems of bandwidth contention, there is the potential for eavesdropping, which makes security planning paramount.

A *service set identifier* (SSID) is a sequence of characters that identifies a wireless LAN. Access points that support multiple SSIDs and multiple VLANs should be used. This will permit different policies and functionality to be assigned to different groups of users e.g. staff, student or visitor. There is little point in using *hidden* SSIDs as a form of security, as any SSID that is not being broadcast can easily be discovered: it is better to treat the SSID as a community identifier. VLANs are employed to create a logical separation of the wireless network and the institution's LAN, and firewalling techniques should be used at the logical point of interconnection to support the access policies in force.

Encryption should be used for any sensitive information, including user IDs and passwords. It is common for an *open*, or insecure and unencrypted, network to be configured for use by visitors with access limited to publicly available services, such as the institution's website or library catalogue. This is also used as the starting point for authentication. For example, the first attempt at web access could be redirected to a page that requests authentication with user ID and password (web based redirection) before permitting access. Once authenticated, other services might be authorised, such as VPN or IMAP, but these should be limited as there is no encryption on the wireless connection. Clearly, the web page requesting authentication should be protected by https/SSL, otherwise passwords will be passed in clear text. Open networks can quickly run out of IP address space as a lot of clients will automatically join them in preference to other networks. This can also cause problems with load on the captive portal machine, for example, if a client joins the network without the users knowledge and it tries to synchronise its data, then it will be constantly

37. <http://www.ja.net/services/publications/factsheets.html>

getting re-directed to the captive portal page, which can cause the captive portal server to run out of resources whilst trying to constantly provide the portal page.

Security is better served if encryption is used on the wireless link. The early encryption protocol, known as WEP, is widely acknowledged as not providing sufficiently robust security. It might be considered adequate for use on a home network but it is not recommended for use on larger, institutional networks. *WiFi Protected Access* (WPA) using the *Temporal Key Integrity Protocol* (TKIP) algorithm should also be avoided as it has been proven to be easily crackable.

For secure connections, institutions should employ *WiFi Protected Access 2* (WPA2). WPA(2) with a *pre-shared key* (PSK), or password, should either be avoided or limited to very small communities as management of the key can become problematic. Rather, institutions are recommended to deploy WPA2 with 802.1X using *extensible authentication protocol* (EAP) and a RADIUS authentication server. This is often referred to as WPA2 Enterprise and is, essentially, a subset of 802.11i, the enhanced security and authentication standard. Reference may be made to the JANET(UK) Technical Sheet³⁸ that provides an introduction to the 802.1X standard. WPA2 with *Temporal Key Integrity Protocol* (TKIP) should not be used, and is no longer ratified by the Wifi Alliance, instead WPA2, with Advanced Encryption Standard (AES), should be deployed.

It is possible to implement, simultaneously, both web based redirect and WPA2/802.1X security, for different levels of service or functionality, by using two different SSIDs.

Institutions that have deployed a wireless network infrastructure should consider joining eduroam³⁹ as a *visited* organisation. eduroam allows mobile users visiting another participating institution to use the user ID and password provided by their home organisation to gain network access. eduroam facilitates a range of network access scenarios, ranging from casual visits and meetings to large conferences and classroom sharing. To operate as a *home* organisation the institution must implement a RADIUS service to authenticate their own users and deploy an *Organisational RADIUS Proxy Server* (ORPS), which may be the same server(s), to link to the *JANET National Radius Proxy Servers* (NRPS). To operate as a *visited* organisation, the institution needs to deploy the ORPS, configure the RADIUS service to implement the authentication mechanism for visitors, configure the SSID *eduroam* on appropriate wireless access points, and permit forwarding of certain IP protocols, for visitors, on institutional firewalls between the wireless network and JANET.

A common problem encountered by many institutions is the connection, by a user, of a wireless access point (or base station) to a network point on the wired network or in a student's bedroom. This is usually done in ignorance of the serious impact that such a connection might have. Such equipment will not be implemented to the institution's standards and will rarely have adequate security. This is likely to allow untraceable access to the internal wired network by anyone in the area. The default configuration often includes operation as a DHCP server, which would cause other users' network services to fail, and it might also advertise itself as the default route for IP traffic, causing serious disruption to traffic on the wired network. The wireless channel being used is also likely to overlap with any being used by the institution and, thereby, degrade the performance of the institution's service. Institutional policies should be in place to permit the identification of any rogue devices, and to allow problem devices to be disconnected or shutdown.

Recommendations

- 11.1 Any deployment of wireless networking must be carefully planned and the expectations of users carefully managed. Currently, wireless networks should be promoted as complementing wired networking rather than an alternative.
- 11.2 Access points that support multiple SSIDs and VLANs should be used, with SSIDs being used as community identifiers and VLANs employed to permit differentiation of services.
- 11.3 Removal of 802.11b support from access points should be considered.
- 11.4 Encryption should be mandated for all sensitive information, including user IDs and password, and institutions should employ WPA2 for secure connections.
- 11.5 For secure access, institutions should deploy WPA2 with 802.1X and a RADIUS authentication server – WPA2 Enterprise.
- 11.6 Institutions should consider joining eduroam.
- 11.7 Procedures should be in place for the identification of any rogue wireless devices. Organisations should consider making a decision on the use of active mitigation technology, including a risk assessment of which access points may be mitigated.

38. <http://www.ja.net/documents/publications/technical-guides/8021x-tg-web.pdf>

39. <http://www.ja.net/services/authentication-and-authorisation/janet-roaming.html>

12 Email

Much useful information is available from the JANET email web page *JANET mail services*⁴⁰, which is recommended reading.

Names and addresses

Email addresses have the form local-part@domain-name. Organisations have almost complete control over the local-part, and very limited control of the domain-name that they use. Most UK academic institutions will have domain names under the ac.uk domain – for example, Camford College might have the name camford.ac.uk. Many organisations choose to use this *master domain* for all their mail addresses, which has the advantage that any internal staff and server transfers are not visible externally, so there is no need to advertise address changes outside the institution. Others choose to insert subdomains; in the case of Camford, this might result in domains law.camford.ac.uk and ee.camford.ac.uk for the Law and Electrical Engineering departments, respectively. However, with this option, if departments merge, split up or are renamed email addresses that include the old subdomain name become anachronistic. Whichever choice is made, there is a strong preference for institutions to have a single mail gateway service (resiliently based on multiple machines) handling all incoming mail to intercept any spam or other malware and to act as a *sorting office*, distributing mail to the various mail systems within the institution. A similar approach should also be adopted for outgoing email.

Rules for the format of the local part are less restrictive. Almost any string of letters and numbers can be used together with some punctuation characters but, in practice, it is best to avoid all punctuation, except the full stop (.) and hyphen (-) characters. Mail addresses are not usually case sensitive. Institutions should choose a format for mail addresses and then stick to it, as changing addresses leads to problems with remote contacts and mailing list subscriptions.

One of the simplest schemes to administer is to use each person's user identifier as the local part of their mail address. This usually results in short addresses which are quick to type but which may not be very informative to outsiders. It also exposes user identifiers to outside view, and some sites prefer to hide this information.

A very common scheme uses personal names in the form *Initial.Initial.Surname* or *Firstname.Lastname*. This can be considered as one of the more user friendly options, though some people dislike the implied informality of the second of these. There can also be a problem with name clashes, especially in larger organisations. This scheme hides user identifiers from public view, though it is common to *accept* mail that is addressed to the user identifier as well as that addressed to the official address.

Many sites also use *role* email names for various institutional officers – for example: Vice-Chancellor@domain and Admissions@domain. This method is strongly encouraged to assist the institution, for example, in meeting its obligations under the Freedom of Information Act in the event that the role holder is absent, as well as to promote internal efficiency by enabling messages to be routed to the most appropriate location. Mail to role addresses may be directed to a dedicated mailbox, to an internal mailing list, or to the individual mailbox(es) of those who perform the role.

Standard addresses

There are several addresses that every mail domain should support and these are described in RFC 2142⁴¹. Mail to the addresses postmaster@domain and abuse@domain should be checked frequently (at least daily during the working week) by a mail system administrator. The address webmaster@domain should also be regularly checked. Some sites direct mail to these addresses to the administrator's own mailbox; others have them set up as mailing lists, and others have *role* mail accounts that are handled by the administrator on duty at the time.

40. <http://www.ja.net/services/mail/index.html>

41. <http://www.ietf.org/rfc/rfc2142.txt>

Mail user policy

The legal and organisational aspects of email are getting a lot of attention. Sites are recommended to have an email policy covering issues such as:

- Appropriate language;
- Avoiding defamation;
- Harassment by email;
- Appropriate use of attachments;
- Appropriate use of distribution lists;
- Disclaimers;
- Avoiding the creation of unauthorised contracts.

The issues are much the same as with paper mail, but the ease of sending email and the perceived informality of the medium have led to problems in some organisations. The following is suggested as guidance to the end user on the use of bulk email (the term bulk email is used in preference to more emotive words such as *spam*); also see *RFC2635, DON'T SPEW – A Set of Guidelines for Mass Unsolicited Mailings and Postings (spam)*⁴².

Guidelines on the use of bulk email

Bulk email is essentially identical email sent to small or large groups of people, irrespective of category, reason, or whether this is done by a single mailing or repeated individual or group mailings, or any other means. Although the use of bulk email can occasionally be in the interests of particular communities within an institution, it can nevertheless present real problems to those who receive it and those who deliver it, but rarely to those who send it.

Unsolicited email is email that the recipient did not request, either explicitly or implicitly (e.g. by joining a particular email distribution list), irrespective of whether it is welcome or not.

Principles for the acceptability of bulk email

It is not acceptable to do anything that is likely to provoke external retaliation against the institution, such as being put on a blocking list, thus risking the continued operation of the institution's core business. Also, unsolicited email sent in bulk outside the institution is likely to be in breach of the JANET acceptable use policies. It is, therefore, normally unacceptable to send bulk email outside the institution's local network. (See below for *Marketing*). Even within the institution, it is not normally acceptable to send bulk unsolicited email unless there is a reason to suppose that a substantial proportion of the recipients either will or should be interested in the email's contents. For example, advising a group of students that a lecture has been postponed.

Although an individual may have assented to certain bulk email implicitly by being on some particular email distribution list, or indeed by virtue of being a student or staff member of an institution, caution is needed in any such assumption. Agreement to receive email that is inappropriate to a distribution list is never implied, and the mere appearance of an individual in some staff or student directory does not imply assent either.

It is not acceptable to send email in such quantity that it would overwhelm the underlying network or email servers, whether the quantity is on account of the number of ultimate targets, the volume of data sent, the volume of data consequentially stored, or anything else. It is the sender's responsibility to verify the logistical acceptability of the proposed mailing. Particular care should be taken where the message includes an attachment – it is normally preferable to give a web link instead. The web, other electronic bulletin boards and discussion forums are usually more appropriate than bulk email and should be preferred wherever possible and reasonable.

Bulk unsolicited email is frequently counterproductive, both to the immediate purpose and to other, perhaps more important, endeavours. It can generate considerable annoyance, is often discarded without being read when it can be recognised as unsolicited, and can cause other, more important email to be lost or overlooked when a large quantity of email is received. Those who would send bulk email often overestimate the importance of their endeavour and underestimate the negative impact of the method. Any sending of bulk email should be preceded by the utmost caution and reluctance.

Rules and guidelines that apply to ordinary individual or personal email apply at least as strongly to bulk email, if not more so, on account of the size and generality of the audience.

42. <http://www.ietf.org/rfc/rfc2635.txt>

Procedural guidelines for sending bulk email

If some particular instance of bulk email is deemed to be appropriate, having considered the above guidelines, the email itself should be constructed and sent following the guidelines below:

- The email should be legible on the most basic of equipment and should not require the recipients to have particular software or hardware platforms that they might not be reasonably expected to have. These days most bulk emails are sent out in html format, but most MUAs (Message User Agents or email clients) will create their own text/plain part for every text/html part of a message. The message should not contain any encoded material, nor should it include any part in any proprietary format. Moreover, it should not include material replicated in different formats (e.g. plain text and HTML).
- The message should be short, perhaps the equivalent of an A4 page at most, and preferably not more than a screenful.
- If the object is to draw the reader's attention to bulkier material or to material inherently in other formats, then appropriate references should be included, e.g. as URLs, to enable the reader to locate it.
- The message as a whole should make it plain that it is a distributed message and should make it clear who the target recipients are, if this is not already clear from the *To:* header field.
- The message must indicate who has sent it and under what authority. Clear instructions must be included for anyone who has received it in error to remove themselves from the list.
- The message headers should be such as to prevent any replies (to the sender) from accidentally being sent to the whole constituency. Expert advice should be sought as to specific methods for doing this.
- Care needs to be taken to avoid exposing the full list of recipients, which could subsequently be misused. If possible, a list server should be used, if not then the list should be entered in the BCC field and the message sent formally to the sender.
- If the constituency is large, it may well be appropriate to take special steps to minimise the logistical impact of such a bulk emailing. Expert advice should be sought.

Relaying policy

Mail systems must be set up to restrict relaying, which is the process of receiving email from *outside* the institution and passing it on to another external organisation. Mail should only be transferred from *outside to inside* or vice versa. Systems that allow *outside to outside* relaying are soon discovered by unscrupulous people who then use them to relay junk mail to millions of recipients. The institution might wish to provide a mail relay for their own mobile users, in which case incoming connections from those users when they are *outside* the institution must be authenticated.

Enforcing an appropriate relaying policy is an essential defence against misuse of resources. More information can be found at the *Fight Spam on the Internet*⁴³ site and *RFC2505, Anti-Spam Recommendations for SMTP MTAs*⁴⁴ also provides advice on good practice.

Blocking email

There are sites that coordinate anti-spam measures by identifying domains that regularly send (or relay) spam email. They inform domains that email from their organisation is being added to a blocking list and if the organisation responds appropriately to prevent the misuse, they will be removed from the block list. An institution can configure their email servers to use such lists to stop email coming from those block listed domains and so provides a relatively easy, but basic, mechanism to control some of the spam email. This does, however, run the risk of rejecting genuine mail from those domains or from domains that might have been incorrectly listed. Other effective measures include open source tools or outsourced services that evaluate the content (but see *Content scanning* below).

43. <http://spam.abuse.net>

44. <http://www.ietf.org/rfc/rfc2505.txt>

There are several Realtime Blackhole Lists (RBLs), the spamhaus list being one of the most used. JANET(UK) has its own copies of the following leading DNS blocklists and whitelists⁴⁵:

- Spamhaus Zen lists; SWL and DWL
- MAPS RBL+
- SURBLs
- URIBLs
- DNSWL.org

RBLs should only be used as a part of the decision as to accept or deny the email, not as a blanket accept or deny rule.

In a technique known as *greylisting*, an institution's mail server will temporarily reject mail from an unrecognised sender with a *try again later* message and make a record of that rejection. A legitimate and properly implemented sending server will try again after a short period, this will be recognised by the institution's server and the mail will now be accepted. This process is effective because mail sending servers responsible for abuse usually do not retry a failed delivery, so a significant proportion of the spam will be rejected early. There is a cost, as legitimate email might also be subject to a delay, and could cause problems for users expecting email, so this technique is usually operated in conjunction with a *whitelist* system, where the incoming mail is checked against a list of known, trusted servers for acceptance, prior to being passed to the greylisting process. Greylisting has the advantage over scanning that the actual message is not examined, although many institutions also use spam filters to identify or delete any spam, which does get through.

Content scanning

An increasing number of sites subject email to content scanning as it passes through their mail systems. Most do this to remove viruses and spam, but some also use it to enforce their appropriate use policies. This can be a very emotive area. Whilst automated filtering of email is generally accepted, if procedures could result in filtered email being viewed by a member of staff, this might be subject to legal constraints – the *Data Protection Act*, *Human Rights Act*, and the *Regulation of Investigatory Powers Act* (see section 4, *Legal and responsible use*). Any site considering content scanning, particularly for reasons other than virus or spam protection, should make sure that it has considered the appropriate legislation when establishing its procedures and has policies in place to justify any manual action, should it arise. These policies should be publicised and users made properly aware of them. Staff with responsibility for monitoring should be made clearly aware of the ethical constraints of the task and that any abuse of power will be regarded as gross misconduct.

Use of email (and SMS) for marketing purposes

Since December 2003, with the implementation of the *Privacy and Electronic Communications Directive*⁴⁶, it is a criminal offence to send unsolicited messages to individuals, unless there is a pre-existing relationship on a related matter. If, despite the guidelines above, your institution still wishes to use electronic media (email, SMS, telephone) for marketing purposes, then you should seek legal advice to ensure that you comply fully with the law. What follows is Good Practice; it is not legal advice.

Even where the communication is legal, you are strongly advised to include in the message your basis for believing that you have the recipient's explicit permission. This might show the date and circumstances when the permission was given. You should include a link to a relevant Privacy policy and you must provide a means to unsubscribe from the list with immediate effect. You should maintain a record showing precisely the evidence for consent, since you may be required by the Information Commissioner to produce this in response to a complaint. You are strongly advised not to trust so called *opt in* lists provided by third parties. Finally, authority to send electronic marketing should be restricted to a strictly limited number of people who are trained in the law and each message should be approved by a Senior Officer of your institution.

In a clarifying *Good Practice Note on Electronic Mail Marketing*⁴⁷, the Information Commissioner makes it clear that the law applies equally to individuals in companies as it does to private individuals.

45. <http://www.ja.net/services/mail/janet-rbl/janet-dnsbls.html>

46. <http://www.opsi.gov.uk/si/si2003/20032426.htm>

47. http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/electronic_mail_marketing_12_06.pdf

Mailing lists

Mailing lists are very useful for discussions and for distributing information. Almost all mail domains support some lists. Issues to consider include:

- List maintenance (how do people get on and off lists);
- Size of acceptable messages;
- Should lists accept mail from outside the organisation.

In the case of the last of the above, if a mailing list is able to accept mail originating anywhere outside the organisation, it can be used as a spam amplifier. The same is true for external mailing lists that do not restrict who may send mail to the list.

Denial of service

Email is now *mission critical* to many organisations and the loss of the service can have a serious effect. Protecting against deliberate denial of service attacks is difficult, but a few simple techniques can help a lot:

- Have separate servers handling internal mail and mail coming in from the internet. If the internet facing machine is attacked, the internal mail has more chance of continuing to work.
- Block the SMTP port (TCP 25) at the site boundary to and from all machines except the known internet facing mail servers. This prevents external crackers from attacking weak or poorly configured internal mail services, and also ensures that all outgoing mail can be logged and sanity checked on a well managed machine. Consider blocking SMTP transactions other than to recognised mail servers at network boundaries within the organisation: this can contain the spread of spam and other mail borne malware.
- Apply all known security measures and patches to the internet facing mail servers and ensure that they are kept up to date.
- Consider applying a limit on the rate at which an individual user or site can send mail through your system or imposing a short delay when accepting incoming connections. This helps to defend against attacks and reduces the impact of inadvertent forwarding loops.

Phishing/Spear Phishing

Phishing is the act of trying to gain sensitive information (usernames, passwords, credit card numbers, etc.) from users by masquerading as an official website/company. Phishing is usually carried out via email with specially crafted messages designed to seem as though they are from an official website/company. Spear phishing is a variant of this, whereby users are specially targeted with messages seemingly from organisations that they are a member of, their work IT Services, etc. These messages usually get users to reply with their details or send the user to a fake website designed to look like the real one, they will often harvest the details and then put up a message saying login failed and then redirect the user to the official website.

Filters on outgoing mail servers can be applied to trap anyone replying to these messages with details, which match the institutions username and password policy. This will only catch the *reply to this email* type phishing scam and not the fake website type ones; these would need an institution to have a secure web gateway type device which denies access to certain types of web content.

Management and monitoring

Mail systems need active management and monitoring. Queues can build up for a number of reasons, and if the cause is not found in good time, the service can be seriously degraded. It should be the rule that each mail system is checked by a competent administrator at least once per day.

Data protection

Mail systems are considered to be databases under the *Data Protection Act*, as they contain information about people who send and receive mail. Further, individual collections of messages fall within the scope of the Act. In this context, you should consider very carefully whether it is necessary to send *out of office* messages outside the organisation.

Disclaimers

Unless your institution routinely includes disclaimers in all its correspondence on all media, there is no special reason to do so specifically for email (the reason that this might be done on FAX and Telex is that a number can be misdialled). There is no legal requirement or basis for extending the practice to email, although it does no harm if the disclaimer is kept very brief. Any disclaimer might need to be omitted from an email that is stating the official position of the institution.

Recommendations

- 12.1 Institutions should ensure that appropriate policies and procedures are in place to meet their legal obligations and covering issues such as avoiding defamation, disclaimers, avoiding creation of unauthorised contracts, and bulk emailing.
- 12.2 For every mail domain supported by an institution, mail to postmaster@domain and abuse@domain should be checked frequently by a mail system administrator.
- 12.3 Mail systems must be set up so as to prevent relaying from outside the domain to outside the domain except when the incoming connection has been properly authenticated as coming from an authorised user.
- 12.4 Institutions should make use of reputable block listing sites for configuring their mail systems to minimise the amount of spam delivered to the institution's users. Consideration should also be given to implementing other techniques such as greylisting or content filtering, including the possibility of outsourcing these functions.
- 12.5 Institutions should take account of relevant legislation, including the *Data Protection Act* and the *Regulation of Investigatory Powers Act*, particularly regarding scanning the content of emails, whether for virus-protection or for other reasons, and should make their users aware of the conditions when their incoming and outgoing emails might be monitored.
- 12.6 Institutions should be aware that email is likely to be mission critical and should take appropriate measures to protect the facility from being completely, or partially, disabled through malicious or accidental action. Institutions should consider installing an emergency back-up facility (e.g. an additional connection to the Regional Network or an ADSL line) to protect against external network failure.

13 The CERT, security policy and practice

As part of Information security functions, institutions can choose to form teams that respond specifically to Information security related incidents. Depending on the type of functions defined for the team, they are referred to as CSIRT (Computer Security Incident Response Team) or CERT (Computer Emergency Response Team). There is not a significant difference in the functions between these teams, as their primary task is to deal with IT security related incidents. In the European Union, security response teams based within a university are generally referred to as CERT and the service providers providing this service refer their team to as CSIRT's. In terms of functions, both are expected to have similarities, such as incident response procedure and readiness to respond quickly and efficiently to security incidents is expected.

JANET(UK) provides the JANET CSIRT⁴⁸ incident response team as one of the core JANET services. JANET CSIRT can better help an institution when there is a primary contact for the team to liaise with; institutions are asked to ensure that they have adequate liaison capacity and to consider whether this should be provided by establishing their own CERT functions (see *Effective Incident Response*⁴⁹).

At present, the JANET CSIRT team is required to respond to fault calls within one hour during normal hours Monday to Friday, with reduced services during evenings and the weekends. The coverage difficulty lies with the hours of availability of local staff. The overall network is less heavily loaded at weekends, so many research groups with heavy data transfer requirements work then. There is also a growing requirement to use network facilities in support of teaching during evenings and at weekends. Additionally, students may require the network to be available for studying or project work from their halls of residence at times other than normal office hours. The availability of the network is vital to these activities. Institutions are recommended to review their requirements in this area in order to ensure that appropriate levels of availability are maintained.

JANET CSIRT coordinates, publishes, and updates information relating to computer and network security on a regular basis, this information is not duplicated here. An institution's CERT staff should inform themselves of the issues relating to computer and network security and should keep themselves updated. Information is available via the JANET CSIRT web homepage⁵⁰.

Implementation of security measures to protect an institution from malicious action occurring via JANET will not fully protect the institution's network if there are other poorly protected means of access. Although an institution's IT Services may provide adequately secure remote access, there may be other departments within the institution with their own access arrangements, e.g. remote login software, such as GoToMyPC or a backup modem connection to their servers, and because of lapse in security best practices, there could easily be a breach in infrastructure security.

These shortcomings could be addressed by having an incident response procedure so an effective path of handling and resolution is available for an IT member of staff in the event of a security event. In general, a security incident is an event contravening with a set Information security policy, Acceptable Usage policy etc. which usually encompasses malware infections, phishing attacks, attacks against the infrastructure e.g. Denial of Service among many others.

A more formal definition would be *"Any real or suspected adverse event in relation to the security of computer of computer networks. Example of such are: intrusion of computer systems via the network occurrence of computer viruses probes of vulnerabilities via the network to a range of computer systems"* (TERENA).

An incident response plan is a key function for a CERT or a security incident respondent. Educational institutions are generally faced with resource limitations and through this process they can achieve better cost-savings and preparedness in the event of security incident. As part of this plan it is crucial that clear definitions are in place as to what would be classified as a security incident. This is necessary, as educational establishments are technologically diverse, and not having a proper scope of a CERT could result in many ordinary events being reported as security incidents.

A distinction is necessary e.g.:

A managed server that has been successfully compromised through an external attack would be classified as a security incident, whereas a user who has only accessed a phishing site/email, but not necessarily replied with any information, such as login credentials etc., would not be classified as a security incident.

Note: that the former event could become a security incident if that user's email account is subsequently abused to send spam email.

48. <http://www.ja.net/documents/services/connections/sponsored-proxy-guidelines-june-07.pdf>

49. <http://www.ja.net/documents/publications/technical-guides/gn-incident-response.pdf>

50. <http://www.ja.net/services/csirt/>

An incident response plan should contain the following:

1. Appropriate tools to assess impact and resources required to handle the identified/reported incident:

This would include assessing the overall business impact, relevant section of people informed of the event, resource required, and have estimated time required to resolve i.e. when services will be restored to normal.

2. Response procedure:

This would be a series of tasks that would cover various areas of information gathering to technical analysis and steps to resolve the issue. To begin with, an immediate course of action must be determined, e.g. isolate the host, gather/sample further traffic etc. If the incident is directly related to a legal investigation the host must be preserved (forensically if required) as evidence.

Further technical analysis can include steps to determine the source of the attack/compromise. If the compromised server was found hosting and distributing malware content then determining if any other hosts within the network were involved. Is it possible to contain and repair the service or does it have to be reinstalled completely (which may require liaising with the backup team to obtain last known good configuration.). A plan is required to ensure that all changes have been tested satisfactorily before the service is brought back online. An audit trail of all activity carried out as part of the investigation should be maintained and should document any decisions by management to proceed when such authority is required.

3. Review response/resolution:

This step should involve a review of the steps that were taken to handle and resolve the security incident. This will help identify areas that could have been addressed in a better way, any lessons learnt, if it was possible to have identified this before the incident occurred and if so, recommend necessary changes to the relevant department.

4. Final Report:

The last stage of an incident response must be a final report that would briefly highlight details about the incident, how it was addressed, i.e. from the logs maintained during the process, and any lessons learnt and room for improvements identified (if any).

This approach should assist in handling incidents in a sustainable way with suitable resource allocation and also provide justification to management for the need for the service, as well as possible financial savings in dealing with future incidents.

By implementing preventative security measures, an institution protects not only itself against malicious and accidental actions, but also other institutions due to the close collaborative nature of the sector.

Recommendations

- 13.1 Institutions should ensure that they appoint a Computer Emergency Response Team (CERT) contact that has adequate internal powers to take actions recommended by JANET CSIRT.
- 13.2 An institution's CERT staff should inform themselves of the issues relating to computer and network security and should keep themselves updated.
- 13.3 An institution should take what preventative action is possible to protect their institution against those with malicious intent.
- 13.4 Institutions should develop a suitable Incident response plan with clear definitions of scope of security incidents that will be handled. The plan can be based on the recommendations made in this document.
- 13.5 Institutions should keep JANET CSIRT informed of security incidents involving their JANET traffic.
- 13.6 Institutions should ensure that their network operations provide adequate local support for CERT activities and out of hours cover sufficient to support their main operations.

14 Business and Community Engagement

Increasingly, institutions are collaborating with other organisations, providing them with on-campus accommodation, services and access to the internet using their campus network; they might create a spin-off company to exploit some developments; or might host web sites for local charitable organisations.

Strategy

Institutional strategy, may highlight Business and Community Engagement with specific commitments, for example, to have unrivalled links with *industry, business and the professions*.

These words seek to capture the wide range of partnerships and engagements that reach far beyond the classic *industry-academia* link. So, for many, this activity spans engagement across many sectors, including public and statutory groups, quangos, NGOs, voluntary and charitable groups.

The benefits to engaging in these activities are varied, but include:

- Collaborative research projects;
- Funding and co-funding of research;
- Facilities, subjects and testing environments for research;
- Graduate recruitment;
- Sandwich year placements and other work experience opportunities;
- Input to curriculum development and enhancements;
- Expert guest lectures on teaching programmes;
- Co-supervision of student projects;
- Recruitment to taught masters and PhD programmes;
- Additional coaching resources available to students;
- Shared capital and operational costs for research facilities;
- Development of spin-out initiatives in sports related areas;
- Joint engagement with manufacturers and suppliers;
- Substantial reputational advantages.

In turn, the JISC Business and Community Engagement Advisory group provides a complete definition of BCE.

“Business and Community Engagement (BCE) is the strategic management, by higher and further education organisations, of relationships with external partners and clients, and of the associated knowledge exchange and workforce development services. The objective is to deliver benefits to the economy and society, resulting in a more highly skilled workforce, a more efficient, dynamic and sustainable economy and a more cohesive, knowledge-enabled society.”

“The scope of engagement includes the commercial sector, the public sector (including charities and trusts) the cultural landscape and the social and civic arena. All organisations undertake BCE across this scope of engagement, but the exact mix and the resulting services deployed depend on organisational strategies.”

JANET Network Context

At the time of writing, in December 2010, JANET(UK) are responding to changes within the sector by reviewing their policies to provide support for this activity. Previous support was provided in the form of Sponsored and Proxy connections⁵¹. With the increasing important of this area of activity, some improvements were identified, and JANET(UK) undertook a comprehensive consultation exercise to best address the needs of the community.

For the latest information on Business and Community Engagement, it is recommended that one visits the JANET(UK) website. During the period of change, any request or uncertainty can be addressed by contacting the JANET Service Desk, who will be happy to advise.

51. <http://www.ja.net/documents/publications/policy/connection-policy.pdf>

For all connections, the host must ensure that users within the third party organisations are aware of, and abide by, the JANET Acceptable Use and security policies, and must have mechanisms in place to deal with failure of the third party to abide by these policies.

There are two important aspects of providing such a service: retaining the status of the JANET network as a private network entity, and complying with the regulatory implications of state aid and competition law.

If the external organisation is one which involves children, e.g. a school or youth club, the host institution should consider both its legal position and possible adverse consequences that might result if the children were able to access unsuitable material, in particular pornography, on the internet.

Local Context

Support for the connections has to be provided by the host institution. The resource that is needed to provide this support should not be underestimated and an appropriate charge should be made internally.

When providing the connections, the local organisation should be clear about the service being provided at what level and any restrictions.

Electronic mail services will certainly be a requirement of many connections, although this is not a service that always needs to be provided by the host in the era of managed and cloud services. Email servers are notoriously liable to attack by hackers and spammers and, therefore, need to be configured carefully and securely. If email services are provided by the host for the external organisation:

- The host institution will probably wish to ensure that the email addresses of the connected organisation's users do not imply that they are members of the host institution;
- The host institution should be aware that their email host could be added to an email blocking-list if the connected organisation engages in spamming;
- The host institution is aware of the email etiquette expected when using the JANET network.

Recommendations

- 14.1 Institutions should carefully consider the legal implications and publicity issues of possible abuse of a network connection by an external organisation.
- 14.2 Institutions should keep a watching brief of Business and Community Engagement activities within JANET(UK) and the sector.
- 14.3 Institutions should consider making realistic charges, covering both direct costs (including any network charges) and indirect costs (including support), to the connected organisation.
- 14.4 Institutions should ensure that their own bandwidth requirements have first call on the access bandwidth allocation so that external organisations are not subsidised at the expense of the host institution.
- 14.5 Institutions should consider what network-related services (such as domain name allocation, nameservers, electronic mail, etc) that they are prepared to provide to the connected organisation.

15 References in order

1. Internet Engineering Taskforce, *RFC 1173: Responsibilities of Host and Network Managers*, <http://www.ietf.org/rfc/rfc1173.txt>
2. and 5. UCISA, *Information Security Toolkit*, <http://www.ucisa.ac.uk/ist>
3. University of Cambridge, *Use and Misuse of Computing Facilities*, <http://www.cam.ac.uk/cs/iss/rules/guidelines.html>
4. UCISA, *Model Regulations for use of Institutional IT Facilities and Systems*, <http://www.ucisa.ac.uk/publications/modelregs.aspx>
6. Internet Engineering Taskforce, *RFC2196, The Site Security Handbook*, <http://www.ietf.org/rfc/rfc2196.txt>
7. JANET(UK), *Terms for the Provision of the JANET Service*, <http://www.ja.net/documents/publications/policy/service-terms.pdf>
8. Office of Public Sector Information, *Statutory Instrument 2002 No. 2013, The Electronic Commerce (EC Directive) Regulations 2002*, <http://www.opsi.gov.uk/si/si2002/20022013.htm>
9. JISC Legal Information Service, *Legal Guidance for ICT Use in Education, Research and External Engagement: RIPA (Part 2) New Codes of Practice (22/01/2010)*, <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/tabid/243/ID/1235/RIPA-Part-2-New-Codes-of-Practice-22012010.aspx>
10. Information Commissioner's Office, *Data Protection: The Employment Practices Code*, http://www.ico.gov.uk/upload/documents/library/data_protection/detailed_specialist_guides/employment_practices_code.pdf
11. JISC Legal Information Service/Higher Education Information Directors Scotland (HEIDS), *Inappropriate use of computers – the technical investigation process*, <http://www.jisclegal.ac.uk/publications/Inappropriateuse.htm>
12. JISC Legal Information Service, *Cybercrime*, <http://www.jisclegal.ac.uk/cybercrime/cybercrime.htm>
13. JISC Legal Information Service, *Legal Risks and Liabilities for IT Services in Further and Higher Education*, <http://www.jisclegal.ac.uk/publications/legalRisks.htm>
14. ACAS, *Advice leaflet – Internet and e-mail policies*, <http://www.acas.org.uk/index.aspx?articleid=808>
15. JISC Legal Information Service, *Intellectual Property Rights Overview*, <http://www.jisclegal.ac.uk/ipr/IntellectualProperty.htm>
16. JISC Legal Information Service, *Data Protection*, <http://www.jisclegal.ac.uk/dataprotection/dataprotection.htm>
17. JISC, *JISC Data Protection Code of Practice for the HE and FE Sectors*, http://www.jisc.ac.uk/publications/generalpublications/2001/pub_dpacop_0101.aspx
18. Information Commissioner's Office, *Information Commissioner's Office home page*, <http://www.ico.gov.uk>
19. Scottish Information Commissioner, *Scottish Information Commissioner's office home page*, <http://www.itspublicknowledge.info>
20. and 29. JANET(UK), *User authentication factsheet*, <http://www.ja.net/documents/publications/factsheets/041-user-authentication.pdf>
21. JISC Legal Information Service, *Internet Service Provider Liability*, <http://www.jisclegal.ac.uk/ispliability/ispliability.htm>
22. JISC, *Improving Network Reliability: Senior Management Briefing Paper 11*, <http://www.jisc.ac.uk/media/documents/publications/smbp11improvingnetwork.rtf>
23. Information Commissioner's Office, *Employment Practices Data Protection Code: Part 3 - monitoring at work*, http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/coi_html/english/employment_practices_code/part_3-monitoring_at_work_2.html
24. and 31. JANET(UK), *Logfiles – technical guide*, <http://www.ja.net/documents/publications/technical-guides/logfiles.pdf>
25. JANET(UK), *Investigating Copyright Complaints factsheet*, <http://www.ja.net/documents/publications/factsheets/077-investigating-copyright-complaints.pdf>

26. JISC Legal Information Service, *Legal Guidance for ICT Use in Education, Research and External Engagement: The Digital Economy Act 2010: Implications for UK Colleges and Universities (23/07/2010)*, <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/tabid/243/ID/1603/The-Digital-Economy-Act-2010-Implications-for-UK-Colleges-and-Universities-23072010.aspx>
27. JISC Legal Information Service, *Legal Guidance for ICT Use in Education, Research and External Engagement: Intellectual Property Rights*, <http://www.jisclegal.ac.uk/ipr/IntellectualProperty.htm>
28. Internet Engineering Taskforce, *Netiquette Guidelines*, <http://www.ietf.org/rfc/rfc1855.txt>
30. JANET(UK), *Using passwords factsheet*, <http://www.ja.net/documents/publications/factsheets/026-using-passwords.pdf>
32. JANET(UK), *JANET NTP Service*, <http://www.ja.net/services/ntp/>
33. and 39. JANET(UK), *JANET Roaming*, <http://www.ja.net/services/authentication-and-authorisation/janet-roaming.html>
34. Internet Engineering Taskforce, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*, <http://www.ietf.org/rfc/rfc3931.txt>
35. Internet Engineering Taskforce, *Securing L2TP using IPsec*, <http://www.ietf.org/rfc/rfc3193.txt>
36. Internet Engineering Taskforce, *Negotiation of NAT-Traversal in the IKE*, <http://www.ietf.org/rfc/rfc3947.txt>
37. JANET(UK), *JANET factsheets*, <http://www.ja.net/services/publications/factsheets.html>
38. JANET(UK), *IEEE 802.1X: Implementation at JANET-Connected Organisations*, <http://www.ja.net/documents/publications/technical-guides/8021x-tg-web.pdf>
40. JANET(UK), *JANET mail services*, <http://www.ja.net/services/mail/index.html>
41. Internet Engineering Taskforce, *Mailbox names for common services, roles and functions*, <http://www.ietf.org/rfc/rfc2142.txt>
42. Internet Engineering Taskforce, *DON'T SPEW - A Set of Guidelines for Mass Unsolicited Mailings and Posting*, <http://www.ietf.org/rfc/rfc2635.txt>
43. *Fight Spam on the Internet!*, <http://spam.abuse.net>
44. Internet Engineering Taskforce, *RFC2505, Anti-Spam Recommendations for SMTP MTAs*, <http://www.ietf.org/rfc/rfc2505.txt>
45. JANET(UK), *The MAPS RBL+ in JANET*, <http://www.ja.net/services/mail/janet-rbl/janet-dnsbls.html>
46. Office of Public Sector Information, *Statutory Instrument 2003 No. 2426 The Privacy and Electronic Communications (EC Directive) Regulations 2003*, <http://www.opsi.gov.uk/si/si2003/20032426.htm>
47. Information Commissioner's Office, *Electronic mail marketing Good Practice note*, http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/charities_and_marketing_12_06.pdf
48. JANET(UK), *Sponsored and Proxy Connections to JANET – Guidelines for Licence Holders*, <http://www.ja.net/documents/services/connections/sponsored-proxy-guidelines-june-07.pdf>
49. JANET(UK), *Effective incident response Guidance Notes*, <http://www.ja.net/documents/publications/technical-guides/gn-incident-response.pdf>
50. JANET(UK), *JANET CSIRT*, <http://www.ja.net/services/csirt/>
51. JANET(UK), *JANET Connection Policy*, <http://www.ja.net/documents/publications/policy/connection-policy.pdf>

16 References by publisher

- ACAS, *Advice leaflet – Internet and e-mail policies*, <http://www.acas.org.uk/index.aspx?articleid=808>
- Fight Spam on the Internet!*, <http://spam.abuse.net>
- Information Commissioner's Office, *Data Protection: The Employment Practices Code*, http://www.ico.gov.uk/upload/documents/library/data_protection/detailed_specialist_guides/employment_practices_code.pdf
- Information Commissioner's Office, *Information Commissioner's Office home page*, <http://www.ico.gov.uk>
- Information Commissioner's Office, *Electronic mail marketing Good Practice note*, http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/charities_and_marketing_12_06.pdf
- Information Commissioner's Office, *Employment Practices Data Protection Code: Part 3 - monitoring at work*, http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/coi_html/english/employment_practices_code/part_3-monitoring_at_work_2.html
- Internet Engineering Taskforce, *DON'T SPEW - A Set of Guidelines for Mass Unsolicited Mailings and Posting*, <http://www.ietf.org/rfc/rfc2635.txt>
- Internet Engineering Taskforce, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*, <http://www.ietf.org/rfc/rfc3931.txt>
- Internet Engineering Taskforce, *Mailbox names for common services, roles and functions*, <http://www.ietf.org/rfc/rfc2142.txt>
- Internet Engineering Taskforce, *Negotiation of NAT-Traversal in the IKE*, <http://www.ietf.org/rfc/rfc3947.txt>
- Internet Engineering Taskforce, *Netiquette Guidelines*, <http://www.ietf.org/rfc/rfc1855.txt>
- Internet Engineering Taskforce, *RFC 1173: Responsibilities of Host and Network Managers*, <http://www.ietf.org/rfc/rfc1173.txt>
- Internet Engineering Taskforce, *RFC2196: The Site Security Handbook*, <http://www.ietf.org/rfc/rfc2196.txt>
- Internet Engineering Taskforce, *RFC2505: Anti-Spam Recommendations for SMTP MTAs*, <http://www.ietf.org/rfc/rfc2505.txt>
- Internet Engineering Taskforce, *Securing L2TP using IPsec*, <http://www.ietf.org/rfc/rfc3193.txt>
- JANET(UK), *JANET CSIRT*, <http://www.ja.net/services/csirt/>
- JANET(UK), *JANET Connection Policy*, <http://www.ja.net/documents/publications/policy/connection-policy.pdf>
- JANET(UK), *Effective incident response Guidance Notes*, <http://www.ja.net/documents/publications/technical-guides/gn-incident-response.pdf>
- JANET(UK), *factsheets*, <http://www.ja.net/services/publications/factsheets.html>
- JANET(UK), *Investigating Copyright Complaints factsheet*, <http://www.ja.net/documents/publications/factsheets/077-investigating-copyright-complaints.pdf>
- JANET(UK), *JANET mail services*, <http://www.ja.net/services/mail/index.html>
- JANET(UK), *JANET NTP Service*, <http://www.ja.net/services/ntp/>
- JANET(UK), *JANET Roaming*, <http://www.ja.net/services/authentication-and-authorisation/janet-roaming.html>
- JANET(UK), *IEEE 802.1X: Implementation at JANET-Connected Organisations*, <http://www.ja.net/documents/publications/technical-guides/8021x-tg-web.pdf>
- JANET(UK), *Logfiles – technical guide*, <http://www.ja.net/documents/publications/technical-guides/logfiles.pdf>
- JANET(UK), *Sponsored and Proxy Connections to JANET – Guidelines for Licence Holders*, <http://www.ja.net/documents/services/connections/sponsored-proxy-guidelines-june-07.pdf>
- JANET(UK), *Terms for the Provision of the JANET Service*, <http://www.ja.net/documents/publications/policy/service-terms.pdf>
- JANET(UK), *The MAPS RBL+ in JANET*, <http://www.ja.net/services/mail/janet-rbl/janet-dnsbls.html>
- JANET(UK), *User authentication factsheet*, <http://www.ja.net/documents/publications/factsheets/041-user-authentication.pdf>
- JANET(UK), *Using passwords factsheet*, <http://www.ja.net/documents/publications/factsheets/026-using-passwords.pdf>
- JISC, *Improving Network Reliability: Senior Management Briefing Paper 11*, <http://www.jisc.ac.uk/media/documents/publications/smbp11improvingnetwork.rtf>

- JISC, *JISC Data Protection Code of Practice for the HE and FE Sectors*, http://www.jisc.ac.uk/publications/generalpublications/2001/pub_dpacop_0101.aspx
- JISC Legal Information Service, *Cybercrime*, <http://www.jisclegal.ac.uk/cybercrime/cybercrime.htm>
- JISC Legal Information Service, *Data Protection*, <http://www.jisclegal.ac.uk/dataprotection/dataprotection.htm>
- JISC Legal Information Service, *Intellectual Property Rights Overview*, <http://www.jisclegal.ac.uk/ipr/IntellectualProperty.htm>
- JISC Legal Information Service, *Internet Service Provider Liability*, <http://www.jisclegal.ac.uk/ispliability/ispliability.htm>
- JISC Legal Information Service, *Legal Guidance for ICT Use in Education, Research and External Engagement: Intellectual Property Rights*, <http://www.jisclegal.ac.uk/ipr/IntellectualProperty.htm>
- JISC Legal Information Service, *Legal Guidance for ICT Use in Education, Research and External Engagement: RIPA (Part 2) New Codes of Practice (22/01/2010)*, <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/tabid/243/ID/1235/RIPA-Part-2-New-Codes-of-Practice-22012010.aspx>
- JISC Legal Information Service, *Legal Guidance for ICT Use in Education, Research and External Engagement: The Digital Economy Act 2010: Implications for UK Colleges and Universities (23/07/2010)*, <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/tabid/243/ID/1603/The-Digital-Economy-Act-2010-Implications-for-UK-Colleges-and-Universities-23072010.aspx>
- JISC Legal Information Service, *Legal Risks and Liabilities for IT Services in Further and Higher Education*, <http://www.jisclegal.ac.uk/publications/legalRisks.htm>
- JISC Legal Information Service/Higher Education Information Directors Scotland (HEIDS), *Inappropriate use of computers – the technical investigation process*, <http://www.jisclegal.ac.uk/publications/Inappropriateuse.htm>
- Office of Public Sector Information, *Statutory Instrument 2002 No. 2013, The Electronic Commerce (EC Directive) Regulations 2002*, <http://www.opsi.gov.uk/si/si2002/20022013.htm>
- Office of Public Sector Information, *Statutory Instrument 2003 No. 2426 The Privacy and Electronic Communications (EC Directive) Regulations 2003*, <http://www.opsi.gov.uk/si/si2003/20032426.htm>
- Scottish Information Commissioner, *Scottish Information Commissioner's office home page*, <http://www.itpublicknowledge.info>
- UCISA, *Information Security Toolkit*, <http://www.ucisa.ac.uk/ist>
- UCISA, *Model Regulations for use of Institutional IT Facilities and Systems*, <http://www.ucisa.ac.uk/publications/modelregs.aspx>
- University of Cambridge, *Use and Misuse of Computing Facilities*, <http://www.cam.ac.uk/cs/iss/rules/guidelines.html>

Acknowledgements

This is the fourth edition of this document; much of the structure and text of the current document has been taken from the previous versions and the work of the original authors and editors is hereby acknowledged. Paul Whitton and Mohamed Imran produced much of the revised text for the fourth edition, which was reviewed by members of the UCISA Networking Group. Peter Tinson and Sue Fells edited this edition; the graphics, design and layout was carried out by Sam Westwood.