NHS/JANET Architectures -
a discussion paper for UKERNA and NHSnet

Andrew Cormack
UKERNA

PB/CERT/002 (02/11)

Contents

1. Introduction

NHS and JANET sites, for very good reasons, generally run entirely separate networks. Likewise at the national level there is no special treatment, with each network regarding the other simply as part of the public Internet. However there are a large number of people, both staff and students, who routinely work on both sides of this divide and who suffer inconvenience and expense as a result of it. For example staff who have legitimate reasons to connect to both networks may have to have duplicated computing facilities to do so and students may be unable to access essential teaching materials if they are not in the 'right' location. This report considers whether there are ways that these people's work might be made simpler while still meeting the operational and security requirements of each network.

This report takes as fundamental principles that any working arrangements should not significantly reduce the security of the networks involved and, furthermore, that the controls to ensure this should be in the hands of the organisations that bear any risk. At this stage the report considers only possible designs for network services, and not the implementation of those designs: where software or hardware packages are mentioned these are simply as examples of a particular approach to a problem.

The report considers three major groups of users and proposes two models to meet their requirements for access to information. For sensitive information, with strict requirements for authentication and needing fine-grained access control, a three-tier thin client/server architecture is proposed. This model incidentally makes access to legacy 'fat' client/server systems straightforward. For lower security applications, such as access to learning resources, diary systems or non-sensitive electronic mail, centrally provided and managed web gateways can provide services to a large number of remote users. Although this study concentrates on particular groups of users, the same facilities may also meet the requirements of other users provided appropriate policies and network connectivity can be provided.

Finally, technical measures alone cannot provide effective security against human ingenuity or carelessness, so it is essential that those who have access to particular networks and data follow the policies and practices set out for those networks and data at all times.

2. Users

This study addresses users who, for the purposes of their work, have access to computing facilities both at NHS and JANET sites. There appear to be three main groups of such users, each group having a particular set of requirements.

1. Students of JANET-connected organisations who spend some or all of their working time at NHS sites.
2. Students in JANET-connected organisations who need to access networked teaching materials or other resources that are hosted within NHS sites.
3. Joint staff of NHS and JANET organisations.

Since the requirements of these groups are different, there is no single solution that will address all their needs. The designs discussed here for each case are entirely independent and could, indeed probably should, be considered separately for implementation. It should also be noted that although this report specifically addresses the needs of a particular set of network users, the generic solutions proposed can be applied to many other different groups of users who have requirements to access facilities in one organisation from another. Hence the same design principles can be used when addressing the needs of other health care workers working inside and outside NHS sites.

2.1 Requirements

The three main groups of users, and their requirements, are as follows:

1. Students of JANET-connected organisations who spend some or all of their working time at NHS sites. These may well require access to networked computing facilities at their home organisation, for example to send and receive e-mail, access controlled resources such as electronic journals, or use administrative facilities. These remote students may indeed be more reliant on electronic contact with their organisation than local students who can have direct, as well as networked, access to services and resources.

2. Students in JANET-connected organisations who need to access networked teaching materials or other resources that are hosted within NHS sites. These are in the inverse situation to those considered above although, of course, the same individual may be in both positions at different times. Medical students may well need, and be entitled, to access online resources but at present are likely to be prohibited by network access controls from doing so.

3. Joint staff of NHS and JANET organisations. Many medical staff have roles in both JANET and NHS organisations and suffer inconvenience if their performance of those roles is limited by their physical location. When they are at the NHS site they are in a similar position to the student group (1) above, although they may need to access more sensitive systems such as finance, student records or other administrative systems and networks within their university or college. When located in JANET organisations, they may need to access sensitive NHS data; at present this requirement is commonly addressed by duplication of computing facilities - both networks and desktop computers - an expensive and inconvenient approach.

3. Suggested Architectures

3.1 Students of JANET-connected organisations who spend some or all of their working time at NHS sites

Access from NHS sites into JANET organisations is not a unique problem and systems already planned or in place for other purposes may well be sufficient to provide it.

From the viewpoint of a workstation connected to the NHS network, JANET is part of the external Internet. If students have access to the Internet there should be no additional risks for the NHS workstation or network in accessing servers at JANET-connected organisations. For the JANET-connected organisation, there may at present be problems in providing off-site access to internal systems. In the past access to such systems has often been restricted to users physically connected to the same Local Area Network (LAN), or with a particular IP address. However there are growing requirements for staff and students to be able to access internal systems from remote locations so facilities will be needed in any case to meet this demand for access from 'anywhere on the Internet'. These same facilities can provide access from NHS systems.

If local security or bandwidth restrictions prevent students accessing the Internet through the NHS central network, then a less flexible alternative is to extend their university or college network into nearby NHS premises and connect it to a cluster of dedicated workstations there. This approach has been used as part of a learning resource centre for students. However the use of such workstations may need to be controlled to ensure that JANET connection and licensing terms are not breached.

Remote access between sites may need to pass through authentication and proxy systems. This will be made much easier if the systems providing remote access are implemented using standard Internet protocols, such as Hypertext Transfer Protocol (HTTP), rather than proprietary systems that may in any case not be suitable for use across the Wide Area Network (WAN). LAN protocols such as Network Basic Input Output System (NetBIOS) should not be relied upon across wide area links unless they have additional protection such as an encrypted virtual private network (VPN).

The architecture to support remote access to JANET servers from the NHS network is therefore likely to consist of two parts. At the NHS end, there is likely to be a proxy gateway permitting outbound access to remote Internet sites, but limiting return traffic to that initiated by local users. At the JANET end, some form of authentication of the remote user will be required. In some cases, with low security requirements, this may be provided on the individual server. Encryption is recommended to protect usernames and passwords so, for example, remote access to e-mail could be provided by a web gateway running secure HTTP over Secure Sockets Layer/Transparent Layer Security (SSL/TLS), known as HTTPS.

Access to more sensitive systems should first require authentication by a dedicated authentication server. This avoids exposing the application server directly to the Internet. Servers that allow remote access should be designed and implemented with care to ensure that they do not accidentally allow access to functions or services that were not intended.

As mentioned above, the server systems to support this type of interworking will also support other uses. It is therefore likely that they will be provided centrally at the main entrance of the Internet into the network or organisation, rather than being

dedicated to NHS-JANET use alone. Traffic could be routed via the public Internet, or via a private peering between the JANET and NHS networks.

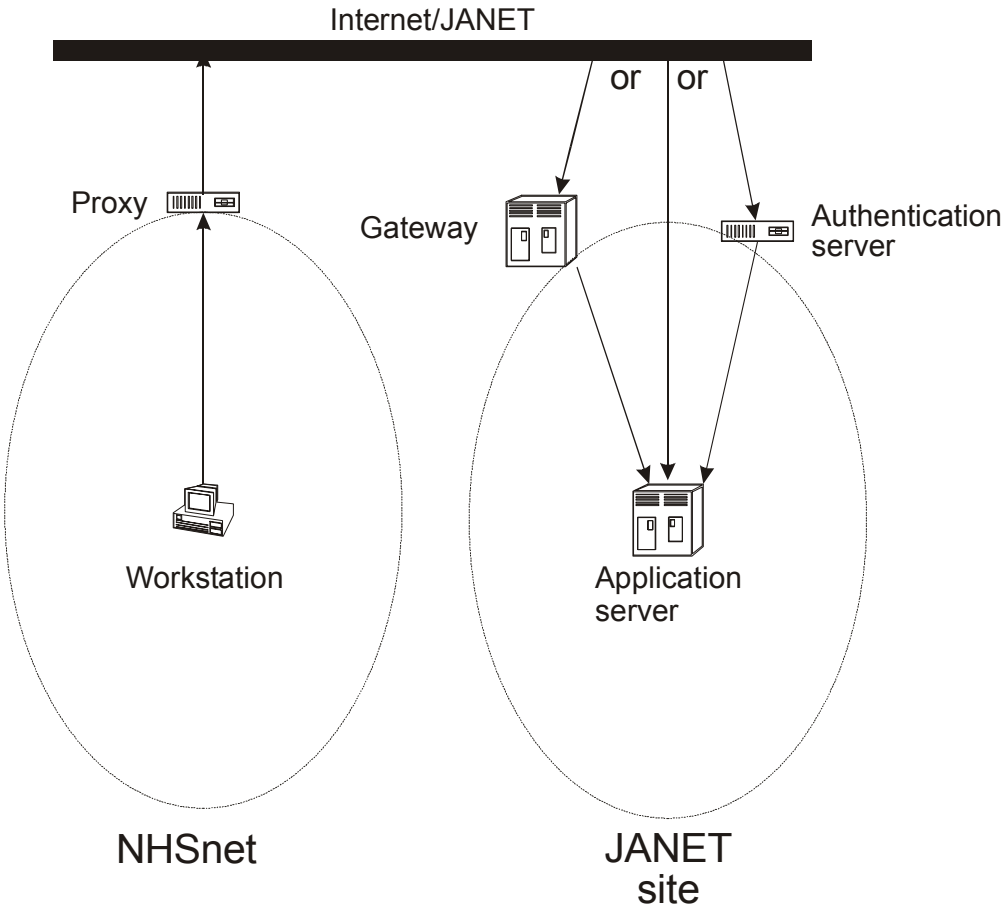

Figure 1: NHS access to JANET

3.1.1 Threats and controls

There are two threats to any server that is exposed to traffic from the Internet. The first is that the server itself may be attacked, using vulnerabilities in its software or operating system. This threat must be addressed by running only essential services on such servers and ensuring that they are configured and maintained in a sufficiently secure manner. Routers or other network control devices may also be used to reduce the exposure to hostile external traffic, if necessary by requiring a user to authenticate to the network before permitting any access to particularly sensitive servers. The impact of any compromise of the server may in some cases be reduced by limiting the access that the server itself (and hence anyone who has gained control of it) has to other systems on the internal network. One network design that is commonly used is to place public servers, or the proxies that grant access to them, in a De-Militarised Zone (DMZ) at the perimeter of the site network.

The second threat is that passwords or other sensitive data may be read off networks or devices while in transit between the client and the server or, more likely, taken from the client machine itself. If reusable passwords are to be relied on for security then they must be protected in transit by using encrypted channels across shared networks (for example using SSL to encrypt web traffic or encrypted VPNs for other protocols). Strong authentication systems using one-time passwords prevent intruders stealing authentication information but on their own do not protect the data exchanged during the subsequent session. Passwords and sensitive information must not be disclosed by users and care must be taken if they are stored on the client machine.

It should not be forgotten that there is the possibility of encountering hostile content anywhere on the Internet. Viruses and trojan horse programs (the latter appear to be benevolent but conceal some less desirable functions) can install themselves on a client computer and allow attackers access not only to current activity on the client machine but also, in many cases, to any information that is accessed from that client in future. A mixture of technical measures (for example running personal firewall and anti-virus software) and good user practice (for example not opening software or documents from unknown sources) are requirements for any client connecting to the Internet. Policies and management systems should be in place to ensure these are followed.

3.2 Students in JANET-connected organisations who need to access networked teaching materials or other resources that are hosted within NHS sites

There are a number of facilities, provided centrally on the NHS network, that staff and students in JANET organisations may need to access. For example the NHS may provide teaching or library resources for health and social care students, while staff may need access to these or other resources that do not involve sensitive material.

This is the reverse of the situation discussed in the previous section, so a similar architecture is appropriate. The same recommendations to use robust standard protocols also apply. As for universities and colleges, a number of systems already exist or are planned to allow NHS users to access central facilities from home or other locations using commercial dial-up Internet Service Providers (ISPs). Adding JANET organisations to the list of locations for clients should not change the threats or risks to such services.

As the NHS internal network has stronger perimeter security than the JANET network, the impact of a compromise to a server within the NHS network is greater. This suggests that direct access to servers on the internal network should not be permitted; but that services should instead be provided on secure gateway machines located close to the protected boundary between the NHS network and the public Internet. For example, access to web-based learning resources should be passed through an external-facing web proxy or content delivery host. Access to e-mail could be provided through a web mail gateway, which passes messages and commands to and from the central e-mail server. These gateways can perform authentication if required, but if static usernames and passwords are used then they should be protected by encrypted tunnels.

As these services are provided for access from the general Internet, traffic from JANET to them could be routed either over the public Internet or over a central private peering between the two networks.

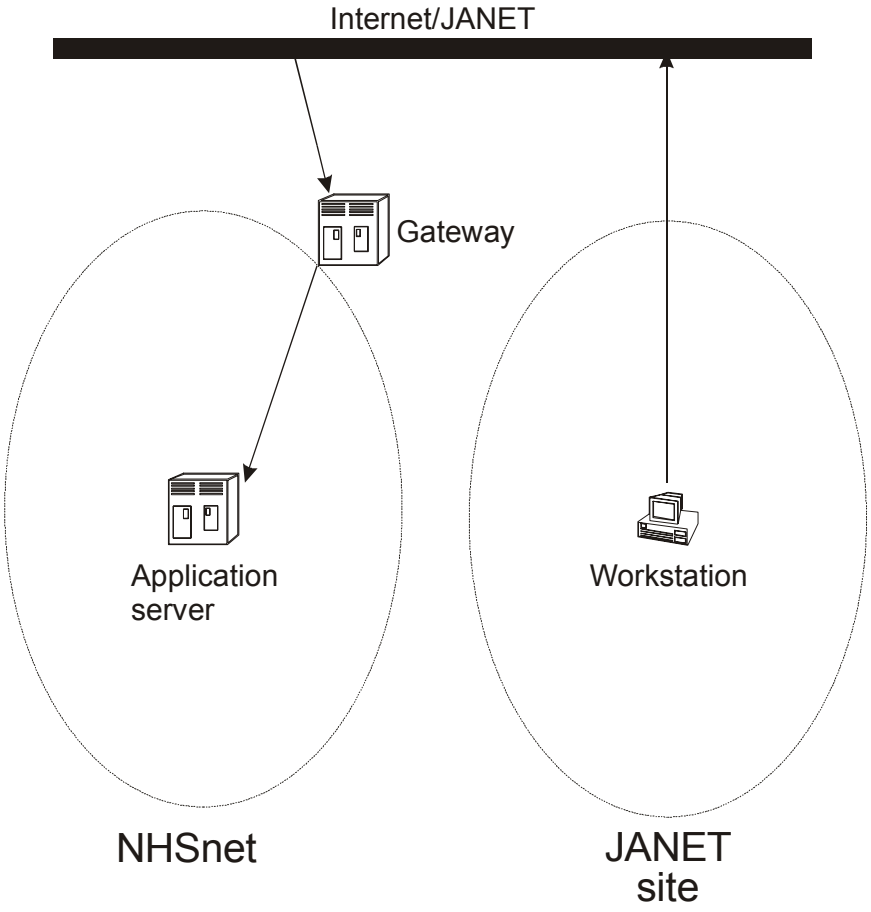

Figure 2: JANET access to NHS service

3.2.1 Threats and controls

As before, the significant threats to security are the theft of passwords and the compromise of a server that is accessible from a public network. Since the potential impact of a server compromise is greater when it is connected to the NHS network, it is recommended that gateway systems be used that, even if compromised, have only limited access to the internal NHS network. In particular, router controls should be in place to ensure that the gateway host can only communicate with a very limited set of internal hosts. In the case of web traffic it may be possible to use content distribution techniques so that the gateway server cannot initiate communication with any internal servers.

3.3 Joint staff of NHS and JANET organisations

Medical staff who are employed by both JANET and NHS organisations may have offices in both organisations. Ideally these individuals would like to be able to perform both roles from both offices, but the different security requirements of the data on the two networks make this difficult to arrange. At present these individuals, and their employers, suffer considerable inconvenience if they wish to work on a particular system belonging to one organisation from the other location. Access to NHS systems from university or college offices often requires dual wiring and duplicate PCs. Access to university or college systems from within NHS hospitals may well be impossible.

As previously discussed, the requirement to access university or college administration systems from off-site does not only arise from NHS sites. Staff may wish to work from other places on the Internet. Any system that is implemented to deal with one off-site requirement is likely to support other locations as well. In contrast, the requirement for staff to access NHS systems from off-site is likely to be limited only to university and college offices and is likely to require an individual solution, rather than borrowing facilities used by others.

The problem here is that such users will expect to be able to access a wide-range of applications, some of which will reside on legacy systems and will not use well-known protocols that might be amenable to firewalling. Simply providing open access between the two networks to allow communication with any application is clearly not acceptable from a security point of view, so a different model is required. The existing, expensive, solution of providing an extension of the hospital network with a PC on the end of it suggests a possible model. Since what users actually need is a workstation connected directly to the hospital network, then a system that creates virtual workstations with displays and keyboards on physical devices located elsewhere on the network may well be satisfactory.

Such remote terminal systems use the same model as the early mainframes, with a physical display and keyboard that happen to be external but appear to the main server to be internal. All applications run on the server and it is only the instructions to change the display or read input devices that pass over an external wire. Modern systems are not limited to simple text displays: typically the remote device is a PC or other desktop system running a small, standard, ‘thin’ client program. The terminal server is a dedicated computer which runs a copy of Microsoft®Windows® for each client, receiving input from the client workstation’s keyboard and mouse and sending display instructions back to the workstation screen, which displays a complete Windows desktop to the user. From the viewpoint of the network, and any application servers connected to it, the terminal server appears to be a client PC directly connected to the network. This client can run any applications that are provided on the server, use any server or network-based authentication systems, and access any data that are authorised by that authentication. Where information security rules permit, the desktop can also be configured to show disks and other resources on the desktop machine, thus allowing file transfer or local printing. However from the client side, the only host that is visible is the terminal server. If that server has two network

interfaces, then there is no need for any physical connection between the two networks and provided the server does not route traffic between its two interfaces then there is no logical connection for individual packets either.

This design, known as a three-tier client/server model, extends the common two-tier server, model by inserting an additional control and interpretation point between the client and the server. Breaking the direct connection between the client and server greatly restricts the possibility for attacks on the application servers, as there is no direct route for hostile traffic to reach the application. The middle tier can implement tight controls on what the client can request: any commands must first pass low-level and authorisation checks on this server before they are passed to the final application for action. The only packets that will be accepted are those that are valid input to the terminal server program. These may, if authorisation permits, then cause commands to be sent over the internal network, but are not themselves passed on. Packets that are not valid input to the terminal server program, or that request unauthorised actions, will be rejected well short of any sensitive data or servers.

The complete system therefore consists of a terminal server host and software, such as Citrix MetaFrame®, located within a hospital with one interface connected, probably through a router to control which internal hosts can be reached, to the hospital network. The other interface is connected, through a firewall or restrictive router, to the university or college network[1]. Medical staff can then use the same desktop system both to access university or college systems and to run the client software that provides them, through the server, with a virtual PC on the hospital network. Authentication may be done using the same credentials as would be presented to a PC within the hospital, or a different authentication process may be chosen for remote users. The communication between the workstation and the server can take place across an encrypted channel if required. This would certainly be recommended if reusable passwords were used to authenticate.

A further advantage of the three-tier model is ease of maintenance. The client workstation only runs a small ‘thin’ client program, which can be easily installed and occasionally updated by the user. All application programs (e-mail, word processing, legacy applications, etc.) are invoked from the client, but run on the server. This means that the server needs to have reasonable memory and processing power - it may be handling a different, resource-hungry application for each client - but conversely means that installation, configuration and upgrades of application programs need only be done on this one machine rather than on each individual client workstation. Furthermore, since this system is not subject to the vagaries of individual users’ hardware and software configurations, applications should run much more reliably and without requiring individual work-arounds. Indeed although most systems

[1] If there is already an approved direct connection between the hospital network and the university or college network then the terminal server needs only a single network interface attached to the DMZ of the firewall protecting that connection. The firewall then provides the function of both internal and external routers. However a direct network connection is not required if only a terminal server is to be used.

[2] Unix is a registered trademark of The Open Group in the United States and other countries.

provide a virtual Windows desktop, many can support client programs that run on PCs, Macintosh or Unix® systems[2].

It would be possible to connect the server only to the hospital network and route packets to it over the JANET and NHS backbones. This is likely to be a complex and slow path over the network, however the terminal server protocol is quite efficient and this may be a viable option for small numbers of users. A more serious problem is the need to maintain a list of all terminal servers on the central NHS firewalls and to permit traffic to these servers from any external Internet address. For reasons of performance and control, therefore, it is recommended that where possible the network connection between the workstation and the server be local, rather than passing over the public Internet and the national NHS backbone. This allows each hospital to have control over access to the system; performance over fast local links is also likely to be much better than using national links to communicate between systems that are in the same geographical location.

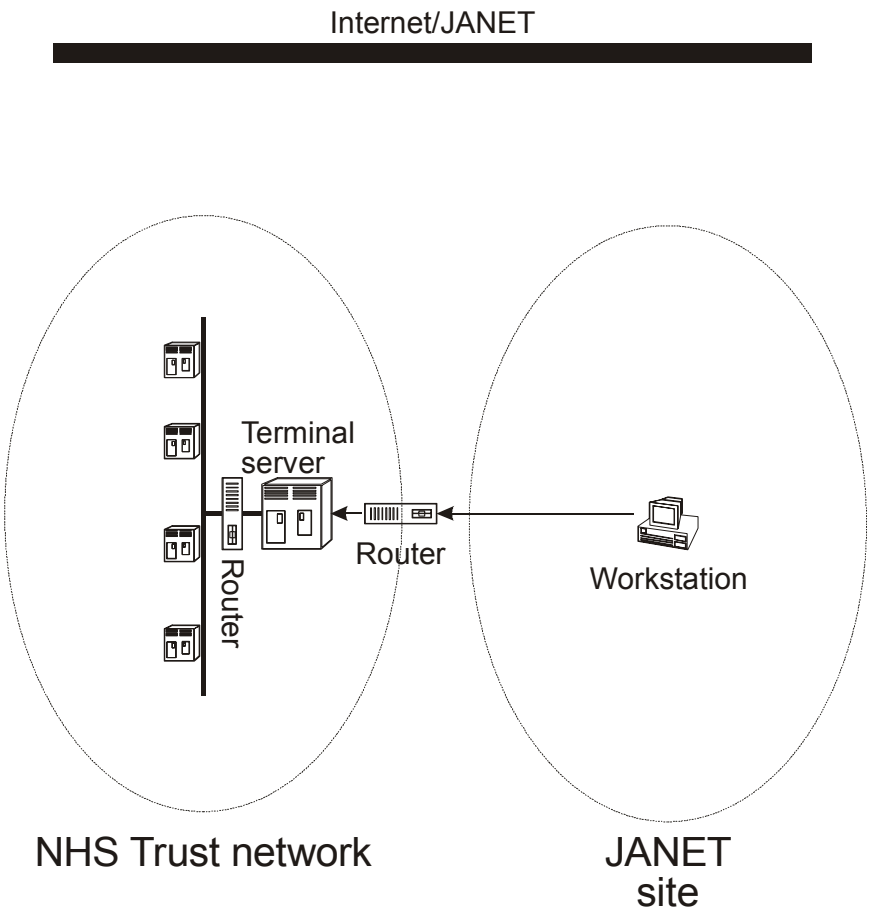

Figure 3: Remote terminal access to sensitive NHS service

3.3.1 Threats and controls

An obvious concern when providing access between two networks is that users on one network might be able to use the link to attack systems on the other. The proposed design greatly limits the scope for such attacks, as there is no direct or logical connection between the two networks. The terminal server is the only place where the two networks are connected to the same device, but it cannot pass traffic through from one network to the other. The only type of traffic that will be accepted by the server from the external network is that comprising requests to connect to the terminal server software, and these will immediately require the user to authenticate using credentials provided by the hospital. The process handling these credentials should be set to lock any account that suffers repeated failed login attempts. This protects against attempts to obtain a password by guessing or brute-force attacks. To further protect the server itself against attack it is recommended that the router connecting it to the external network should only permit connections using the terminal server protocol, and only from IP addresses belonging to the local university or college. Although it would be possible to restrict to the individual IP addresses of recognised staff workstations, in practice the maintenance of such detailed lists of addresses has been found to be more trouble than it is worth. However if the medical school has its own allocated subset of IP addresses then restricting to these is probably a good idea. There remains the risk that a bug in the terminal server program might be attacked from a trusted workstation, however the reputable software packages of this kind have a good record as regards security vulnerabilities. A restrictive router between the terminal server and the hospital network may be used to reduce the number of internal hosts exposed to direct attack in case the terminal server is compromised. It should not need to be stated that the server host must be configured and maintained securely and must not run any unnecessary software or configurations.

The other points of attack in this design are the network between the client and the server, and the client itself. Most products allow the network connection to be encrypted, which prevents interception by other devices on the network. In any case the risk is somewhat reduced as what is passed over the network is no more than instructions to build and alter a screen display, plus keystrokes and mouse movements. If encryption is not used, then the login process must use some form of strong authentication so that credentials cannot simply be replayed.

The possibility of theft of data from client workstations exists wherever these are used to view sensitive information, whether in hospitals or anywhere else. This risk is best addressed by policies and good user behaviour. No technical measure can prevent information displayed on a screen being seen by others, or stop a member of staff storing or transmitting the information insecurely. However the terminal server software may provide additional configuration options to reduce the opportunities for bad practice: blocking access to a floppy disk drive from the virtual PC, for example. Another growing risk is from viruses and other attack tools that install back doors onto infected PCs. These allow intruders to connect to the infected PC and view the contents of the screen or disk. Again the best defence is a combination of technical measures, running anti-virus and personal firewall software for example, and good practice by users. No user should open an unexpected e-mail attachment, or run or

install software whose origin is suspect; these basic precautions should be followed especially carefully by users who have access to sensitive data.

4. Policies

Any technical security measure can be defeated by carelessness or malice on the part of its users. Good passwords are useless as proof of identity if they are written down or shared with others. Clear and comprehensive policies and good practice guides will be needed and must be followed by users of any system. JANET and the NHS network already have user policies, as do many universities and colleges. Part of any implementation project should be a review of the applicable policies to check that they meet the needs of users and of the networks and information they are there to protect.

Since users in any of the structures described in this document will be using at least two different networks, they must abide by the policies of all the networks used. Furthermore, since actions taken on a computer at one time can have security consequences at others, they must abide by these policies at all times, not just on those occasions when they happen to be using both networks. Otherwise the user may install an untrusted program when connected to an open network, which then gives intruders access when they later connect the same machine to a sensitive network. This type of problem has already been encountered by businesses that allow staff to login from home; careless multiple use of the same home PC can result in the business firewall being completely by-passed through that PC.

It is therefore important, that as well as ensuring that policies for the use of the individual networks are defined, that a joint policy is defined so that the users of both networks are aware of the policies when accessing one network from another.

When security incidents occur that affect both networks the network security teams need to collaborate to resolve any security issues.

5. References

A case study of a hospital using virtual terminals to permit access from anywhere on the Internet can be found at: **http://www.citrix.com/press/news/profiles/ourlady.htm**

Microsoft and Windows are either trademarks or registered trademarks of the Microsoft Corporation in the United States and/or other countries.
Citrix MetaFrame is a registered trademark or trademark of Citrix Systems, Inc. in the United States and other jurisdictions.
Macintosh is a trademark registered in the United States and other countries.
Unix is a registered trademark of The Open Group in the United States and other countries.