



# Logfiles

**Technical Guide**

---

## UKERNA Technical Guides

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

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

JANET Customer Service

UKERNA

Atlas Centre, Chilton, Didcot

Oxfordshire, OX11 0QS

Tel: 0870 850 2212

Fax: 0870 850 2213

E-mail: [service@janet.ac.uk](mailto:service@janet.ac.uk)

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

**<http://www.ja.net/services/publications/technical-guides/>**



# Contents

<b>1</b>	<b>Introduction .....</b>	<b>5</b>
1.1	<b>Motivation .....</b>	<b>5</b>
<b>2</b>	<b>Using Logfiles .....</b>	<b>7</b>
2.1	<b>Privacy and Legal Issues .....</b>	<b>7</b>
2.2	<b>Data Retention.....</b>	<b>7</b>
2.2.1	Data Preservation.....	8
2.3	<b>Access by Others.....</b>	<b>9</b>
2.3.1	Regulation of Investigatory Powers Act 2000.....	9
2.3.2	Other Statutory Notices .....	11
2.3.3	Police and Criminal Evidence Act 1984.....	11
2.3.4	Data Protection Act 1998.....	12
<b>3</b>	<b>Tracing Misuse .....</b>	<b>13</b>
3.1	<b>Clients .....</b>	<b>13</b>
3.1.1	Summary of Logs .....	13
3.1.2	Federated Authentication Systems.....	14
3.2	<b>Intermediaries .....</b>	<b>14</b>
3.2.1	Proxies and Caches .....	14
3.2.2	E-mail and News .....	15
3.2.3	Network Address Translation .....	15
3.2.4	Gateway Servers.....	15
3.2.5	Summary of Logs .....	16
<b>4</b>	<b>Examples .....</b>	<b>17</b>
4.1	<b>Attempted Break-in.....</b>	<b>17</b>
4.2	<b>Inappropriate E-mail .....</b>	<b>17</b>
4.3	<b>Abuse of Webmail Service .....</b>	<b>18</b>
4.4	<b>Denial of Service Attack with a Web Server Intermediary</b>	<b>19</b>
<b>5</b>	<b>Identifying Attacks .....</b>	<b>21</b>
5.1	<b>Authentication Logs .....</b>	<b>21</b>
5.2	<b>Service Logs.....</b>	<b>21</b>
5.3	<b>Summary of Logs.....</b>	<b>22</b>
<b>6</b>	<b>Implementation.....</b>	<b>23</b>
6.1	<b>Central Logging.....</b>	<b>23</b>
6.2	<b>Timestamps .....</b>	<b>23</b>
6.3	<b>Automated Processing of Logs.....</b>	<b>24</b>
6.4	<b>Graphing Activity .....</b>	<b>24</b>
<b>7</b>	<b>Next Steps.....</b>	<b>25</b>
7.1	<b>Network Flows.....</b>	<b>25</b>
7.2	<b>Intrusion Detection .....</b>	<b>25</b>
<b>8</b>	<b>References .....</b>	<b>26</b>
8.1	<b>Information and Guidelines on Logfiles .....</b>	<b>26</b>

<b>8.2</b>	<b>Data Retention.....</b>	<b>26</b>
<b>8.3</b>	<b>Access to Data .....</b>	<b>26</b>
<b>8.4</b>	<b>Analysing and Processing Logfiles .....</b>	<b>27</b>

# 1 Introduction

Logfiles are one of the most useful tools in detecting and investigating problems with computer systems. Logs can provide information about system faults and misuse as well as early warnings of problems. This Technical Guide discusses the logs that should be kept, the conditions under which they should be held and some of the uses to which they may be put. It is concerned mainly with logging of activity on computers, whether they act as clients, servers or proxies. Direct logging of activity on networks is not covered.

## 1.1 Motivation

Without collecting and analysing logfiles, it is impossible to know what is happening on a computer system. There will be no indication of faults and misuse and when they finally result in complaints from users, there will be no evidence to show the cause of the problem or how it can be cured. Failure to keep logfiles therefore leads rapidly to an unreliable system on which users will naturally be unwilling to rely for any critical function. Reliable systems can only be achieved if their performance is recorded and action taken to prevent or remedy problems. Logfiles also provide information about the usage of a service, and allow upgrades or alternative provision to be planned and installed before the load on the existing system becomes a critical problem.

As well as these internal pressures to deal with problems, there are also likely to be external pressures. Wide Area Networks, such as JANET, are shared resources and a problem on one computer or site can soon affect others. For example, a fault that causes excessive network traffic is likely to cause congestion for others as their traffic competes for the finite bandwidth and routing resources available. In cases of misuse it is common for an individual at one site to attack systems or users at others. If reasonable requests to deal with the problem are not satisfied then the responsible site is likely, at best, to suffer a tarnished reputation in the eyes of its peers. The JANET community, and the policies that support it, require its members to behave responsibly and not to cause unnecessary problems for others or harm the good reputation of the JANET network.

In extreme cases failure to deal with problems, whether arising from the lack of logfiles or unwillingness to use them, may even lead to legal cases. Service providers have paid large damages to individuals or companies harmed by the actions of their users. At present we are not aware of any cases where educational organisations have been held liable for the computing activities of their students or staff, but solicitors have expressed their view that courts might indeed find against the organisation. Another area where legal action might arise is in negligence: it has been suggested that if an organisation had been warned of a problem but did not deal with it, then subsequent victims might have a valid claim against the organisation.

In the case of faults, the best that logging can offer is the early detection and resolution of problems. However, in cases of misuse there is good reason to believe that a publicised practice of recording and analysing logfiles and dealing with those who misuse the system may itself be an effective deterrent. Logfiles can therefore act as a preventive measure, reducing the number of problems experienced by users and system owners.

Logfiles enable an organisation to improve its service to its own users and maintain a good reputation with others. In the near future, it appears likely that logfiles and a process to use them may be essential to defend against threats of legal action. It is important to note, however, that simply keeping logs is unlikely to be sufficient. It is also important to have processes for checking them, analysing the information they contain, and dealing promptly and effectively with problems.



---

## 2 Using Logfiles

### 2.1 Privacy and Legal Issues

Any comprehensive programme of logging will capture information about the activities of individual users. In some cases this has the potential to intrude on the privacy of those individuals. Users and system administrators must be clear that the sole purpose of logfiles is to provide a better service to legitimate users, by providing computers and networks that are fit for their intended purpose and which work as reliably as possible. Article 8 of the European Convention on Human Rights states that ‘Everyone has the right to respect for his private and family life, his home and his correspondence’: all users and administrators must respect this.

Most logfiles contain personal data, so are also subject to the provisions of the Data Protection Act. Users must be informed what information will be recorded and what it may be used for (as noted above, the very fact of notification may well discourage misuse of the system). Logfiles must also be protected from unauthorised access, use or modification. The Act requires that personal data should not be kept for longer than necessary; guidelines for the interpretation of ‘necessary’ are discussed in Section 2.2, on data retention, below.

The Data Protection Act allows data subjects to request copies of information held about them. Such Subject Access Requests (SARs) can be one of the hardest and most costly parts of the Act to comply with: however, the same processes that allow logfiles to be used to investigate problems should provide most of the apparatus required to deal with SARs.

Various Acts of Parliament make a useful distinction between traffic data and content data. Traffic data (sometimes referred to as communications data) is information about the existence of communications. For example, records that a particular user logged on to a workstation at a certain time, sent e-mails to a number of other, recorded, e-mail addresses and then logged out, would comprise traffic data about that session of use. The texts of the e-mails would, however, constitute content data. The law appears to be treating traffic data as less likely to involve a breach of privacy than content: for example the rules for police access to traffic data require a lower level of authorisation than for content data.

Most logfiles will contain only traffic data. However, there are a number of electronic systems where the distinction between traffic and content data is unclear. It is considered that the subject lines of e-mail are content, rather than traffic; similarly in the case of web requests the identity of the server from which a page was requested is traffic data, whereas the rest of the URL, which may well specify exactly what the user saw, is considered to be content.

The third type of data that is essential in identifying the individuals responsible for cases of misuse is the identity of the real world individuals who own, and should be responsible for, each login account or other online identity. This information is known by ISPs (Internet Service Providers) as subscriber data, but in universities and colleges it is likely to form part of student records and staff personnel files. As information relating to an identifiable individual it is, of course, subject to the Data Protection Act.

### 2.2 Data Retention

Problems are rarely detected the instant they occur so, to be useful, logfiles must be kept for some period of time. However, logfiles can grow very large, so shortage of storage space may put an upper limit on what this time may be. Even if logs can be physically stored, there is little point in keeping them for so long that the quantity of information prevents convenient searching. Where logfiles contain personal data, the Data Protection Act’s Fifth Principle also requires that they not be kept for longer than necessary for the specific purpose for which they were collected. The European Directive 2002/58/EC on Privacy and Electronic Communications, which applies Data Protection law to electronic

communications networks, states that traffic data must be anonymised or destroyed once it is no longer required, but identifies the provision of value added services and investigation of unauthorised use as legitimate reasons to collect and retain this data. Collecting and keeping traffic data for as long as necessary to investigate misuse of computers and the network is therefore acceptable.

A number of Codes of Practice have been written in an attempt to establish a reasonable balance between usefulness of logs on the one hand and privacy and practicality on the other. Following a recognised Code of Practice should be a good defence against accusations of keeping either too few records or too many. For some time, the Code of Practice most relevant to computer and network logging has been that produced by the London Internet Exchange (LINX) in 1999, which is available online at:

**[https://www.linx.net/www\\_public/community\\_involvement/bcp/traceabilitybcp\\_v1](https://www.linx.net/www_public/community_involvement/bcp/traceabilitybcp_v1)**

The LINX document was prepared and is maintained by members of the Exchange, who include many of the major ISPs in the UK. The document was also reviewed by the Data Protection Commissioner, responsible for ensuring compliance with the Data Protection Acts. The document recommends that traffic data should be retained for a minimum of three months to allow misuse to be traced, but that to comply with the Data Protection Act it should not be kept for more than six months except where it relates to a known case of misuse. If an investigation is in progress then data relating to it may be kept until the investigation is complete. The same minimum retention time is recommended for subscriber data. However, users of university and college computers will usually be students or staff. Both of these legal relationships involve much longer retention periods to comply with education and employment law, so information about these users' identities will normally and legitimately be held for much longer than six months.

In 2004 the Home Office published a further Code of Practice under the Anti-Terrorism, Crime and Security Act 2001 (ATCSA). This voluntary Code establishes a new Data Protection Act purpose – protecting national security. Providers of public networks may use this purpose to retain data for longer than might have been necessary for investigating misuse: up to 12 months for subscriber data and six months for most traffic data. The Code sets a much lower limit of four days for web cache logs, apparently in recognition of the immense rate at which these can accumulate. The Act allows the Home Secretary to make the Code mandatory on particular networks or groups of networks, or to extend its scope. However, no plans have been announced to do so. As JANET is a private network, it and most university and college networks are not covered by the present ATCSA code. Organisations providing these networks should therefore ensure they have legitimate reason if they extend their data retention periods beyond what is needed for investigation of normal misuse.

In 2005 proposed European legislation on data retention was discussed by both the European Council of Ministers and the European Commission. Early drafts of both the Council and Commission proposals appear to make retention of some types of communications data mandatory; however, the details and proposed retention times vary between the two documents. It appears that either text, if adopted, would involve the storage of more data than the UK's existing ATCSA code. However, both documents apply only to public networks, so should not directly affect JANET or most of its connected sites. Were any European law to be agreed and passed there then would be a period of further discussion (typically up to two years) before it is implemented as UK law.

### **2.2.1 Data Preservation**

On a small number of occasions, following major terrorist or other criminal incidents, the UK police have asked the providers of communications networks (including JANET sites) to preserve logs and other relevant files in case they contain information relevant to the investigation. The purpose of these requests is to prevent existing information from being overwritten or deleted, not to cause additional information to be collected. There is no requirement to comply, but such exercises have protected useful information for the police

in the past and are considered helpful. In practice, unless the police request contains more specific instructions, the most usual response is to take a backup of main server logs and to reserve this along with a recent set of backup tapes that are not reused until the police investigation is completed.

Such data preservation is permitted under the Data Protection Act 1998, where section 29 allows processing of data for the purposes of the prevention and detection of crime and the apprehension or prosecution of offenders. Section 28 provides similar permission where necessary for the purpose of safeguarding national security. Data preserved under either section may be exempted from the normal subject access provisions of the Act where disclosure might be harmful to the purpose.

The preserved data should be kept by the organisation in a secure place: if the police find they need access to it then they will use one of the legal mechanisms described in the next section.

## 2.3 Access by Others

Evidence from logfiles may be useful to the police and other investigating authorities in cases where unlawful acts have been committed. A number of different Acts of Parliament include provisions under which such authorities may request or require such evidence to be provided to them. This section attempts to summarise the provisions likely to be encountered by universities and colleges (for that reason, provisions that only apply to public networks have been omitted). However, it does not constitute formal legal advice. The definitive source of information is the original Acts and Codes of Practice: web addresses are listed in Section 8.3. Since there may be a legal requirement to comply promptly with some of these notices, organisations should consider instituting standard procedures for responding to them. They may also wish to discuss these procedures and any requests with their lawyers.

Requirements to produce (or collect) data	<ul style="list-style-type: none"> <li>Regulation of Investigatory Powers Act 2000 (Part I Chapter II)</li> </ul>
	<ul style="list-style-type: none"> <li>Other Acts including: Social Securities Fraud Act 2001 Consumer Protection Act 1987</li> </ul>
Requirement to produce or give access to data	<ul style="list-style-type: none"> <li>Police and Criminal Evidence Act 1984</li> </ul>
Requests to produce data	<ul style="list-style-type: none"> <li>Data Protection Act 1998</li> </ul>

In the case of the police, access to data will normally be obtained by a notice under the Regulation of Investigatory Powers Act 2000. If the information sought is not communications data then a request under section 29 of the Data Protection Act 1998 will normally be used. A Production Order under Schedule 1 of the Police and Criminal Evidence Act 1984 will only be used where neither of these routes has been successful, or where the voluntary request under the Data Protection Act 1998 would not be appropriate to the investigation.

### 2.3.1 Regulation of Investigatory Powers Act 2000

Part I Chapter II, and in particular §22, of this Act deals with the disclosure of communications data to law enforcement and other public bodies. This came into force in January 2004 and is now the normal process for all access to communications data, replacing §29(3) of the Data Protection Act 1998.

Communications data is information about traffic on a network, but **not** the contents of that traffic. §21(4) of the Act provides a full definition of Communications Data, separating it into three types:

- (a) Information forming part of a communication, that is needed by the system to deliver the communication from its source to its destination. For example, source and destination addresses and routing information.
- (b) Other information concerning the use of the system by individual users. For example, times when individual users were logged on and the IP addresses they were allocated.
- (c) Other information about the users of the system, referred to elsewhere as subscriber data. For example, the identity of the owner of a login name or e-mail address.

Logfiles may contain any or all of these types of communications data. Some logfiles will also contain information that is not communications data such as the subject lines of e-mails or full URLs of web requests (only the identity of the web server is communications data), which must not be disclosed under §22. Responding to a §22 notice is therefore likely to require making edited versions of logfiles with these unauthorised types of information removed.

The Act permits any designated authority to issue a notice to a communications provider requiring either that existing communications data be disclosed or that particular communications data be collected. Communications providers are widely defined and would certainly include any university or college providing Internet access to its members. A provider receiving such a notice must act on it, otherwise it may itself be committing an offence. The Regulation of Investigatory Powers Act makes the authority that issues a notice responsible for ensuring that it is proportionate: the communications provider releasing the information is not required, or entitled, to make any judgement on this. The purposes for which a notice can be served include interests of national security, detecting crime and preventing crime or disorder, national economic well-being, public safety, protection of public health, assessment of taxes and duties, preventing death, and preventing or mitigating injury to an individual's physical or mental health.

To be allowed to issue notices under §22, an authority must be designated by the Home Secretary. The initial list of authorities was published as The Regulation of Investigatory Powers (Communications Data) Order 2003 (Statutory Instrument 2003 No. 3172). To the law enforcement authorities included in the Act (listed in Schedule 1 of the Regulations) this adds the emergency services, central and local government departments, the NHS and others with powers to investigate compliance with particular laws (listed in Schedules 2 and 3). Many of these authorities do not have powers to access the whole range of communications data set out in §21(4) and above – many are limited to the subscriber data of type (c) and some are restricted to particular types of communication services – and in some cases a more senior officer is required to authorise notices for the more intrusive types of data. The Schedules to the Regulations set out these arrangements in detail.

For some time, police forces have had designated Single Points of Contact (SPoCs) for dealing with the communications industry. Officers staffing the SPoCs have been specifically trained both in the legal requirements of handling data and in what is likely to be practical for network operators to provide. SPoCs have been useful to ensure that the law is used properly and that the evidence obtained is suitable for the investigation and subsequent prosecution. The Home Office has therefore granted the new authorities powers under §22 to ensure that their staff have equivalent training and work in a similar way as the police SPoCs. The Home Office is maintaining a register of individuals designated to exercise the powers on behalf of each authority, and every §22 notice must be approved by one of these designated persons before it is served on a communications provider.

The process of issuing notices should be further standardised by a Code of Practice. A draft Code was published in 2001: it is hoped that this will be finalised and come into force soon. A standard form requiring disclosure of communications data has been published and should be used for all notices. Notices that are received by JANET sites should be checked to confirm that they come from a designated authority, request data which that authority is entitled to receive, and have been issued by the appropriate designated person or SPoC. JANET-CERT has access to the Home Office register and can confirm that notices have been approved by the correct designated person. Notices that appear incorrect should not be

acted upon. There have been reports that individuals have attempted in the past to use other statutory powers (see next section) to gain access to information they did not have authority to see and the Home Office has asked for reports of any attempts to abuse the §22 powers in this way.

It is strongly recommended that any organisation likely to receive statutory notices under the Regulation of Investigatory Powers Act 2000 or other statutes (see below) should designate and train a person or office to deal with the notices, and that all enquiries regarding notices should be directed to that person or office. Legal advice is likely to be helpful when setting up these procedures.

### **2.3.2 Other Statutory Notices**

The Regulation of Investigatory Powers Act 2000 (RIPA) is just one of a number of pieces of legislation that create rights for designated authorities to obtain information for particular purposes. These include the Police and Criminal Evidence Act 1984 (law enforcement), the Consumer Protection Act 1987 (trading standards), and the Social Security Fraud Act 2001 (benefits agency), as well as court orders and police warrants. The Home Office intends that all access to communications data will eventually be done under RIPA powers: however, it is likely to be some time before the other powers stop being used.

When presented with a valid statutory notice by a person entitled to issue that notice, it will normally be an offence not to provide the required information. However, anyone receiving such a notice must check both that the notice is valid and that the person is entitled to use it. This will normally involve checks with appropriate third parties.

The LINX has published a Best Current Practice document on privacy, which contains useful guidelines on dealing with statutory notices. This is available at:

[https://www.linx.net/www\\_public/community\\_involvement/bcp/bcp\\_userprivacy\\_v1](https://www.linx.net/www_public/community_involvement/bcp/bcp_userprivacy_v1)

### **2.3.3 Police and Criminal Evidence Act 1984**

Schedule 1 of the Police and Criminal Evidence Act 1984 (PACE) enables a police constable to ask a judge to make an order requiring a person to either produce or give the constable access to information that the person holds or has access to. An order will only be granted if there are reasonable grounds for believing that the information will be of substantial value in investigating an indictable offence, and could be used as evidence. Furthermore all other means of obtaining the material must have either been tried or found inappropriate. The judge will then decide whether it is in the public interest to make the order that the information be produced.

A production order may be served on an individual, a partnership or corporate body, and may be delivered either by hand or post. The person or body on whom the order is served must then either produce the information, or give access to it, within a fixed period, typically seven days from the issue of the order. Failing to comply with an order, or tampering with the information once the order has been served, is a contempt of court: a serious criminal offence.

PACE production orders are used as a last resort, generally where information is not accessible by any other power. They may also be used in place of Data Protection Act 1998 requests (see below) in cases where it is desirable to have a judge rule on the proportionality of disclosure before it occurs, rather than after as is the case with the Data Protection Act process. PACE production orders should be simple to deal with: the recipient must comply with the order or commit a serious criminal offence.

### **2.3.4 Data Protection Act 1998**

Various sections of the Data Protection Act permit data controllers to disclose personal data without breaching their obligations under the Act. In particular §29(3) permits this where the information is required for the prevention, detection or prosecution of crime, and §28 applies where the disclosure is required in the interests of national security. In all cases, disclosure is voluntary and the data controller must consider whether the breach of privacy is proportionate to the stated reason of why the information is needed.

Although these provisions are still in force, they have been superseded for communications data by the powers under the Regulation of Investigatory Powers Act described previously and should no longer be used for this type of information. Where other types of information are concerned, the request for disclosure should be made in writing on a standard form, giving enough information about the purpose for the assessment of proportionality to be made. In the case of a request from the police, this form should always be issued by the force's Single Point of Contact (see above).

As discussed in the previous section, a PACE production order may be preferable in some circumstances to a Data Protection Act request as it allows this assessment of proportionality to be made by a judge rather than the recipient of the request.

## 3 Tracing Misuse

### 3.1 Clients

Most cases of misuse will be reported as originating from one or more electronic identities, for example e-mail or IP addresses. Such identities are public, and can be seen by anyone on the Internet, but in most cases it will only be the local site that can relate these electronic identities to the individual responsible person. Ensuring that the actions of these electronic addresses can be assigned to responsible individuals is therefore fundamental to any attempt to track down network misuse. Since most electronic identities can be forged with various degrees of technical difficulty (e-mail addresses are easiest, IP addresses used for UDP packets slightly more difficult and IP addresses used in TCP connections significantly harder) it is also important to collect, as a matter of routine, sufficient reliable information to be able to prove when a forgery has taken place and thereby remove blame from an innocent individual.

Individuals are usually identified to computers by login name or e-mail address, which are not usually the identifiers that are used in complaints, so a conversion process will usually be needed to identify the individual who may be the subject of a complaint.

In the simplest case a reported IP address will be the fixed address of a workstation in a private office. The owner of the office will normally be responsible for activity by that IP address and only a record of the ownerships of allocated addresses will be needed

The situation is more complicated where a number of different users may use the same computer, either because the computer is a system that supports multiple users simultaneously, or because it is a public workstation that may be used by different people at different times of day. In the latter case, it should be possible to identify a single login account that was logged in on the workstation at a particular time. To achieve this it is essential to have login records that can be searched by workstation address and time. A record of the ownership of login accounts and e-mail addresses is, of course, a basic management tool. It is recommended that all users be made formally responsible, and accountable, for the activities of accounts allocated to them.

The next level of complication arises when different client computers use a single IP address at different times. This occurs whenever a pool of IP addresses is shared between a number of computers, for example in dial-up, mobile or fixed networks where addresses are allocated temporarily to active computers (the BOOTP or DHCP protocols are commonly used to manage addresses in these situations). Here it is essential to have logs of which client computer was allocated each IP address and the times when the use of the addresses began and ended. Once the workstation has been identified, records of logins and times can once again be used to identify the responsible person.

#### 3.1.1 Summary of Logs

Type of System	Logs required
Single-owner workstation, fixed IP	<ul style="list-style-type: none"> <li>• IP -&gt; owner</li> </ul>
Shared workstation, fixed IP	<ul style="list-style-type: none"> <li>• IP + time -&gt; login login -&gt; owner</li> </ul>
Dynamically configured workstation	<ul style="list-style-type: none"> <li>• IP + time -&gt; workstation</li> <li>• workstation+time -&gt; login</li> <li>• login -&gt; owner</li> </ul>

### 3.1.2 Federated Authentication Systems

Federated Authentication Systems are a relatively recent development where an organisation providing a service to a user relies on another organisation (typically the user's home organisation) to authenticate the user, rather than maintaining its own local database of usernames and passwords. Examples of federated authentication systems include Shibboleth (<http://www.ja.net/development/aa/shib/>), used by publishers to authenticate access to online resources, and eduroam (<http://www.eduroam.org>), used by education organisations across Europe to authenticate visitors from other organisations and provide them with network access.

Some federated systems are designed to preserve the anonymity of the user. In these cases the organisation providing a service is merely told 'yes, this is one of our users', accompanied, if necessary, by attributes that may be required by the service such as whether the user is a student or member of staff. Such systems add an additional step to the process of tracing access through logs, since the organisation that gives the authenticated user access to its service may not itself be able to link the service provided to the user it was provided to (the step resulting in a 'login' name in the previous table). Instead this requires two steps: the service provider needs to identify the home organisation and the home organisation needs to identify the user they authenticated.

Details of the logs that are needed will vary between different federated authentication systems and should form part of the federation agreement. In general, service providers need to record at least the home site that provided the authentication for a particular service request, together with other identifying material such as the service that was requested and any identifier that was provided by the home site. Home sites that are providing authentication to external organisations need to record at least the source and time of the authentication request, the same identifying information for the request, and the local user who was authenticated.

## 3.2 Intermediaries

The types of client logging discussed above deal with the situation where a direct TCP/IP connection exists between the client machine and the server. However there are also a number of services and configurations where some other machine is involved as an intermediary between the client and the server. Additional logs are needed in these cases as the end server will see the activity as originating with the intermediary, while the client logs will only show a connection to the intermediary and not to the end server where the alleged misuse occurred. Intermediary logs are required to link these two sets of information.

The most obvious examples of intermediaries are proxies and caches; store and forward systems such as e-mail servers also act as intermediaries. There are also other systems, particularly those that act as gateways between different Internet services, that may act as intermediaries and therefore need to record appropriate logs.

Some intermediary systems handle requests for very large numbers of clients and servers so that a simple timestamp may not be sufficient to identify a communication uniquely. Logs on these systems will often need to record additional details of each transaction, such as URLs for web requests or subjects of e-mail messages, to allow a particular communication to be identified. Complaints where these details are not included are likely to be very hard to investigate.

### 3.2.1 Proxies and Caches

Proxies are systems explicitly designed to act as intermediaries. Clients make requests to a proxy and the proxy may send that request to a server on behalf of the client. Proxies are usually designed to support one or a small number of protocols, for example HTTP and

---

FTP, and, unlike gateways, use the same protocol for the requests they receive and send. A caching proxy may respond directly to some requests as an alternative to passing them on.

As the server's records will show the request coming from the IP address of the proxy, the proxy must itself retain a log of the client IP address or authenticated user on behalf of which each request was made. A busy proxy will often make many requests each second to a popular server so the time of the request and identity of the server will not be sufficient to identify an individual client request uniquely. It is therefore normal for this type of proxy to retain additional information about each request, for example the web URL, that will allow the responsible client to be linked to the information recorded by the server. Where the protocol and local policy permit, a great deal of investigation time can be saved if the proxy includes the client address in each request it passes on. If this is visible to the end-user, or recorded in the logs on the final server, then there will be no need to search through large volumes of proxy logs.

### **3.2.2 E-mail and News**

Store and forward systems, such as news and electronic mail, differ from proxies in that the transaction with the client is completed before the message is passed to another server. However they still act as intermediaries so should retain a log of the client from which each message was received and the server or other destination to which it was passed. In theory each mail or news message includes as part of its content a full record of its origin and path: however, this information is relatively easy for a malicious user to forge. Trustworthy logs kept by servers are important tools in detecting this kind of forgery. This may be especially important if messages are forged so as to appear the responsibility of an innocent party. As with proxies it is common to record the message subject or ID to ensure correct identification of a particular message.

### **3.2.3 Network Address Translation**

Network Address Translation (NAT) is another type of intermediary, but one that works at the network rather than the application layer. Clients of an address translation system usually have private IP addresses. Clients send any packets for external destinations to the NAT system, which rewrites them with a public source IP address and forwards them to their destination. The NAT system must of course remember the state of each communication so that when it receives response packets it can rewrite their destination addresses and send them to the correct client.

NAT systems can use a variety of strategies to allocate external addresses to internal clients. Some create static mappings between external and internal addresses for each client; more commonly the system will behave like a proxy, using one or a small pool of external addresses for all communications.

Address translation systems have the same basic logging requirement as other intermediaries: to be able to relate a request made by the translation system to the client that invoked it. However the complexity of mapping used by some systems can make this a challenge. Attention to traceability requirements must therefore be included at the design and implementation stages of any address translation system.

### **3.2.4 Gateway Servers**

The final class of intermediaries is gateway systems that take a request from a client in one form and use it to generate a request in another form to a server. The most familiar gateway at present is probably a webmail server, which takes an HTTP form submission and uses it to generate an SMTP request. The command sent by the gateway will not necessarily contain any information originating from the client, so it is particularly important that gateway servers keep reliable records of their activities. Some gateways may add client information

to their output - for example webmail servers often include the client IP address as a header in the SMTP mail - but it is important to know how much of this information can be relied upon, and how much is under the control of a potentially malicious user. Of course if the gateway itself is compromised then nothing about either its output or its logs can be trusted. Some systems can be made to act as accidental gateways, for example a badly configured web cache may allow e-mail forgery. Such systems, and others with inadequate logging, are a hazard to the Internet as they provide abusers with complete anonymity for unauthorised or illegal activity.

### 3.2.5 Summary of Logs

Type of System	Logs required
Proxy server	<ul style="list-style-type: none"><li>• destination IP address + time + request -&gt; client/user</li></ul>
Mail or news server	<ul style="list-style-type: none"><li>• destination IP address + time + request -&gt; client/user</li></ul>
NAT/PAT server/SOCKS proxy	<ul style="list-style-type: none"><li>• source &amp; destination IP address (+port) + time -&gt; client/user</li></ul>
Gateway server	<ul style="list-style-type: none"><li>• server IP address + time + request -&gt; client/user</li></ul>

## 4 Examples

The following examples show some of the types of information that are available to the victims of computer misuse. Real examples have been used with names and addresses modified to protect the sites involved. These are typical of the evidence that may be sent to a site to complain about the activities of its users. In each case the receiving site will need to use additional logs relating to its clients and intermediaries to understand and investigate the origin of the misuse.

### 4.1 Attempted Break-in

The following entries were recorded by syslog on a UNIX® system called `victim`, whose clock is known to be synchronised to a reliable UK time source (see section 6.2):

```
Jul  4 19:08:11 4Q: victim telnetd[338556]: connection from
attacker.camford.ac.uk
Jul  4 19:08:12 0F: victim telnetd[338556]: ignored attempt to
setenv(_RLD, ^?D^X^\    ^?D^X^^
^D^P^?^?$^B^Cs#^?^B^T#d~^H#e~^P/d~^P/~^T#`~^O^C^?^?L/bin/
sh%32614c%11$hn%86
000c%12$hn)
```

This shows an apparently normal telnet connection from the host `attacker.camford.ac.uk`, during which the attacker attempts to overflow a buffer in the telnet server program. This is clearly an attack, whose intent is to obtain a command shell (`/bin/sh`) with root privilege, so the victim site would expect Camford to investigate. Provided `attacker.camford.ac.uk` is an end-user machine, this should simply be a case of identifying the person who was logged on to that machine at 19:08 on 4 July. If `attacker` is an intermediary, for example a network address translation system, then its logs will need to be used to identify the internal host that was the source of the activity. In practice it is most common for this type of attack to come from a host that has itself been compromised, often using the same attack, so the file system and logs on `attacker` will also need to be checked for signs of malicious activity. This also means that times and logs on the attacking machine may have been modified so user records from a central authentication server (if available) may be a more reliable source.

It is worth noting that `victim` was one of a number of similar machines in its department that were subject to this attack. `victim` had been patched, so the attack on it was unsuccessful. The other machines were compromised, and no trace of the attack was left in their syslog files.

### 4.2 Inappropriate E-mail

A complaint was received from a Microsoft network subscriber who had received an offensive e-mail:

```
From: heidi32396@camford.ac.uk <heidi32396@camford.ac.uk>
To: user@msn.com <user@msn.com>
Subject: Teens & Hot Horny Housewives
```

This was already suspicious as the username given in the From field of the e-mail is not of a form used by Camford and there is no user of that name. The full headers from the e-mail were obtained from the complainant, and showed clearly that the e-mail had not originated from a JANET site:

```
Received: from cpimssmtpa02.msn.com - 207.46.181.107 by
email.msn.com with
Microsoft SMTPSVC;
    Fri, 11 May 2001 13:09:34 -0700
Received: from njkkkkkkk.com ([38.31.27.7]) by cpimssmtpa02.msn.com
```

```
with
Microsoft SMTPSVC(5.0.2195.3225);
  Fri, 11 May 2001 12:56:53 -0700
Message-ID: <NApRVdL7bndr-.YGar-xZdAA18Fi2SQJtihM@njkkkkkkk.com>
From: heidi32396@camford.ac.uk <heidi32396@camford.ac.uk>
Bcc:
To: user@msn.com <user@msn.com>
Subject: Teens & Hot Horny Housewives
Date: Sun, 05 Mar 2000 20:34:33 -0400 (EDT)
MIME-Version: 1.0
Content-Type: text/plain; charset="US-ASCII"
Content-Transfer-Encoding: 7bit
Return-Path: heidi32396@camford.ac.uk
```

The Received headers created within the `msn.com` domain indicate that the message in fact originated from a customer of an American ISP, and that references to Camford had been forged to conceal the true origin of the message. The logfiles of the mail server at Camford further confirmed that no message had been sent to the recipient e-mail address from that site.

The message may have been forged by the ISP's customer, or may have been inserted into the Internet through a badly configured proxy or other system at the customer site. Both techniques are all too common ways to generate volume advertising by abusing services provided and paid for by others.

## 4.3 Abuse of Webmail Service

A customer of an international ISP received an offensive e-mail from an address at `hotmail.com`. Hotmail is a webmail gateway, so is effectively an intermediary. Like many other webmail systems, Hotmail adds headers to the e-mail it sends that include the IP address of the host that submitted the web request to Hotmail that caused it to generate the message. On inspection of the full headers of the offensive mail message the following information was found.

```
Received: from mail pickup service by hotmail.com with Microsoft
SMTPSVC;
  Thu, 12 Apr 2001 08:12:26 -0800
Received: from 192.251.0.8 by lw3fd.law3.hotmail.msn.com with HTTP;
  Thu, 12 Apr 2001 16:12:21 GMT
X-Originating-IP: [192.251.0.8]
Date: Sun, 1 Apr 2001 10:00:00
```

The X-Originating-IP header and lines above it were written by Hotmail and are usually reliable; note that the Date header has been forged by the creator of the offensive mail. Including reliable information in the outgoing message means that in most cases it will not be necessary to search through the logs on the Hotmail intermediary. However, these logs should still be kept, as there have been attempts to forge or conceal this X-Originating-IP information.

The IP address `192.251.0.8` belongs to a site cache, so the logs on this cache must be checked to determine which local host was responsible for the request. This involved searching for the Hotmail host named in the last of the received lines: `lw3fd.law3.hotmail.msn.com`. The following entry was found, but note that there is a 10 second difference in the time stamps. In this case the correct sender was identified but even this time difference, which was due to a failure to synchronise the system clocks to an international standard, could have prevented or cast doubt on the identification of the offender.

```
Thu Apr 12 16:12:31 2001 6913 babel.comp.camford.ac.uk TCP_MISS/200
12339 POST http://lw3fd.law3.hotmail.msn.com/cgi-bin/premail/4284 -
DIRECT/209.185.240.250 text/html
```

`babel.comp.camford.ac.uk` is a single-user workstation and its login records showed the identity of the account that was logged in at the time of the offensive posting.

---

badguy pts/2 babel.comp Thu Apr 12 12:08-17:23 (05:14)

## 4.4 Denial of Service Attack with a Web Server Intermediary

The site `intermediary.ac.uk` observed an unusually large traffic load on its link to the JANET network. At the same time the web server of `victim.com` suffered a denial of service attack, receiving a large number of packets from `www.intermediary.ac.uk`. On further investigation, the following request was found in the web logfile on `www.intermediary.ac.uk`.

```
2001-06-28 09:34:37 webcache.attacker.ac.uk - www.intermediary.ac.uk
80 GET /scripts/../../../../winnt/system32/cmd.exe /c+ping+-v+ping%20-
n+2000+-l+65500+-w+0+www.victim.com 132 502
```

This indicates that the web server received a request from a host called `webcache.attacker.ac.uk`. It is a reasonable guess that this host is itself a proxy. The filename requested contains a series of `../` entries which attempt to move the program out of the initial `/scripts` directory and indeed out of the area normally containing web scripts or files. This should not be valid and should be rejected as an illegal request, but the web server program had a well known bug, known as a directory-traversal vulnerability, which let it accept and service requests of this type rather than returning an error message. The directory traversal is used to move to the Windows system directory and to run the `'ping'` command with parameters to make it generate 2000 packets, each 64K bytes in size, as fast as possible. These caused both the unusual traffic flow and the denial of service attack.

Because the web server logs are available it is possible to identify the system `webcache.attacker.ac.uk` from which the request came. However, as this system is itself a proxy, the attack must be traced back by checking the logs on that proxy for a request made to `www.intermediary.ac.uk` at that time, and containing the same request string. The cache logs should identify the client machine responsible and from its login records the offending user can be found.

This final example illustrates the range of logs that can often be needed to trace activity back to its source. Not all the computers through which an attack passes will themselves be compromised; they may be performing quite correctly or just offering a service that has been used in an unauthorised way. Indeed, even though the web server in this case could have been broken in to using the same vulnerability, it was not necessary to do this to make it participate in a denial of service attack that was disruptive both to the target and to the intermediary site.



## 5 Identifying Attacks

The remaining group of systems whose logfiles are likely to be of interest is servers. Whilst logs from clients and intermediaries will usually indicate attacks against other sites, logs from servers will normally be used to detect attacks, or attempted attacks, either on the servers themselves or on other local systems. Public servers such as web or mail systems are likely to be the most exposed to hostile activity on the Internet so these should always be configured to keep good and secure logs. Internal servers should also record logs as these may be subject to attack from within the organisation, or may be used by malicious local users to practice before attacking systems elsewhere. Detecting and preventing such activity at an early stage by recording and monitoring server logs can save the organisation a great deal of trouble.

Attacks against servers generally have one of two intentions. One is to gain access, presumably unauthorised, to the information or services supplied by that particular server; students might well be motivated to try to gain access to the server that contains their examination results, for example. The second aim is to make a server perform some function for which it is not intended, for example to make it act as an intermediary in another attack, as shown in Section 4.4. Traces of these two types of attack will often appear in different sets of logfiles. It is therefore essential to ensure that both types of logs are recorded and checked regularly for suspicious activity. All logs should record both successful and unsuccessful attempts to use the system: a successful access by an unauthorised person to log in to a system is unlikely to stand out in a logfile, but it is likely to be preceded by a large number of unsuccessful attempts which should alert the operator to the problem.

### 5.1 Authentication Logs

Systems that require authentication should always keep a record of the users who authenticated successfully, and also of failed attempts to authenticate. Where a number of consecutive failures cause an account to be locked out, this must be recorded. In most cases the system should take additional measures to alert the operator to this event. Authentication failures may be due to mishaps – genuine users can mistype or forget their passwords – but any patterns of failures should be investigated. Authentication logs should also be checked for any unexpected periods of silence, as these may indicate that an intruder has been able to tamper with the logs to conceal evidence of their activities. Entries in authentication logs should always be associated with an accurate time; where a single authentication gives access to a session, rather than a single transaction, it can also be helpful to record the time when the session ended.

### 5.2 Service Logs

Servers that are accessible to untrusted users should also retain logs of the requests made to them. For example, public web servers should usually record the URLs requested by their clients. The time and the IP address from which the request came should also be retained. As with authentication logs, unusual events are often a sign of problems. These may include periods of unusually low or high activity, though web servers in particular can see unexpected surges in legitimate requests. A common way to attack a server is to present it with unexpected input: very long requests, or those containing unusual characters, should be investigated, as should any request containing the name of a command interpreter, such as `/bin/csh` or `cmd.exe`. Service logs often cannot show whether an attack was successful – even a request that failed as far as the service is concerned may have achieved its malicious purpose before it was rejected.

### 5.3 Summary of Logs

Type of System	Logs required
Authentication service	<ul style="list-style-type: none"><li>• userid + time + login success/fail + source IP (userid + time + logout) details of locked accounts</li></ul>
Information service	<ul style="list-style-type: none"><li>• source IP + time + request + result</li></ul>

---

## 6 Implementation

The ways to enable and configure logging will vary from one computer and software system to another, and should be covered in the system documentation. This section cannot deal with such detailed instructions, but identifies a number of common topics that have been found to be useful in many different circumstances.

### 6.1 Central Logging

One of the uses of logfiles is in the investigation of attacks on computer systems. However, a successful attack will often give the intruder complete control of the system, including the ability to delete or modify files. Tampering with logfiles to remove evidence of the break-in is normally one of the intruder's first priorities. The risk of finding deleted or corrupted evidence can be greatly reduced by holding logs on a different computer. A successful intruder may still have the ability to add records to the logfile but is much less likely to be able to rewrite history. Having logs on a dedicated central system can also make it much easier to deal with logs from a large number of different computers, as these will automatically be gathered into one place.

The most commonly used system for writing remote logfiles is the syslog service, which was originally written for UNIX® but is now available for most other operating systems. Syslog allows messages of a standard form to be both written to a local file and transmitted over the network to one or more central logging hosts. Syslog sends each message separately, so there should be little delay between the local and remote copies of the logfile being updated. Messages are automatically timestamped by the service, but this still relies on the system clocks being synchronised to some common standard (see the next section).

Two potential problems need to be borne in mind when using the syslog service over a network. The first is that messages are sent over the network using the User Datagram Protocol (UDP), so there is no guarantee that they will arrive at their destination. On a Local Area Network with little congestion this is not normally a problem, but messages may be lost if there is a high traffic load on the network. Syslog over UDP may not be sufficiently reliable to be used over a Wide Area Network. For this reason it is good practice to store logs locally on the system that generates them, as well as sending them to a central logging host. The other problem is that use of remote logging can itself lead to network congestion if a very large number of error messages are generated. For example a denial of service attack on a network will be made much more effective if the target systems are trying to report each attack over the same network. Avoiding this problem requires careful design of the logging system. One option is to summarise batches of repeated log messages into a single message including the number of repeats in a particular period of time.

Highly reliable central logging systems can be built by using separate network connections to carry logging messages. In such a system, the critical computers generating logs will have dedicated network links to the central logging host. These links should not be used for normal traffic, nor connected to the production network. Such systems can also protect against sophisticated attacks where a denial of service attack against the logging system is used to conceal an intrusion into vital production services. The ultimate in tamper-resistance is to write logs immediately to an unerasable form of storage such as a write-once, read-many CD-R or DVD-R drive.

### 6.2 Timestamps

Most incidents involve more than one computer so it is common for an investigation to have to deal with logfiles from many different systems, possibly at different sites or even in different countries. Entries in logfiles that refer to the same event are most commonly matched by comparing their timestamps. To ensure that the times from different logs can be compared within a site, or with a complaint made from the other side of the world, it

is essential that the times of different computers making the logs be synchronised to an international standard. The Network Time Protocol (NTP) is the common way to do this across computer networks and countries. UKERNA provides a central NTP service, linked to a number of atomic clocks keeping standard international time, which JANET sites can and should join. For more details of this service, see :

**<http://www.ja.net/services/network-services/ntp/>**

International incidents often occur across different time zones so that the numeric values of time may not be directly comparable. All systems should be set up to record the time zone against which they are logging, and whether daylight saving has been applied. Unfortunately there are a number of different formats for recording this information. Indeed some time zone names are not unique, so correlating logs is often harder than it should be. When reporting an incident, you should always include the time zone with respect to Coordinated Universal Time (UTC, for example '16:19:00 +0100' for British Summer Time) and whether timestamps are synchronised to an international standard. Without this a great deal of effort can be wasted.

### 6.3 Automated Processing of Logs

Logfiles can grow very large, and routinely scanning them by eye may not be possible. There are a number of computer programs that can help to monitor logs, and many sites have written their own scripts for this purpose. Such programs aim to identify patterns in the logfiles, and they can be very effective at identifying common, known problems. However a program is unlikely ever to be as good as a human at spotting new or unusual patterns. A compromise solution may be to use programs to filter out entries that are known to be harmless (though even these should be checked occasionally) and well-known problem patterns, and then to scan the remaining information by eye. As new patterns are identified they can be added to the known-good and known-bad filters. This process of tuning can be time-consuming, but is the most effective way to extract information from the logs. Raw logs should still be kept, subject to the issues relating to Data Retention in Section 2.2, as it can be useful to review them when new patterns are identified. It is very common for hostile activity to take some time to be noticed, and reassuring at that stage to be able to review older logs to determine when it actually began.

### 6.4 Graphing Activity

Humans are quite good at spotting patterns in textual information, but they are extremely good at finding them in graphical presentations. A highly efficient way to monitor the health of any system is to graph some appropriate measures of its performance. Graphs of network traffic levels such as those provided by the JANET Netsight system:

**<http://www.ja.net/services/network-services/netsight/>**

are fairly commonly used. Less frequent, but very useful, are graphs of, for example, number of requests to a web server, number of failed and successful logins to a network, or system idle time or memory usage. Regular patterns in these graphs will quickly become familiar to the operator. Once the system is well known, unexpected changes will often be spotted without conscious thought. Detailed investigation can then be done using the original logfiles.

---

## 7 Next Steps

This document has dealt only with the logs that can be recorded by individual computers and other systems. A great deal of useful information and early warnings can also be obtained by looking at computers and networks in combination. Any organisation that is concerned to protect its own systems and reputation should also be developing systems to monitor these systems. Two examples are given below of what can, and should, be done.

### 7.1 Network Flows

Networks can be characterised very effectively by knowing what flows of traffic are taking place along them. Flows can be classified by their source and destination IP addresses, or groups of addresses, and ports, along with the volume of traffic making up the flow. For example it would be quite normal in an open-access workstation room to see large numbers of packets coming from the HTTP ports of external machines into the workstation room – web browsing is a legitimate use for these types of systems – but much less normal to see traffic flowing out from the HTTP port on the workstations. The latter situation may indicate that someone has set up a web server on a workstation, either with or without authority, or that systems have been compromised and are being controlled by a remote intruder. Thus simple flow information can be effective in detecting both security problems and breaches of policy.

Network flows are often the only way to trace denial of service attacks, since these commonly use forged addresses to conceal their origin. If addresses are forged then an attack can be hard to trace at the IP level; instead information taken from routers and switches will be needed to determine which ports or interfaces are carrying the traffic. Such information can rarely be gathered from a single central point but needs effective, secure reporting and management to be set up on the network devices when they are deployed.

### 7.2 Intrusion Detection

Network flow monitoring examines where packets are going to and from; network Intrusion Detection Systems (IDS) examine the content of packets or search for patterns in time. For example an IDS might be configured to check HTTP packets for the commands used by well-known attack tools, or could detect that a particular IP address had sent packets to all the addresses in a range. IDS can be very effective at warning of known problems, though they are less good at identifying new, suspicious activity. Some IDS packages can take action to respond to a detected threat, either by blocking the hostile traffic or by targeting a response at the apparent source. However these options run the risk of denying service unnecessarily (any system, computer or human, will occasionally make mistakes) or of taking reprisals against an innocent party.

## 8 References

### 8.1 Information and Guidelines on Logfiles

LINX Best Current Practice – Traceability

[https://www.linx.net/www\\_public/community\\_involvement/bcp/traceabilitybcp\\_v1](https://www.linx.net/www_public/community_involvement/bcp/traceabilitybcp_v1)

Information Commissioner’s Employee Monitoring code

<http://www.informationcommissioner.gov.uk/eventual.aspx?id=437>

### 8.2 Data Retention

LINX Best Current Practice – Traceability

[https://www.linx.net/www\\_public/community\\_involvement/bcp/traceabilitybcp\\_v1](https://www.linx.net/www_public/community_involvement/bcp/traceabilitybcp_v1)

Data Protection Act 1998

<http://www.opsi.gov.uk/acts/acts1998/19980029.htm>

Directive 2002/58/EC on Privacy and Electronic Communications  
(to search, enter year=2002, number=58. Select ‘Directive’ radio button.)

[http://www.europa.eu.int/eur-lex/en/search/search\\_lif.html](http://www.europa.eu.int/eur-lex/en/search/search_lif.html)

Anti-Terrorism, Crime and Security Act 2001 (ATCSA)  
Text of the Act (see Section 2 on Data Retention)

<http://www.opsi.gov.uk/acts/acts2001/20010024.htm>

Home Office Page containing codes of practice etc.

<http://www.homeoffice.gov.uk/terrorism/reports/legisguidance.html>

Retention of Communications Data Code of Practice

<http://www.opsi.gov.uk/si/si2003/draft/5b.pdf>

### 8.3 Access to Data

LINX Best Current Practice – User Privacy

[https://www.linx.net/www\\_public/community\\_involvement/bcp/bcp\\_userprivacy\\_v1](https://www.linx.net/www_public/community_involvement/bcp/bcp_userprivacy_v1)

Regulation of Investigatory Powers Act 2000 (RIPA)  
Text of the Act

<http://www.opsi.gov.uk/acts/acts2000/20000023.htm>

Home Office Page (includes links to draft and final codes of practice)

<http://www.homeoffice.gov.uk/crimpol/crimreduc/regulation/>

The Regulation of Investigatory Powers (Communications Data) Order 2003 (Statutory Instrument 2003 No. 3172)

<http://www.opsi.gov.uk/si/si2003/20033172.htm>

RIPA draft Code of Practice (draft published in 2001)

<http://www.homeoffice.gov.uk/docs/pcdcpc.html>

RIPA section 22(4) form requiring disclosure of Communications Data

<http://www.homeoffice.gov.uk/crimpol/crimreduc/regulation/forms.html>

## 8.4 Analysing and Processing Logfiles

Log analysis resources

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

## Tell us what you think

Guidance Notes are produced and published by UKERNA for use within the JANET Community.

We welcome your comments on all aspects of this document and on any other UKERNA publication.

Please direct feedback to JANET Customer Service, at the address below, or e-mail us directly at:

documentation@janet.ac.uk

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

For further information please contact:

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

Tel: 0870 850 2212  
Fax: 0870 850 2213  
E-mail: [service@janet.ac.uk](mailto:service@janet.ac.uk)

Copyright:

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

Trademarks:

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

Disclaimer:

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

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

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

Availability:

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

This document is also available electronically from: [http://www.ja.net/documents/gn\\_logfiles.pdf](http://www.ja.net/documents/gn_logfiles.pdf)



