Logging network activity

Why?

To trace use of Janet, both legitimate and otherwise, helping to investigate and learn from security incidents. Whilst some network events are of a continuing nature, others may only occur sporadically and unexpectedly, and logging of activity can help us understand what took place in the past. Many networks now use Network Address Translation (*NAT*) or proxy devices that obscure the source of of a connection to the external world, which can prevent the timely investigation of serious incidents.

Any well managed network will maintain logs of particular events of interest. It is consistent with the intention of the Janet AUP and Security Policy that you should also keep enough information to trace use of Janet to specific computers or individuals in your network at the request of Janet (normally Janet CSIRT) or Law Enforcement Agencies (*LEAs*).

Questions

Your logs should help you answer questions about activity on your network. Questions from both within and outside your organization (eg from Janet CSIRT) that you may eventually face include:

- Who sent this e-mail?
- Can you confirm that a user of your network read this Web page on 23rd April?
- Are you able to confirm that a particular IP address downloaded this file?
- Please confirm can you isolate the source of this port scanning and make it safe.
- When was the last time that this computer was connected to the network?
- Who installed this unauthorised software on this computer?

What?

When Janet CSIRT ask you to investigate something, we will nearly always tell you which of your external IP addresses was concerned and the time. Usually we also know the remote IP address and destination port with which your address was communicating, or a URL or e-mail address identifying the remote entity. We may have other details such as the TCP or UDP ports used or the volume of traffic.

From the above information you should, through your logs, normally be able to identify the computer involved and, if appropriate, the individual user. There may then be a separate obligation on you to take some action, which will **not** normally include passing details of any user identified to Janet CSIRT. Where the police or another LEA ask you to trace some activity, they may be entitled to such information on production of a proper warrant or notice.

Which logs?

Logs you should consider generating and keeping include, but are not limited to:

- **Routers and firewalls**
  A router can export network flow records, which normally need processing to extract useful information, to a collector system. They can also keep records of packets matching access control lists. Whilst continuous logging of all traffic is recommended, often a specific access control list can be useful in answering a question. For instance, it would be possible to log all attempts to send e-mail directly from client computers, in support of a policy that they should not do so.
  Where a VPN (*Virtual Private Network*) is implemented in a router or firewall, the device can capture records of use, attempted use and authentication.
- **E-mail gateways**
  In many cases the header of an outgoing e-mail message includes enough information to trace the message back to a particular computer in your network, and (possibly in combination with other logs) to an individual user. There are, however, good reasons why an e-mail server might not put every detail into outgoing messages; you might not want to expose usernames or the private addresses of client computers.
  If your e-mail server does not include enough information in the message header, you should ensure that it logs it, and that you keep the log files concerned. Where the only information available is a report of e-mail abuse such as the header of a UBE message (*Unsolicited Bulk E-mail*), some details are suppressed by the reporting mechanisms and again may need to be recovered from your logs.
- **NAT gateways**
  Any NAT device (*Network Address Translation*) should log each new translation, so that given the public IP address in your network, source port for some traffic flow, and the time, you can recover the internal source IP for that traffic.
- **Proxy servers**
  For the same reasons as NAT, Web proxies should log each URL with the time and the internal IP address requesting it, so that any request can be traced to its source.
  Logging is likely to be complementary to that from NAT gateways. Similar considerations apply to proxies for other kinds of service.
- **Shared workstations**
  It is desirable that there is a record of who was logged in to any computer at any time. If a computer is used by more then one person (such as those in labs and classrooms) then it is essential. Typically the records will be generated by an authentication, login or RADIUS server.
  (Failing that, it should still be possible to record which class was using the lab at the time and perhaps a seating plan).
- **Network Access / 802.1x**
  Many sites are now using 802.1x and other Network Access Control technologies to authenticate network layer access to their networks. Records of authentication can be very useful in tracing use of public networks such as wireless networks back to an individual rather than a specific computer.
  If access to the network is not determined by user credentials, low-level information such as the MAC address of each computer or device connecting is important for tracing; logging from a switch or DHCP server of port allocations and leases may also be useful for internal management of the network. This is particularly evident in wireless

networks.

It is a feature of many of the devices above that they present an interface to the outside world behind which behavior in your network is hidden. The information logged will often resolve internally what is ambiguous externally.

It is important to make sure that any devices purchased are not only able to cope with expected loads and traffic, but are also able to continue to produce adequate logs at these levels. The requirements of the Janet AUP and Security Policy should not be considered secondary to system performance.

How much?

On the face of it the requirements above imply that the more you log, the better. At the same time it is important that the coverage is not excessive. To keep a lot of data for a long time is expensive, it can also be difficult to search for trace details in a large or poorly structured mass of data, and you are at risk of holding personal data without reasonable justification. Don?t log so much that you are overwhelmed; and do impose some structure on what you keep (with sensible filenames, consistent formats etc). It may be useful to age your log files gradually, so that you keep raw data close at hand for a short time in case of immediate operational need, and then archive it to another location in a standard format for long term storage. Otherwise you should document the procedures for responding to likely requests, so that even if the process is not quick, you have some estimate of the time it will take.
The UNIX tool *logrotate* is particularly useful for this role, as it can rotate and compress log files according to size or age. As most log files are repetative compression can save large volumes of storage at the expense of access and search speed.

When?

It is essential that you are able to match timestamps in your log information with the details supplied to you; so you should ensure that all devices generating logging have their clocks synchronised to external references. NTP (*Network Time Protocol*) is the normal way to maintain accurate time on at least one server with an external connection; other internal devices may use NTP, SNTP or a proprietary protocol locally, perhaps referring to this single internal server, to keep their clocks correctly set. You may find the Janet network time service useful.

Where?

You should consider having one or more dedicated logging servers, to which most or all computers and other devices in your network can send their log information. Management of the logged information is then much simpler.

Many machines (including those with UNIX-derived operating systems and most network devices and appliances) will be able to use a central *syslog* server, but you may need different servers for systems with proprietary applications or platforms.

Logging servers have distinctive security requirements. They must be continuously available and they should run as few other services as possible, to reduce the risk to critical data from failures and vulnerabilities in applications. Access to logged information (especially write access) must be carefully controlled, authenticated and audited, both in terms of the

individuals authorised to use the data and of limiting the range of addresses from which the log servers are reachable.
Bearing in mind that backups of log servers can result in short-term data being unwittingly retained for longer than intended, you should ensure that adequate backup arrangements are in place.

Who?

UK legislation imposes various constraints on the disclosure, non-disclosure, retention and destruction of logged information. One major consideration is that much information directly or indirectly identifies individual people and is subject to Data Protection legislation.

Read access to the logs should be limited to people in appropriate roles. The ability to modify or destroy logged information must be even more closely restricted. It may be appropriate to provide simple tools for initial searches by authorised administrative staff. Where the same logs are routinely used for both service management and administrative tracing it may be difficult to separate the levels of privilege provided, and procedures for all users of the logs should be clearly documented.

The police and other LEAs are not automatically entitled to the information you have. UK legislation specifies the notices and warrants which authorise or require you to release information, and you must insist on having such authority. You must also verify the identity of any officer or agency to whom you disclose information. If you have problems doing this, please contact CSIRT who are able to identify authorized LEA officers.

How long?

Janet CSIRT will not expect you to retain logs for more than three months, although you may wish to keep the information for longer (and certainly LEAs would like you to). Equally, you may not have the resources to keep all of it for as long as that, and for some things one month may be enough; but anything as short as two weeks would be considered irresponsible. Whatever you decide, it should be included in a written policy which also specifies how log information is to be kept secure and securely destroyed, and the purposes for which it is kept and may be used. One issue for the policy is the retention of log files on backup media, whose default lifetime may be longer than that intended for the information; again, you should document the position.

As noted above, UK legislation imposes various constraints on the disclosure, non-disclosure, retention and destruction of logged information. The above advice does not in any way alter the application of such legislation.

Requirements of Law Enforcement Agencies

LEAs may need information as evidence, in which case the chain of custody becomes very important. Where possible your procedures should anticipate that possibility. You may, for instance, be asked to demonstrate that the records you have made available could not have been altered.

Sometimes LEAs are looking not for evidence but intelligence, where the formal requirement is less severe; enquiries by Janet CSIRT are almost always of this simpler kind. The higher standard of proof for evidence may be useful if disciplinary procedures with your users

become difficult.

Supplementary sources of information

When there is a need to hold an individual accountable for some use or misuse of your network, sources of data other than logs from the network itself may also be valuable, such as CCTV images or records of physical access from card entry systems. This applies particularly to substantially open access computers such as those in teaching rooms or library areas. If one user fails to protect their account by logging out properly, or it is necessary for some reason to use shared accounts, such out-of-band information may be the best available and may be essential.

More information

- The Janet network time service [1]
- LINX guidance on traceability [2]
- Janet Logfiles Technical guide [3]

Links
[1] http://community.ja.net/library/janet-services-documentation/network-time-service
[2] https://www.linx.net/good/bcp/traceability-bcp-v1_0.html
[3] http://community.ja.net/library/janet-services-documentation/logfiles-technical-guide