

Case studies

Overview

A questionnaire was distributed to site security contacts in the summer of 2002 to find out the typical provision for incident response at Janet sites. The questionnaire looked at: staffing of incident response; what authority the incident responders had; what services were offered; and what contacts and relationships, both formal and informal, team members had to help them in their work. Thirteen responses were received. Clearly those who responded were self-selecting so may represent the more active organisations, but they include all sectors of the community and range in size from less than a thousand users to nearly thirty thousand and so represent a reasonable cross-section of Janet sites.

Almost all sites provided the core services of announcements and incident handling, though more provided co-ordination than hands-on incident response. All could call on local system and network administrators for assistance. Most had a written statement of authority and were permitted to disconnect a problem system from the local network, though fewer mentioned the power to suspend user accounts when these appeared to be the source of problems.

At almost all sites, the individuals responsible for computer and network security did incident response as part of their job. Most had other responsibilities. A few organisations had one or two full-time incident response staff, but usually the work was shared out among a group. All incident response teams had a published e-mail address for reporting problems and nearly all had a telephone number for incident reporting and more general helpdesk functions. Only one in three teams provided out-of-hours cover by having staff on call.

Most sites had contact details for the individuals responsible for servers (in many cases these were the same people who responded to incidents). Contacts in departments and sponsored connections were also common where these ran their own networks or computers. However a list of contact details for every system connected to the network was still an unfulfilled ambition for all but the smallest sites, where a single person appeared to be responsible for everything.

Most teams offered additional services beyond the basic list. Education and incident tracing were common, and many teams were involved more or less formally in providing input to organisational processes by performing consultancy and risk analysis. A high proportion of the respondents not only had the expected relationships with technical staff but also with policy makers and senior management. In these organisations incident response is not just about fixing current problems but gives additional benefits to the organisation by informing planning and policy processes.

Three sites were asked to act as case studies and provided the information for the following sections. Particular tools are mentioned only as examples, not recommendations. In some

cases these may have been superseded. The [Janet CSIRT website](#) [1] contains up to date lists of security tools.

Case Study 1

At a medium-sized college, effective use of technical measures has reduced incidents to the level where the most common reported problem is students accessing inappropriate material. The site connects to Janet through a router implementing a default-deny policy. This greatly reduces the exposure to external attack. Desktop workstations and servers are centrally managed so updates to software can be implemented quickly and easily. This reduces the number of vulnerabilities present within the network. Configurations are discussed with users to ensure that they meet their requirements. Network flows are measured using Multi Router Traffic Grapher (MRTG) and a central logging server collects error messages from all network devices and generates e-mail alerts for serious problems. This often allows problems to be detected and resolved before they become an issue for users.

All problem reports, including security incidents and exception reports from logfiles, should be directed to a central helpdesk, by phone, e-mail or in person, however not all users are yet used to contacting the helpdesk rather than individual technicians directly. Problems are allocated by the helpdesk to an appropriate person among the six IT staff, either technical or management depending on the nature of the problem. Any of these staff may therefore be involved in handling a security incident. A formal job tracking and escalation system is being implemented that is expected to lead to more efficient communications, better use of staff effort and a faster response for users. In resolving problems, staff have access to data about computers, software and users through the software management system as well as to address allocation, web proxy and event logs.

Inappropriate use of the network leads to formal procedures against the users responsible including suspension of network accounts. Procedures are also being developed for other breaches of the organisation's policy though to date these have seldom been required.

The small number of technical security breaches means that when these occur they can be investigated in depth with the aim of identifying the full scope of the attack and any other systems affected. Thanks to the central software configuration management this should allow other vulnerable systems to be identified and remedial and preventive measures to be deployed across the organisation.

Future plans include improving the logging and detection of network and traffic problems by implementing a full firewall and possibly intrusion detection systems as resources permit. The existing logging system may also be improved. These should allow a secure and reliable service to be provided with even less disruption to legitimate users.

Case Study 2

The second case study is a medium-sized research-led university. Here too a default-deny policy on the site's connection to the Internet has greatly reduced the number of incidents. In addition to the central computing services, a number of university departments run their own systems that are not directly under central control. The incident response team therefore combines the roles of incident response for central systems with incident co-ordination for those managed by others.

The university has one full-time network security officer, with three other members of the computing service available for incident response work. The three staff do not have time allocated to incident response work, but provide cover for absence and busy periods. Most incidents are reported by e-mail to a distribution list, which ensures that all four of the incident response staff receive the report. Telephone reporting is to the network security officer, with the phone numbers of other staff published internally as backup. Work is allocated informally among the team, whose offices are close to one another.

The main aim of the incident response team is to remove problems as soon as possible and prevent recurrence. For a rapid response, systems may be disconnected from the internal network by disabling their local router or switch port, or external traffic may be blocked at the site router. Where local accounts are the source of the problem, these may be disabled on central servers. Lack of time usually prevents investigations doing more than identifying the method of attack and planning how to prevent it in future. Diagnosis may involve records from the central network management systems or logs from the central logging host, which central systems are configured to use. If problems occur on systems managed by departments these are usually blocked from the network and the department asked to resolve the problem before re-connection is authorised. If a department does not have the necessary skills to resolve a problem, incident response staff may be available to provide assistance. Occasional talks on specific security issues are given to departmental system administrators.

The level of security incidents has not yet required a more formal incident tracking system than is provided by the incident e-mail list, though the university does have an existing job tracking system that could be adopted if necessary.

One area where the university has already implemented a full tracking process is for investigations that may have legal implications, in particular where there are requests to access user filestore or complaints about published content on servers. Before these cases are investigated, approval must be obtained in writing from the appropriate Head of Department. Senior administrative staff of the university may also be involved at an early stage, for example when decisions to remove content or provide information to law enforcement agencies are required. The computing service departmental secretary keeps written records of these procedures in case of future challenges.

The incident response team also provides some pro-active services, in particular distributing security advisories from other sources. Where necessary an introduction is added to these to explain them in the context of the university. If vulnerabilities affect software that is used to provide public services then the owners of these servers are contacted shortly after the advisory is distributed to confirm that they are acting on it. If necessary, services can be blocked at the site router until a server has been secured. On occasion, the incident response team have used network and wireless scanners to identify potential weaknesses in the site network and any problems identified by these methods are followed up. A traffic graphing

system, implemented using MRTG for network management purposes, has already proved very effective in identifying compromised systems by highlighting unusually large traffic flows originating from workstation systems.

Future plans include improving the facilities for reporting by telephone, as well as finding time to make better use of an intrusion detection system to identify problems more quickly.

Case Study 3

The final case study is a large university whose policy permits unrestricted Internet access to most computers within the organisation. Priority is therefore given to detecting any hostile traffic as early as possible and responding to limit the extent of any damage that may be caused. One full-time and one part-time member of staff spend a significant amount of time on incident response work, though this is not their only job. They can call on a number of members of the central network and systems team for additional assistance and also on departmental system administrators when there is a problem with a system that is not centrally managed. The network management team assign one duty person to respond immediately to problems. Problem reports arrive by e-mail to an incident reporting address or by phone, either direct to the incident response staff or via the site switchboard or helpdesk. The person who receives a report will normally manage that incident; if they need help then they may communicate verbally or by e-mail with the others in the team.

All incidents are recorded in a ticketing database with a record of the evidence and steps taken to resolve it. This can be displayed or printed in the form of a job sheet for each incident, including any e-mail messages sent and received, which allows any member of the team to work on the incident knowing what has already been done. A standard header includes the start and end times of the incident and the identity of the person responsible for the next remedial action. The database can also generate standard e-mails to report problems to other teams or to the owners of internal systems. Reports of network probing are also entered in a standard form into an Excel spreadsheet; this can be sorted to determine whether a particular IP address or range has a history of such activity and what action is therefore appropriate.

Prompt analysis of network flows is particularly useful in detecting problems early. The bulk of this process has been automated by a program that analyses central router logs to look for internal or external Internet Protocol (IP) addresses that have attempted to make connections to large numbers of other hosts or ports in a particular time period. This is a very unlikely pattern in legitimate use, but is common when a system is scanning systems or networks for vulnerabilities. When this pattern is detected an alert is sent to the incident response team by e-mail. Further information on network traffic can be obtained from flow logs collected by Network Traffic Meter (NeTraMet). Key systems also run ZoneAlarm® or Snort as intrusion detection systems and these logs are also inspected routinely for attacks that evaded the flow monitor checks. When investigating particular incidents logs from applications such as web servers and individual systems may also be used to learn as much as possible about the attack.

If a local system is identified as the source of an incident, the person or department responsible for it will be contacted and asked to resolve the problem. If the system appears to be a threat to local or remote networks then the duty network manager will be asked to block its traffic at the appropriate router until the owner has confirmed that the problem has been

identified and removed from the system. The incident response team may probe the system using a network scanning tool such as SAINT™ to determine what vulnerabilities may be present. They may also help the system owner to find and resolve the cause of the problem.

If an external system is reported as attacking the university, this will usually be reported to the abuse contact for that network and also to Janet-CERT. If the system is particularly persistent or appears to be targeting the university rather than scanning randomly, then the external site may be contacted by telephone. Where this appears necessary to protect the university systems, router blocks may be installed to prevent return traffic to the hostile address.

As well as these response services the team also distribute advisories and technology updates within the university and perform some security audits. These are areas they would like to develop as time permits.

Source URL: <https://community-stg.jisc.ac.uk/library/janet-services-documentation/case-studies>

Links

[1] <http://community.ja.net/library/janet-services-documentation/janet-csirt>