Tools

Document Reference: GEN-DOC-005 - *Please see here [1] for additional document control information.*

3. Tools

As stated, the following sections describe a number of tools and techniques that can be used to reduce the risk of misuse of the network. They are presented here in the context of providing network access for guests, though many of them can also be used for local users. None of the tools can make misuse impossible: each section describes which risks can be reduced by a particular tool and which risks may remain or be increased.

As discussed above, a successful system for guest access is likely to contain a balance of preventive measures – those that constrain use of the network to make deliberate or accidental misuse more difficult – and reactive measures that increase the likelihood of identifying the person responsible for the misuse and, perhaps, holding them to account. Guest users present a particular problem in this latter respect: since many will have no formal link with the organisation, enforcing sanctions on them is likely to be more difficult and less effective than for local users. Any guest access facility is therefore likely to require at least some preventive measures with reactive measures being used to back these up. As stated in section 2.2, the appropriate balance between these, and the choice of measures, will vary depending on the number and nature of guests and the services that the organisation wishes to make available to them. A wide range of examples are contained in the case studies in the next section.

3.1 Supported Equipment

One of the most significant decisions is whether to allow guests to connect their own computers (many will have either a laptop or an Internet-enabled PDA or mobile phone) or whether to provide fixed terminals for them to use. The latter might include, for example, fixed kiosks in libraries providing controlled access to catalogues or databases licensed for walk-in use, as discussed in the HAERVI study, or possibly shared use of workstations in open-access rooms. Each has advantages and disadvantages.

Providing fixed terminals allows all guests to connect to the network whether or not they have their own computers. It also allows the organisation to control the configuration of the computer and how it checks for misuse. With most modern operating systems the organisation can, for example, ensure that up-to-date virus detection software is running and that no new software can be installed. However fixed terminals also limit the number of guests that can be supported and the locations where they can connect. For those used to their own computers, an organisation-provided working environment may be unfamiliar and may prevent

the use of software such as virtual private networks or authentication tokens that some guests need to gain access to systems at their own organisation. Finally, as on any shared computer, measures must be taken to prevent the activities of one user affecting those who follow – if one user infects a computer with a virus or spyware then this may remain on the computer to harm subsequent users. This is a particular risk where guests share a computer with local users: by persisting, a virus introduced by one user may later gain additional access under the permissions granted to another. Allowing a guest to share a computer used by a single local user is particularly risky since such computers may well have sensitive or critical information on their local hard disks which it will be very hard to prevent the guest from accessing.

Providing facilities for guests to connect their own computers, often using wireless networks, has its own benefits and risks. Those users who have their own equipment can use their familiar working environment and tools, and the issues around shared equipment do not arise. Wireless networks can be located more flexibly than fixed terminals, though providing complete wireless coverage across a campus is likely to be expensive and their restricted bandwidth will eventually limit the number of simultaneous users who can receive an acceptable service. However guests without their own computers will be excluded and the organisation has no control over the equipment that may attempt to connect to the network. Guest-owned computers may well introduce viruses or security vulnerabilities and may interact in unexpected ways with the existing infrastructure. Guests may expect the organisation to provide support for their equipment, at least to get it connected to the network, even though it may be completely unfamiliar, brand new or many years out of date. Providing support may be technically difficult or impossible, especially if guests do not have administrator rights on their own computer. This, or licence issues, may prevent the installation or use of the organisation’s standard software on the guest’s equipment.

If the organisation performs proactive scans or tests of equipment connected to its networks it is important to ensure that guests are aware of this and any impact it may have on their computer. To avoid accusations of unauthorised access or modifications to the guest’s computer it may be best to include authority to run such tests in the policy that guests accept when they connect their equipment to the network. Many organisations also remind guests of the need to protect their own computers (for example by installing patches, anti-virus and firewall software) since firewall and other protective arrangements may be different from those at the home organisation.

3.2 Physical Controls

In some cases it will be possible to physically restrict the ability to connect to the network. Wired networks have always implicitly used the fact that connections can only be made where there is a socket to exclude unauthorised users simply by preventing them entering rooms in which there are live sockets. This can also be used to manage some types of guest access; for example when staying in residences a guest may be told that they will be held responsible for any traffic to or from thenetwork socket in their room.

Physical controls can also be used for wireless networks, though they are significantly less effective and should always be supported by other control techniques such as port and bandwidth restrictions, monitoring and incident response. Obtaining a wireless network signal outside a building is usually much easier than getting a network cable through a brick wall! However, for some buildings a wireless survey around the outside of the building may indicate

that the signal is sufficiently well contained for this to be a helpful control measure in a wider plan. Locating access points away from external walls and managing transmission strengths make it more likely that this will be achieved, but unless a building is actually designed to contain radio frequency transmissions (such protection is often referred to by the military term Tempest and is very expensive) it is always likely that sufficient signal will escape to allow a determined attacker to connect by using a more sensitive aerial. Of course, even if it is possible to limit the spread of the wireless signal beyond a building this is of little benefit unless physical access to the building is also controlled.

3.3 MAC Addresses

Every network access device, such as an Ethernet or wireless LAN adaptor, is allocated a unique Media Access Control (MAC) address, often referred to as an Ethernet address, by its manufacturer. Since a MAC address should uniquely identify a particular piece of network hardware, it is sometimes suggested that a list of MAC addresses can be used to control which devices are permitted to connect to a network. The ability to do this is included in many wireless access points and network switches. For a home network, where there is a relatively small and unchanging set of devices permitted to connect, controlling access through this mechanism may be a useful precaution but for guest access it may be more trouble than it is worth.

To implement this first requires a MAC address for every authorised guest device to be obtained and configured into all access points and network edge switches where the guest may connect. Doing this in a secure fashion for a large number of guests is a considerable organisational and technical challenge and may result in a significant delay for the guest in obtaining a network connection. Determining the MAC address of a particular device is not always easy. Second, many devices, such as PCMCIA wireless cards, are designed to be swapped between different machines so the link between a MAC address and a particular computer may well change. Even desktop computers often have their Ethernet cards changed if a fault is reported. Finally, many devices allow their MAC address to be changed by relatively simple software; a fact that may be exploited by a malicious user who may well be able to monitor traffic on a network (particularly a wireless one), identify a MAC address that is being granted access and change their own address to be the same, thus defeating the intended protection measure.

3.4 SSIDs (wireless networks only)

Wireless networks are identified by a Service Set Identifier (SSID); for example it is a requirement that networks providing eduroam use the SSID ‘eduroam’.Where an organisation provides a number of different wireless services (for example one for local users, one for guests and one for other visitors), different SSIDs may be used to let users select the appropriate service. However, since there is nothing to stop users choosing the ‘wrong’ SSID, some additional check of their entitlement to use the service they select is likely to be required.

It is also possible to configure access points not to advertise their SSID; however this is another measure that is better suited to a home network than a guest facility. With a ‘hidden’ SSID some means will be required to tell authorised guests what the SSID is and they may then have difficulty configuring their computers to use it. There may also be performance

penalties. Nor does a hidden SSID provide much protection against misuse: a passive attacker can simply wait until an authorised user connects and discover the SSID from their traffic, while an active attacker can forge sufficient management frames to make the access point reveal itself.

3.5 Shared Secrets

A shared secret is something that is known to authorised guests, but not to other visitors, which the guest can use to prove their entitlement to use a service. Although these have some similarities with ‘hidden’ SSIDs, they have two advantages in that they can be made easier for guests to enter and harder for nonguests to discover.

The simplest form of shared secret is a password or other piece of information that is given to every guest. For example some conference facilities give conference organisers a code that they can then share with all delegates. However, if every guest uses the same code there is no way to distinguish individual guests should it be necessary to disable a particular guest user or contact them to deal with misconfigurations or misuse. For this reason it is often better to have a separate code for each guest so that individuals can be distinguished. Some conferences include individual codes in their delegate packs (one printed the codes on delegate badges, which is probably not a good way to keep a secret!); others issue codes on request, recording which code is issued to which delegate. Shared secrets, particularly those that are known by many people, should be changed frequently to stop guests continuing to use the facility after they have left, or sharing the code with others who were never entitled to use it. For the same reason, predictable algorithms such as ‘today’s date’ should also be avoided.

There are two common ways to implement shared secrets. On wireless networks, several encryption methods support pre-shared key (PSK) options. When guests first connect to a network that uses one of these, their wireless client software will generally prompt them to enter the shared secret and then offer to store it for future use. The other approach, which can be used on both wired and wireless networks, is to block all packets until the user starts a web browser, then re-direct the request to a server that invites users to enter the shared secret. This technique is commonly used to control access to commercial networks in hotels and elsewhere, and a wide range of products and open source implementations is available. On wireless networks, however, there is a risk that a fake access point virtually indistinguishable from a real one can be set up and used to collect secrets (see our factsheet ‘Safe Use of Web Re-Direct Networks’ (PB/INFO/058a)). Particular care is needed if a Web Redirect system is used to access to systems and files that require logon passwords to be entered as well as just the network.

To ensure that secrets stay secret, it is important to use technical measures to protect them from eavesdropping. Wireless technologies should always protect their pre-shared keys using encryption but there has been a history of problems with poorly designed, and poorly implemented, systems. In particular, the original WEP encryption should never be relied upon (see our factsheet ‘WEP, WPA or Other’ (PB/INFO/071)). Cracking WPA PSK (a simplified form of WPA suitable for home wireless networking) is known to be at least theoretically possible, especially for short secrets. The successor to WPA, WPA2 PSK is currently believed to be secure against anything other than a brute-force attack of guessing all possible secrets, though it may not be supported by all devices. It is good practice to choose secrets that are

hard to guess and to change them frequently. Where secrets are entered through a web page, SSL (i.e. an https: URL) should always be used to ensure that they are strongly encrypted as they pass across the network.

3.6 IP Addresses

When guests connect their own device to a wired or wireless network, the network will normally provide the device with an IP address and other configuration information using the DHCP protocol. Since IP addresses can be used to control what guests can do and, in particular, which resources and systems they can connect to, it is useful to assign a different IP address range to distinguish devices used by guests from those used by local users. Kiosk devices and open access workstations may also be placed in the ‘guest’ address range when appropriate, though for workstations used by both guest and local users this may require technologies such as IEEE 802.1X that can assign network configurations after determining the user’s identity and authorisation, rather than before.

Depending on the organisation’s existing address allocation and the facilities it wishes to provide, guest addresses may either be taken from a separate pool within the public IP address space or from the private RFC1918 addresses (e.g. 192.168.x.x, which is commonly used by commercial network providers). In the latter case, providing access to Janet and the public Internet will require either proxies for all supported protocols or else a NAT/PAT device to convert the assigned private addresses to public ones.

If guest IP addresses are distinguishable from those of local users they can, for example, be configured into router Access Control Lists to block direct access by guests to parts of the organisation’s LAN to which they should not have access, or into web and other servers to prevent unauthorised access to documents or services intended only for local users. If guest IP addresses cannot be distinguished then there is a risk that both local and remote resources may grant them unintended access. If guest machines are placed into the local IP address range then it is particularly important to ensure that the policy for connecting them includes the right to subject them to any of the organisation’s normal monitoring and scanning processes, with no liability for the consequences. Some monitoring systems might be confused by the presence of an unrecognised device or configuration, while scanning a device of an unknown type may cause it to crash or respond aggressively to the scanning network.

3.7 VLANs

As well as separating guest users at the IP level of the network (layer 3) it may also be useful to separate them at the network layer (layer 2) by using Virtual Local Area Networks (VLANs). This separates their network traffic so that it can be managed separately from that of local users; it also makes it easier to contain the impact of any deliberate or accidental unusual configurations, virus infections, etc. Some network switches allow VLANs to be configured so that guest machines can only send traffic to and from the VLAN router, not to other hosts that may happen to be connected to the same VLAN. This provides further protection for individual guest machines since they need not worry about other devices connected to the network at the same time.

VLANs can be used to implement fine-grained controls, for example by placing different types of guest into different VLANs, each with an appropriate level of control of what the guest is permitted to do. VLANs can be assigned using IEEE 802.1X technology by authenticating either the device or its user (a default VLAN can be used for guests who do not authenticate); alternatively many wireless access points allow a different VLAN to be assigned to each SSID that is supported (see section 3.4 above). For example guests who connect and successfully authenticate to the ‘eduroam’ SSID might be assigned a less restricted VLAN, since by authenticating them their home organisation has accepted responsibility, under the Janet eduroam Policy, for their activities (see section 3.11 below).

3.8 Port Restrictions

Depending on the confidence the organisation has in its guests’ good behaviour, and the range of services it wishes to make available to them, it may be appropriate to reduce the opportunities for them to cause harm by using a router or firewall to limit the TCP and UDP ports that they can make connections to and from. Port numbers do not completely guarantee what service is running but they can still be a useful control measure. This is likely to be much easier to implement if guests are connected to a separate LAN or VLAN, of course. Where such a guest-(V)LAN is implemented, two different sets of port restrictions can be implemented – one between the guest-LAN and the Internet and one between the guest-LAN and the organisation’s internal LAN.

Restrictions between the guest-LAN and the Internet are most often used to control the harm that a guest can cause to external sites or users. This should still be a concern, even though it does not directly harm the organisation, because any harmful traffic will be traced back to the host organisation (being associated with an IP address registered to that organisation) and the organisation’s reputation may be damaged. This can become a technical issue because external e-mail and other services may block access from IP address ranges that they regard as potentially harmful, so an organisation may find itself unable to send e-mail because of the activities of a guest or a guest’s computer. Jisc's eduroam service has identified a list of ports and services that aim to achieve a balance between the desire to provide services to guests and the risk that those services will be abused. The full list of ports, which provides web browsing, file transfer, e-mail, VPN and terminal services, can be found in the eduroam Technical Specification; the controls available to address the remaining risks are in the factsheet on eduroam Security Measures (PB/INFO/067). Note that this list was developed for a service where guest users are authenticated and where their home organisation has agreed to deal with any misuse, so guest facilities without those protections may wish to provide fewer services.

As discussed in the factsheet ‘Connecting Wired and Wireless Networks’ (PB/INFO/068), filtering is recommended between guest and local LANs, especially if the guest LAN offers wireless, rather than wired, connections. Indeed, provided DHCP, routing and any authentication and logging services are accessible from the guest LAN, there may be no need to allow packets to pass between these two LANs at all. Internal networks are often assumed, whether correctly or not, to be accessible only to local users, so allowing guests to access them may well break security assumptions. One facility that may be useful for travelling guests is printing; however since every sheet of paper and toner has a cost associated with it for the host organisation they may wish to have a human controlling access to this. Asking

guests to e-mail their documents to a local user, or to copy them to a USB memory stick, may well be more acceptable, though both mechanisms need to be protected against the introduction of viruses.

3.9 Bandwidth Management

Some types of malicious or undesirable activity by guests or their computers can be reduced by reducing the bandwidth available to them. Even a simple limit on the throughput of the router connecting the guest LAN should prevent a guest from saturating the organisation's own access link, though it may still allow significant amounts of spam or DDoS traffic to be sent, apparently from the organisation. More complex rules involving rate or volume limits per-port, per-destination or per-user are supported by some routers and traffic management devices, though remember that too low a limit may also prevent a guest from downloading a security patch or anti-virus update!

3.10 Network Access Control/Trusted Network Connect

Various commercial and open source systems are now available that allow some aspects of a computer's security precautions to be checked before it is granted full network access. Rather than either granting or denying full connection, some of these systems include an intermediate option of connecting an insecure computer to a subnet from where it can only see the patch and anti-virus services it may need to update itself. A number of terms are used to refer to such systems, including Network Admission Control (NAC) and Trusted Network Connect (TNC).

Although these systems may appear attractive for guest facilities they often assume that the organisation has standard hardware and operating system configurations, or that a client can be installed on the computer to be checked. Guests bringing their own devices are unlikely to satisfy either of these requirements. A standard is being developed for communication between a client and the device that checks security posture and authorises connection, and eventually compatible clients may be present on most laptops and other portable devices, but at present most systems are proprietary.

Organisations using NAC or TNC will probably always need to set policies for their guest networks that are more tolerant of different or uncooperative devices. Insisting on a particular operating system or anti-virus product may be possible on the local LAN but it makes little sense to offer a 'guest' facility that, in practice, refuses to connect anything that is not a device owned and configured by the host organisation.

3.11 User Authentication

The previous sections have described tools that reduce the likelihood of misuse occurring, either by preventing problem systems or users from gaining connectivity in the first place or by limiting the connectivity that is available so as to make it less open to misuse. Another technique is to deploy systems that allow a particular user to be identified so that they can be traced and the problem resolved individually. Being able to find a user permits a range of responses such as informing them of the need to secure their computer, telling them that a particular activity is not permitted under local policies, or holding them accountable in some

way for any harm they have done. Having said this, accountability may be less effective for a guest than a local student or member of staff. Even a complete ban on using the network may have little effect unless the user had planned to be a regular guest. Eduroam, described later in this section, provides a solution to this problem for guests from other Janet-connected, or international peer, organisations. Even without accountability, the ability to identify a particular user can be a useful tool in providing protection for the organisation since it allows a problem to be contained by disabling just the problem user without having to turn off guest access across a whole area or site. Even if (as will often be the case with wireless access) it is not possible to locate a particular guest physically, disabling their network connection will often cause them to visit a helpdesk to have their connectivity restored, at which point the problem can be addressed. For this to work, disabling an account must also terminate any established network connections. Merely preventing the guest logging in next time they visit will be much less effective!

Identifying an individual user requires that each guest must have their own username and password. As the case studies below indicate, there are many ways that these credentials can be given to guests. However there are essentially three different ways to manage usernames, each method having its own requirements and benefits: a pool of static guest accounts, accounts for individual guests, or the guest's home account. Each of these is now considered in turn.

3.11.1 Guest Account Pool

With this option the organisation maintains a fixed pool of guest accounts, each of which may be loaned temporarily to a particular guest if authorised by the guest's host or event organiser. A record should be kept of the authorising person since they should be best able to find the guest if there is a problem. This option requires the least electronic infrastructure: the guest accounts are created at one time and an individual guest only needs to be given the current password for the account they are loaned (this may, as in some of the case studies, be a paper-based process). Pool account credentials should only be valid for a short, fixed period so that it is clear when they can be reissued to a different guest. It may not be necessary to allow the guest user to change the password and this should certainly be avoided if it may result in the account lifetime being extended.

To ensure that there is no conflict between successive users of the account, the password must be changed and any associated storage or configuration changes deleted before the account is next issued to another guest. If storage is provided, the guest should be informed that any information will be deleted at the fixed time when the pool account expires.

Pool accounts are often used for training courses or guests who are only present for a well-defined, short period, and who will generally only require access to a limited range of facilities (for example 'Internet access' or 'training room login'). Pool accounts can be used for guests from anywhere and any type of equipment and network access.

3.11.2 Individual Guest Account

When guests have a longer-lasting relationship with the organisation, it may be appropriate to create an individual account for a particular guest. This may be done centrally on the request of the host member of staff, or Identity Management systems may be configured to allow staff to create their own 'sponsored guest' accounts. In either case a record should be kept of the

link between the guest and their sponsor. Sponsors must be aware of their responsibilities to authorise or create guest accounts in accordance with Janet and local policies. Individual guest accounts allow individual settings of lifetime, storage, facilities accessible etc. to provide anything from simple Janet/Internet access through the organisation's wireless network facilities to something close to a full local user account. Organisations need to ensure that they do not set up accounts in breach of licence conditions for software or other resources, for example by giving non-members of the organisation access to resources licensed only for members (see section 3.6 above for more details). They should also be aware of any blame or liability they may attract as a result of the guest's activities. The guest should be informed of what privileges the account does, and does not, give them and what the conditions of use are.

Individual accounts may be used for a wide range of requirements but are especially suitable where there is a longer-term collaboration between the host organisation and the guest, lasting across a number of visits. At the other extreme, Internet access during conferences may use either individual accounts or pool accounts, depending on the duration of the conference and the flexibility of provision required.

3.11.3 Home Account (eduroam®)

The eduroam service, and the international eduroam® federation of which it is a part, have been designed to provide simple and trustworthy network access to guests from other organisations that are members of eduroam®. For these users, the eduroam infrastructure can replace the need to create or issue individual guest accounts. When using eduroam, guests login with credentials based on their existing home username and password. These credentials are forwarded securely using RADIUS technology to the home organisation (the visited organisation cannot see them) for it to determine whether this is a user for whom it is prepared to take responsibility. The home organisation then returns a response indicating whether or not the guest should be granted network access.

Under the eduroam Policy, by returning a response indicating that its user should be connected, the home organisation promises that it will assist in investigating any misuse and will deal with such misuse as a breach of its own Acceptable Use Policy. This means that, unlike the other ways of providing guest access, eduroam can provide effective sanctions for misuse even after the guest has left the host organisation, and there is no need for a local sponsor to take responsibility for the guest.

The main purpose of eduroam is to enable member organisations to offer each other's users the ability to connect to their home organisation and the Internet. At present it is less suitable for providing individually controlled access to local facilities, though there have been some proposals to develop this area. For this type of facility, and for guests from outside the eduroam® members, individual or pool guest accounts are likely to be more suitable.

3.12 Usage Monitoring

As with any network, early detection and containment of problems is an excellent way to reduce their impact. Network flow monitoring and intrusion detection systems can be as useful on the guest network as on the internal network. Indeed, since the flows on the guest network will normally be smaller and simpler than those on the internal network such tools may even

be more effective and may detect attacks against the organisation that would otherwise be lost in the network ‘noise’. Flows on the guest and local networks should be collected and analysed separately since their expected patterns are different. If the guest network looks the same as ‘normal’ local network traffic then this probably indicates a problem! If router or firewall rules are used to restrict access to and from the guest network then any hits on these rules should be monitored, as they may indicate either attempts to misuse the guest network or a guest who needs a service that is not available.

Trends in use should also be monitored over time to ensure that the guest facility continues to provide the services required by guests. Types of connectivity, capacity and services expected have changed considerably in the past and user expectations will continue to develop alongside changes in technology and habits. These changes may have nothing to do with education or research: for example conference delegates developed an interest in local printing facilities when airlines offered the possibility of checking in online, though this may decline again as airports deploy systems that can read boarding passes off the screens of laptops, PDAs and mobile phones.

3.13 Incident Response

Monitoring the network to detect problems will be much less effective if the organisation cannot respond quickly and effectively when problems occur. If misuse of guest facilities is not resolved it can harm other guests, local users and the organisation’s reputation in the Internet community. For example many sites with effective perimeter defences have discovered that their major source of worms and viruses was guest laptops. Once carried around the perimeter, these could spread both within and outside the organisation. In other cases, guest laptops have advertised themselves as DHCP servers, boot servers or high-priority routers, any of which can be highly disruptive for the organisation’s core operations.

Earlier sections have discussed some of the technical controls that may be needed to contain and resolve a problem. Ideally a guest facility should provide the ability to disconnect a single guest user or device immediately; it should also be possible to close down the guest facility if it appears to present an imminent threat to the organisation or to the external network.

Policies and procedures are needed to decide when and how these controls will be used. Since a problem may need to be contained rapidly (ideally within minutes of being detected), these should be developed and agreed in advance. For almost all organisations the guest facility will be less critical than the internal LAN or the external link to Janet so a policy that reduces the guest service where necessary to protect either of those should be acceptable. The policy should also state what support or alternative connectivity is provided to a guest who inadvertently causes a problem, remembering that not all guests will have the authority or ability to manage their own devices. In some cases the only option will be to tell them to have the device checked by their home organisation.

The policy should also ensure that incidents are used to update the organisation’s assessment of the risks presented by its guest facility and the controls that are required. Incidents that are resolved quickly and effectively with little disruption generally indicate that the system is working well; however, incidents may also reveal that either the risk assessment or the chosen control measures, both of which inevitably rely on estimates of a changing threat, need to be revised.

Source URL: https://community-stg.jisc.ac.uk/library/janet-policies/tools

Links
[1] https://community.ja.net/library/janet-policies/network-access-guests-technical-guide