The Use of Firewalls in an Academic Environment

JTAP-631

Tim Chown
Jon Read
David DeRoure

Department of Electronics and Computer Science
University of Southampton

April 2000

1 Contents

2 Executive Summary

The Internet has evolved to the state where it is now an everyday part of personal and business life. At the time of writing, there are an estimated 300 million people using the Internet at least once a week, and some 70,000,000 hosts permanently connected to the Internet.

The rapid growth of the Internet enables new methods for Universities to perform distance teaching and research, to conduct business, and to attract new students and staff. But it also brings the danger of an increased online population from which “hacking” or other undesirable activity may originate.

The main findings of our one-year JISC/JTAP project are as follows:

1. The adoption of default deny firewalls on JANET sites is very limited. The typical firewall policy involves a handful of protocols being filtered at a site’s JANET point of presence and perhaps a small firewall protecting financial systems. We estimate that less than 10% of all sites have a default deny firewall/security policy, for inbound or outbound traffic.

2. We were able to successfully deploy an inbound default deny firewall at a site with 1,300 users on over 600 live hosts. The solution we chose was Check Point Firewall-1 on a Sun/Solaris platform. In doing so we determined a wide range of selection criteria that should be applicable to assist sites considering their own deployments.

3. User education and awareness are critical issues. The traffic passing through the point at which a firewall may be inserted can be monitored to see which services users are accessing. However, it is important to consult the users of a network to find what their feelings are and to see what they perceive their requirements to be.

4. The introduction of a firewall or set of firewalls is only a means by which to enforce a security policy. That policy must be determined by the institution concerned. An ongoing risk assessment exercise is one method by which to maintain and refine a security policy.

5. Despite the intuitive feeling that a department of computer scientists would include many people running unusual applications over the Internet on unusual IP port numbers, we discovered that the volume of such traffic is much lower than expected. After the initial move to a default deny firewall configuration, the rate of additional long-term service-enabling requests was very low, of the order of at most five per month, often less. Temporary requests for the purposes of software demonstrations were more common, but still not of a level to cause a severe administrative overhead. It was very important that service requests were processed quickly and fairly.

6. There are new network services becoming popular where simple packet-based filtering is not able to perform the desired “firewall” function, e.g. video streaming protocols and chat clients like ICQ. In such instances application level gateways such as SOCKS5 appear to offer a good solution.

7. The cost to an institution of a breach of security is very hard to evaluate. Those sites adopting a more rigorous security policy have typically been exposed to a major incident

from which they have learnt that cost the hard way. This has made the passing of a firewall policy through such university committees that much easier.

8. Failing to deploy a firewall system can have indirect repercussions. If a site is found to be open to "spam e-mail relay" abuse, it may be added to one of a number of blacklists. Many sites make use of such black lists when filtering for junk e-mail or for general network traffic, so becoming blacklisted can be a major problem.

9. While the deployment of a firewall has immediate costs in terms of new hardware and/or software and staff training, that cost can likely be recovered by the notable reduction in systems staff time spent pursuing problems caused by security breaches.

10. A firewall or set of firewalls is only one risk reduction measure. A site security policy should also encompass areas such as secure access to data, e.g. via secure shell (ssh) or secure socket layer (SSL), and authenticated access to data beyond plain usernames and passwords, e.g. via X.509 or PGP software certificates, or via physical tokens such as SecurID.

11. Firewall technology is improving rapidly. Some can now do content filtering in silicon. Established products such as Firewall-1 allow content-based vectoring whereby WWW, FTP and e-mail data can be passed to a process on another machine for filtering (e.g. automatic non-intrusive virus checking of inbound e-mail). However, the basic security principles (as reported in this document) remain the same.

12. The Data Protection Act 1998 requires all personal data to be held under a reasonable level of security. That may include data encryption but also the selective blocking of access to certain network servers or subnetworks. It is a site's responsibility to decide what information needs protecting, and to what level the site will protect that information.

13. Many police forces throughout the UK are now addressing IT issues. The popular press tends to cover the more sensational pornography-related cases. While these may be a cause of embarrassment to a site, there is also a corporate liability under which sites should be able to identify individuals carrying out illegal acts over a network. This is difficult from non-authenticated access points. While in some cases no law may have been broken by the site from which an incident originated, if the site can be shown to have acted negligently in a civil court it may be liable for a fine or compensation. Thus some method of outbound authentication may be required.

Recommendations for institutions and the JISC can be found at the conclusion of the report.

3 Introduction

This report is the result of ongoing work at the Department of Electronics and Computer Science (hereafter referred to as ECS) at the University of Southampton into methods by which secure and “safe” network access can be offered to staff and students. The report includes reference to ECS as well as to Southampton University Computing Services (SUCS), who manage the computing facilities on the rest of the campus.

The goal of the work has been to minimise the required change of habits for the users as a result of the new security policies. The project began with the brief to investigate methods for secure transient access to networks (e.g. via dialup and ad hoc laptop PC connections). However, initial informal surveys of other institutions led us to shift the focus slightly; as a result the report is now focused on the general issues behind a default deny firewall installation.

A firewall is now an important component in enforcing a network security policy in any organisation that wishes to connect its computers to the Internet. At the time of writing there are over 90 different commercial firewall solutions available, and it is often a difficult and confusing process to narrow down the choices to those that are suitable for your organisation. It is therefore one purpose of this report to provide a framework and methodology for evaluating the available solutions. We also report on technical deployment issues as observed at our site.

The report includes reference to social issues in a firewall deployment. The notion of “academic freedom” is one that has a strong bearing in any University environment, and it is a sentiment that requires respect from anyone attempting to draw up and enforce a security policy.

3.1 Acknowledgements

We are grateful to the Joint Information Systems Committee (JISC) for the one year of funding which enabled this project to take place.

The assistance of external steering advisors was also most helpful, including Mark Hindess (Bath University), Tish Roberts (JISC) and Colin Bird (IBM).

The Networks and Distributed Systems (formerly Systems and Networks) staff and postgraduates have shared their experiences and wisdom throughout the project. Thanks are due in particular to Mark Thompson and Kevin Page.

Members of the ECS systems team have been most helpful in the work, both from a technical viewpoint and in the dissemination of information to the 1,200 or so users registered for the 600-700 hosts in the Department. Julian Field’s work in deploying a GNAT-based firewall for an ECS spin-off commercial unit is reported in a separate case study. Our thanks also go to Jon Hallett, Chris Gutteridge, Al Riddoch and Lance Draper. Also the work of Ben Crosby was most valuable; Ben helped pioneer our firewall project, but he has since departed ECS.

Finally, we also thank Alex Walker and Simon Lane from Southampton University Computing Services (SUCS) for their input and assistance in the project.

4 An Introduction To Firewall Concepts

Before we consider the issues involved in deploying a firewall in an academic environment, we first need to briefly describe what a firewall is, and the general concepts that lie behind the use of firewall systems. Later sections go into the fine details of functionality and deployment.

4.1 What is a Firewall?

The Internet is a network of computer networks. It has evolved from the interconnection of networks around the globe. Interconnection is a good thing; it allows the free exchange of information via the Web, e-mail and file transfer. But it also carries a price, namely the risk that your Internet connection may be used by "hackers" (or as some would rather call them "crackers") to gain unauthorised access to your local network. Availability of computing facilities can also be targeted by Denial of Service (DoS) attacks.

A firewall is a system that implements and enforces an access control (or security) policy between two networks; it usually guards an internal private network from an external public one, isolating an intranet from the Internet. Essentially a firewall connects two or more networks but only allows specified forms of traffic to flow between them. The firewall is a means by which a security policy can be enforced.

A security policy defines general security principles for a site. In general, it will state what standards, guidelines and practices should be adhered to. It need not go into specific detail, but may specify policies such as "e-mail may only be delivered into the site to e-mail servers maintained by authorised systems support staff". The trick is choosing the right policy for the right environment. Some degree of flexibility is required such that a site's users can continue to work and exchange information with remote sites.

The "bible" for those wishing to deploy a firewall is Brent Chapman's "Building Internet Firewalls" [CHAP]. This book is essential reading. There are many freely available documents on the Internet acting as "beginner's guides" to firewalls, e.g. the "Keeping Your Site Secure" guide [WACK], the Firewall.Com portal [FWC], the Security Portal [SPORT], and the Firewalls Frequently Asked Questions (FAQ) list [FWFAQ].

4.2 Types of Firewall

There have historically been two main types of firewall; application layer and network layer:

1. **Application layer firewalls** implement a proxy server for each service required. A proxy is a server that enables connections between a client and server, such that the client talks to the proxy, and the proxy to the server on behalf of the client. They prevent traffic from passing directly between networks, and as the proxies are often implemented for a specific protocol they are able to perform sophisticated logging and auditing of the data passing through them. A disadvantage of application layer firewalls is that a proxy must exist for each protocol that you wish to pass through the firewall; if one does not exist then that protocol cannot be used. Some protocols, such as SMTP for e-mail, are natural proxies. Others, such as FTP for file transfer, are not.

2. **Network layer firewalls** make decisions on whether to allow or disallow individual Internet Protocol (IP) packets to pass between the networks. IP is the protocol by which almost all data is routed around the Internet. IP connections rely on a unique source and destination IP address for the communicating hosts. TCP layer port numbers (the "application layer endpoints") are also readily available to a network layer firewall. For example, port 25 is the agreed port number for SMTP e-mail transfer. The firewall can make filtering decisions based on the IP and port number values. This type of firewall can be very flexible. However the added complexity increases the risk of security holes through misconfiguration.

Modern firewall architectures tend to lie somewhere between these two firewall types. "Stateful inspection" techniques allow network layer firewalls to parse IP packets more fully (by looking inside the packet to the embedded TCP layer data) and to keep track of individual connections. In doing so they allow comprehensive logging and auditing to occur. Additionally, many firewall solutions provide application proxies for some protocols, while handling others through a packet filtering system.

Some firewalls, such as Sunscreen EFS3 [SUNEFS3], can operate in a stealth mode. In doing so they present no targetable IP address to internal or external hosts. The firewall acts similarly to a layer 2 switch – it does not route packets but can filter based on IP addresses and interfaces.

TCP layer proxies also exist for relaying connections between an internal and external network. An example of such a proxy is SOCKS [NECSOCK], currently at version 5 [RFC1928]. SOCKS can be used in situations where enabling full access to a host inside your network is undesirable. For example, in the case of the ICQ chat system, by acting as a "smart" relay the external host will interact with the (secure) SOCKS gateway and not the (relatively vulnerable) client host. This allows the client to receive data without the need to open up permanent "holes" in the firewall.

The very latest firewalls offer layer 4 (transport) and 5 (application) filtering or switching abilities, e.g. the ability to switch data based on Web browser cookie content. However, the principles of firewall deployments described in this report remain the same. Protocols such as FTP, HTTP (Web) and SMTP (e-mail) may be intercepted by a firewall (e.g. Firewall-1) and "vectored" intelligently to a separate process (on a separate server) for filtering. The Content Vectoring Protocol (CVP) [CVP] is Check Point's open protocol by which application layer content can be passed to a server for processing; one typical use is e-mail virus scanning by a co-operating product such as InterScan VirusWall [ISVW].

4.3 Modes of operation

There are two very distinct and different modes for network firewalls to operate in.

1. **Default allow firewalls** allow all traffic in and out of a site. Some specified services may be blocked on the firewall, but all others can freely pass through.

2. **Default deny firewalls** block all traffic in or out of a site (though commonly they only block inbound, rather than outbound, traffic). Only named services are allowed to pass through the firewall.

The "textbook" recommendation is to run default deny, but such a policy would intuitively be in stark contrast to the notion of "academic freedom". As a result, the vast majority of UK

Universities appear to run in default allow mode (see Section 13 on UKERNA activities), with perhaps a handful of known "problem" services blocked at the point of entry to their network.

The problem with a default allow firewall is that as and when new security vulnerabilities are reported, the firewall administrator has to play an ongoing "catch up" game with the potential attackers. This can consume valuable time.

Firewalls are unlikely to be able to tackle all potential external threats. We have already mentioned that a firewall can be used as a means to identify inbound SMTP traffic and redirect it through a virus-scanning host. Denial of service (DoS) attacks are also not uncommon. Given that many security policies do allow some connection types into the internal network, it is desirable to be able to detect when these services are being attacked. An intrusion detection system (IDS) will typically look inside the TCP data element of an IP packet for certain data sequences that may be used in an attack, and may also spot DoS incidents such as the SYN, SMURF and Fraggle attacks [CERT]. This report does not focus on the use of such systems, though we will report on them as part of our ongoing Secure Internet Protocols JTAP project [JTAP032].

4.4 Where should a Firewall be situated?

Most networks will have a single point of presence on the network through which they connect to the Internet. For a campus site, that is typically the router (backbone edge node or BEN) through which they attach to JANET. For a department, it may be their link to the campus network, though in many cases a department may be spread across many buildings.

In the case where a router is used, it may be possible to run a firewall on that router (e.g. on a Cisco router running IOS and a Firewall-1 module, or within a Cisco router itself [CISCOR]). It may be possible to also run an IDS on the router. If a router solution is not possible, or if the router is unable to meet the processing and logging/management requirements of a high-capacity firewall, then a separate dedicated system can be used.

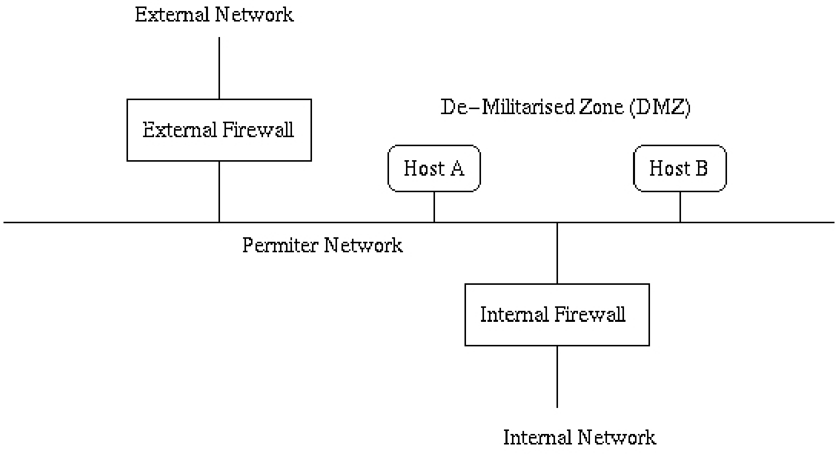

A traditional DMZ topology

The textbook solution is to have an exterior router/firewall acting as screening firewall with an internal router/firewall protecting the internal network. The perimeter network that lies between these two devices is often referred to as the De-Militarized Zone or DMZ. The benefit of a DMZ is that hosts on it (often called "bastion" hosts) can be placed in a "no man's land" where they can be accessed from external or internal hosts. However, if a bastion host is compromised, the knock-on threat to the internal network is minimised (in comparison to the implications were the breached host inside the internal network).

Rather than have two separate firewall devices, it is possible to collapse exterior and interior routers into a single system if that system has multiple interfaces and you can specify rules across each combination of interfaces. It may be beneficial to have multiple DMZ networks. The main danger of such a collapsed system is the potential effect of a denial of service (DoS) attack on the external interface of the device (in a traditional perimeter network system the internal device would not be affected).

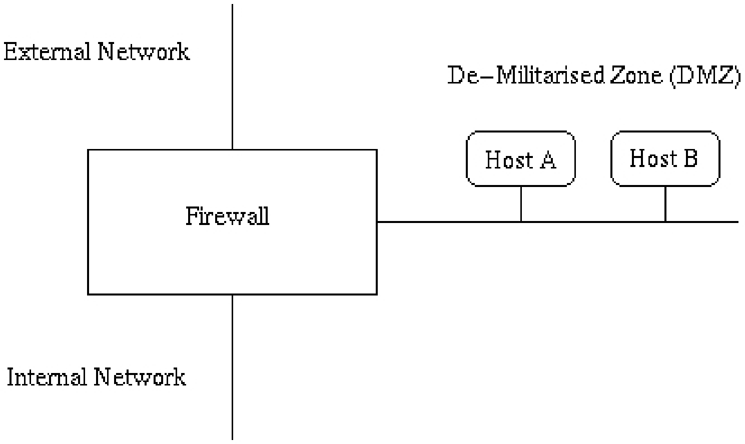

A collapsed DMZ topology

This is the solution adopted at the authors’ site. We will discuss uses of the DMZ(s) in Section 8 of this report.

4.4.1 JANET-level firewalls

It would in theory be possible to do firewall or filtering operations at a JANET interchange point, i.e. at the border between JANET and other national or international networks. The problem is that such filtering would be subject to requests from any one of hundreds of individual sites to enable certain services. A list of “approved” mail servers might stretch to 2,000 IP addresses, and would be subject to a heavy administration load. It may be possible for JANET to choose to block certain known dangers (e.g. BackOrifice). However, for reasons of performance, scaling and manageability, we believe the firewall and security policies should be left to the individual sites (or perhaps the regional networks or MANs) attached to JANET.

5 General Firewall-related Issues

The deployment of a firewall system is touched upon by consideration of many issues, including general "best practice" standards already in existence and UK Law applicable to matters related to computer security.

We also consider recent trends in security incidents on JANET, and the general growth of the Internet host and user population.

5.1 The Computer Misuse Act and Corporate Liability

There are two main legal aspects to computer security.

The first, the Computer Misuse Act (1990) [CMA] is "an Act to make provision for securing computer material against unauthorised access or modification; and for connected purposes." The Act makes it an offence to make unauthorised access to computer material, to make an unauthorised access with intent to commit or facilitate commission of further offences, or to make unauthorised modification of computer material.

The Computer Misuse Act is the Act by which prosecutions for "hacking" are generally made.

The recent Networkshop 2000 event [NWS28] highlighted another issue of legal importance, namely an institution's corporate liability when computer hacking incidents originate from that site. Dave Reid of the Lothian Borders Police Force explained the issue at the event. If the institution is unable to show that it took reasonable steps in authenticating users for access to its computer systems, and as a result cannot identify individuals commiting offences, it may find at the very least a civil case for negligence being brought against it.

The implications are that services like public non-authenticated "docking stations" may be used on a campus, but if they are to be used to gain access to external facilities, some authentication service should be employed (perhaps at a campus firewall) to grant that access. The strength of that authentication remains an issue – students are notorious for sharing computer account details, so a user name and password is not necessarily a guarantee of a computer user's identity.

We found the "Introduction to Computer Law" book by Bainbridge to be useful [ICL].

5.2 Data Protection Act 1998

Another legal minefield that has arisen in the past couple of years, and which came into effect as of March 1st 2000 is the new Data Protection Act 1998 [DPA98]. The Act gives extended rights to individuals to have access to information held about them, to know why information is being held and for what purposes. The data controller also has increased responsibilities to ensure that (sensitive) personal data is held securely, and if transmitted is done so with due care and respect for the individual's right to privacy.

The pertinent part of the Act relating to personal data is the Seventh Principle which states that "Appropriate technical and organisational measures shall be taken against unauthorised or unlawful processing of personal data and against accidental loss or destruction of, or damage to, personal data.” At present there is little in the way of test cases under the new Act. Will encryption of personal data be required as an “appropriate technical measure” to prevent unauthorised access to that data? And if data is encrypted, will a key escrow policy be required to ensure that if the encryption key is lost some “responsible party” can return a copy of the key to the user?

Computer systems managers should now also be aware that they are also personally liable for offences under the DPA 1998, with fines ranging up to £5,000.

5.3 DTI/BSI/ISF Standards

Standards and recommendations for security already exist in the form of guidelines from the Department of Trade and Industry [DTISEC] and the British Standards Institute [BSI].

The DTI advice for computer security is presented as a set of tasks that should be performed by companies connecting to the Internet. The advice in general appears relatively good.

1. *Identify your business needs*
   This involves defining which services you require, e.g. e-mail, Web, information exchange (file transfer) and e-commerce, where they run (intranet or extranet), and who should have access to which information.

2. *Assess the risks*
   Here you assess the value of your information, the harm that its exposure could cause, and the likelihood of a security breach, given existing security controls in place.

3. *Select your connection*
   In the case of a business, the choice here is wide, but for HE institutions the choice is almost universally restricted to a JANET service.   The service may also be provided under a Metropolitan Area Network (MAN) agreement between co-operating Universities.   In some instances, Universities may seek commercial Internet connections, perhaps for network resilience or to service commercial spin-off companies.

4. *Develop a security policy*
   The policy should state the services that can be used, who authorises connections, who is responsible for security, what standards, guidelines and practices should be followed, and the users' responsibilities. Such a policy should be regularly reviewed.

5. *Implement security controls*
   These may include encryption, digital certificates, firewalls and staff awareness and training.

6. *Monitor and maintain effective security controls*
   This implies some iteration in the risk assessment exercise. As the DTI advice states, the risks can be managed, providing that they are understood.

The British Standards Institute [BSI] has recently updated BS 7799 on Information Security Management [BS77991, BS77992] from its previous (and now withdrawn) 1995 version. These standards are available for a small fee (£50-110, depending on membership status) from the BSI. The DTI recommends adoption of BS 7799, saying that a company doing so

1. practises what it preaches in terms of adopting best practice in information security management;

2. provides business partners in the UK and overseas with a “yardstick” by which to judge it;

3. enhances its prospects for tendering to other secure organisations and of winning contracts.

and that customers should insist on BS 7799 when undertaking Internet trading.

Consultants such as Gamma Secure Systems [GAMMA] recommend that companies at least inspect the BSI standards. BS 7799 provides over 120 security guidelines structured under 10 major headings to enable readers to identify the security controls which are appropriate to their particular business or specific area of responsibility. As well as giving detailed security controls for computers and networks, BS 7799 also provides guidance on security policies, staff security awareness, business continuity planning, and legal requirements.

JTAP carries reports by the Information Security Forum [JTAPISF] on its Web site. These do not directly address firewalls, but do feature papers on BS7799 (though apparently the 1995 version), general best practice, risk analysis, as well as specific issues such as a Windows 2000 Security Checklist released in February 2000. The papers date back to 1991 or so, but many appear to still be relevant today.

5.4 Growth of the Internet

At the time of writing, there is an estimated 300 million people using the Internet [MIDS, NUA] at least once a week, and some 70,000,000 hosts permanently connected [ISC, MAPPA, NSZ] to the Internet. The number of hosts has doubled since July 1998 [NSZ].

The true number of hosts connected to the Internet can not be known for sure. The biggest obstacle to “sizing” surveys is the wide-spread use of Network Address Translation (NAT) devices, which masquerade multiple Internet devices behind single Internet Protocol (IP) addresses.

Estimates for the number of regular Internet users as of January 2000 vary from 280,000,000 up to 470,000,000, depending on the source quoted. Whatever the number, the demographics generally put the bulk of those users in the US and Europe, with Asia and other regions relatively sparsely connected.

The growth of the ICQ Chat system [ICQ] is another barometer of Internet growth. There have been some 70,000,000 ICQ registrations since the service began. Even given some duplications where ICQ accounts have been restarted, the ICQ user base has ballooned from the few hundreds of thousands it stood at just two years ago.

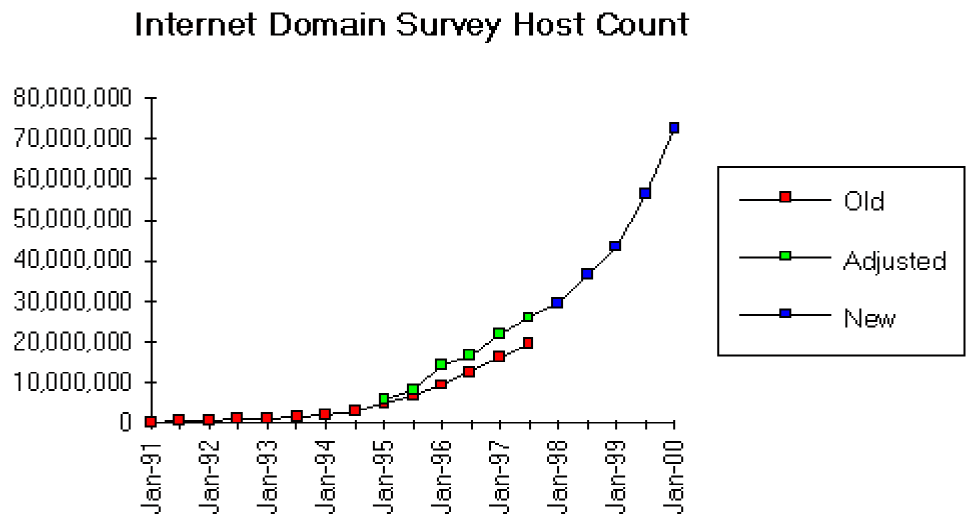

Source: Internet Software Consortium (www.isc.org)

The growth in the number of connected hosts is occurring across commercial, domestic and academic sites. HE is a very large and varied user group with 1.5 million students and 200,000 staff. With FE sites joining this will reportedly increase numbers to 4.5 million; it was stated at Networkshop 28 [NWS28] that some 430 new Further Education sites would be connected to JANET in the coming 12 months. This growth represents additional sites that may either be subject to “attack” over the Internet or that may be the source of such an attack.

The growth of the Internet, together with imminent voice-data-video convergence to packetised IP services, will lead to the adoption of new Internet protocols. The most important of these will be IPv6 [V6F], which we discuss later in this report. It is very important that security measures and policies are abstracted as much as possible from specific technologies, so that they can be adapted to encompass those new technologies when they arrive.

In short, the Internet is expanding in scope rapidly, and dependence on it is growing. Failure to recognise that fact may prove costly to an institution. The cost of being offline for a few days as the result of an attack may be much higher than intuitively thought.

5.5 Security Incidents on JANET

The JANET-CERT team collects monthly statistics of security incidents on the JANET network [JCSTATS]. These are very usefully categorised into a number of incident types. While month-by-month incident levels fluctuate, there is a general pattern upward. The Web site lists incidents for 1997-99; we also have figures for 1994-96 direct from JANET-CERT.

For overall incidents the counts for the three years over which statistics have been reported are:

Year	Incidents
1994	173
1995	344
1996	507

1997	806
1998	1594
1999	1712

The 1998 figure includes a very large number of probing (mscan) attacks in July. These appear atypical and do skew the figures a little. A more accurate indication of severe incidents comes from the counts of root compromises over the same period. Such a compromise means the attacker has gained complete control of the target host.

Year	**Root compromises**
1994	6
1995	29
1996	43
1997	51
1998	96
1999	174

Just one such incident can potentially cause a number of days of downtime for a core or critical computer system (e.g. a campus file server or set of file servers). The problem is that once a host has been compromised it, and systems that trust it, will have to be rebuilt from scratch if confidence in system integrity is to be restored. A disaster recovery policy is required in the event that the security policy (or the implementation of it) fails.

Network scans and probes (such as from mscan, ISS or SAINT, as discussed in Section 8) are precursors to a potentially more severe attack. Such probing is designed to find weaknesses in your defences. The probes may not occur in one single rapid session; the more wily attackers will spread probes out over a period of time, and only look for a small set of vulnerabilities.

These statistics indicate that there is a greater need than ever before to take precautions against potential incidents. Arguably more people may now be reporting incidents given JANET-CERT's recent rise in profile, but we can only deduce trends from what is reported. It is quite possible that sites with no security measures may be oblivious to attacks, in particular those being relayed or staged through them.

5.6 Impact of an Attack

As reported in the 1999 JTAP Security Workshop [JTAP043], there are three areas of impact:

- **Financial loss**.
  This may be via lost income, or possibly from fines/compensation imposed by a court.

- **Loss of reputation**.
  If embarrassing material is revealed, or perhaps from a forged e-mail.

- **Denial of access to resource**.
  If a key piece of network or server equipment has been rendered unusable.

In each case, assessing the exact cost is very difficult. We live in an Internet-driven academic environment where resources need to be accessed 24 hours a day, seven days a week. A researcher in California or Japan may be many hours behind or ahead of the local time.

Information has to be available to existing or potential project partners whenever it is required. People now assume (rightly or wrongly) that e-mail is an acceptable way to reliably communicate, and may not fax or post copies of letters in the event that the e-mail system at the recipient’s site is forced offline.

6 Evaluating a Firewall Solution

There are many yardsticks against which a firewall system can be measured. We discuss these in this section, with the aim of generalising the selection criteria for a firewall system, rather than discussing specifics of individual systems. This will hopefully make the criteria relevant for those considering a firewall deployment in the future.

6.1 Cost

When putting together a firewall system, it is worth remembering that the cost involved does not end with the hardware and software purchases; there may also be costs involved concerning the installation of the system and the training of the firewall administrators. Additionally, there are likely to be ongoing costs associated with external support and internal administration of the firewall, as regular policy reviews and security audits are recommended. There is little point in operating a firewall that has not been recently updated with the latest security fixes and patches (it is worth noting that even a market leader such as Firewall-1 v4.0 has had five service packs released for it).

It is also prudent to investigate exactly which features are included as standard with a firewall product, and which are only available at extra cost, or even not at all.

The purchase cost will vary from nothing for the likes of Linux ipchains, to tens of thousands of pounds for high-throughput silicon firewalls.

Educational discounts are available on some products, e.g. through CHEST [CHEST], where ESOFT [ESOFT] resell Firewall-1. ECS originally bought Firewall-1 from CenturyCom [CCOM], but then discovered that CenturyCom had stopped their educational discount. We thus renewed our maintenance and support contract through ESOFT. However, consider what may happen if the ESOFT deal with CHEST is discontinued – ongoing support costs may rise as a result.

6.2 Functionality

The functionality of a firewall solution is perhaps the most important criteria for evaluation; does a chosen solution fully meet your current and predicted requirements? You should have a growth plan for your network. Will your firewall still work three years from now? You might want to run the same software product but on different or upgraded physical hardware (e.g. to meet rising bandwidth usage). The lifecycles of network equipment for Internet connections are fairly short, so you should make sure that the basic architecture that you put in place is likely to be viable in the long term.

Note that every feature the firewall has that is not being used adds an extra unnecessary risk, in that it can be targeted by a potential attacker. This might be unused functionality on the operating system of a software firewall (e.g. Windows NT or Solaris), or unused hardware, such as a hard disk in a floppy-based firewall such as the GNAT Box.

General firewall references on the net that have already been mentioned carry useful pointers to factors to consider. Two other good references are “How to Pick an Internet Firewall” [FPICK] and the Great Circle Associates site [GCA], originally set up by firewall guru Brent Chapman of "Building Internet Firewalls" [CHAP] fame.

Factors and features to consider include:

- **De-Militarised Zone**
  A DMZ, also known as a perimeter network, is a third network added between the internal and external networks, or alternatively an extra independent interface (or more) on your firewall host. Services that you wish to be made available to external users may be located on the DMZ. If these services are compromised by an attacker, they will not have access to your internal network, because you will have another firewall, or rules applicable across another interface, to protect your internal network.

- **Virtual Private Networking**
  Virtual Private Networks, or VPNs, are a low-cost alternative to dedicated leased lines for connecting two or more sites. In essence, a VPN works by creating encrypted virtual channels over a public network, such as the Internet. Modern firewall solutions often include support for creating VPNs, either built-in or as an associated product, e.g. SecuRemote [SECU] for Firewall-1. Client software should be available for mobile users, enabling them to securely connect to the internal network from anywhere. There are several standards for VPN product compatability (eg. IPsec, S/WAN), and it is worth investigating if remote sites you wish to work with are running firewall products, and if so which standards are supported by them. It may be advantageous to have an integrated firewall and VPN solution, as a separate VPN solution opens a second avenue of attack into your network. However, public domain VPN technology can be used, as discussed later in this report.

- **Network Address Translation**
  A firewall may hide the IP addresses of machines on the internal network, by keeping track of connections from a machine to the outside and rewriting packets on the fly. NAT is invariably used where a site has more hosts than available IP addresses, or where it wishes to masquerade multiple hosts behind one IP address for administrative or other reasons. NAT helps to protect machines on the Internet network from being discovered and targeted by attackers, but it also breaks the end-to-end security model and transparency of the Internet; this is one reason why IPv6 [V6F], with its much greater IP address space is an attractive protocol. A department using NAT on a campus network introduces a second level of network administration, which is both a cost and a potential insecurity.

- **Media**
  Does the firewall system under evaluation support the media interfaces required, e.g. 10 or 100Mbit Ethernet, quad Ethernet cards, Gigabit Ethernet, ATM or FDDI?

- **Filtering**
  - If a firewall performs stateful inspection of packets (e.g. SMTP, FTP or HTTP), or uses a proxy system, which protocols does it cover?
  - Can it filter by time of day?

- **Number of interfaces**
  It is worth checking that the system under evaluation supports the number of network interfaces that are required; most firewalls should be able to perform filtering between more than two networks, e.g. for a "collapsed" DMZ configuration.

- **Transparency**
  How transparent is the firewall to the end user, in both outgoing and incoming directions? Does the user need special software or configurations in order to perform their tasks?

- **Authentication**
  - Does it support standard passwords, S/Key, RADIUS, TACACS or SecurID?
  - The firewall typically does not take the place of the vendor's authentication server. Rather, it forwards requests from the user to the authentication server, and, depending on the authentication result, either allows or disallows the connection.

- **Content Control**
  Does the firewall have the ability to control the content of the data that passes through it? For example, firewalls often have the ability to provide access control and enforce policy for web browsing, and may also scan for possibly malicious content such as Java applets, ActiveX controls, or even viruses in e-mail attachments. Firewalls may alternatively be able to redirect content to another server for processing or filtering.

- **Denial-of-service (DoS) attack detection**
  - This is a more recent development, and is typified by intrusion detection systems (IDS) such as Check Point's RealSecure [RSEC]. Which DoS attacks does a firewall support? Or will you run a separate IDS?
  - Can the IDS modify firewall rules on the fly to react to DoS or other attacks?
  - Can "dangerous" live connections be manually killed if detected?

- **Reporting**
  Reports are one of the most important aspects of a firewall's functionality; the firewall may be preventing someone attacking your internal network, but does it provide you with enough information about the attacks for investigation and maintenance? The following are some points to be considered.
  - How much detail of events do the reports give, and can the detail level be tailored? Is it possible to change the level of detail by event type, so that important events are recorded in detail but not obscured by huge numbers of unimportant ones?
  - Can the reports be securely logged to a remote machine, printer or other device? If your firewall is compromised, the logs may be altered or erased by the attacker to cover their tracks.
  - Are comprehensive log analysis tools included as standard with the firewall product, or is one available at extra cost? A good log analysis program can save a lot of time in identifying and tracing attacks. Can the logs be exported to plain text or an open format for processing by your own scripts?
  - What sort of intrusion detection capabilities does the firewall have, and will it be able to alert you under specified conditions? For example, you may wish to be alerted to an ongoing attack by visual and audible signals, by pager or by phone, or even by email for less serious cases to be investigated later. Can you silence individual alerts if necessary, leaving the others active?
  - How easy-to-read are the reports (no cryptic error messages or warnings), and in what formats can they be generated? (e.g. plain text, HTML...)

- **User interface**
  - Where can the system be managed from? (console, internal, external, dialup,…)

    - Is it terminal-based, command line, a good GUI, or Web browser based?
    - Is there a remote configuration tool (if so, how secure is it?)
    - Can you run a single console to manage multiple firewalls?
    - What SNMP management is offered, if any?
  - How secure is access to the user interface, and the system the interface runs on? (It is worth noting that independent of the firewall itself, a dedicated secure room may be warranted for the firewall management point. Dedicated fibre optic cabling may also be desirable.)
  - What is the level of ease of configuration (there is probably less risk of introduction of security holes through mis-configuration on a simpler system)
  - The interface should allow an unsophisticated user to build a simple configuration in line with policy, but also allow an expert user to fine-tune the configuration as necessary.
  - What log management system exists?

- **Firewall security**
  - How secure is the firewall platform itself?
  - Is it running on a hardened kernel and operating system?

- **Firewall architecture**
  - Does it offer proxies for control of some applications, e.g. ICQ?
  - Does it offer packet filtering for speed or where proxies aren't needed?
  - Are there hooks for third-party or add-on systems (authentication, VPN etc)?
  - Is there significant "freeware" support from a large user community?
  - Is there standards adherence (e.g. Internet Key Exchange protocol, IPsec...)?
  - If the firewall is a software product, be sure that the hardware that you plan on using is supported and sufficient. These criteria vary depending on the firewall product.
  - How many systems are in the architecture? This affects maintenance time and cost and the importance of a centralised maintenance station.

- **Platform architecture**
  - Can the hardware be upgraded while keeping the same firewall software, in the event of faster processing or greater throughput being required?
  - Is the hardware proprietary or open?
  - Would you run your firewall on NT, or should you demand a Unix version? A Unix (e.g. Solaris) version may be more robust, but may cost more.
  - Is the existence of and support for the OS guaranteed for the foreseeable future?
  - Can security patches be applied to the OS independently of the firewall package?

- **Fault tolerance**
  - If the firewall goes down or is compromised, can a backup system take over automatically?

- **Performance**
  - Does it run in silicon or software? While a software version may be slower, it may be more readily upgradable.
  - What is the maximum packet forwarding rate?
  - What is the VPN encryption overhead?
  - What is the stateful inspection overhead?
  - Can it handle large rule sets and host or protocol object lists?

- Can it load-balance on multiple firewall interfaces, or between co-operative firewalls?

6.3 Training, Support and Documentation

- **Documentation**

  - How comprehensive is the documentation? Some trial or beta versions of firewalls come with manuals that are very similar to the full release versions, so you can judge the quality well.
  - Is the documentation printed, on the Web, or on CD?
  - Are there easy-to-follow tutorials?

- **Technical support**

  - Where is it based? In the UK?
  - Availability
    - What is the response time?
    - What about support on the hardware platform and its OS, if separate?
    - Is it 24-hour support, 7 days a week?
    - Will support be on-site, or available by phone, fax, or e-mail?
    - Is it from the vendor or an independent consultant or reseller? Beware “box shifters” who don’t understand their product.
    - Is support included in package, or extra? If extra, how much does it cost?
  - What is the upgrade and patch availability now? Future support will likely be of the same quality.

6.4 Miscellaneous features

Other factors include:

- Has the firewall been subject to third party certification, e.g. by ICSA.net [ICSA]?
- Are there good product reviews in reputable publications?
- Can you find reference sites that will vouch for the product’s reliability and performance?
- What about the firewall author company credentials
  - Number of years in business (overall, and on the security and firewall side)
  - Size of installed user base
- Can you obtain a demonstration version for evaluation for a suitable period, e.g. 4-8 weeks?
- Make sure the firewall will integrate with your existing network configuration.

In the next section, we present a summary of these criteria in a “checklist” format.

6.5 Firewall Evaluation Checklist

This checklist summary is a concise version of the notes in the previous section.

Functionality	
Interfaces	
Maximum number	
Media supported	
10/100Mbit Ethernet	
Gigabit Ethernet	
Quad Ethernet card	
ATM	
FDDI	
Token ring	
Filtering	
Protocols covered by stateful inspection or proxies?	
At which layers can it filter?	
By source	
By destination	
By time of day	
Content control	
By site	
Virus scanning	
Content vectoring	
Java	
JavaScript	
ActiveX	
Authentication	
Passwords	
S/Key	
RADIUS	
SecurID	
TACACS	
Others	
VPN support	
Third party/freeware integration?	
NAT support	
Transparency	
Incoming	
Outgoing	
Intrusion detection	
Manually terminate a connection?	
Can it detect many DoS attacks?	
Can it react to a DoS attack?	
Reporting	
Level of detail	
by service and event type	
Remote logging supported?	
Are log analysis tools included?	
Can product alert you to attacks?	
How readable are the reports?	

User interface	
Remote	
Internal	
External	
Dialup	
Interface	
Console	
Browser-based	
Configuration tool	
How secure is the method for remote configuration?	
Is the interface easy to use?	
Does the interface allow an expert to fine-tune the settings?	
Fault tolerance	
Firewall security	
Firewall architecture	
Freeware support from users?	
Multi-system interaction	
Standards adherence	
Established underlying OS?	
Hooks for third party products	
Platform architecture	
Is the hardware/OS proprietary?	
Is the future of the OS good?	
Is there a robust Unix version?	
Can hardware be upgraded?	
Performance	
In silicon or software?	
Packet forwarding rate	
VPN encryption overhead	
Stateful inspection overhead	
Load balancing possible?	
Documentation	
Is the documentation comprehensive?	
Are printed manuals supplied?	
Are tutorials included?	
Technical Support	
Are support engineers available locally?	
How quickly can they respond?	
Is support available 24 hours a day?	
Availability	
On-site	
Phone	
Fax	
Email	
Support from vendor?	
Support supplied by third party?	
Is support included in the purchase cost?	
Support for hardware platform and OS	
Are patches made available regularly?	
Are upgrades part of the package?	

Cost	
Hardware	
Software	
Non-bundled extras, e.g. VPN support	
Educational (CHEST) discount offered?	
Installation	
Training	
Administration	
Support	
Upgrades	
Miscellaneous	
Independent certification	
Good product reviews in publications	
Company credentials	
How many years has the company been in the firewall business?	
How large is the product’s installed user base?	
Can the vendor supply reference accounts for contact?	
Is a demonstration version available?	
Will the firewall be easy to integrate into your existing network configuration?	

7 Introducing a Firewall

It is possible to deploy a firewall in default deny mode in a matter of hours. But, if the firewall is to form part of a site security policy that your users respect and want to work with, the approach should be more gentle and should be performed over a reasonable period of time (i.e. months rather than days).

In this section we describe the process we went through within ECS in what we called "Project Eclipse" (named after the 1999 summer of the total solar eclipse), from initial non-intrusive hardware deployment through to a full (inbound) default deny policy.

The feedback from our user base of some 1,300 staff and students has been very positive. The more vocal opponents of firewall systems at the outset of the project have reported that the policy adopted has worked much better than they had feared. Academic institutions often have the most difficulty in setting up a firewall solution due to the notions of "academic freedom"; users are used to being able to do what they want, when they want. In theory, they tend to want to experiment with a large variety of different and sometimes obscure network features. They may have a high likelihood of resentment to the firewall (and will try to circumvent it) if they do not feel involved in the deployment process. As we report later though, the number of users who do run "obscure" network protocols is not as great as might be intuitively expected.

7.1.1 Drivers for change

The driver for our own firewall installation was a security incident in the summer of 1998 in which the Southampton campus servers were hacked from a site in France. While the hacker was caught – he was reportedly working from Watford using Paris as a staging post, and was traced by an impressive line of co-operation involving multiple police forces – the damage was significant. Due to the topology of the network, all Departments were off the Internet for 5 days while the central servers were rebuilt. The chronicle of events is detailed on the University Web site [HACK].

As a result of the incident, our University convened the SOTON-CERT committee [SOTCERT], through which reviews of security practice have been enacted and through which the campus is also moving to a default deny firewall environment. ECS chose to act more quickly, though this was made easier by the move to a new building and a new network backbone technology.

It is also worth noting here that a University campus network should be designed where possible such that incidents (be they security-related or weather-inflicted) should be able to be isolated on the affected parts of the network. This may require internal firewalls for parts of the network, or it may involve a simple "pulling of the plug". Areas of the network can then be "quarantined" while the remainder can continue to operate.

7.2 Initial Non-filtering Deployment

ECS occupied a new building in the summer of 1998. At this point in time the Department's network topology changed somewhat. One of the design choices in rolling out the new ECS network was that we would have a single point of entry to that network. This meant running a private fibre to a remote building, rather than using our campus network to carry traffic to this remote group. But the private link meant that we could place that building behind our single point of presence, and thus behind any firewall deployed at that point.

The choice of firewall, as already discussed, depends on many criteria. ECS did not have a high-end Cisco router as its network gateway to the campus and beyond. Traffic volumes to and from the campus were reasonably high, in part due to NFS traffic exchanged between ECS file servers and public workstations on campus. The campus ran one flavour of ATM (3Com LANE) while ECS ran another (Newbridge MPOA). In terms of performance and functionality, a smaller solution such as GNAT Box [GNAT] was not sufficient, but nor were we of a size where firewalling in silicon (as possible on high-end Fore ATM switches) would be economical.

Our evaluations of firewalls, undertaken through the criteria mentioned in a previous section, homed in on software-based solutions that would run on a Unix platform. The main contenders were Firewall-1, Sunscreen, Raptor and Gauntlet. A number of independent reviews can be found on such products (e.g. PC Magazine [PCMFW1]), and these almost universally recommend Firewall-1. Sunscreen has similar functionality, but at the time of purchase (late 1998) its support for ATM interfaces was, based on our own tests, poor. We chose to run a Unix firewall for the additional stability offered. The UKERNA Risk Reduction Workshop [URISK] made a recommendation that critical security services should not be run on a Windows NT system. While a Linux version of Firewall-1 is en route, we chose the Solaris platform.

Having chosen Firewall-1 [CKPFW1] as the software solution, we sought a supplier. At the time we were buying, CHEST struck up a deal on Firewall-1 [CHESTFW1] through ESOFT [ESOFT], with an academic price of around £6,450 agreed. We in fact bought an enterprise licence through CenturyCom [CCOM], who were very competitive. We also bought an annual support contract for around £1,300. The basic licence does not include optional extras such as VPN security, which are extra. While this is not a cheap firewall solution, we decided that the investment in a fully-featured system, with a pedigree and good interoperability with other products, would be worth the cost.

7.3 Performance, Deployment and Cost Issues

Firewall performance is an important issue. A firewall not only has to cope with a certain average load, but it should also not drop packets during peaks of activity.

The deployment of a campus firewall as opposed to a departmental firewall may pose an interesting difference of requirements. A campus connection to JANET is most likely running at a slower data rate than an internal intra-campus network link to a department; in our own case the JANET link is up to 34Mbps, against a 100Mbps or 155Mbps ATM link to the Department. However, the rule set for a campus will almost certainly be more complex, because there are more hosts wishing to gain access to external services. The ability of a firewall to process a complex rule set is as important as its ability to process packets at the required rate.

7.3.1 Performance

It is perfectly possible for a site to measure its packet throughput on its external network link. This is a useful exercise in assessing performance requirements, though it is quite likely that functionality will rate higher as a requirement. In ECS’ case the data throughput varied from a background “trickle” of some 1,000 packets per second (less than 1Mbits) through to peaks of around 50Mbits during heavy file transfers (we had observed data rates through our external ATM network link of up to 6MB/sec to campus servers). With many sites deploying Gigabit Ethernet, those rates will only rise.

7.3.2 Hardware

We sought advice on hardware specification from other sites running Firewall-1; as a result we bought a 333MHz Sun Ultra 10 with 512MB of RAM. Such a system can currently be bought for around £3,500 at academic prices. Additional cost (around £1,500) came from the required ATM card and a quad Fast Ethernet (QFE) card. The QFE card would enable us to run a "collapsed" multiple DMZ topology, though we did not plan to make use of the DMZ until later in the deployment cycle. We have not had any performance issues with this specification.

Hardware support costs may vary with the supplier. A Platinum (2 hour) support contract from Sun is not overly expensive for an Ultra 10; ours was around £500 p.a. at the time of purchase. A more expensive alternative is to buy and hold spares; if this method is chosen one should be wary of issues such as firewall licences that use a host's unique ID.

7.3.3 Other costs

The cost in hardware and software is easy to assess; in our case approaching £15,000 for the Sun and Firewall-1. Staff time and training is less readily identifiable. A Firewall-1 training course costs some £1,200 or more; we have not sent any staff on such a course, but we have heard that the basic and advanced examinations are worth sitting to gain a useful certification. Staff time in deploying a firewall depends on how much evaluation and pre-deployment time is invested. While this project ran for over a year (the majority funded by the JISC), were a solution chosen rapidly and research and deployment focused and driven from existing reports and work, the staff time would likely amount to 2-3 man months. There would also be post-install support as an ongoing task, though we have found that that is relatively minimal (2-3 hours a week at most). More time can be spent educating and consulting users (as we did on this project), but the level to which that is done is likely to vary between institutions. We certainly recommend devoting as much time as possible, particularly just before and after the default deny changeover.

If a dedicated room is to be used for firewall management, the costs of such space should be included.

7.3.4 Installation disruption

A final point to consider is the installation downtime. Our firewall was deployed as part of a new network install, but the installation could have been done in well under an hour by pre-building the Solaris box and re-plugging the network live into a "default allow" IP forwarding system. Case Study A, described later, saw a downtime of less than a minute. There are many ways to make the transition quick and painless. Running an interior routing protocol such as OSPF between two routes that lead to the network egress point, one of which is firewalled, one of which is not, would allow a rapid changeover. Downtimes of many hours or even days should not be required.

7.4 Running a Watching Brief

Our initial installation saw Firewall-1 merely forwarding all packets between networks. No filtering was performed. The system allows you to log connections based on whatever criteria you choose without filtering. This basic "monitoring" mode is a very useful tool in assessing your required firewall policy.

It is possible to monitor all traffic entering and leaving a firewall, and thus to observe which protocols are being used by people on the internal and external networks. However, a site's users will most likely want to feel consulted about their network requirements, and thus a "social" survey of the users is a wise action to undertake. They may also know of future requirements that cannot yet be observed. Consultation forms part of the process of a phased deployment, along with education and dissemination.

7.4.1 User survey

We chose to issue a questionnaire to all staff and postgraduate students. We could have included undergraduate students, but did not for two reasons. First, the consultation happened to fall outside of term time (the beginning of the summer holiday), and second, we felt that teaching staff could reply on behalf of the students, knowing what tools were required on their taught courses. With hindsight, a student consultation may have been a good idea. We would ideally also run the survey in peak term time (e.g. March or October). We did however chat to the students informally, and they were consulted at later stages in the process. One should also not forget to ensure that all support staff understand the firewall deployment principles, and that no one member of staff alone holds specific operational knowledge.

Before we sent out the questionnaire, our University Computing Services (SUCS) had already applied some inbound packet filtering at the campus firewall. Connections to internal POP, IMAP and DNS servers (both zone transfers on TCP port 53 and client traffic on UDP port 53) were restricted to a limited set of hosts, in response to the summer incident which reportedly stemmed from an IMAP vulnerability.

The questionnaire is reproduced in Appendix A (with minor modifications), with a summary of responses listed in Appendix B (with some host names modified). Note that the level of technical knowledge in the responses varied greatly. The most interesting property of the survey was that very few users stated that they were running "obscure" services on "obscure" ports. This was our first clue that a default deny policy, at least on inbound data traffic, may be workable, and workable without a heavy administrative overhead.

We asked what services people accessed outbound from ECS even though we only intended to apply the default deny policy inbound. We felt this would raise the feeling of consultation. Of the inbound services declared, the main ones of interest were VNC (which allows remote control of a Windows PC), ssh (the secure shell – sessions are encrypted but other protocols can be tunneled inside the session), X11 (Unix windowing system) and telnet and e-mail access. Given that our e-mail was run from a small number of well-managed servers, the main "problem" we (apparently) faced was educating users that they would not be able to telnet to their desktop (though the number of such requests were small). Restricting remote telnet access would also allow us to check users of ECS resources, i.e. we could ensure they possessed a valid ECS account.

In hindsight, questionnaires geared to certain user categories (e.g. secretaries) would have been a good idea.

Perhaps the most useful point made in any reply was that the firewall policy must allow new services to be opened up at short notice, i.e. within 24 hours or ideally same day. While one might hope that users wanting a service enabled would plan ahead, invariably they will not, and emergency requirements can never be discounted.

7.4.2 Understanding the traffic

Interpreting observed network traffic is sometimes not straight-forward. What service runs on port 4500 for example? While common services are “well-known”, e.g. SMTP port 25, telnet port 23, POP port 110, others are less so. To understand the data flows, it is necessary to understand the port numbers in use. There is a very useful list of officially assigned port numbers, maintained by The Internet Assigned Numbers Authority [IANA]; the list of recognised port numbers [IANAPORT] is available online, and runs to over 100 pages of invaluable A4 text.

Whilst logging and watching the traffic, and surveying your users, it is also important to be pro-active in dissemination. We thus held open seminars on the proposed firewall system, and how it would affect users. Slides from one such seminar can be seen in Appendix C. The post-seminar discussions were very useful. One of these led to the suggestion that the firewall policy should insist on certain services being blocked to all bar systems staff-support machines, but that other services should be “default allowed” into the “default deny” regime unless there is a known security issue with the service. We thus accepted that the firewall deployment would be a risk reduction measure, not a risk elimination exercise (though no firewall can ever claim to achieve such a status).

Part of the education process also involved explaining why incoming user modems could not be allowed on desks; unmanaged back doors through the firewall would not be permitted.

7.5 Evolving default allow rules

Given that we were monitoring traffic in and out through the firewall, and we were consulting our users, how did we go about building the rule sets with a view to a final default deny inbound regime?

The questionnaire had given us an insight into the services our users perceived as important to them, although of course we know which protocols they really were using from our network monitoring statistics.

7.5.1 Classifying traffic

We felt the best course of action was to build a set of firewall rules that would allow us to model and classify the traffic passing through the firewall, without actually blocking any of that traffic. The initial deployment would be transparent to the users. By classifying the traffic, we mean that we could add rules to the firewall rule set that would be matched by packet flows, e.g. a rule to classify inbound e-mail to a set of e-mail servers, which could then be logged as an event under that particular rule number. Our aim was to classify as much traffic as possible, creating groups of hosts which would fall into categories such as ‘incoming-smtp-servers’, ‘login-servers’ or ‘ftp-servers’. We accepted that initially we may have 20 or more of each, where ideally we would want less, but to get a grasp of the existing network flows, we felt this was the best way to proceed.

What sort of services did we commonly see? The most obvious were:

FTP	File transfer
NNTP	Usenet news
SMTP	E-mail delivery server
POP	E-mail access

IMAP	E-mail access
HTTP	World Wide Web
TELNET	Remote login

How do the firewall rules work? In Firewall-1, you specify rules as a triple of source IP address, destination IP address and a service. The source and destination may be any host or group of hosts. The service may be any source or destination port, or group of ports, or a particular packet pattern specifiable via a (INSPECT) script, or it might be a remote resource (e.g. a pointer to another filtering server). In a simple case, that might be a connection from Host A to Host B on port 25 (SMTP). If the triple matches, then an action can be defined, normally to accept or reject the connection. All rules are applied in the order they are listed in the rule set (this set is called the "firewall policy" on Firewall-1).

Note that the initial connection is the item that is filtered upon, so if an inbound SMTP TCP connection is allowed, subsequent traffic back out to the calling client will be allowed. Also, in Firewall-1, rules apply across all pairs of interfaces, so the source and destination addresses are matched against both interfaces. This is usually meaningless in a two-interface firewall host, but will be much more significant in a three or more interface host, e.g. in a collapsed DMZ topology. Our initial pre-default deny topology had no DMZ.

In our discussion we do not mention specific hosts, for (we hope) obvious reasons.

We were quickly able to draw up a rule set of some 20 or so rules matching on the basic protocols, with the final rule being the "catch-all" rule which would be matched by any connection not matching any of the rules above it. The rules took the form of

- If the inbound connection on a service (e.g. FTP) is to an "approved" server (e.g. one of the 'ftp-servers' group) then allow the connection.

- Otherwise deny the connection for that service.

Or put another way:

Source	**Destination**	**Service**	**Action**
Any	ftp-servers	ftp	accept
Any	ECS	ftp	reject

As time went by, we would refine the rule set such that the catch-all rule would see less and less data. The catch-all rule would allow traffic matching it, but would be "flicked" to deny traffic to make the transition to default deny.

7.6 Adding the first inbound deny rules

We realised that we would have to begin introducing deny rules at some point. We decided that we would monitor the standard services (telnet, ssh, rlogin, rsh, HTTP, FTP, SMTP, POP, IMAP, NNTP, NFS, etc) for three months before introducing any deny rules. Having monitored a service for that period of time we would be confident that we had identified all internal hosts running those services, such that if we blocked those services to other internal hosts, the chances of blocking an important service would be minimised.

Three months from the initial firewall advertisement to our users we, with notification, added an inbound block on the common protocols to all but the groups of identified hosts running those services. In making that step, we would be able to limit the potential for new internal hosts to offer those services. We decided that if users did ask for such services to be enabled, that we would "default allow" those requests, provided they were not for one of seven key protocols: rlogin, telnet, ssh, rsh, POP, IMAP and SMTP. These protocols were the ones that, by our own security policy, would only be allowed to run to hosts maintained by systems support staff. The immediate beneficial effects were that

1. Any user accounts in the Department had to be created on systems that were only in the control of support staff; thus, for example, a postgraduate student could not create an account on his desktop Linux PC for a friend in America.

2. Users could not run an "ISP" like service offering e-mail accounts to non-Department members.

3. Inbound e-mail would only be able to be received by systems-supported e-mail servers, ones properly configured to reject "spam-relay" attacks. Spam-relay is a method generally employed by commercial users to route massive e-mailshots through a third party host. The danger is that the victim can be added to e-mail blacklists such as ORBS or RBL, and as a result the host (or network) may be blacklisted for e-mail from legitimate users, or worse, blacklisted for network traffic.

Note that secure shell (ssh) is an interesting protocol. It is beneficial in that it allows secure encrypted sessions, yet "weak" in that users can tunnel other protocols inside ssh. This poses something of a dilemma for firewall administrators. It is also worth noting that protocols can be tunneled in other services, such as HTTP. The latter is a technique that can be used, for example, to access Quake (a multi-player PC game) servers.

7.6.1 Consulting the users

At this point we still had some users running these "systems supported only" services. Part of our dissemination and education process involved going to the users running the services that we had made host groups for (e.g. 'ftp-servers') and asking them why they ran (for example) FTP to their desktop when they could FTP from a systems-supported infrastructure login server. In many cases the users were very happy to run the services from central servers rather than their own. This allowed us to whittle down the size of these host groups, slowly but surely. We were also able to advise and help users install the latest (and most secure) versions of services where they insisted on running them on their own workstations.

At this time we also began blocking well-known security threats on services that the logs showed our users were not actively running. This included services that JANET-CERT [JCERT] advisories recommended sites should block, e.g. Back Orifice [BO]. For example we put global inbound firewall blocks on the following services:

Service	Port number	TCP/UDP
Back Orifice	31337	UDP
Net Bus	12345 12346 20034	TCP
SNMP	161	UDP

	162	
finger	79	TCP
TFTP	69	UDP
NetBIOS	137 138 139	TCP/UDP

We had no complaints from users about any of these blocks, bar one postgraduate who had his own software running on port 12345 (a reasonable “random” port number to pick), but who was happy to move his service to another port. It is not generally necessary, for example, to have SNMP (network management protocol), TFTP (a simple file transfer protocol) or finger probes available to external users.

We chose to block Back Orifice outbound as well to prevent any of our own students (or staff) abusing any other sites.

7.6.2 An evolved Firewall-1 rule set

Combining security blocks with our classification of hosts into allowed service groups, we evolved the following (approximate) rule set by three months into our firewall policy building process.

RULE	SOURCE	DESTINATION	SERVICES	ACTION
1	Any	ECS	BackOrifice Netbus	reject
2	ECS	Any	BackOrifice Netbus	reject
3	Any	ECS	finger	reject
4	Any	ECS	SNMP	reject
5	Any	ECS	TFTP	reject
6	Any	ECS	lpd	reject
7	Any	ECS	nbname nbsession nbdatagram	reject
8	Any	dns-servers	DNS-TCP	accept
9	Any	ECS	DNS-TCP	reject
10	sucs-nis-servers	nis-servers	NIS	accept
11	Any	ECS	NIS	reject
12	Any	mail-access-servers	pop-2 pop-3 imap	accept
13	Any	incoming-smtp-servers	smtp	accept
14	Any	ECS	pop-2 pop-3 imap	reject
15	Any	ECS	smtp	reject
16	Any	www-servers	http https	accept
17	ecml-clients	ecml-server	https	accept
18	honours-clients	honours-server	https	accept

19	Any	ECS	http https	reject
20	NatGallery	iip-server	telnet ftp	accept
21	Any	iip-server	telnet ftp	reject
22	Any	ftp-servers	ftp	accept
23	Any	ECS	ftp tftp	reject
24	Any	login-servers	shell login telnet exec ssh	accept
25	SUCS-CLUSTERS	ug-login-servers	telnet ssh	accept
26	Any	ECS	shell login telnet exec ssh	reject
27	SUCS-CLUSTERS	UG-VLAN	X11	accept
28	cardlock-mmc-server	cardlock	Cardlocks	accept
29	Any	cardlock	Any	reject
30	cvs-clients	cvs-server	cvs	accept
31	SUCS-CLUSTERS	cad-licence-server	cad-licence	accept
32	sucs-mbone-relay	mbone-relay	mbone	accept
33	SUCS-CLUSTERS	ug-nfs-servers	NFS-SGI	accept
34	Any	ECS	NFS-SGI	reject
35	sucs-webcache	ecs-webcache	squid-3130	accept
36	Any	news-server	nntp	accept
37	Any	ECS	nntp	reject
38	Any	Any	Any	accept

These rules included some special services, such as a web cache peering, access to a software licence server from external hosts, imaging access from the National Gallery, access to our cardlock PC from University maintenance and control staff, and access to services such as a CVS repository.

7.6.3 Initial technical issues

We had two SSL-enabled Web servers handling our ECML (requisitions) and Honours (financial system) database front-ends; these we chose to only make available to named hosts outside the Department over HTTPS (the secure, encrypted HTTP protocol). Further discussion of SSL occurs later in this report in Section 9. This meant that any external access to sensitive information would be relatively secure.

At this stage we were logging all accepted and dropped connections for each rule, bar the last one (the outbound connections). We also configured Firewall-1 to detect the SYN denial of service attacks and to detect IP spoofing (a packet appearing on an external interface with a forged internal source IP address). Both are useful features of Firewall-1 and many similar firewall products.

The DNS server traffic was limited to our DNS servers, and our DNS servers protected locally to only accept DNS zone transfers from approved hosts (external secondary DNS servers). We applied similar protection for our NIS servers; NIS can be protected such that servers only reply to requests from certain IP subnets. We blocked inbound NTP traffic, as we were not offering any network time services.

7.6.4 Blocking SMTP

Restricting SMTP traffic to the `incoming-smtp-servers' meant that we had some level of protection against spam-relay e-mail attacks, given that our allowed SMTP servers were correctly configured to prevent spam-relay abuse. We chose not to block SMTP outbound as we knew that some staff and students were using "friendly" external mail relays from mail clients on campus, and we were happy that if anti spam-relay measures were in place, we would be unlikely to generate spam e-mail internally.

However, when you do block SMTP inbound, but not outbound, you need to be careful to cater for mail that goes out but which might not then be received back. Consider a user sending out an e-mail from host abc.ecs.soton.ac.uk with a From:' line that refers to the full host name and not just the ecs.soton.ac.uk domain. If the recipient tries to reply to the full domain name, the firewall will likely block the connection. Thus users either have to ensure their e-mail client is properly configured, or they have to use a conformant 'Reply-To:' line, or the domain's DNS will need an MX entry for every host pointing at the real mail server to use when talking to hosts under that domain. When the latter technique is used, the mail server will need to be able to understand mail targeted at hosts under the main domain.

The patterns of SMTP entries in our early logs were quite revealing. When we did still have a small number of open SMTP relays during the initial rule set evolution process, we could see clear evidence that once one external "spammer" had successfully attacked a machine, there would be a significant number (20-30) of repeat attacks from other spammers within a day or two of the first incident. The spammers were obviously trading lists of open relays. However, our stance on anti-spam (as being undertaken through a separate JISC JTAP project) appears to have dramatically reduced the number of attempted SMTP relay attacks through our site.

7.6.5 The portmapper

The portmapper (TCP/UDP port 111) is an interesting protocol because it enables other protocols via Sun RPC (remote procedure calls). One such protocol is NFS. It is unlikely that sites will want to make use of remote NFS mounts, but many archive sites do enable NFS access to their files, so users may request NFS mounts of certain external services. Whether internal services need to be made NFS-available is another issue; most likely client workstations may be mis-configured and exporting NFS file systems to the "world", so a default block on NFS exports from a site is generally wise. We did not observe any NFS exports during our monitoring period, bar our arranged NFS exports of student file store to some campus workstations. Such an export is a considered risk (one risk being the modification of trusted access files such as *.rhosts*). Firewall-1 allows individual RPC services to be enabled, and additional ones to be defined, but such services should be handled with care.

So, at this stage we had employed blocks on some “dangerous” traffic, and we had contained access to common services (Web, FTP, etc) to known groups of servers. We had added rules to match the most common inbound IP traffic, and were seeing only a small proportion of traffic fall through into the “catch-all” final rule.

7.7 Defining and establishing a security policy

Getting a security policy drawn up, approved, and accepted is not necessarily an easy task. At our own site, our security policy has evolved as part of the Department computer use regulations. These are subject to approval by Departmental Board meetings at which representatives of all Department members are present.

Our Department is subject to University Calendar statutes, which must be honoured. For example, Charter/Statutes Section 30, Construction 1(a) [page 99] states the University should *“ensure that academic staff have freedom within the law to question and test received wisdom, and to put forward new ideas and controversial or unpopular opinions, without placing themselves in jeopardy of losing their jobs or privileges.”* This sort of statute implies freedom to make available and access information, and it is most likely that many universities will have similar statutes.

However, a university will also wish to safeguard its computer systems and data, and in the light of the Data Protection Act 1998 it has a legal obligation to take reasonable steps to do so where (sensitive) personal data is concerned. Similarly the university will have a corporate liability to ensure that if offences are committed that originate from its systems, it is not seen to be negligent in offering non-authenticated access to such systems.

At Southampton, the hacking incident of the summer of 1998 has enabled Computing Services Committee and then Senate to pass measures to allow a default deny security policy to be operated by Computing Services.

7.7.1 What questions should a firewall policy answer?

It is perhaps not prudent to publish the precise details of the Department’s firewall policy here, but the questions that should be considered in formulating the policy should include:

- Who is responsible for implementing the policy?

- Who should be approached when disputes over policy occur?

- Who approves the policy and the updates to it?

- How is the policy reviewed – by whom, by what process, and how frequently?

- Which data services can be enabled through the firewall on request?

- What is the process for registering new service requests?

- What is the maximum turnaround time on new service requests?

- Will such services be monitored for appropriate use, and if so by whom?

- Which data services may only be run to systems staff-supported hosts? These may typically include rsh, rexec, rlogin, telnet, ssh, pop, imap and smtp.
- Which standards and practices should be followed? One might include BS 7799.
- Who has access to what data?
- Should secure communication channels be used, and if so where?
- Should authentication be used, if so where and to what strength?
- If a host is believed to have been compromised, what action can be taken and by whom?
- What are the users' responsibilities?
- How is the policy disseminated, by whom and to whom?

Such a policy, expressed in a document or via regulations, should form part of a complete data security policy, which should probably also include paper-based systems.

7.8 Going default deny

Our timeframe for moving from a default allow environment to a default deny one under Project Eclipse spanned many months.

Date	Activity
August 1998	Firewall host installed. No services blocked. Standard services monitored.
January 1999	Notification to users of intention to go default deny. Blocks imposed on standard services to non-registered hosts and on known "security threats". Firewall seminars held.
June 1999	Firewall rule set evolved to a state where the "catch-all" rule matches were minimised. Ongoing refinement of rules for specific service requests. Default rule still set to "allow". Date for default deny announced.
September 1st 1999	Default rule switched to "deny". Further rule refinements via service requests from users.

The key statistic to observe during the evolution of the rule set and the migration to default deny was the number of connections which were logging under the final default allow (and then deny) rule.

Once we were confident that all the major services were catered for in our rule set, and that users had been consulted about the traffic falling into the final catch-all default allow rule, we were in a position to go ahead with the change to an inbound default deny policy. This was achieved by simply changing the catch-all rule action from “allow” to “reject”.

7.9 Six months on

Our firewall is now at the state where, as a risk reduction measure, we have a balance between ease of use and a reasonable level of security and protection for our 1,300 Department members.

There will always be trade-offs in policy between security and ease of use, e.g.

- We allow X11 connections from some public clusters on campus.
- We allow most forms of inbound ICMP traffic.
- We allow logins to some hosts from outside, though where you can go from there is limited.
- We have some relatively open hosts in our DMZ.

Such decisions have to be made by each site based on that individual site’s needs.

7.9.1 Summary of five weeks of blocked firewall events

We studied a summary of our firewall log for a five week (35 day) period starting in February 2000 and running into March 2000. The intention was to generate overview statistics on what types of traffic were being blocked the most under the default deny regime. We generate daily log files from our firewall which are processed with our own scripts. The issue of logging is discussed in a later section. We did not at this time log all inbound firewall traffic, e.g. the campus workstation NFS traffic was not logged, so the total inbound connection count is less than the true figure.

Service	Connections	% of Blocked
All logged inbound	1,297,745	-
Rejected inbound	174,510	100.0
HTTP	73,175	41.9
SMTP	38,664	22.2
External probes (20+)	13,843	7.9
ICQ	8,115	4.7
NetBIOS (name)	6,086	3.5
NTP	4,678	2.7
33434 (traceroute)	3,908	2.2
3128 (unused webcache)	3,772	2.2
finger	2,795	1.6
FTP	2,011	1.2
NNTP	1,828	1.0
Socks	1,403	0.8
6970	1,205	0.7
32775	1,183	0.7

telnet	1,044	0.6
6446	988	0.6
6666 (Mud server)	543	0.3
7777	506	0.3
Napster (6699)	496	0.3
sunrpc	267	0.2
8000 (web server)	170	0.1
NetBIOS (session)	101	0.0
ssh	22	0.0
IMAP	5	0.0
Other to dialup hosts	1,070	0.6
Other	6,632	3.8

The statistics showed that we were logging around 250,000 “interesting” inbound connections per week, of which around 25,000 were connections that were blocked by the firewall. Or put another way, 3,500 connections per day, the vast majority of which we are very confident are unwanted connections, are being blocked. The average data rate for traffic into our Department over the period was observed to be approximately 3Mbits/sec, with peaks at up to 50Mbits/sec.

The most significant levels of blocking are occurring with Web (HTTP) and e-mail (SMTP) traffic. In the case of HTTP, the vast majority is to redundant web servers, and a lot of the hits to these come from “web crawlers”. Our users make much wider use of our centralised Web server(s) now, rather than attempting to run personal “vanilla” Web services from their desktop machines. Unfortunately there is direct no equivalent of an MX record (a service redirection from one host to another) for HTTP, though firewall content-vectoring may be able to offer some benefit in this area.

The vast majority of the e-mail blocking appears to be related to “random” attempts to send or relay spam through our site, in particular e-mail being targeted at or through our main Web server host. The use of MX records should reduce the chances of genuine e-mail being blocked; if all e-mail to a domain is hidden behind a single mail domain (e.g. username@ecs.soton.ac.uk) it is important the users understand that, especially those configuring their own machines (whether or not the MX technique is applied).

The level of access to dial-up hosts was not insignificant; clearly there is a need to consider protecting those who choose to dial in for access. Southampton is unusual in that it has a historical free local off-peak calls arrangement via Videotron as were (Cable and Wireless as they then became, and NTL as they are now). This means some students like to occupy dial-up lines for lengthy periods, and presumably run services off their connections.

7.9.2 Unsolicited probes against our network

External probes very rarely now sweep all ports on all hosts in a domain; attackers know this will significantly raise the chances of them being detected. It is more likely that you will see activity such as the following (a real event from our log, from the same source IP, to one of our hosts in February 2000):

Destination	Port number/service
x.ecs.soton.ac.uk	6699 (napster)
x.ecs.soton.ac.uk	6700
x.ecs.soton.ac.uk	http

x.ecs.soton.ac.uk	ftp
x.ecs.soton.ac.uk	telnet

In this case a very selective probe to just one host was made. And of course it only takes one insecure firewall hole to let an attacker in.

In the five week analysis period above, we had one significant probing attack on one of our dial-up lines (through our Portmaster primary rate dial-up service), i.e. an attacker was attempting to probe a wide range of ports (some 1,700 or more) on a host connected via dial-up.

Probes on our network as a whole were numerous, with at least 20 being noted by their pattern, with more that may have been probes (but which were in any event blocked). Ports probed included a general sweep, a small range (e.g. 33459-33470) and specific ports on many hosts, e.g. the POP2 port, or port 1243.

We observed two instances of IP spoofing being used against our network in the five week period.

The firewall is an invaluable tool to hide information by blocking external scanning of a network, as well as protect against the majority of potential attackers.

7.9.3 Services that are hard to filter

There were a number of services which we initially found very hard to filter effectively using Firewall-1, and which we feel most firewall products will have problems with. These are services that a user initiates an outbound connection to on one TCP port, and where the server then tries to initiate a connection back to the client on a new TCP port. Firewall-1 can match up inbound TCP traffic on the same connection as the outbound call, and it can match up UDP replies with UDP packets sent out (up to a set timeout). But it has little built-in support for the services which specifically require new ports to be opened up inbound in response to a service request outbound.

We discuss these services further in Section 8 of the report. A typical example is the ICQ chat service, a facility that is very popular with staff and students alike.

7.9.4 Maintenance on Firewall-1

We were originally running Firewall-1 4.0 under evaluation on Solaris 2.7 with co-operation direct from Check Point. When we decided that Firewall-1 was the product we wished to run with permanently we bought the licence from CenturyCom [CCOM], along with one year of technical and upgrade support. We learnt that CenturyCom have since abandoned educational pricing on Firewall-1, leaving ESOFT [ESOFT] as the best educational distributor via their CHEST agreement.

Our future support contracts will be taken out through ESOFT, who at the time of writing are attempting to negotiate academic (or CHEST) deals for SecurID products and the Trend Interscan VirusWall e-mail and FTP/HTTP virus scanner.

7.9.5 Outbound filtering

We have discussed inbound filtering as the focus of this firewall report. There are some areas where we experimented with and used outbound filtering. Some of these were defensive measures to deter our own students (and staff) from abusing other sites (c.f. the corporate liability issue) and others were to help keep the use of resources “fair”.

For example:

1. Napster [NAP] is a system that can be used to share files over the Internet. Its current hot use is for gathering collections of music files, in particular MP3 format files. As of early 2000, our Computing Services determined that 10% of its IP traffic off campus was Napster. While Napster can be used legitimately for knowledge sharing (e.g. there is talk that the Human Genome Project is considering the use of Napster to share gene information), the University has chosen to block Napster traffic [NAPB], and will enable the service to individual hosts on request. However, there are other methods for MP3 files to be shared, e.g. Gnutella [GTELLA]. While this sort of filtering is contentious, it does deter University members from acquiring illegal MP3 files, and will reduce bandwidth charges significantly.

2. Back Orifice. As previously discussed, we blocked internal users from being able to attempt exploit the Back Orifice vulnerability in Windows systems on hosts off-site. This had no negative impact and the measure is still in place.

3. We have some PC’s which are dedicated to CD writing or image scanning. To encourage their “proper” use we chose to prevent Web browsing from these machines to external hosts. This policy has been accepted and remains in place.

The police have made clear their stance on a site’s responsibility to ensure its users behave responsibly. There may thus be more pressure to run firewall filters that either increase the level of authentication for users connecting off site from “transient” access points (e.g. docking stations), and/or that prevent the exploitation of security vulnerabilities on other sites.

In terms of saving on bandwidth, SUCS is considering using the firewall to redirect all Web (HTTP) traffic through a site-level Web cache. This would have the effect of reducing bandwidth usage (and charges for bandwidth, given the cache peers with the JANET Web cache) off site. By using a Firewall-1 redirect the users would not have to change any settings on their workstations. The firewall host would run an interior routing protocol that would enable it to allow HTTP traffic through directly if both Web cache hosts failed. While this may not be a popular policy with users, the realities of bandwidth charges may force the issue on to universities. Presumably those users or departments wishing to retain a direct Web connection could choose to do so if they picked up the bandwidth cost.

7.9.6 Interaction with Computing Services

The introduction by our Computing Services of a campus level firewall poses the dilemma of how to manage the dual systems. The options are:

1. Allowing Computing Services to manage the ECS firewall remotely. The Firewall-1 management GUI enables remote management very easily. This feature is handy, for example, for managing a collection of firewalls, or routers (such as those from Cisco) running Firewall-1 modules (which can be relatively inexpensive).

2. Running both firewalls as separate administrative entities, with the ECS firewall offering additional protection for ECS from attacks from campus hosts. ECS could install outbound filters to help improve security of core non-user login campus servers (by denying direct

access to them from inside ECS). In this scenario, the campus can in some ways be treated as a rather large DMZ by the Department.

3. SUCS pass all IP traffic direct from the JANET point of presence to ECS for it to filter as it sees fit. This would imply some level of trust between Computing Services and the Department, which may exist at Southampton but not at all universities.

While ECS and SUCS work closely at a technical level, problems could arise where there may be policy differences, e.g. should ECS choose to declare ICQ as a valuable educational tool that should be enabled, while the SUCS view were that it should not be allowed.

The department in a tiered firewall hierarchy is likely to want access to the logs connections to its hosts that are being rejected by the campus firewall.

There is also the issue of how users submit service requests. Do they submit one to the Department and one to the University? Or do they submit one to the Department that then makes a request to the University if the service is enabled through the ECS firewall? The crux of the issue is that the user is able to have the service enabled quickly and efficiently.

Where universities and their departments each run firewall services, it is vital that they co-operate and work together as closely as possible. One method by which to achieve this is by setting up a campus-wide CERT team or committee, at which such issues are openly discussed. Southampton has its own CERT team [SOTCERT].

8 Operational Firewall Issues

In this section we discuss the "day to day" issues that have arisen since we deployed our firewall. This includes notes on testing the firewall, what servers and services can be placed in a DMZ, how to analyse log files, and discussion of technical issues.

8.1 Firewall Testing

Once a firewall has been installed, it is a wise precaution to test whether it's performing the way it is intended to. The usual method to do this is by "port scanning" the hosts behind the firewall to see which services can be seen from the external network. Scanning should be done on a regular basis. It should be scheduled, or at least advertised, so that users on workstations on the internal network will be prepared for the scan. We have not experienced any hosts crashing from the scans we have performed, but the possibility should not be overlooked.

Scanning can be arranged through an external "professional" company, but this can be very expensive. It is probably better, or better value, to arrange for the scan to come from a topologically adjacent University site, or another department at the same site in the case of a departmental firewall. Where a De-Militarised Zone (DMZ) exists, scanning should be performed across all combinations of interfaces. It is most likely that the location the test scan is done from will have much better (faster) connectivity to your site than an attacker, so a local scan will be the best test of resilience and compliance.

The scan should look at all ports (0-65535). If a restricted scan is performed, consider which ports you must inspect – the IANA Port Number listing [IANAPORT] is very useful in this respect, as it is for understanding the function of open ports.

It is possible to buy commercial scanning software, such as one by Internet Security Systems [ISS], but these are typically very expensive, with little or no academic discount. We were, for example, quoted £15,000 for a licence to run the ISS scanner on a Class B network (a typical 16-bit campus subnetwork). In contrast, the SAINT (Security Administrators' Integrated Network Tool) vulnerability tester, from World Wide Digital Security Inc [WWDSI], is available for free, and was the best scanning package we could find. It is open source, well maintained, and very easy to use. SAINT is a more advanced version of the more well-known SATAN (Security Administrator's Tool for Analyzing Networks) software. Another commonly used probing tool is mscan [MSCAN].

A scan to a large department network may generate many millions of firewall log entries; it is important that your firewall can cope with such "extreme" logging conditions. We found that our own version of Firewall-1 seemed to cope without problems.

SAINT generates some very useful and easily configurable vulnerability reports. Shrewd network administrators should run SAINT inside their firewalls to help protect their servers and systems from internal attack.

8.2 Technical Issues

One of the most important aspects of firewall and security policy maintenance is to keep up to speed on the latest security developments. To that end, sources such as CERT [CERT] are invaluable. For technical assistance with Firewall-1, we found the Phoneboy site [PHONE] to be excellent; this collection of Firewall-1 information and utilities is a must-read.

8.2.1 Getting the more complex services running

There are many services which are desirable to have running but which need some special consideration. In this section we present some tips and pointers for a few such services. Our notes refer to Firewall-1, but the information should be fairly generic:

1. **Mbone traffic**. The mbone uses IP in IP tunnels, which can be detected and matched by a Firewall-1 INSPECT script that looks for IP protocol 4, i.e. ip_p = 4. You may also find that you need to match against the Internet Group Management Protocol (IGMP) as well, which is IP protocol 2, i.e. ip_p = 2. Since your mbone feed will likely come from one host, you can improve security by requiring the connection to come from that host.

2. **IPv6 tunnels**. Where IPv6 is tunneled in IPv4 (the current version of IP), IP protocol 41 is generally used, for which an INSPECT script should match on ip_p = 41.

3. **ICQ**. This chat software requires connections to be opened from the ICQ server(s) to the client host. You must thus either trust source IPs and source ports (very bad, an an attacker can spoof the IP in an attack), or use a proxy service instead [ICQP]. The SOCKS proxy is one such solution, and is discussed below.

4. **RealPlayer**. This video streaming protocol should be supported by most firewalls; you can check out a list of the ones that are at the Real site [REALF].

5. **NetMeeting**. The Microsoft site carries information on getting this videoconferencing tool running through a firewall [NETMF].

If your firewall product of choice doesn’t support a given protocol, the most likely solution is that you will need to run a proxy server.

8.2.2 Proxy servers – SOCKS

The SOCKS proxy is the result of standardisation processes ongoing in the IETF [IETF], in particular in the Autheticated Firewall Traversal (AFT) Group. In short, the client application communicates with the proxy, and the proxy communicates with the remote server. If the proxy is run on a more “open” part of the network (e.g. in a DMZ), it can trust the remote server more than the client is able to. Direct communications from server back to client are not possible, thus protecting the client.   Note that client applications must be configured by the user(s) to use the proxy service, e.g. an ICQ user must enable the SOCKS proxy option.   It is however also possible to forcibly redirect protocols via a proxy (e.g. forcing all HTTP through a Web cache).

Sites such as COAST (Computer Operations, Audit, and Security Tools) [COAST] and CERIAS (Center for Education and Research in Information Assurance and Security) [CERIAS] carry more information on the SOCKS proxy.   There is a free reference version of SOCKS5 available to academia from the NEC Networking Systems Laboratory [NECSOCK], though not as many clients support this version of the protocol yet as support SOCKS4.   The ICQ client, for example, supports both, but this is not typical of all applications.

The NEC Laboratory have also made a commercial version of the proxy called e-Border [EBOR], for which a trial version is available. The focus of current SOCKS and e-Border development at NEC is currently on streamed multimedia applications. The rise in popularity of services such as IP telephony, as exemplified by services such as Dialpad [DIALPAD], may make the use of proxies more common in firewall deployments, whether as part of an integrated firewall or as a standalone proxy server.

8.3 Use of a De-Militarised Zone (DMZ)

A DMZ offers a “buffer” network between your internal and external network. Hosts in the DMZ can be protected by firewall rules between the external and DMZ interfaces (in the case where the DMZ is one interface on a firewall), and the internal network can be protected from the DMZ by further rules.

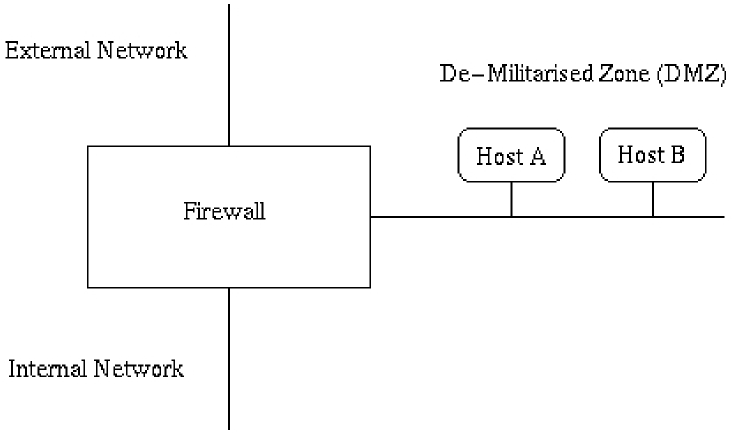

The firewall topology at ECS

In the case of Firewall-1, rules apply across all combinations of source and destination objects, and thus with three paths (external to DMZ, DMZ to internal, and external to internal), you have to be aware that any rule may match against any path. At the same time you must recognise that the rules are applied in the order they are listed in. Adding a DMZ adds a lot of value to your firewall-protected network, but it also adds some complexity.

Our own firewall now runs in conjunction with a DMZ. The DMZ was added just before we went default deny, and we have been migrating services to the DMZ as it has become practical to do so. Those wishing to deploy a DMZ from the outset may wish to consider adding some of the services listed here to their DMZ from an early stage of planning.

Our current or near-future DMZ services include:

1. **Dialup.** Our Portmaster dialup service is located in our DMZ. The dialup device requires an authenticating server, which can be placed either in the DMZ or internally, assuming the traffic is enabled. All users connecting via dialup are thus assigned an IP address from our DMZ IP subnet.

2. **SOCKS server**. This proxy is a recent addition and handles our ICQ traffic. We are running with the NEC reference version [NECSOCK].

3. **External Web servers.** We have some Web services which are aimed at an external audience only, and which may have administrators logging in remotely. We locate these in our DMZ.

4. **IRC/talkd.** To run an IRC server you can enable a hole in the firewall for the IRC traffic, or (if the Direct Client Connection mode is a concern) the IRC server can run from the DMZ. The same applies to the (Unix) talkd.

5. **Wireless Access Point.** We have our "public" 802.11 Breezecom wireless access point(s) located in our DMZ. Any users wishing to use wireless access in our seminar rooms will be served an IP address in our DMZ subnet by the DHCP server in the DMZ.

6. **Docking Station Points**. These include seminar room network points, and other (colour-coded) points in the building. Visitors and students using these points will get a DMZ IP from our DMZ DHCP server. This implies (for user convenience) that IP addresses are also served by DHCP for users connecting from their normal network locations (e.g. a member of staff bringing a laptop from an office to a seminar room). Hosts on these non-authenticated points should be challenged for authentication before traffic leaves the site (e.g. via an authenticating proxy or a challenge-response on the firewall).

7. **External DNS server(s)**. It is generally a good idea, but a notable additional expense, to host a slave DNS server in the DMZ. This DNS server is advertised to the outside world as your primary DNS server, meaning that all external DNS requests are resolved through it. Your internal master DNS server is thus protected from external access. To be more robust, you may want two such servers.

8. **External e-mail server(s)**. Similarly, the same can be said for e-mail. The MX servers advertised to the external Internet should sit in the DMZ, with the "real" internal servers not accessible from external hosts bar the DMZ "relays". This is in effect another proxy DMZ service.

9. **Hosts for specific research projects.** If the project staff require root or administrator access on their hosts, and perhaps access from outside the network, then it may be wise to connect the host in the DMZ. This will depend on where the host is situated and the physical network topology.

10. **Student-run hosts**. We have some student-run hosts which require root access for the students; we place these inside our DMZ.

Servers that are intended for internal use only, such as an InterScan VirusWall server or a SecurID ACE server, are best kept on the internal network.

8.4 Issues with Firewall-1

Our experience with Firewall-1 has been positive. Of the few quibbles we have, the following are worthy of note:

1. The Solaris OpenLook interface has a limit to the number of host or service objects it can display in a pop-up menu. If you want to have more than 120 or so host or network objects, the pop-up window simply won't appear. You either have to uncheck the "show in menu" option for the host, or buy the Motif GUI licence (an extra £600). This is somewhat poor.

2. When removing objects from the hosts or services lists, Firewall-1 will not check which rules or groups they are present in. This can cause problems when the rule sets are reloaded (the reload process may even fail). In a handful of instances, manual rule file editing may be required.

3. If a host is renumbered, you have to load up and change its Firewall-1 definition. All host to IP mappings are static. This may be a good security measure, but it is a chore if you renumber a couple of Class C size subnets with a number of firewall-listed hosts resident in them.

4. If the firewall is presented with a port number for a non-RPC service that it needs to check against an RPC rule, the port number may not be cached and the firewall will then query the portmapper of the target machine. We were seeing portmapper requests on a wide variety of machines from the firewall until we took care over our rule set construction and ordering. If a private IP address is used for the firewall link, this can also cause some problems.

In general however, the firewall policy and rule set management functionality is very good.

8.5 Behind the Firewall

A firewall only reduces the risk of an external attack on the network. Good practice should still be followed on the internal network, e.g. with respect to applying security and other patches, using sensible passwords, restricting access to hosts (e.g. via tcp_wrappers on Unix machines), acting on CERT warnings and advisories, etc. Back-door routes into the network should be restricted, e.g. inbound personal office modems, or dual-homed links to commercial leased lines. Such good "defence in depth" practice is beyond the scope of this report.

Online sites of interest include the Security Portal [SPORT] and Security Focus [SFOCUS].

Mailing lists to subscribe to include *uk-security@mailbase.ac.uk* (general security information and alerts) and *firewall-admin@mailbase.ac.uk* (set up as part of this project).

Other useful resources include the monthly Crypto-Gram newsletter [CGRAM] by Bruce Schneier and the Computer Incident Advisory Center [CIAC].

Good practice can extend to physical network design, e.g. using Ethernet switches instead of hubs to reduce the likelihood of network packet snooping being successful (though it is worth noting that switches can be duped into sending "private" traffic down the wrong interfaces).

8.6 Analysing Firewall Logs

The logs from a firewall are the only means by which attempted attacks can be detected, and from which statistics on the firewall actions can be generated. These notes are geared up to Firewall-1, but should have general relevance.

The Firewall-1 logs can be exported automatically (e.g. via a Unix cron job) on a daily basis to a text file format. The daily log file will vary in size depending on your logging level and network activity. A 100MB plain text file per day may not be uncommon (though it may compress to less than one tenth of that size). Given the file is plain text, you can devise home-grown software (typically Perl scripts) to generate statistics and to look for "dangerous" patterns in the logged entries. Firewall-1 generates two logs, an activity log and an accounting log, so you can use Firewall-1 to generate (for example) Web bandwidth used per IP subnet (with appropriate logging rules).

However, Firewall-1 has a very large user community, and this has led to excellent sites such as Phoneboy [PHONE], from which free utilities such as *fwrule42.pl* can be downloaded. This indispensable tool generates an HTML page from your firewall configuration file(s), letting you view all your rules, objects and policies from a (secure) remote Web server. Another good free package is *fwlogsum* [FWL], which generates overview reports.

The commercial log analysis tools include the Check Point Firewall-1 Reporting software. This is only available on the Windows NT platform, but the price is not unreasonable at around £900 to academia. The functionality is good, but you may find you can do everything you need to with your own scripts from the (free to generate) plain text files.

Another commercial offering is the WebTrends Firewall Suite, which costs $1500. However, we found that the trial version couldn't handle large log files, so we didn't pursue the product further. The package crashed after (very) slowly reading one 300,000-event log file. It does reportedly generate Excel output though, which may be useful.

At present, we generate daily log and accounting files. The Data Protection Act 1998 [DPA98] raises the question of how logs are kept, since personal workstation activity could be seen as personal data. The DPA 1998 also states that data should be kept for no longer than is necessary, and that the subject should be informed of the data gathering process, and what the data is being used for.

One potentially useful service that we are considering implementing at present is to offer to e-mail a user a log of blocked connections to their workstation on a daily or weekly basis. Users would then see (some of) the data being collected, and potentially be able to react to suspicious activity. Of course, it may be that someone may argue that the IP or host address of the connecting machine in the log may also be personal data…. Interesting times ahead…

8.7 Intrusion and Virus Detection Systems

Given that some firewall "holes" exist, e.g. for Web and e-mail traffic to enter the internal network, consideration needs to be given to associated risks and dangers.

An Intrusion Detection System (IDS) will look for patterns of network traffic that might indicate a Denial of Service (DoS) attack is underway, or it may look for certain sequences of data in Web or other traffic where that sequence indicates that a security hole is being exploited. A simpler IDS might be a module that has 50 or so "signatures" that runs on a Cisco router. Firewall-1 has an anti-SYN attack mechanism built in, called SYNDefender.

An IDS can run on a host or in tandem with a firewall. Check Point have RealSecure as the "partner" product to run alongside Firewall-1. RealSecure runs the Check Point Suspicious Activity Monitoring Protocol (SAMP) and is geared to detect DoS attacks including SYN floods, Smurf, ping floods and the “ping of death” [CERT]. However, RealSecure is not a cheap product, costing a similar amount to Firewall-1 (around £6,000). If funds are limited, investment in a firewall has to be seen as more useful than investment in an anti-DoS tool; recovering from a DoS attack is in general a lot easier than from a severe hacking incident.

Check Point’s OPSEC (Open Platform for Secure Enterprise connectivity) integrates and manages all aspects of network security through an open, extensible management framework. One of the strengths of Firewall-1 is its Content Vectoring Protocol [CVP] support. This allows Firewall-1 to delegate inspection of an application layer connection (SMTP, HTTP or FTP) to another process or host, and have the result returned to the firewall. The CVP API is open, so third party developers can make products that support CVP. Defining CVP objects in Firewall-1 is relatively simple.

One such product is Trend InterScan VirusWall [ISVW]. While many products that scan for viruses in e-mail run on NT only (e.g. MIMEsweeper, F-secure products, or Norton), VirusWall runs on Solaris, and it supports CVP. Thus, rather than having a virus-checker run a one-to-one filter by being added as a site’s primary MX host, VirusWall can check for e-mail viruses on any number of mail servers on a site.

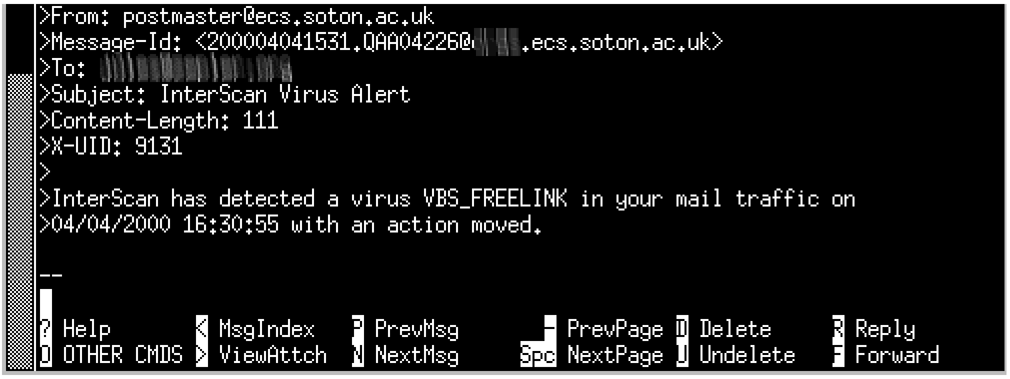

```
>From: postmaster@ecs.soton.ac.uk
>Message-Id: <200004041531.QAA04226@[redacted].ecs.soton.ac.uk>
>To: [redacted]
>Subject: InterScan Virus Alert
>Content-Length: 111
>X-UID: 9131
>
>InterScan has detected a virus VBS_FREELINK in your mail traffic on
>04/04/2000 16:30:55 with an action moved.


--

? Help        < MsgIndex   P PrevMsg      - PrevPage D Delete     R Reply
O OTHER CMDS  > ViewAttch  N NextMsg    Spc NextPage U Undelete   F Forward
```

VirusWall alerts a user to a virus that was in a mail addressed to them.

VirusWall has the additional advantage that it can use CVP to also redirect FTP or HTTP traffic through the same virus-scanning module. Thus one host can virus scan all the major vehicles by which viruses can enter a site. It is also smart enough to understand some 20 or so compression and archiving formats. We have run the trial version successfully with a number of viruses, and (as advertised) viruses embedded in .tar, .tar.gz, .tar.Z and .zip archives are detected, with notification to the sender and recipient. The management is done via a simple Web interface, and new virus patterns can be automatically downloaded daily.

VirusWall looks to be a good product, but, and it’s a big but, the pricing is excessively high. While a 250 user licence is some £3,000 or so, the cost ramps up with a 20,000 user licence costing £65,000. At the time of writing ESOFT (who have the CHEST Firewall-1 agreement) are looking to strike an academic deal on VirusWall, perhaps with CHEST, preferably with a flat rate cost as per Firewall-1 and RealSecure. Without such a deal it is unlikely to be good value in comparison to the cost of a site licence for a host-based virus checker (which a site may also wish to run).

8.8 User Issues

The most obvious “user problem” is that users will dislike what the firewall represents. In our experience, we have found that this is almost universally not the case, but that fact may in part be due to the way the firewall was introduced.

The other issue is the fact that when a firewall is known to be in place, many observed “network” problems will be blamed on it. Some of these may be justified, the majority most likely will not, but all have to be handled considerately. Comments we received included:

- “Does newsgroup reading have anything to do with firewall default deny policy? I’ve noticed that I cannot access any newsgroup with Netscape recently.”

- “Hi, I was wondering whether the firewall is stopping access?”

- “So I was wondering if some new firewall thing has got in the way. Or is this some other glitch? I know the service is working elsewhere on campus.”

- “I traced the port to 1558 but don't know if its tcp or udp - I think its a straight one to one so tcp? netstat didn't seem to show it... ergh...”

- “The curious thing is that I never had any symptoms like these prior to September 1st.”

It took around a month for the firewall doubts to settle down, by which time users were probably more accustomed to the way the firewall worked. In one instance it turned out that a problem was actually with a remote site who had just introduced a default deny inbound firewall policy of their own (and the remote user had no idea that a firewall had been put in place at their site).

9 Secure Transient Firewall Access Methods

The original JTAP bid and specification for this project had a greater focus on secure transient access methods through firewalls. However, having discovered just how few universities run a default deny inbound firewall (we estimate less than 10%) that focus shifted to more general firewall deployment issues in academia.

This section describes our experiences and recommendations for secure access methods. The issue of secure (encrypted) sessions must be understood to be different from the issue of authentication (proving who a user is) and authorisation (granting privileges based on who a user is).

We will report in more detail on these issues in our JTAP 659 "Secure Internet Protocols" project, due for completion in July 2000 (see our interim report [JTAP032]).

9.1 User Authentication

One method of authorisation relies on a simple username and password combination, e.g. Unix passwords or the S/Key one-time passwords. The current trend in authentication is to centralise the authentication service, such that users can have a single password for an entire domain, e.g. via RADIUS [RFC2138], TACACS+ or an ACE Server [SECID]. The username/password model is still widely used, but it of course does not guarantee that the person supplying the username and password is the user who was issued with the combination originally.

Another approach is to issue digital certificates to users (e.g. X.509 or PGP [PGPI]). These will typically need to be unlocked with a pass-phrase to be used. Certificates have a private and a public part. Public certificates can be shared via a Public Key Infrastructure (PKI) model, or by a directory services such as LDAP, or perhaps by a dedicated public key server (as per PGP). But again, it is possible for a certificate to be lent to another user, or to be compromised.

SOCKS5 offers the facility to authorise users as proxy connections are established. This is one method by which users on otherwise non-authenticated systems (e.g. docking stations) can be authenticated before sending any traffic off site.

Encrypted sessions between hosts, or Virtual Private Networks (VPNs) can be set up when a user is carrying an appropriate private key. The remote server can assert the user's identity and initiate an encrypted session. Commercial encryption products include the DataFellows F-Secure VPN tools and Check Point's SecuRemote (which is designed to integrate directly into Firewall-1).

Much thought has to be given to the ease of use for the end user, who may be not be very computer literate. The solution has to be simple to operate, and ideally transparent to the user.

Here we first report on a SecurID evaluation, and then on public domain (free) tools for enabling secure data sessions. We foresee strong authentication as the next logical step after initially securing network access (via firewall and proxy technology) and ensuring connections are secure (encrypted) where appropriate.

9.1.1 SecurID tokens

We evaluated the ACE server v3.3 from Security Dynamics/RSA on our Unix systems. The system manifests itself as a key fob or credit card sized token that displays a 6-digit number which changes every 60 seconds. The tokens are in sync with the ACE software server, such that any user authenticating themselves via securID can do so by keying in their current token number after a 4-digit PIN. Traditional username and password techniques can be used in combination with SecurID.

The ACE server management software is a little primitive in style, but allows user profiles to be set up, and tokens to be assigned to users. The token may be configured to only be valid for certain hosts on the network. Hosts that are to be SecurID authenticated require client software installed (which talks to the ACE server) and require the user's local shell to be replaced by the SecurID shell (which then invokes the real shell after a successful authentication).

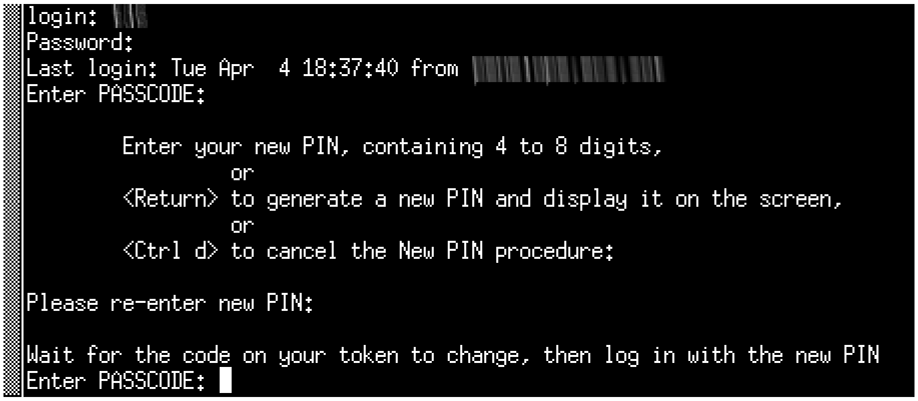

```
login: [redacted]
Password:
Last login: Tue Apr  4 18:37:40 from [redacted]
Enter PASSCODE:

        Enter your new PIN, containing 4 to 8 digits,
                 or
        <Return> to generate a new PIN and display it on the screen,
                 or
        <Ctrl d> to cancel the New PIN procedure:

Please re-enter new PIN:

Wait for the code on your token to change, then log in with the new PIN
Enter PASSCODE:
```

Logging in for the first time with a new SecurID token

There are a number of potential drawbacks to the SecurID system, including:

- If the fob or card has no value to the holder, they may lend it to others. While systems staff would be unlikely to do this, students probably would. Ideally the SecurID token would be incorporated into something of value, such as a mobile phone or a cash-carrying smart card.

- The lifetime of a token may not match a typical student 3-year cycle; extending the lifetime has a cost. The administrative and financial costs of token replacement need to be minimised.

- The credit card token is a little bulky – the key fob is more practical. However, we have had key fobs break after less than a year's use.

- Instructions on Firewall-1 challenge-response integration (from RSA or CheckPoint) are poor.

- Having to type a 10-digit code for every login can become a chore. Authentication comes at a price to convenience.

- The anti-replay attack measure, which prevents two successful admissions on the same PIN and passcode, also prevents the user from being able to open multiple telnet windows in quick succession.

- The system is proprietary; the issue of ease of integration into user applications is not clear, e.g. if the SecurID client "agent" were required to do a web-based authentication for access to a web server area. ACE server v3.3 (which we evaluated) includes code samples for integration, but the object files which must be linked in only run on selected supported platforms.

- The PIN and username-password combination can still be snooped for a telnet session, or grabbed from a vulnerable X11 session or a watched user. If the token is lost, the system is then vulnerable.

- The OS support for the system is not complete; this affects what the system can be used for. RSA have dropped support for Irix [ACEAG], which is the operating system that our own file servers and the University's file servers run. There is also no support for Linux or FreeBSD, which are operating systems of choice in many universities. This is obviously a major drawback. Also, new OS versions, such as Solaris 8, are not supported even under the latest ACE server 4.0 at the time of writing. This illustrates a major weakness of such a proprietary solution.

- How many users can be authenticated per second on the server? We have not been able to run tests for more than 10 users. It may be possible to "drown" the server in authentication requests.

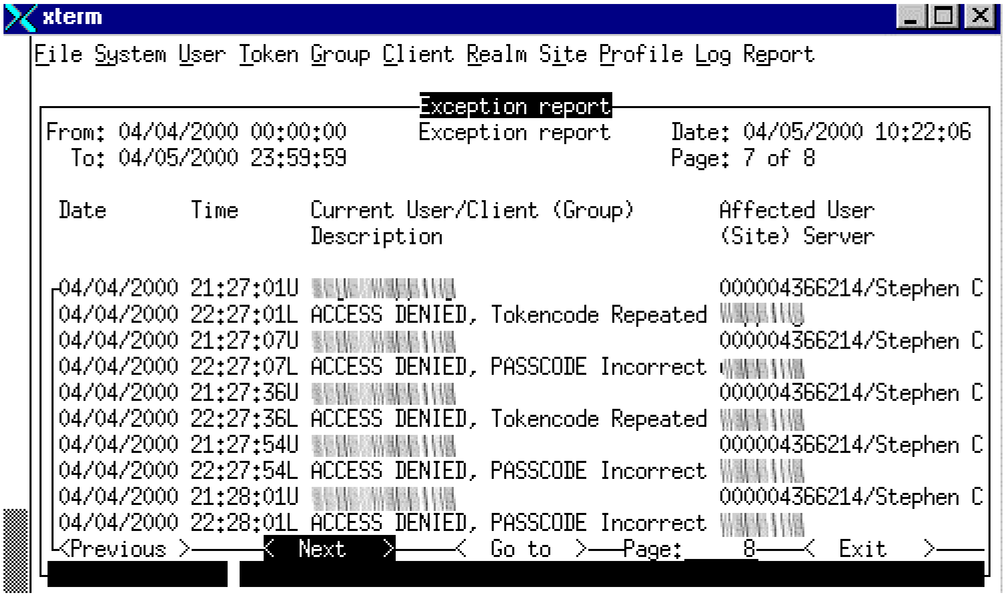

```
File System User Token Group Client Realm Site Profile Log Report

                               Exception report
From: 04/04/2000 00:00:00      Exception report     Date: 04/05/2000 10:22:06
  To: 04/05/2000 23:59:59                           Page: 7 of 8

 Date       Time      Current User/Client (Group)       Affected User
                      Description                       (Site) Server

 04/04/2000 21:27:01U [redacted]                        000004366214/Stephen C
 04/04/2000 22:27:01L ACCESS DENIED, Tokencode Repeated [redacted]
 04/04/2000 21:27:07U [redacted]                        000004366214/Stephen C
 04/04/2000 22:27:07L ACCESS DENIED, PASSCODE Incorrect [redacted]
 04/04/2000 21:27:36U [redacted]                        000004366214/Stephen C
 04/04/2000 22:27:36L ACCESS DENIED, Tokencode Repeated [redacted]
 04/04/2000 21:27:54U [redacted]                        000004366214/Stephen C
 04/04/2000 22:27:54L ACCESS DENIED, PASSCODE Incorrect [redacted]
 04/04/2000 21:28:01U [redacted]                        000004366214/Stephen C
 04/04/2000 22:28:01L ACCESS DENIED, PASSCODE Incorrect [redacted]
 <Previous >      <  Next   >     <  Go to  >  Page:     8   <  Exit   >
```

The SecurID management menu – checking login failures.

For all its potential flaws, the SecurID system is still attractive. It is used by companies such as BT, for example. The cost is high: at commercial rates over £20,000 for a 500 user ACE server licence with £4,000 p.a. maintenance. The fobs are £50 each. Clearly this does not scale effectively to a campus-wide solution for over 10,000 people. Academic rates are currently being negotiated, e.g. by ESOFT and Chernikeef, but there are no confirmed prices at the time of writing. It would thus seem that SecurID only lends itself to small deployments, e.g.

1. A small team of systems administrators wanting to restrict access to a small cluster of infrastructure machines and servers. However, with Irix, Linux and FreeBSD not supported, this may not be practical.

2. A small team of financial administrators seeking to restrict access to servers holding sensitive financial data, with perhaps Web-based SecurID authentication required to view Web-based accounts driven from Windows NT systems.

One useful property of SecurID tokens is that they can be shared between domains, e.g. the same token can be imported to a university ACE server and a BT ACE server if a staff member is doing work with both their university and BT.

The SecurID token is in effect an authenticating certificate. If it is lost along with its PIN and user information it is no more or less secure than alternative certificate-based systems.

As an alternative, it would be possible to use secure shell (ssh) in a mode where access were only granted if the user authenticated with a private key unlocked by a passphrase. The problem in such a case is making the private key available to the user wherever they are. Carrying a floppy disk around is feasible but not too practical. Carrying the private key on a laptop is also possible. The ssh tools are generally freely available.

One-time passwords, such as those offered by S/Key, are another option.

9.2 Making ssh available over the Web

We will report more fully on ssh and friends in our JTAP 659 project report; here we cite some early findings an recommendations.

It is possible to use ssh [SSH] to just encrypt sessions (i.e. leaving the authentication to the Unix or other system at the remote end where the ssh daemon is running). This at least prevents passwords being carried plain text over public networks. One way to make ssh available to remote users is via a client they carry on a floppy disk or portable. Putty [PUTTY] is one such client (for ssh version 1), being lightweight (around 200KB for the Windows version) and only requiring the executable to run.

A more interesting method is to use a Java ssh client that can be served from a Web page at the user's home site. If the Web site is also SSL-protected, the session is doubly secure. One such client is MindTerm [MINDT], which is freely available with source code under the GNU Public Licence. MindTerm supports colour and mouse modes, though it is also only an ssh version 1 client at present.

A version of the SSHv2 server has recently been made available for non-commercial use [SSH2].

9.3 SSL for secure web servers

While ssh enables secure use of passwords and data in transit in a login session, encryption of access to a Web server generally requires different technology. The most common method of achieving this is to run an SSL certificate (e.g. from Thawte [THAWTE] or Verisign [VERI]) to a Web server running with SSL support (e.g. Apache with OpenSSL [OSSL]).

An SSL (X.509-based) certificate costs as little as $100 from a “trusted” certificate authority (CA) such as Thawte. With support for Thawte and Verisign certificates built into common Web browsers – i.e. MS Internet Explorer and Netscape – the encryption is simple for the user to work with. Self-signed certificates can be used, but these force users to jump through a series of potentially confusing pop-up menu “hoops”.

We bought $1,000 worth of Thawte certificates, for user/e-mail and Web/SSL use. By buying in bulk ($1,000 or more) Thawte allows you to manage your own certificates. One certificate applies to one host name only (e.g. www.ecs.soton.ac.uk), so it is economic to put your secure services on one domain host rather than splitting them across several (so long as you are confident the host is secure on all fronts). We use SSL to protect Web pages where users enter passwords, and to protect Web servers holding sensitive data, e.g. Department finances.

Note that Thawte were recently acquired by competitors Verisign. This has not affected our use of our Thawte certificates.

9.4 SSL for secure e-mail access

While ssh can encrypt login sessions over a public network, it does not protect the other main remote “transient” user requirement, e-mail access. Typically remote e-mail access will be presented via POP or IMAP. We deploy IMAP in preference to POP within ECS for its more natural support of users accessing e-mail from a variety of locations and for the way it manages remote folders. Unlike most POP-based mail access implementations, it also has the advantage that it will not fetch large attachments unless requested by the user.

There are a number of Web-based IMAP e-mail clients available. IMP [IMP] is a very good free client. Users can access their e-mail remotely via a Web client that talks to IMP running on their home Web server; that server then communicates with the user’s IMAP mail server. The advantage with this system is that access to the IMAP server is restricted to the Web server and internal hosts, and thus IMAP vulnerabilities will be less easily exploited externally.

To secure the mail session, we run SSL on the Web server. This ensures that passwords and all e-mails viewed in the session by the user are encrypted as they is transported over the network. We are running WebIMAP [WEBI] as our server in ECS; this has a small cost (around $200 to academia), but we were impressed enough with the authors’ other well-known product, the DNEWS news server, that we chose to run with WebIMAP. It has proven to be very easy to customise and install, and so far it looks to be a good product.

Experiences of this pilot will be reported under JTAP 659.

9.5 Use of Transport Layer Encapsulation for Secure Access through a Firewall

There are other ways to secure communications to remote users, which may not be so intuitively easy to deploy, but which we are investigating under JTAP 659. In this scenario, we use stunnel [STUNNEL] for remote communications.

In the scenario, a user outside of their home network requires secure use of internal services, such as e-mail reading or internal USENET discussion groups from an insecure, remote Foreign Network. The home network is secured with a firewall at the network’s connection point to the Internet, where the site is employing a default deny policy only allowing incoming connections on pre-configured ports.

It is assumed that the remote machine that the user is working on has either supported applications installed or the user has the ability to run a service redirector application. Current software that is freely available and that uses this Transport Layer Security [TLS]/SSL technique for secure access to e-mail and news includes Outlook Express on Windows and Solaris, and Netscape Communicator on Windows and most Unix platforms.

To realise this scenario, a staging machine is configured outside the secured home network, in a DMZ for example, to offer tunnel endpoints for the remote applications to talk securely to the core services in the home network. This requires the firewall to be configured such that the secured services from the remote network to the staging machine, and the insecure services (the normal SMTP/NNTP connection) from the staging host to the home network are permitted through.

Clients on the remote machine configure their TLS/SSL-enabled application to access the secured services on the staging host. If their application does not have this functionality for secure communication, the remote machine needs to run an external relay that forwards communication securely.

9.5.1 Security

The tools used in our experiments permit three levels of authorisation at the point of setting up the secure communication channel: None; Cert status; and Cert comparison. All three modes rely on the encapsulated protocol for authorisation, should it be required of the application.
In the first instance, where just TLS is required without authority, neither side of the tunnel carries out authentication checks. The staging host's X.509 certificate is only used to seed the channel encryption for the session.

The second mode will only permit the client to communicate with the staging host (and thus the core service in the home network) if the X.509 certificate presented in the link set-up stage of the connection has characteristics that match the status of the staging host. For example, a common strategy used here is to match the domain name encoded in the certificate against the name of the target machine.

The third mode compares the certificate offered by the staging host with that held locally by the remote machine, only permitting communication if the two match. This shared-secret style approach is suitable for scenarios where the remote user has been issued with a disk containing a site-certificate, or is using their personal laptop (with the certificate resident) on a remote network.

9.5.2 Hazards

If the encapsulated protocol does not offer any form of access authorisation, (e.g. standard SMTP) then the service is only as secure as the TLS/SSL communication channel. If either mode 1 or 2 above are employed anyone would, for example, be able to route mail through the home network as if they were physically on-site.

If the staging machine is compromised - a higher risk for a host in a DMZ than for a host in a protected network - the internal services that are routed through it are then vulnerable to Man-in-the-Middle, Service Spoofing and Denial of Service attacks. As with all DMZ services, the status of the tunneling software has to be closely monitored as to verify its serviceability. This is also a reason for not running the tunnel endpoints on the services machine itself.

9.5.3 Example configuration (and current test service)

The current test configuration at ECS offers secure USENET news for internal newsgroup browsing, IMAP for mail retrieval and SMTP for mail delivery. Our staging host is currently a SPARC Ultra-5 running production Solaris 8 code, with the OpenSSL [OSSL] libraries installed for X.509 certificate management and TLS/SSL applications.

The stunnel package [STUNNEL] by Michal Trojnara is a freely available reference implementation of TLS/SSL tunneling for Windows and Unix platforms where OpenSSL has been installed.

The Staging host

Our staging machine sits in our DMZ network and has three tunnel endpoints configured for SNEWS, SIMAP and SSMTP [IANA], redirecting incoming requests to our production (unsecured) services.

The configure script in this case is:

```
/path/to/stunnel -d snews -r nntp.ecs.soton.ac.uk:nntp
/path/to/stunnel -d simap -r imap.ecs.soton.ac.uk:imap
/path/to/stunnel -n ssmtp -d ssmtp -r smtp.ecs.soton.ac.uk:smtp
```

The Firewall machine

The firewall between the ECS network and the outside world needs an additional rule to permit the routing of the tunneled protocols through it:

```
Src: ANY
Dst: staging-host
Services: snews (tcp/563), simap (tcp/993), ssmtp (tcp/465)
Rule: Accept, with log
```

To allow access between the staging host in our DMZ and the internal ECS network servers, we also require the following three rules in the firewall configuration:

```
Src: staging-host
Dst: nntp.ecs.soton.ac.uk
Services: nntp (tcp/119)
Rule: Accept, with log

Src: staging-host
Dst: imap.ecs.soton.ac.uk
Services: imap (tcp/143)
Rule: Accept, with log

Src: staging-host
Dst: smtp.ecs.soton.ac.uk
Services: smtp (tcp/25)
Rule: Accept, with log
```

Client remote machine running a TLS/SSL enabled application

For example, in Netscape Communicator 4.52's preferences, under *Mail & Newsgroups* preference tree:

In the *Mail Servers* pane:

> Outgoing Mail Server; set server to "`staging-host:465`"
> Select "`always`" for *Use Secure Socket Layer (SSL) for outgoing messages*
>
> Add Incoming Mail Server; set server name to "`staging-host`", of type "`IMAP`"
> Under *IMAP* tab, select "`Use secure connection (SSL)`"
> Check IMAP server directory under *Advanced* tab

In the *Newsgroups Servers* pane:

> Add new Newsgroups Server, set server name to "`staging-host`"
> Check *Support encrypted connections (SSL)*
> Ensure port number is set to **563**

This will configure the client to use the secure sessions.

Client on remote machine not using TLS/SSL enabled application

Existing applications that aren't SSL-aware can be used securely by employing an encapsulating proxy process on the remote machine. In our experiments, this has been done using the *stunnel* package on the client running in "client mode" to encapsulate and forward packets to the relevant *stunnel* process on the staging host. The following script brings up the client-side tunnels to the staging host on the home network:

```
/path/to/stunnel -c -d 40119 -r ssl.ecs.soton.ac.uk:snews
/path/to/stunnel -c -d 40143 -r ssl.ecs.soton.ac.uk:simap
/path/to/stunnel -c -n ssmtp -d 40025 -r
  ssl.ecs.soton.ac.uk:ssmtp
```

And then, to demonstrate the tunnel is working, we configure a Netscape session not using its TLS/SSL capabilities, pointing instead to the local tunnel endpoints:

In the *Mail Servers* pane of Netscape Communicator's preferences dialog:

> Outgoing Mail Server; set server to "`localhost:40025`"
> Select "`always`" for *Use Secure Socket Layer (SSL) for outgoing messages*
>
> Add Incoming Mail Server; set server name to "`localhost:40993`", of type "`IMAP`"
> Under *IMAP* tab, select "`Use secure connection (SSL)`"
> Check IMAP server directory under *Advanced* tab

In the *Newsgroups Servers* pane:

Add new Newsgroups Server, set server name to "`localhost`"
Check *Support encrypted connections (SSL)*
Change port number to **40563**

Note that the port numbers used are above the 0-1024 range, enabling user processes to create the socket endpoints. If the configuration were required such that the tunnel endpoints replaced the actual service ports on the remote client machine (e.g. *stunnel* listening on the real localhost port 25, forwarding to the staging host securely), then the process invocations would need to be run as the Super User (*Administrator* on Windows, *root* on Unix).

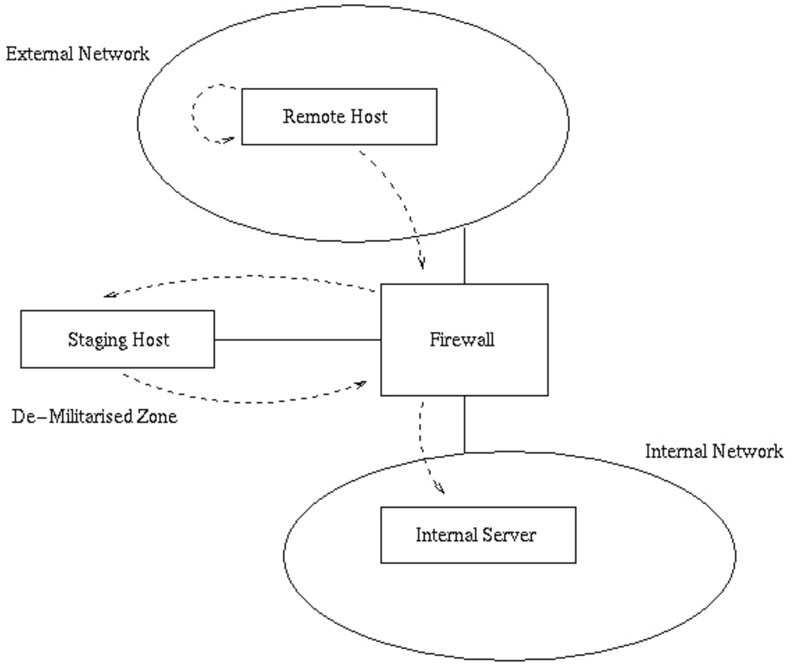

The data flows for a secure external connection using stunnel with client encapsulation.

9.5.4 Summary

As a mechanism for enabling secure access to data across an insecure network, the TLS/SSL mechanism presented here is fairly successful. However, if unsupported software is used on the remote client machine, the extra load on the user to bring up a remote tunnel endpoint could be too high for the general or casual user.

This technique is best suited to applications where user authentication is done as part of the application protocol (e.g. IMAP), as the existing TLS/SSL implementations do not cover this granularity of authentication.

An alternative technique, similar in complexity to the one presented here is to use Secure Shell Port Redirection. Running ssh with redirection is one of the focuses of our JISC Secure Internet Issues in UK HEI project, JTAP-659.

By the nature of TLS/SSL services, only connection-based (TCP) services can be secured in this manner. Whilst it is technically feasible that connectionless services could be secured over a

TLS/SSL link, the overhead in repeatedly setting-up and tearing-down secure connections is probably too great to render the encapsulated service usable in practice.

10 Case Study A: Commercial Spin-off Firewall

During the course of the JTAP project, ECS was involved in the deployment of two smaller-scale firewall systems. We report on these here through three different authors as the case studies may be relevant for sites looking to firewall just a small part of their network. A typical instance might be an administrative or financial unit. In such cases, a full-blown Firewall-1 solution may be excessive. Our questionnaire (see Appendix, question 4 responses) showed that some users believe that a site firewall is sufficient to protect their sensitive data, but in many cases firewalls within a site will be very wise precautions.

In this first example, we consider the installation of a firewall for a small spin-off unit from the Department that carries out commercial activity on a commercial leased line. The network is not connected to JANET. We refer to the site as the "Unit".

10.1 Requirement

The Unit has a 9-bit IP subnet containing Web, Mail, DNS and SQL servers running Windows NT and Linux. This network needed to be protected from external attack while still allowing the commercial customers full access to services they have paid for.

Due to the small scale of the Unit, a low-cost product was desirable.

The internal (protected) network is an Ethernet network running over 10BaseT cabling. The external network is 10BaseT to a Cisco 2501 router owned by a commercial Internet Service Provider. This in turn is connected to a 256Kbit leased line.

10.2 Choice and Survey of Firewalls

We considered firewalls operating as

1. static packet filters,
2. stateful inspection systems (dynamic packet filters) and
3. proxy servers

Most high-end firewall products combine all 3 modes of operation.

Static packet filters alone are not considered to provide strong enough protection and are usually confined to routers. Proxy servers are the safest but have a very high overhead. There is also the major problem of implementing a new proxy for every new service required by a customer.

Stateful inspection systems have a much lower overhead than a proxy server and are usually also faster. However, they also provide much stronger protection than static packet filters. The firewall remembers the state of every connection and bases its decision on allowing/denying the packet on all the OSI level data in the packet. It can therefore effectively screen all packets, including high port numbers and UDP traffic.

Most low to mid-range products appear to either run on Windows NT or a stripped-down version of Unix. The stability and security of Windows NT and its ability to survive long up-times are still questionable, so this class of products (e.g. CommandView [CMV], at £629 for a 25-host licence at the time of investigation, and CyberGuard [CBG], at £967) was dismissed.

Some systems use their own proprietary hardware and software (e.g. BIGfire [BIGF], at £8,000) but these were dismissed on the basis of cost.

Other firewalls are software products, based on a stripped-down, hardened variant of Unix, the most popular choice appearing to be BSD. Examples include BorderWare Firewall Server [BWARE] (£1,920 for 25 hosts), Firewall-1 [CKPFW1] (£6,000 through the CHEST academic agreement, but more for a commercial copy) and GNAT Box [GNAT] (£995). The author considered building a "free" public domain solution using Linux and its "IP Chains" firewall kernel modules, but this was dismissed as requiring far too much time to setup and maintain. Ongoing ease-of-maintenance is an important issue.

The final product chosen was GNAT Box from Global Technology Associates Limited. This product can be bought as a hardware and software combination, but in this case the software-only version was used.

The GNAT Box is a surprisingly capable piece of software, especially given its size: the entire firewall including configuration fits on a single 3½" floppy disk. It supports a wide variety of network interface cards, from old 10Base2 cards through to FDDI and 1000BaseSX Gigabit cards, while also supporting dialup connections using PPP via modem or ISDN. Administration of the firewall can either be done remotely using the supplier's own Microsoft Windows program or a WWW interface, or done locally on the firewall console.

The list of software features is very extensive and beyond the scope of this document. See the GNAT Box Web site [GNAT] for more information.

10.3 GNAT Box Hardware Requirements and Installation

GNAT Box can run on very low-cost hardware. To support around 32,500 simultaneous connections, it requires no more than a PC containing a floppy disk drive, 32 Mbytes RAM and 2 network cards. It does not use a hard disk, so these are best removed to improve the reliability of the system. The display console is entirely text-based, so an old VGA card is quite sufficient. It also makes little use of the CPU, so an old Pentium or even a 80486 CPU will work (though a Pentium system is recommended).

As the GNAT Box fits on a floppy disk, more than one image can be created. One disk was created containing a simple "allow everything" firewall rule set, and another created containing a more useful rule set. The firewall was initially setup using the "allow everything" rule set, so that the potential for disturbing customers could be kept to a minimum. Liasing with the ISP by phone, the firewall was plugged in at the same moment that their router was re-configured to support it. As a result none of the customers ever noticed its installation, and the actual down-time was under one minute.

The rules were constructed by adding a rule for each service that was known to be provided to customers, keeping an "allow everything but log it" rule active. This "catch-all" method is similar to that used in the deployment of the main ECS firewall (described earlier in this report). The logs were analysed every couple of days to see if any rules had been forgotten, and over the next 2 weeks the rule set was refined to a point where the last rule was changed to "deny everything but log it".

Some disallowed traffic appeared to be very common (e.g. various Microsoft Windows attacks, Back Orifice probing) and was generating a large amount of logs. So further rules were added so this traffic was blocked but not logged (the choice of logging depends on what action you might wish to take based on log entries). GNAT Box is generally set up to log remotely via the Unix syslog mechanism (if it has no hard disk, it has to do so), so the less logging required, the better for network performance. After a further week or so, the amount of disallowed traffic being logged had dropped to an acceptable rate.

With time, various rules have been added to enable staff outside the Unit to use specific machines to develop software for Unit customers, but otherwise the ruleset has remained static.

10.4 Conclusions

The GNAT Box product has proved itself to be a very capable firewall. It costs very little to deploy (£1000 and an old PC) and is easy to set up and manage. The installation involved less than one minute of down-time, and at no point have any of the commercial customers complained of any problems with accessing their servers.

The firewall is monitored and managed remotely, and the system’s console has not been used since the day it was installed. There have been no stability problems in the 3 months it has been in operation, and it has not required rebooting or otherwise interrupting the network service.

This deployment is somewhat different to the academic scenario. The smaller scale and scope of the firewall allowed a more rapid progression through the rule set evolution before a default deny environment was enforced.

11 Case Study B: Internal Firewall

In this case study we consider the use of internal firewalls for small departmental LANs. These notes refer to installations undertaken at Southampton.

A single firewall at the external connection to a campus network does not prevent parts of that network from attack by users who have access inside the firewall. This potential threat might include tens of thousands of users. This is one of the reasons for internal firewalls, which also provide additional layers of protection in the event of a security breach of outer firewalls. As a general principle, internal firewalls can be more restrictive than external ones. Internal firewalls provide a level of defence between individual machines and the campus level, protecting parts of the network, perhaps just a few offices.

The ultimate internal firewall is of course the machine itself - individual machines should be configured with a view to security, only making essential network services available and monitoring access. As well as protecting the machine, this can also help reduce internal problems by preventing machines being used as a staging post in attacks. Techniques such as TCP wrappers support this approach by providing access control when TCP services are requested.

Consider a small department with ten members of staff, wishing to install workstations to provide access to a local server as well as campus and Internet services. The local network is to carry confidential data. We have investigated several installations of this type, categorised as follows:

1. The server has two network adapters, one to the campus network and one to the private network and acts as a router, configured defensively to control access.

2. A machine is dedicated to being the firewall, as above but without other services running.

3. The private LAN is connected via routers (e.g. a satellite site using ISDN) which can be configured to control access.

Option (1) is inappropriate if a high degree of protection is required for the server but in fact it is a common configuration, providing improved security at the cost of an extra network card - arguably it should not be described as a firewall. Indeed, sometimes external access to the server is desirable, such as for a Web server. This is discussed in the following section.

Options (1) and (2) are characteristic of internal firewalls in that they are host-based solutions. Internal firewalls are typically LAN-to-LAN and can therefore be implemented on a PC with two network adapters, i.e. off-the-shelf hardware which is familiar to systems administrators of small networks. The LAN-to-LAN interconnectivity can also sometimes imply performance requirements which exceed those of firewall solutions designed for Internet connectivity to service providers using lower bandwidth leased lines. While firewall performance is an issue with larger departmental LANs, our 10-client scenario can readily be supported by a PC-based solution.

In all three cases the configuration follows standard firewall practice for ‘screening routers’, controlling access by protocol types, Internet Protocol (IP) addresses and port numbers. Configurations can be far simpler than campus firewalls because there are fewer services on the private network so external access to the private network can be highly restricted. A typical starting point is to give the private users full access to the campus but deny all access from users on the campus side. Restricting outbound traffic might be a better security policy but introduces an ongoing support task. The trade off between security and support is discussed later.

In addition to protecting small departmental LANs, sometimes an internal firewall is used to control access from a ‘public’ area such as a seminar or meeting room. This requires a set of rules which permit only certain traffic such as DNS traffic to specific nameservers, telnet and FTP, HTTP (possibly just to campus hosts and a proxy), SMTP and IMAP to specific mail servers, SSH and specific ICMP packets. Note that if the private LAN has any form of external access (such as dial-up) then this amounts to the same situation.

When a department is split across a site, internal firewalls could be introduced between each part and the campus; however, the users on the two private LANs will wish to function as if on a single private LAN. This involves creation of a Virtual Private Network (VPN). While a number of solutions exist off-the-shelf, the extra costs of configuration and maintenance appear to be a deterrent to introducing internal firewalls in this situation. An alternative is to run a private network link (fibre or otherwise) between the remote locations.

11.1 Servers as screening routers

One way of connecting a small department is simply to extend the LAN to all the machines. Traditionally this introduced problems of traffic management and repeater counts. The advent of Ethernet switches has dealt with both these problems and is often now the preferred solution, either providing the private LAN with a switched connection, or using a local switch to interconnect workstations, server and campus. The latter solution also reduces risk from eavesdropping locally (but does not eliminate it, due to possible spoofing attacks on Ethernet switches).

By introducing a second network adapter to a host and then using it as a router we are able to create a basic internal firewall, provided the routing capabilities include packet filtering. This host should ideally be a dedicated firewall machine, on the principle that firewalls should provide no other services. In practice a beneficial degree of security may be achieved by dual-homing the server itself, though this should be seen as sensible practice in the interest of security rather than a secure solution. Although early TCP/IP routing support on PC platforms was far from perfect, Novell (since 3.12) and NT 4.0 provide satisfactory solutions. 'Multi Protocol Routers' are available for both platforms, providing IP and IPX routing in particular, and filtering capabilities are available. Like any Unix system, Linux provides TCP/IP routing as standard, and extensive filtering capabilities with the *ipchains* package.

The disadvantage of introducing a TCP router is that it requires a subnet to be allocated to the private network, and campus router tables then need to be maintained accordingly. Apart from the support effort involved, allocating subnets in this way can lead to inefficient use of the available Internet address space; e.g. a private LAN with a server/router and two clients needs a subnet of 8 addresses. Our 10 client scenario requires 16 addresses (as would have just 6 clients) and would require additional space when it needs to grow beyond 13 hosts plus the server/router.

We see here a trade-off between security and support cost. The cost of implementing additional security is minimal (i.e. the cost of a network card), but it introduces an ongoing support requirement, especially as the network evolves; in fact, a poorly administered internal firewall could itself represent a security risk as with any routing device. There are also overheads to the two layers of network administration. If every private network on campus results in a different routing solution, the support task scales badly. Furthermore, the ease of use varies considerably across the various low-cost solutions and none of them attempts to integrate smoothly with other firewall management software. Hence they require specialist support, which might only be available at installation time, leading to an insecurity as the firewall might not be properly maintained as the network evolves.

For reasons of support, switched solutions are typically preferred over routed solutions except where there are special requirements, and in the latter case there may be budget to purchase a standard firewall which integrates with other firewall management software on campus. Implementation policies and guidelines for internal firewalls across a campus could assist with this problem, with recommended router configurations for a small set of standard solutions.

A useful compromise has emerged with the increasing availability of solutions which hide a private network address behind one IP address, such as SOCKS, masquerading and NAT. These do not require changes to the campus router configuration and nor do they cause inefficient use of address space. However, such a method does still suffer from multiple layers of administration and the potential for abuse if managed poorly. Worse, the use of NAT breaks the IP layer transparency, and prevents true end-to-end security. Some protocols embed IP information in the data stream (e.g. FTP) and these are not well-suited to NAT. The state information that a NAT router must hold (the mapping of internal to external addresses and ports) is also a burden.

11.2 Practical experience

This discussion of internal firewalling draws from studies of several installations, listed below with some commentary about practical experiences:

- A network of 8 machines and printers using a Novell server (initially acting as a router using Netware 3.11). This system has a number of external users and relies on Novell security to control access. Novell’s approach to firewalls (known as *BorderManager*) explicitly supports creation of internal firewalls; we have reviewed the possible upgrade of this system using *Firewall for NT 3.0* (previously ’Netroad’- a trial version is available from Novell as a download).

- An expanding network of machines from various vendors connected to the campus network by a Sun using *socks 4*. This is a standard Unix-based solution as discussed earlier in this report.

- 12 hosts on a Windows NT 4.0 network, connected using *Multi Protocol Router* (MPR) and subsequently *RRAS* Routing and Remote Access Service (previously known as ’Steelhead’, which is available from Microsoft as a download, also with Windows 2000). RRAS supports IP and IPX packet filtering (and RADIUS). We have reviewed the possible upgrade of this system using *Microsoft Proxy Server 2.0* (which can be purchased from Microsoft) to provide application layer security.

- An industrially-sponsored project with specific data security requirements, with an off-the-shelf CISCO firewall.

- A satellite site interconnected with ISDN using Ascend Pipeline routers. These were configured as a screening router, with the dual motivation of security and minimizing traffic.

- Homeworkers using linux and FreeBSD with PPP over dialup, then supporting the local network with *ipchains* and masquerading.

11.3 Summary

We have drawn a number of conclusions from our experiences with smaller-scale firewalls:

1. An external firewall should not be seen as a complete solution, as many attacks are internal to an organisation.

2. Simple PC-based firewalls can be set up at little cost, but their scope is limited.

3. Where a high degree of security is required, dedicated firewall devices should be used.

4. There are initial and ongoing support tasks associated with an internal firewall that impact at both local and campus levels.

5. To encourage use of internal firewalls, a rationalised approach across campus is required to reduce support costs.

12 Case Study C: A Campus Firewall

This paper was authored by Alex Walker of SUCS in November 1999 as part of the Southampton campus firewall procurement project.

This paper describes the hardware choice for the Super JANet Firewall, and proposes a software solution required to complete the procurement.

12.1 Hardware Choice

The choice of hardware platform was not as straightforward as one might have originally thought. There are a large variety of "Firewall Boxes" out there to choose from, all with various features and benefits. They generally fall into one of three categories.

- A PC
- A UNIX machine
- A "Black Box"

It was felt that a PC did not offer the hardware reliability and scalability required for a network component as crucial as the JANet Firewall, and investigations did not proceed much beyond this. Whilst Novell BorderManager was briefly investigated, its limitation of not supporting ATM adapters finally ruled it out.

Many vendors offered solutions based on a "black box" of some description. All had varying costs and specifications. The only one that came close as a contender was the Nokia IP range of firewalls. These are marketed as being routers, and indeed supported all the functionality one would expect from a router (routing protocols, hot-swap cards etc). The Nokia IP440 and 650 were investigated in detail, and suitable suppliers contacted. It was established that ATM LANE was supported as well as Fast Ethernet. The other major benefit was that it ran Checkpoint's excellent Firewall-1 software. The cost of the hardware was in the region of £12k and would have been within budget. The supported dual-redundant configuration was also investigated and was attractive.

Another vendor of Firewall-1 embedded hardware is Fore Systems. They had an Ethernet switch product called the ESX-2400. This was capable of firewalling at Gigabit speeds, but with a price tag of over £40k excluding software or maintenance, it was too expensive, and would have been overkill for our requirements.

The third hardware platform was a UNIX machine. The limited availability of software for SGI machines meant that they were not an option. The only suitable software was "Gauntlet", but this was many years behind other firewall products in terms of development, and its future was considered to be uncertain.

This left Sun hardware. A suitable machine was deemed to be of the following specification:

- Sun Enterprise E250
- Dual 400Mhz Processor
- 512MB RAM

- 4x9GB SCSI Disks (Mirrored Striped)
- Suitable Network Interfaces (QFE Cards, ATM Cards).

The Sun solution described above was decided to be the best hardware solution for the following reasons:

- Flexibility over the choice of firewall software to run on it
- Good value for money
- Assured future of the product and O.S.
- Hardware is easily maintained and upgraded with standard components
- Existing experience within SUCS of this platform
- Same firewall platform as chosen by ECS. This gives some collaboration benefits.

The Nokia platform was not chosen for the following reasons:

- Intel based (Pentium II Processor), therefore based on a BSD UNIX O.S. There are some potential concerns with using this for a firewall platform because BSD is open source (and thus hackers may study the source for vulnerabilities).
- Proprietary hardware (Nokia CompactPCI Cards)
- Scalability was achievable, but expensive compared to Sun hardware

12.2 Software Choice

With the hardware decided upon as being a Sun E250, this left two software choices – SunScreen EFS 3.0 from Sun Microsystems and Firewall-1 from Check Point.

12.2.1 Firewall-1

Firewall-1 is developed by Check Point in Israel, and is now on version 4.0. Supported platforms include (but are not limited to) Microsoft Windows NT 4.0, Solaris 2.6 and the Nokia platform. Check Point announced only this month (5th November 1999) that they would be porting Firewall-1 to Linux. Compared to SunScreen, it has the following benefits:

- A mature product. Has been through many revisions and appears to be almost bug-free.
- Easy to use. We had Firewall-1 up and running a simple policy in just an afternoon.
- Has a good GUI and Log/Status Viewer
- GUI available for Solaris, Win32, AIX and others.
- Is capable of viewing and terminating active connections
- Is very well documented. Hundreds of pages of quality manuals on the CD.
- Plugin modules available for Secure VPNs (Virtual Private Networks) and Floodgate-1 (Bandwidth control/management)
- Supports a wide variety of Authentication methods, and can alter a users/machines firewall ruleset based on that authentication
- Has a wide userbase (some 60% of Firewalls are based on Firewall-1)
- Incorporates SYNDefender – Detects, prevents, and informs on SYN flooding attacks.

Firewall-1 has the following disadvantages:

- Expensive – The required license is about £6.5k after educational discount – Four times the cost of SunScreen EFS
- Expensive ongoing costs – about £1.3k p.a. – About ten times the cost of SunScreen EFS.
- Being widely used and known, some consider it a threat to be using this firewall product.
- Solaris is not the primary development platform (the Nokia platform is) hence Firewall-1 is “ported” to Solaris.
- Version tested does not officially support Solaris 7, but only Solaris 2.6.

12.2.2 SunScreen EFS 3.0

SunScreen EFS 3.0 is the migration of two previous Sun products; EFS 2.0 and SPF-200. EFS 2.0 was their previous firewall product, and SPF-200 was a stealthing firewall product. Although Sun have entered this market later than Checkpoint, their product appears on paper to be very capable and has some tempting advantages over the rival.

- Aggressive education pricing policy – 90% discount from list price
- A “Stealthing Mode”. This is where the firewall acts as a bridge instead of a router. It hence has no IP address visible to the outside world, rendering it “virtually unhackable”
- Sun software on a Sun O.S. and Sun hardware. Benefits from having all one vendor/supplier.
- The GUI uses the JavaVM – hence any platform that Netscape Navigator has been ported to could support the GUI.

During our encounters testing SunScreen, we came across many problems and bugs. In fact, it took us almost a month to get SunScreen working as a firewall. We found the speed and capability of Sun's software technical support to be somewhat lacking. We failed to get the product operating in Stealthing mode at all, with Sun telling us our problem was a known bug to which no one seemed to be able to tell us the solution. Below are summarised some of the other disadvantages.

- Very slow GUI. Even on a new Ultra 10 333Mhz the GUI was irritatingly slow to use.
- Very poor documentation. It suffers from being the amalgamation of two products. Sun’s own bug database for this product lists poor documentation as something to be looked at for the next release.
- Immature product. The feeling is that this product could be much better in another one or two releases. Lacks some of the advanced features of other products that could be of use to us.
- Bandwidth management is a separate product, not integrated.
- Has a number of bugs in both the GUI and installation process.
- Requires SKIP encryption on the client and server. The installation “out of the box” failed to configure this component properly and required “manual intervention”. We also failed to make the Win32 SKIP client work correctly. This left us with the Solaris-only GUI.

12.3 Software Recommendations

Following evaluation of these two software products over some three months, I feel the only product that is suitable for our requirements is Check Point’s Firewall-1. It is a mature, easy to use, well-supported and well-documented product.

13 Reports on related UKERNA/JISC activities

There has been a notable increase in activity within UKERNA, JANET and the JISC over the last two years with respect to issues of security and authentication. The JISC have set up a new authentication and security committee, the JISC Committee on Authentication and Security (JCAS). The JANET-CERT [JCERT] unit has become more visible and pro-active in its work. UKERNA have also held a number of security-related events, including a Risk Reduction Workshop and a "Protecting Your Network" Workshop.

We survey these activities here, noting events where ECS delegates attended or presented.

13.1 UKERNA Activities

13.1.1 Quarterly Reports

The UKERNA quarterly reports include reference to security incidents; these are tallied and made available by JANET-CERT, but only where such incidents are reported of course.

The trend for sharply rising severe security incidents was reported earlier in this document. In particular reported root compromise incidents are almost doubling each year.

13.1.2 "Protecting Your Network" Workshop (January 19, 2000)

This event was attended by approximately 150 delegates, including one from ECS. We used the workshop as a method to conduct further informal discussions with managers from other sites on firewall issues. A show of hands at the event revealed that only some 8-10 people at the event were working behind a default deny firewall (or were admitting to it). The most common reason for not considering a default deny firewall seemed to be resignation to the fact that university committees would not allow such a policy to be put in place.

The talks at the workshop did not feature any site that had gone to default deny on anything but small financial-type networks. Some sites were using filter rules to block some known threats on their border routers, others were using firewalls on only their administrative networks. Huw Gulliver of Cardiff reported they were only protecting their admin network with Firewall-1, and that an external supplier had set up their firewall. Their system was running on Windows NT (a Pentium II PC), not a Unix platform.

The event featured some discussion of intrusion detection systems and vulnerability scanners. The IDS systems discussed were all commercial, with no reference to any public domain systems. While one scanning talk was an overview with little apparent practical experience to report on, Ben Harris of Cambridge described a home-grown system which scanned all Cambridge hosts (some 30,000 or more) and which made the results public amongst the owners of the scanned networks. Ben reported that open X11 servers and SNMP devices were one of the most common findings.

The clear message from this event is that there is little work being carried out or reported on the use of default deny firewall systems in large academic networks.

13.1.3 SuperJANET 4

The new SuperJANET network is due to roll out in the second quarter of 2001. At the time of writing UKERNA has not yet announced the “winning” tender, but it will be interesting to see what security techniques the new technology may enable or prevent. For example, if multi-Gigabit bandwidth networks are deployed, the firewalls will need to keep up with the demands on throughput.

13.1.4 Networkshop 27 (1999)

Jon Read from ECS attended the 1999 Networkshop, but reported that there were no presentations on firewall issues.

13.1.5 Networkshop 28 (March 2000)

There were five talks at Networkshop 28 that were of relevance to firewall deployments.

- Filtering, an Open Discussion, chaired by Andrew Cormack of CERT
- Risk Analysis in Practice, by Roland Trice, of ULCC and UKERNA
- Computer Crime Investigation, by Dave Reid of the Lothian Police
- The Cisco IOS Firewall, by Bob Lawrence of UCL
- Connecting Student Residences to the Network, by Ian Campbell of Exeter.

The slides from the talks will be made available on the UKERNA Web site soon. The number of security-related presentations illustrates the perceived importance of the topic by UKERNA. While the filtering discussion didn’t come to any notable conclusions, it did discuss a number of topical issues, such as how to deal with the Napster MP3 “piracy” problem [NAP, NAPB].

The Police presentation was perhaps the most enlightening in that it made clear the need for a site to take responsibility for the actions of its users when those users commit a crime from a host at that site. The site must be able to show that it has acted responsibly, e.g. in only allowing external connections from authenticated hosts. Negligence can lead to civil law suits, Dave Reid reported.

Wiring up 900 live network connections in student halls, as Exeter have done, raises many security issues. The talk did not go into depth on how the students are protected from each other, but did state that the site had used private (non-routable) IP addresses and NAT to ensure that only a handful of services (including e-mail, Web) could be run from the rooms in the halls.

Tim Chown of ECS also gave a presentation on IPv6 at the event; IPv6 is discussed briefly later in this report.

13.1.6 Risk Reduction Workshop (1999-2000)

The UKERNA Risk Reduction workshop considered a number of areas of risk, namely:

- Capacity

- Complexity
- The environment
- Security
- Single points of failure

A representative of ECS sat on the two security risk reduction meetings at ULCC in which some 40-50 security-related risks were identified and assessed for both likelihood of occurrence and their potential impact. The threats recognised by the group included hacking (of hosts or routers), DNS spoofing, denial of service attacks, blacklisting via ORBS or RBL, and lack of staff expertise.

The overview report from all groups stated that site security was very much down to individual institutions to deal with (with help from the likes of JANET-CERT). UKERNA would consider risks for its backbone network and services.

13.2 JISC/JTAP Activities

13.2.1 Report of JTAP Security Workshop (June 1999)

This small, focused workshop was attended by Dave DeRoure of ECS. The two-day event led to a report [JTAP-043] summarised by Tom Franklin of the JISC JTAP Programme.

To cite the report, its key recommendations were:

1. The most important single recommendation was that senior management need to understand how appropriate use of information security can aid the business processes within their institutions.

2. Senior management need to understand the responsibilities associated with ownership of information. These include privacy, accuracy and availability.

3. HEIs should use BS 7799 as a means of understanding the security risks and threats in their business and the costs and benefits associated with addressing the main risks.

4. JISC should undertake a review of the key legislative and regulatory frameworks and provide advice to institutions on the measures that they need to have in place to meet them. This will include the Data Protection Act and the e-commerce bill.

5. JISC should undertake a study of the current situation prevailing within HEIs and over JANET so that the scale of the problem can be understood and to ensure that any solutions will meet the needs of institutions.

6. There is a need for JISC working with the institutions to determine the minimum levels of security which institutions must have in place order to be part of the community. This could form a baseline service level agreement.

7. There is a shortage of experience and expertise within the community which JISC can address by commissioning the production of guidelines, templates, examples best practice and training materials.

The report also observed that there have been changes in the type of traffic travelling over JANET with more financial data, commercial and confidential data and personal data (including medical data). This data is not always recognised as confidential by those using it.

The issues raised by the JTAP Leach Reports [JTAP015, JTAP016, JTAP017] of 1997 were revisited by John Leach, who was present at the meeting. His thoughts were reproduced in JTAP-043:

The task for JISC:

- promote a strategic approach
- maintain business focus (i.e. not technology led)
- institutions must identify their key business goals
- security must be seen to bring value/benefit to those business goals if it is to be accorded any priority.

Business goals:

- financial accountability
- reduce operating cost
- prevent unlawful use of resources
- good standards of control

Recommendations:

- adopt business risk management
- produce a baseline security for devolved IT
- develop a structured approach to network access
- harmonise local access
- harmonise national access
- use secure internet protocols

The way forward:

- the way forward is NOT Public Key Infrastructure / Certificate Authorities, smartcards etc. until these meet business goals
- revisit business goals
- identify how information strategy can benefit the HEI
- link key security goals to business goals

Two of the meeting's recommendations have since been acted upon under the new JCAS Programme, by a call for investigations into legal aspects of the DPA 1998 and for reports on best practice for secure communication technology.

13.2.2 JISC Authentication Event (Nov 2, 1999)

Tim Chown of ECS gave an overview of authentication techniques at this seminar. The topics ranged from general principles through to specific use of smart card systems and the current status of Athens. The Athens system enables access to bibliographic data for some 1,500,000 UK academic users, yet the technology it uses is (by most commercial standards) relatively simplistic, e.g. widespread trust by source IP address. A new generation of Athens is being planned.

One of the more interesting presentations was on the use of "Smartcards for digital certificates and public key encryption" by Graham Phillips, a Research Fellow at the University of

Cambridge Clinical School. The Cambridge system will reportedly be made available for testing by other sites soon.

13.2.3 JTAP-032 “Secure Internet Protocols”

This report [JTAP032] is an interim report from our own project due to complete in July 2000. Some of its findings and recommendations have already been reported in this document. Current focus is on user trials of software.

13.2.4 JTAP-040 “Blocking Spam Relaying and Junk Mail”

The “Blocking Spam Relaying and Junk Mail” report by Janusz Lukasiak of the University of Manchester makes a reference to firewall use:

> “The firewall, as mentioned earlier can be either a dedicated unit, or a set of filtering rules in a router located on the boundary between 'inside' and 'outside', as defined in the policy document. Incoming (JANET to local network) SMTP connections should be allowed only to secure mail routers. If outgoing email is forced to pass through mail routers, a similar set of rules should control outgoing SMTP connectivity. The exact format of filter rules obviously depends on the type of firewall or router. Managers of local networks with multiple connections to JANET will obviously have to ensure that the same level of protection is implemented on each connection.”

ECS has just completed a report on the issue of junk e-mail (working heading JTAP-635) which will be made available shortly. As reported earlier in this report, it is important for a site to take measures to prevent it becoming blacklisted by other sites on the Internet.

14 Future Firewalls and IPv6

There are a number of challenges facing firewall designers either now or in the near future. These include:

1. **Speed of operation**. As networks run faster, the processing requirements grow. SuperJANET 4 will offer 2.5Gbit links to sites from 2001, with a core backbone running to a projected 80Gbits by 2005. A campus firewall will thus need to run at Gigabit speeds. This may imply firewalls running in silicon not software.

2. **More application layer processing**. Rather than IP-layer filtering, more application layer filtering will be required, beyond the current HTTP, FTP and SMTP. Firewalls will need to understand and handle new protocols that may have special requirements, e.g. new multimedia streaming applications.

3. **Co-operative firewalls**. As performance issues become more important, it is likely that we will see co-operating firewall "farms" deployed. Firewalls may share the "proxying" load for applications where such an approach is required, or where stateful handling of communications is beneficial (e.g. application layer virus scanning).

4. **The always-on Internet**. More Internet devices will become "always-on", be they mobile IP-based devices or devices in the home. Staff and students on campus will wish to communicate securely but freely with these devices. Staff may also work from home more frequently as technologies like ADSL and cable modems offer productive bandwidth to households.

Firewall technology is improving rapidly, but at the same time the goalposts are moving…

14.1 IPv6

The next generation Internet Protocol, IPv6 [V6F], is primarily aimed at offering a bigger address space to enable new IP-based applications, such as IP-based cellular phone devices. IPv6 offers 128-bit addresses instead of IPv4's 32-bit addresses. The onset of pervasive computing, bringing new innovative IP devices both at work and in the home, will demand the address space that IPv6 offers. The use of Network Address Translation (NAT) is a stopgap measure that will offer sites with more hosts than IP addresses a method to get connected, but NAT won't scale beyond the immediate future.

IPv6 will make backbone IP routing more efficient through hierarchical aggregated addressing principles. It will make IP device management simpler through advanced auto-configuration options. IPv6 is in experimental use today on the majority of the major academic networks in the world (e.g. Internet 2, Renater, DFN, WIDE, ACOnet, SURFnet), and has support in products from the major router and OS vendors, including Cisco (IOS), Microsoft (Windows 2000) and Sun Microsystems (Solaris 8).

While production use of IPv6 over JANET may be a few years away, it is prudent to start preparing now by running test networks and developing trial applications. ECS is participating in UK trials with Lancaster and UCL on a project called Bermuda, and is participating in IPv6 trials at a European level within the Quantum Test Programme (QTPv6). Part of this work [V6ECS] includes evaluation of IPv6 firewall products that are currently very scarce. The two products that do offer IPv6 firewall functionality at present are FreeBSD 4.0's *ipfw* [FBSD] and the *ipfilter* [IPFW] package, both available for free. There is a rumour that Firewall-1 v5.0 will support IPv6, but it is only a rumour at present.

One of the challenges that IPv6 will pose firewall designers lies in its inherent ability for networks to automatically renumber. IPv6 router renumbering allows network prefix changes to be propagated to subnetworks. Extensions to DNS (made available in BIND9) abstract the network and host part of a DNS entry allowing DNS updates in reaction to router renumbering to be done at minimal cost. The same problem impacts on firewalls, which "compile" their rule sets into IP number-based filters. If the underlying IP numbers change, as they would with router renumbering, the firewall would need to react to (authenticated and secure) updates.

IPv6's stateless auto-configuration also poses a firewall problem, in that hosts will join and leave networks ad hoc. Firewalls will need to understand the same sort of (secure) dynamic updates as are sent to DNS servers. However, most services will run from permanently attached servers with fixed addresses.

IPv6 is a technology that will be used in production networks in the near future; firewall manufacturers have as yet been slow to react to the recent acceleration in IPv6 support from the router and OS developers.

For more information on IPv6 there is a mailbase list *ipv6-users@mailbase.ac.uk*.

15 Recommendations

This report makes the following recommendations to institutions and to the JISC:

1. All HEI sites should reconsider the adoption of a default deny firewall policy. A surprisingly small percentage of UK universities appear to be running default deny firewall configurations. Many run a small set of filters on network border routers, and perhaps on administrative network routers. But that will not be enough to deter all but the most casual of attackers.

2. The Data Protection Act 1998 brings legal responsibilities to institutions to take reasonable steps in securing sensitive personal data. This includes access to data networks, but also data in transit and data stored on a computer system. Network managers should note that they are also personally liable under the DPA 1998, with fines ranging up to £5,000. Sites should ensure they are aware of the legal implications of the new Act, and the JISC may wish to invest in a study for the benefit of all HEIs.

3. The cost to an institution of being offline from the Internet is not easy to assess. The cost may depend on whether internal and/or external services are affected. It may be useful for the JISC to commission a short investigation into evaluating the outage costs for a range of representative sites as a means to alert institutions to the potential implications. Such an exercise may offer valuable "ammunition" with which to get a default deny-oriented security policy approved at a site (though the easiest method to gain approval most likely remains to be the victim of a major incident…).

4. The deployment of firewall solutions is becoming a relatively well-understood problem. The use of secure shell (ssh), secure socket layer (SSL) and secure tunnel (stunnel) methods allow encrypted communications between hosts. However the issue of authenticated access, and authenticated peer-to-peer communications, requires needs particular further investigation. A site security policy should include each of these aspects of security. Further studies into best practice are required here.

5. Sites seeking to deploy a default deny (inbound) security policy should consider the benefits of a pre-installation programme of user consultation, awareness and education. By gaining the respect and support of the user community, the switch to (what is perceived as) a more restrictive working environment should be a far more successful and beneficial exercise. The change should be phased over a reasonable period of time.

6. Institutions have a corporate liability for actions performed by users on their computer systems; if incidents affecting remote sites cannot be traced back to authenticated individuals (however strong that authentication might be) then an institution may be liable to action in at least a civil court. Sites should ensure they are aware of their responsibilities in this area. JISC may wish to fund a study for the benefit of all HE sites.

7. A firewall as a network point of entry sees all traffic entering or leaving the network. As such it may be able to log all IP traffic for the purposes of accounting and/or billing. Sites wishing to pass UKERNA bandwidth charges onto departments may wish to consider making use of this firewall feature.

8. The growth of the Internet, together with imminent voice-data-video convergence to packetised IP services, will lead to the adoption of new Internet protocols (at various network layers). One example of such protocols is IPv6 [V6F]. It is very important that security measures and policies are abstracted as much as possible from specific technologies, so that they can be adapted to encompass those new technologies when they arrive.

9. A campus network should designed in such a way as to isolate effect of a security breach or an “act of God”. Many university departments are dependent on their Internet connection for successful and profitable research and teaching programmes. It is no longer acceptable, as it may have been five years ago, for a site, or a department at a site, to be offline for a few days due to an incident of some sort. By planning network deployments carefully, it should be possible to maintain a service to the majority of a network when an incident occurs. Sites should study their existing network topologies for possible improvement in “fault tolerance”.

16 Contact Information

If you would like to discuss the content of this report with the authors, or visit the site for more information and demonstrations, please contact:

Dr Tim Chown
Department of Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
UK

E-mail: tjc@ecs.soton.ac.uk
Phone: +44 (0)23 8059 3257

We have also set up a mailbase discussion list *firewall-admin@mailbase.ac.uk*. Visit the mailbase site at http://www.mailbase.ac.uk for more information.

The master copy of this report is held at the JISC/JTAP site at http://www.jtap.ac.uk.

Other information may be found under the Networks and Distributed Systems web pages within the IAM group web presence at http://www.iam.ecs.soton.ac.uk.

17 References

[ACEAG] *RSA ACE Agents*, http://www.rsasecurity.com/products/securid/rsaaceagents.html

[BIGF] *BIGfire Firewall*, http://www.portcullis-security.com

[BO] *Back Orifice Vulnerability*, http://www.cert.org/vul_notes/VN-98.07.backorifice.html

[BS77991] *BS 7799-1:1999 Information security management. Code of practice for information security management*, 1999

[BS77992] *BS 7799-2:1999 Information security management. Specification for information security management systems*, 1999

[BSI] *British Standards Institute*, http://www.bsi.org.uk

[BWARE] *BorderWare Firewall Server*, http://www.borderware.com

[CBG] *CyberGuard Firewall for Windows NT*, http://www.cyberguard.com

[CCOM] *CenturyCom*, http://www.centurycom.co.uk

[CERIAS] *Center for Education and Research in Information Assurance and Security*, http://www.cerias.purdue.edu

[CERT] *CERT Coordination Center*, http://www.cert.org

[CGRAM] *Crypto-Gram*, http://www.counterpane.com/crypto-gram.html

[CHAP] *Building Internet Firewalls*, D. Brent Chapman and Elizabeth D. Zwicky, O'Reilly, ISBN 1-56592-124-0

[CHESTFW1] *CHEST Firewall-1 Agreement*, http://www.chest.ac.uk/software/firewall-1/contents.html

[CIAC] Computer Incident Advisory Center, http://ciac.llnl.gov/ciac

[CISCOR] *Improving Security on Cisco Routers*, http://www.cisco.com/warp/public/707/21.html

[CKPFW1] *Check Point Firewall-1*, http://www.checkpoint.com/products/firewall-1/index.html

[CMA] *The Computer Misuse Act (1990)*, The Stationery Office Ltd, ISBN 0 10 541890 0, http://www.hmso.gov.uk/acts/summary/01990018.htm

[CMV] *Commandview Firewall for NT*, http://www.mis-cds.com, http://www.elronsoftware.com

[CVP] *Check Point Content Vectoring Protocol (CVP)*, Open API Sepcification, http://www.checkpoint.com/cvpopenspec/index.html

[DIALPAD] *Dialpad IP telephony*, http://www.dialpad.com

[DPA98] *Data Protection Act 1998*, http://www.hmso.gov.uk/acts/acts1998/19980029.htm, 1998.

[DTISEC] *DTI Communications and Information Industries: Security*, http://www.dti.gov.uk/cii/security.html

[EBOR] *e-Border proxy*, http://www.eborder.nec.com

[ESOFT] *ESOFT*, http://www.esoft.co.uk

[FBSD] *FreeBSD*, http://www.freebsd.org

[FPICK] *How to Pick an Internet Firewall*, http://www.clark.net/pub/mjr/pubs/pick/index.htm

[FWC] *Firewall.Com*, http://www.firewall.com

[FWFAQ] *The Firewalls FAQ*, http://www.interhack.net/pubs/fwfaq

[FWL] *fwlogsum*, http://www.ginini.com.au/tools/fw1

[GAMMA] *Gamma Secure Systems Limited*, http://www.gammassl.co.uk

[GCA] *Great Circle Associates*, http://www.greatcircle.com

[GNAT] *GNAT Box*, http://www.gnatbox.com, http://www.globaltech.co.uk

[GTELLA] *Guntella*, www.allskin.com/gnutella

[HACK] *Chronicle of 1998 Southampton Hacking Incident*, http://www.soton.ac.uk/~cert/incident/inc1event.html

[IANA] *The Internet Assigned Numbers Authority, Protocol Numbers and Assignment Services*, http://www.iana.org/numbers.html

[IANAPORT] *The Internet Assigned Numbers Authority, Port Number Allocations*, http://www.isi.edu/in-notes/iana/assignments/port-numbers

[ICL] *Introduction to Computer Law*, D. Bainbridge, Pitman Publishing, ISBN 0 273 61940 3.

[ICQ] *ICQ Chat*, ICQ Inc., http://www.icq.com

[ICQP] *ICQ through a SOCKS5* Proxy, http://www.icq.com/firewall/socks5.html

[ICSA] *ICSA.net*, Firewall Certification, http://www.icsa.net

[IETF] *The Internet Engineering Task Force*, http://www.ietf.org

[IMP] *IMP E-mail client*, http://web.horde.org/imp

[IPF] *IP Filter*, http://cheops.anu.edu.au/~avalon/ip-filter.html

[ISCDS] *The Internet Software Consortium Domain Survey*, http://www.isc.org/ds/

[ISS] *Internet Security Systems*, http://www.iss.net

[ISVW] *InterScan VirusWall*, Trend Micro, http://www.antivirus.com/cpfw1.htm

[JCERT] *JANET CERT*, http://www.ja.net/CERT/cert.html

[JCSTATS] *JANET-CERT Report Statistics*, 1997-1999, http://www.ja.net/CERT/JANET-CERT/monthly_reports.html

[JTAP015] *Findings from the first stage of the Study into the Requirements for Authentication, Authorisation and Privacy in Higher Education*, John Leach, Trusted Information Systems, February 1998, http://www.jtap.ac.uk/reports/htm/jtap-015-1.html

[JTAP016] *Recommended Security Solutions: Results of the Study into the Requirements for Authentication, Authorisation and Privacy in Higher Education*, John Leach, Trusted Information Systems, February 1998, http://www.jtap.ac.uk/reports/htm/jtap-016-1.html

[JTAP017] *Recommended Actions for JISC: Results of the Study into the Requirements for Authentication, Authorisation and Privacy in Higher Education*, John Leach, Trusted Information Systems, February 1998, http://www.jtap.ac.uk/reports/htm/jtap-017.html

[JTAP032] *Working Paper on Secure Internet Issues for the HE Community*, Tim Chown, David DeRoure, Julian Field, Mark Thompson, Jon Read, University of Southampton, June 1999, http://www.jtap.ac.uk/reports/htm/jtap-032.html

[JTAP040] *Blocking Spam Relaying and Junk Mail*, Janusz Lukasiak, University of Manchester, October 1999, http://www.jtap.ac.uk/reports/htm/jtap-040.html

[JTAP043] *Report of JTAP Security Workshop*, Tom Franklin, JISC, 21-22 June 1999, http://www.jtap.ac.uk/reports/htm/jtap-043.html

[JTAPISF] *JTAP Information Security Forum*, http://www.jtap.ac.uk/reports/isf/

[MAPPA] *Mappa Mundi Cyber Geography Research*, http://mappa.mundi.net

[MIDS] *Matrix Information and Directory Services*, http://www.mids.org

[MINDT] *Mindterm Java ssh client*, http://www.mindbright.se/english

[MSCAN] *mscan*, http://www.ja.net/CERT/JANET-CERT/mscan.html

[NAP] *Napster*, http://www.napster.com

[NAPB] *Napster filtering*, http://www.soton.net/napsterfiltering.html

[NECSOCK] *SOCKS at the NEC Networking Systems Laboratory*, http://www.socks.nec.com

[NETMF] *NetMeeting and Firewalls*, http://support.microsoft.com/support/kb/articles/Q158/6/23.asp

[NSZ] *Telcordia Netsizer*, http://www.netsizer.com

[NUA] *Nua Internet Surveys*, http://www.nua.ie

[NWS28] *UKERNA Networkshop 28*, held at Heriot-Watt University, March 21-23, 2000.

[OSSL] *OpenSSL: TLS/SSL libraries*, incorporating X.509 certificate generation/management for a wide variety of Unix and Win32 platforms, http://www.openssl.org

[PCMFW1] *PC Magazine review of Checkpoint Firewall-1*, August 1999, http://www.checkpoint.com/corporate/articles/pcmag-080699.html

[PGPI] *PGP Security Inc*, http://www.pgp.com

[PHONE] *PhoneBoy Productions, Firewall-1*, http://www.phoneboy.com/fw1

[PUTTY] *Putty ssh client*, http://www.chiark.greenend.org.uk/~sgtatham/putty

[REALF] *Real.com Firewall Support*, http://service.real.com/firewall/index.html

[RFC1928] *SOCKS Protocol Version 5*, M. Leech, M. Ganis, Y.Lee, R. Kuris, D. Koblas, L. Jones, March 1996, http://www.ietf.org/rfc/rfc1928.txt

[RFC2138] *Remote Authentication Dial In User Service (RADIUS*), C. Rigney, A. Rubens, W. Simpson, S. Willens, April 1997, http://www.ietf.org/rfc/rfc2138.txt

[RFC2251] *Lightweight Directory Access Protocol (v3)*, M. Wahl, T. Howes, S. Kille, December 1997, http://www.ietf.org/rfc/rfc2251.txt

[RSEC] *RealSecure*, Intrusion Detection System, http://www.checkpoint.com/products/firewall-1/realsecure.html

[SECID] *RSA SecurID*, http://www.rsasecurity.com/products/securid

[SECU] *SecuRemote*, VPN-1 Secure Network Architecture, http://www.checkpoint.com/products/vpn1/securemoteds.html

[SFOCUS] *Security Focus*, http://www.securityfocus.com

[SOTCERT] *University of Southampton CERT Team*, http://www.soton.ac.uk/~cert/

[SPORT] *Security Portal*, http://www.securityportal.com

[SSH] *Secure Shell*, http://www.ssh.org

[SSH2] *Secure Shell server v2*, http://www.ssh.com/about/press/release01032000.html

[STUNNEL] *TLS/SSL tunneling software*, http://mike.daewoo.com.pl/computer/stunnel

[SUNEFS3] *Sunscreen Secure Net 3.0*, http://www.sun.com/software/securenet/index.html

[THAWTE] *Thawte PKI*, http://www.thawte.com

[TLS] *The TLS Protocol v1.0*, RFC2246, T. Dierks, C. Allen, January 1999, http://www.ietf.org/rfc/rfc2138.txt

[UKERNAQ4] *UKERNA Quarterly Report to the JANET Community*, Winter 1999, http://www.ja.net/documents/quarterly.html

[URISK] *UKERNA Risk Reduction Workshop*, London, Spring 2000.

[V6ECS} *IPv6 at ECS*, http://www.ipv6.ecs.soton.ac.uk

[V6F] *The IPv6 Forum*, http://www.ipv6forum.com

[VERI] *Verisign PKI*, http://www.verisign.com

[WACK] *Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls*, NIST Special Publication 800-10, John P. Wack and Lisa J. Carnahan, 1995, http://csrc.ncsl.nist.gov/nistpubs/800-10

[WEBI] *WebIMAP E-mail client*, http://www.netwinsite.com/webimap/index.htm

[WEBT] *WebTrends*, http://www.webtrends.com/security/products.htm

[WWDSI] *World Wide Digital Security, Inc*, SAINT, http://www.wwdsi.com/saint/index.html

Appendices

17.1 Appendix A - ECS Firewall Questionnaire

The ECS network is connected to the campus network and hence the Internet by a machine which is running a "firewall". As part of a JISC-funded project we are investigating best practice for running this firewall.

A firewall can be used to restrict access to hosts and services, usually in the interests of security - the necessity for such measures is illustrated by the Southampton University Computing Services (SUCS) hacking incidents. The firewall currently in place implements some very rudimentary precautions against external hacking. This email is to find out about user requirements and opinion, as part of establishing the best way to operate the firewall in the future. More information concerning firewalls in general can be found in the Internet Firewalls FAQ, which can be viewed at http://www.interhack.net/pubs/fwfaq/.

We welcome your response to the questionnaire, which is available on the Web at <URL>. We will also send the questionnaire out by email, and if you prefer you can respond directly to that message or print it out and return it to me (Jon Read) via internal mail.

Please respond for yourself and (separately if you wish) for any classes you teach or people you supervise who will not be returning questionnaires themselves. E.g If you teach a class that requires access to a service (other than WWW) at another site, please tell us. If your requirements are likely to change in the future, please tell us what they are likely to be as best as you can.

We will make available a summary of the outcome of this survey. If you have any queries or wish to discuss the firewall further, please contact me (Jon Read) or Tim Chown. Your local member of ECS systems staff will be able to assist you should you have any questions when responding to these questions.

Thank you in advance for the information you provide.

1. When you are working *inside* ECS (e.g. in your office or laboratory), which network services outside the department do you access? Please note that SUCS is considered to be outside of ECS.

   You do NOT need to include:

   - sending email from your ECS account (i.e. sending mail from your ecs.soton.ac.uk address)
   - receiving email on your ECS account (i.e. mail addressed or redirected to you at ecs.soton.ac.uk)
   - accessing World Wide Web sites
   - accessing ftp sites, anonymous or otherwise
   - reading or sending to USENET news groups using ECS news servers
   - telnet access to remote sites

   Please list the services (or client programs) you use and the sites you access. Here are some examples:

- "MICS Student records system"
- "access to file store on SUCS fileservers, using NFS"
- "Eudora to read or send mail using an Internet Service Provider"

2. When you are working from Internet sites *outside* ECS (e.g. from SUCS, using an Internet Service Provider other than ECS dialup, or from a company or other university), which network services inside ECS do you access?

   You do NOT need to include access to official ECS World Wide Web servers (those registered with names beginning www.)

   Please list the services you use and how you are connected (in particular whether it is a fixed location). Here are some examples:

- "email with pine by running telnet to host A from SUCS workstations"
- "read and send email from home with Eudora using ECS mail servers, connected via BT Internet"
- "telnet to host B from SUCS dialup"
- "running X windows clients on research machines from my desk at Foo Ltd"
- "telnet access to your machine for collaborators on research projects"

3. (For ECS dialup users) If you use the ECS dialup service, do you need to access any network services outside ECS? Which ones?

4. Do you manage machines or data that require special security measures? If so, what requirements do they have?

5. The firewall is there to offer protection without obstructing your work. Do you have any other comments that will help us get this balance right for you?

6. Would you like your name to be added to a mailing list for discussion of the ECS firewall? We will then keep you posted about the JISC firewalls project.

17.2 Appendix B - Responses to questionnaire

1. When you are working *inside* ECS (e.g. in your office or laboratory), which network services outside the department do you access? Please note that SUCS is considered to be outside of ECS.

    - BIDS
    - CDDBP
    - CUSeeMe
    - CVS access to remote sites
    - games (eg. Quake)
    - Gopher
    - IRC to Host Z
    - NFS
    - NNTP access to remote sites
    - ping, traceroute, whois, dns etc
    - POP/IMAP from external mail servers
    - Publish to external webservers using Frontpage
    - rc5des to distributed.net
    - rsh/rlogin
    - slogin to remote sites
    - SMB
    - SMTP to external mail servers
    - SQL
    - ssh
    - SUCS (campus) filestore
    - SUCS SP2 (login server) "access"
    - VNC
    - X
    - Z39.50 servers
    - "Connections on various high-number ports (TCP and UDP) for experimental remote services, such as DCC DRAW, long-distance SANE, diagnostics..."
    - "I would mount remote Appletalk volumes in Psychology and elsewhere if I could. :-(I hope to be able to in the future."

2. When you are working from Internet sites *outside* ECS (e.g. from SUCS, using an Internet Service Provider other than ECS dialup, or from a company or other university), which network services inside ECS do you access? You do NOT need to include access to official ECS World Wide Web servers (those registered with names beginning www.) Please list the services you use and how you are connected (in particular whether it is a fixed location).

    - Access to Host G port 8888 example IIP server
    - email via IMAP from any external host
    - Eudora (POP3/IMAP)
    - FTP (client-server, server-client-server)
    - FTP access to work machine from any external host
    - ftp to any host from SUCS
    - HTTP + FTP access to work machine
    - incoming ARA via ECS and SUCS modems

- NFS to mount directories on work machine
- read and send email from e.g. conference with Eudora using ECS mail servers (can read, can't send as SMTP host is blocked).
- remote sound (within X session) from lab to SUCS
- secure-VNC
- slogin to work machine(s)
- ssh
- telnet to any host from SUCS workstations
- telnet to my research group machines from abroad
- telnet to Host A from SUCS dialup
- telnet to Host A from SUCS workstations (to read mail using elm)
- telnet to Host B from SUCS workstations (writing WWW pages)
- telnet to Hosts C and D from any external host
- telnet/rlogin to Host E from foreign ISP
- telnet/rlogin to work machine from foreign ISP
- VNC to work machine
- WWW access to handin.ecs.soton.ac.uk from any external host
- X sessions to Host F from anywhere
- X sessions to work machine
- Z39.50 access to work machine
- "Connections on various high-number ports (TCP and UDP) for experimental remote services, such as DCC DRAW, long-distance SANE, diagnostics..."
- "I currently do not have the ability to access my incoming email remotely, despite being able to access Host A and all other mail directories. This recently caused me considerable difficulty during a conference in Germany. Moral: keep it simple and get it right!"
- "I mount my office PC disks on my home PC using the Win95 "map network disk" facility, both using ECS Dialup and SUCS dialup."
- "I would provide Appletalk access to collaborators if I could. :-(I hope to be able to in the future."
- "Web access to my Prograph-based web-server on Host J."
- I run occasionally the Java Webserver on Host P. I use it for a project together with the universities of Oldenburg and Zurich. The project partners have to test prototype Java Servlet implementations infrequently.
- I run PCAnywhere, a remote administration tool, to my desktop.
- I need to have access while traveling around the planet as I often do for academic/industrial purposes. Any machine that allows me to get to Host U via rlogin will do.
- I am not sure if i am affected by the implications of firewalling; Anyway, i am running a FTP and WWW server on Host R.
- I would like the following services on Host R: Web server - port 80, Secure Web server - port 443, Ftp server, Telnet, Imap, SMTP, ssh - port 22.
- A service I make regular use of is PPTP so that I can work on the FOO windows network from outside. Will this be affected?

3. (For ECS dialup users) If you use the ECS dialup service, do you need to access any network services outside ECS? Which ones?

- ssh to Host B

- WWW, FTP, News etc
- telnet to SUCS machines

4. Do you manage machines or data that require special security measures? If so, what requirements do they have?

    - no telnet access, no ftp access, ssh only
    - "Hosts D and E need telnet access from anywhere.
      handin.ecs needs WWW access from anywhere.
      Hosts G and H need NFS access from SUCS.
      Hosts I and J need SMTP access to anywhere.
      Hosts K and L need SMTP access from SUCS.
      Hosts M and N need SMTP access from anywhere.
      No hosts in network N should have SMTP access from outside ECS."
    - "NT workstations in network I often have driver shares enabled with minimal security measures. It would be useful if these machines were protected from outside access."
    - "NDAs with Microsoft, CCLRC etc.
      UK Terbulence repository."
    - "The security I need is protection from network failure, hard disk failure and operating system incompatibilities - all of which have hindered my research at ECS. Currently this is achieved by writing CD archives (when the machine and network are functioning, and I am available) and keeping the acquisition machine in the lab (Trouble in 1/1049) switched off as much as possible."
    - "Our machines have information private to the company and we have considered setting up our own firewall for them for this reason but if they are behind an ECS firewall I would feel this was not really necessary."

5. The firewall is there to offer protection without obstructing your work. Do you have any other comments that will help us get this balance right for you?

    - "Blocking by content is necessary to stop people (like me!) running ftp and web servers on non-standard ports without authorisation"
    - "Access to Z39.50 gateways (and pseudo HTTP servers running on unusual ports) is vital to my work, however they can usually be configured, so most of the ones I run locally could be limited to a range of port numbers if neccessary, however external ones are run on any port you have ever heard of."
    - "Thanks for carrying out this survey. I imagine a lot of people with answers similar to the above (i.e. we don't use the network much) will not bother to reply; I hope you'll take this into account."
    - "The questions above do not apply to me at all, but rather than not submit a reply to the questionaire, I thought that I would just mention this."
    - "The firewall currently passes individual machine name IP addresses to 3rd parties. For security reasons I believe the firewall should substitute machine IP addresses for a firewall IP address or addresses. This would allow the Intranet to remain anonymous to external sites and also allow the firewall to limit external access of internal machine."
    - "Any firewall can be breached, given co-operation on both sides of the firewall, since by its nature the firewall allows some data through, and this can be leveraged as much as necessary. (*) So it is vital that firewall security is done with the

understanding and support of all the users *inside* the firewall.
It is useful to have some commonly abused outbound services (such as the web servers which are all-too-often included silently with new software and hardware) blocked at the firewall to improve internal security.
(*) For this reason attempts to use a firewall for political purposes inside an organisation are futile."

- "Flexibility. We need to be able to make changes within 24 hours. We also need access *at all times* to a connection outside the firewall for test and development."
- "Its important to me to be able to access ECS services viva my BT Internet account from home (because its more reliable and never engaged!)"
- "I would prefer a reliable basic service to an fancy intermittent one."
- "Should I have a firewall on my own machine? I have seen it suggested that everyone should have firewall software on their machines these days."

17.3 Appendix C – Firewall Presentation to ECS Department Users

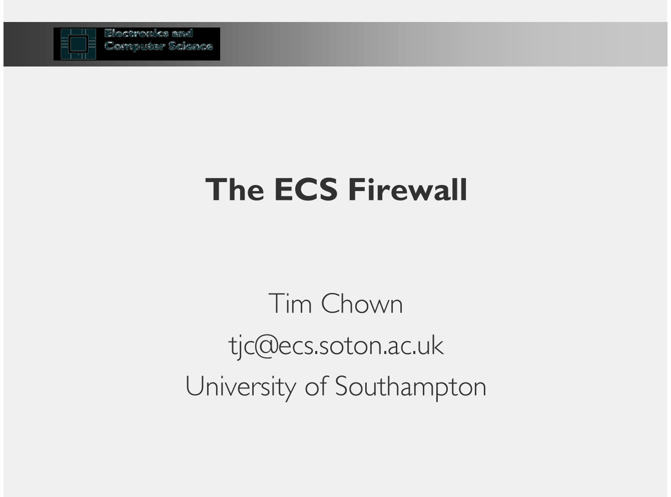

The ECS Firewall

Tim Chown
tjc@ecs.soton.ac.uk
University of Southampton

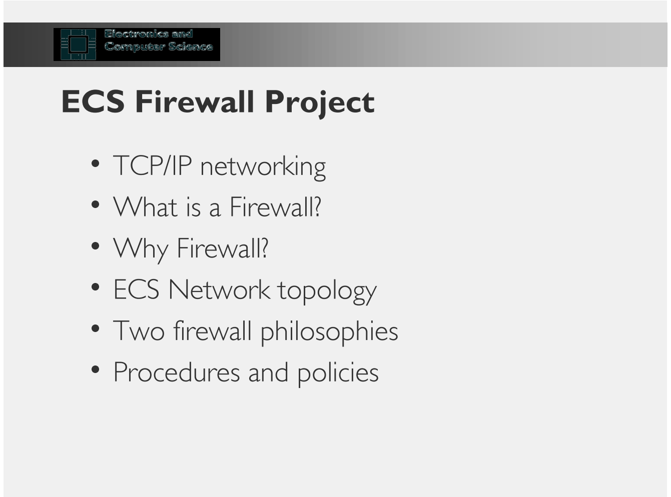

ECS Firewall Project

- TCP/IP networking
- What is a Firewall?
- Why Firewall?
- ECS Network topology
- Two firewall philosophies
- Procedures and policies

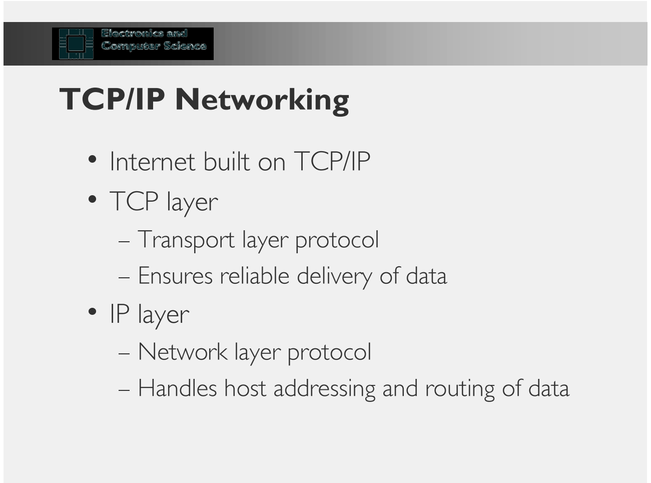

TCP/IP Networking

- Internet built on TCP/IP
- TCP layer
  - Transport layer protocol
  - Ensures reliable delivery of data
- IP layer
  - Network layer protocol
  - Handles host addressing and routing of data

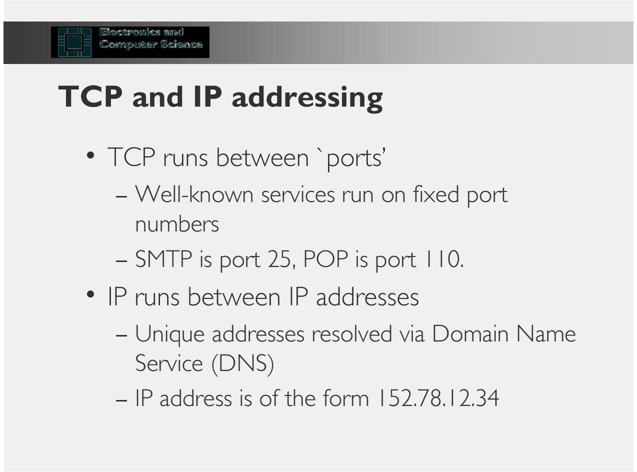

TCP and IP addressing

- TCP runs between `ports’
  - Well-known services run on fixed port numbers
  - SMTP is port 25, POP is port 110.
- IP runs between IP addresses
  - Unique addresses resolved via Domain Name Service (DNS)
  - IP address is of the form 152.78.12.34

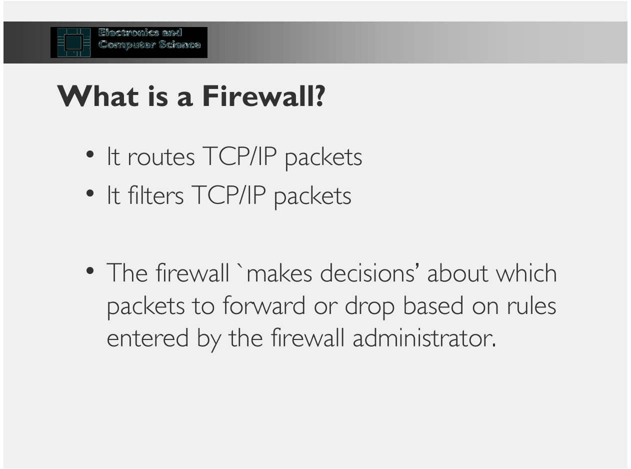

What is a Firewall?

- It routes TCP/IP packets
- It filters TCP/IP packets

- The firewall `makes decisions’ about which packets to forward or drop based on rules entered by the firewall administrator.

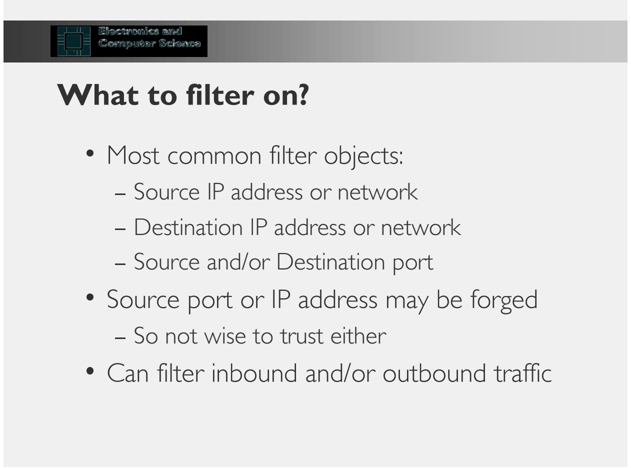

What to filter on?

- Most common filter objects:
  - Source IP address or network
  - Destination IP address or network
  - Source and/or Destination port
- Source port or IP address may be forged
  - So not wise to trust either
- Can filter inbound and/or outbound traffic

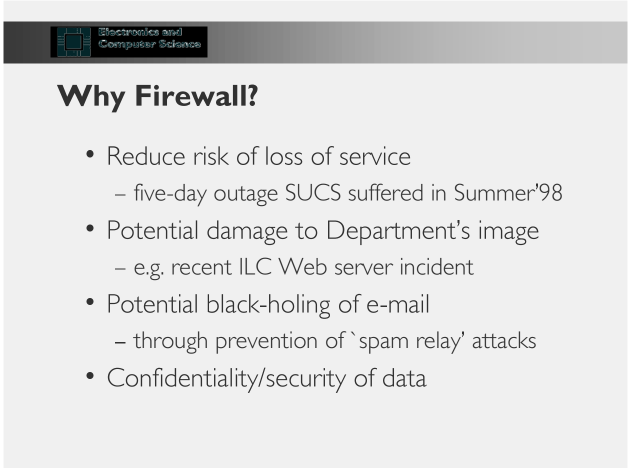

Why Firewall?

- Reduce risk of loss of service
  - five-day outage SUCS suffered in Summer'98
- Potential damage to Department's image
  - e.g. recent ILC Web server incident
- Potential black-holing of e-mail
  - through prevention of `spam relay' attacks
- Confidentiality/security of data

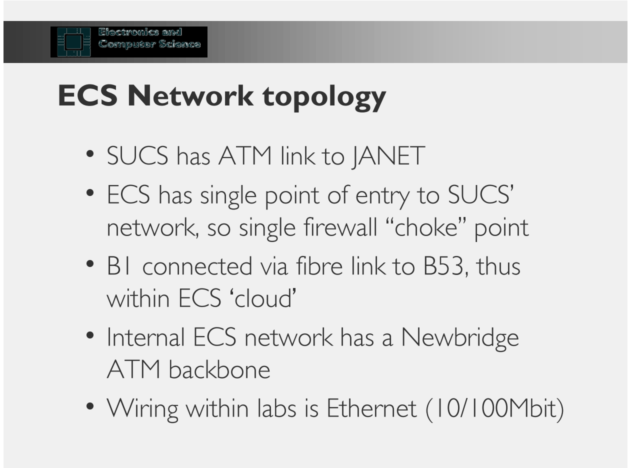

ECS Network topology

- SUCS has ATM link to JANET
- ECS has single point of entry to SUCS' network, so single firewall "choke" point
- B1 connected via fibre link to B53, thus within ECS 'cloud'
- Internal ECS network has a Newbridge ATM backbone
- Wiring within labs is Ethernet (10/100Mbit)

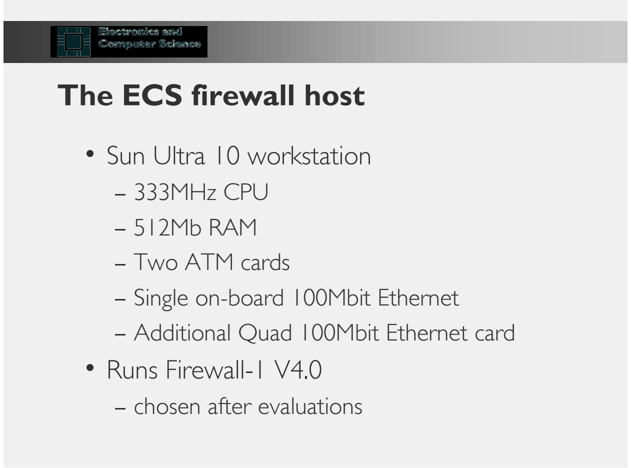

The ECS firewall host

- Sun Ultra 10 workstation
  - 333MHz CPU
  - 512Mb RAM
  - Two ATM cards
  - Single on-board 100Mbit Ethernet
  - Additional Quad 100Mbit Ethernet card
- Runs Firewall-1 V4.0
  - chosen after evaluations

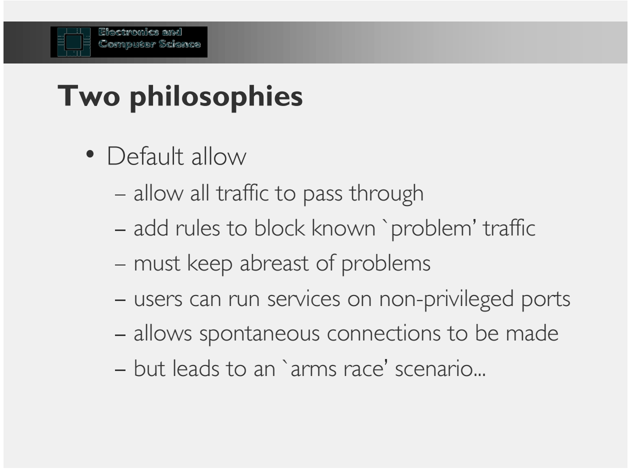

Two philosophies

- Default allow
  - allow all traffic to pass through
  - add rules to block known `problem’ traffic
  - must keep abreast of problems
  - users can run services on non-privileged ports
  - allows spontaneous connections to be made
  - but leads to an `arms race’ scenario...

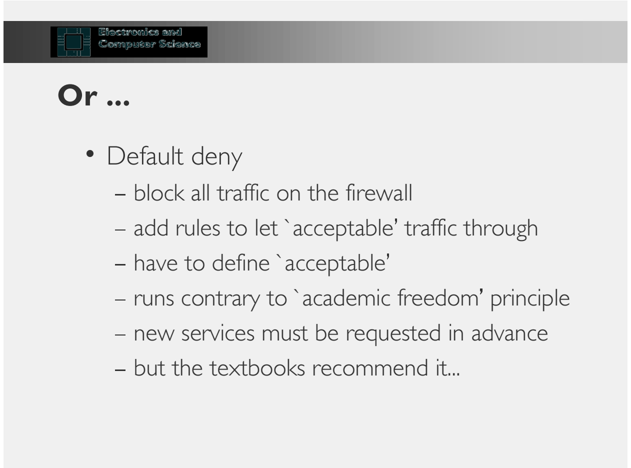

Or ...

- Default deny
  - block all traffic on the firewall
  - add rules to let `acceptable’ traffic through
  - have to define `acceptable’
  - runs contrary to `academic freedom’ principle
  - new services must be requested in advance
  - but the textbooks recommend it...

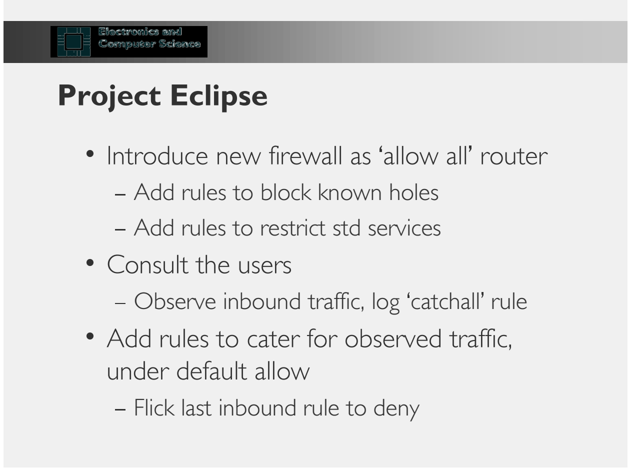

Project Eclipse

- Introduce new firewall as ‘allow all’ router
  - Add rules to block known holes
  - Add rules to restrict std services
- Consult the users
  - Observe inbound traffic, log ‘catchall’ rule
- Add rules to cater for observed traffic, under default allow
  - Flick last inbound rule to deny

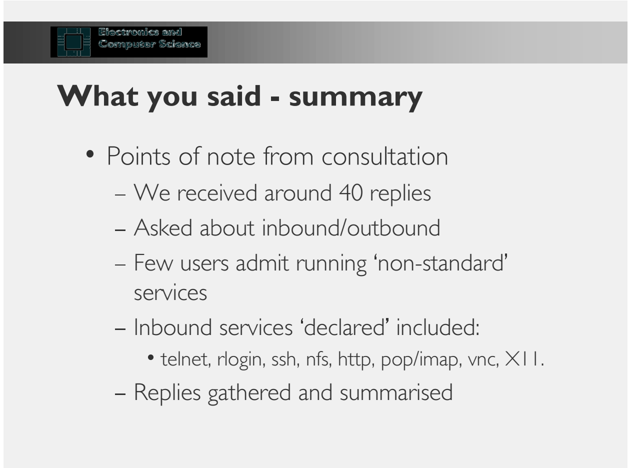

What you said - summary

- Points of note from consultation
  - We received around 40 replies
  - Asked about inbound/outbound
  - Few users admit running ‘non-standard’ services
  - Inbound services ‘declared’ included:
    - telnet, rlogin, ssh, nfs, http, pop/imap, vnc, X11.
  - Replies gathered and summarised

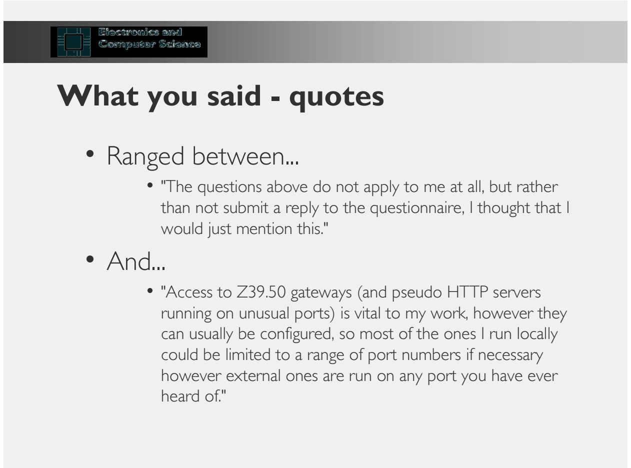

What you said - quotes

- Ranged between...
    - "The questions above do not apply to me at all, but rather than not submit a reply to the questionnaire, I thought that I would just mention this."
- And...
    - "Access to Z39.50 gateways (and pseudo HTTP servers running on unusual ports) is vital to my work, however they can usually be configured, so most of the ones I run locally could be limited to a range of port numbers if necessary however external ones are run on any port you have ever heard of."

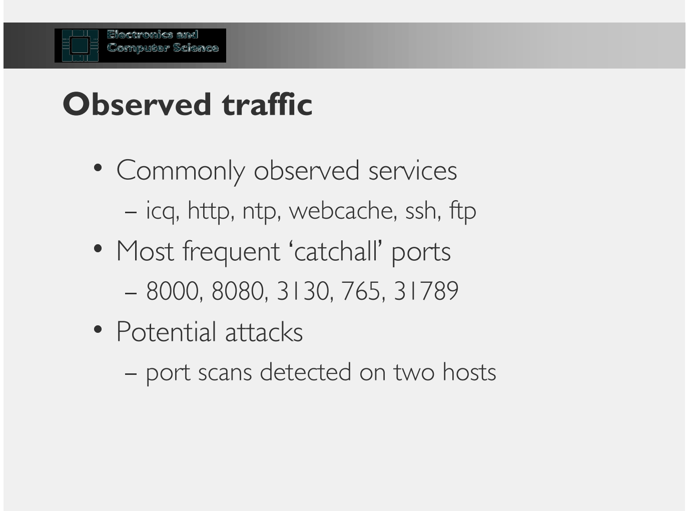

Observed traffic

- Commonly observed services
  - icq, http, ntp, webcache, ssh, ftp
- Most frequent ‘catchall’ ports
  - 8000, 8080, 3130, 765, 31789
- Potential attacks
  - port scans detected on two hosts

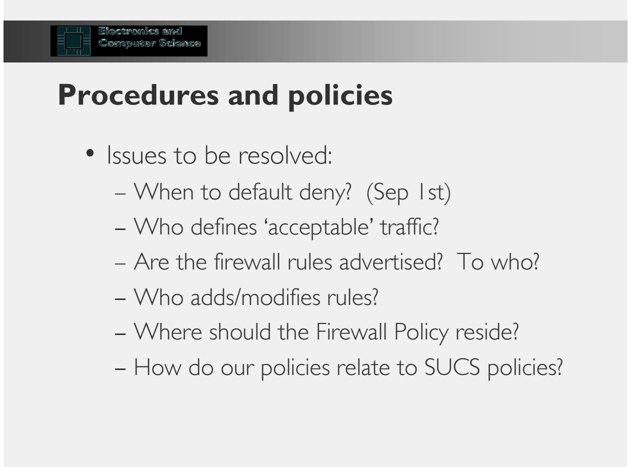

Procedures and policies

- Issues to be resolved:
  - When to default deny? (Sep 1st)
  - Who defines ‘acceptable’ traffic?
  - Are the firewall rules advertised? To who?
  - Who adds/modifies rules?
  - Where should the Firewall Policy reside?
  - How do our policies relate to SUCS policies?

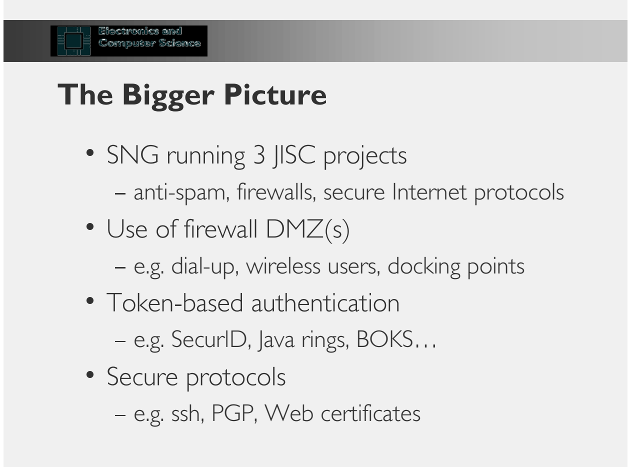

The Bigger Picture

- SNG running 3 JISC projects
  - anti-spam, firewalls, secure Internet protocols
- Use of firewall DMZ(s)
  - e.g. dial-up, wireless users, docking points
- Token-based authentication
  - e.g. SecurID, Java rings, BOKS…
- Secure protocols
  - e.g. ssh, PGP, Web certificates

17.4 Appendix D – Firewall Service Application Form

This form is to be completed by any student or member of staff who wishes to run an inbound IP service which passes through the ECS IP network firewall and which would otherwise be blocked by the ECS `Default Deny' firewall policy.

To be completed by applicant:

Description of service (assistance is available from systems staff, call x24494 if needed):

ECS host requiring external access (name/IP)

Host hardware and operating system

Purpose of service (nature and protocol)

Service characteristics (port, TCP/UDP etc)

External hosts accessing the service (name/IP)

Special requirements (e.g. time of day)

I will endeavour to ensure that any known security-related patches are applied to the service. I also accept responsibility for access offered through this service. I accept that the service may be blocked without notice if systems staff suspect a breach of security, and give permission for my personal details to be held for the purposes of maintenance and operation of the firewall and its rulesets.

Name Staff/Student number

Group Project name (if applicable)

I also confirm that the above IP service will be used in accordance with the regulations and codes of practice applicable to the Department as specified at *http://www.ecs.soton.ac.uk/support/ecscop/*

Date of last modification of version of regulations studied

Signed Date

Tutor name (if student/postgrad) Tutor signature (if student/postgrad)

To be completed by Systems and Networks Manager:

Expiry date on service Special notes

Request approved by Date

17.5 Appendix E - Firewall Vendors and Products

Actane Controller	Actane	http://www.actane.com/
AFS 2000 Internet Device	Internet Devices	http://www.internetdevices.com/AFSHOME.htm
Altavista Firewall	Altavista	http://altavista.software.digital.com/firewall/
Ascend Secure Access Firewall	Ascend	http://www.ascend.com/318.html
aVirt Gateway	aVirt Gateway Solutions	http://www.avirt.com/gateway.html
BIGfire	Biodata	http://www.biodata.com.sg/
BorderManager	Novell Consulting	http://www.novell.com/bordermanager
BorderWare	BorderWare Technologies	http://www.borderware.com/
Brimstone	SOS Corporation	http://www.soscorp.com/products/Brimstone.html
Netwall	Bull	http://www.bullsoft.com/
Centri	Cisco	http://www.cisco.com/centri/
Cisco IOS Firewall	Cisco	http://www.cisco.com/warp/public/732/net_foundation/fire_ds.htm
Cisco PIX	Cisco	http://www.cisco.com/pix/
Citadel	Citadel Data Security	http://www.cdsec.com/
Conclave	Internet Dynamics	http://www.interdyn.com/
CONNECT:Firewall	Sterling Commerce	http://www.sterlingcommerce.com/
ConSeal PC	Signal 9 Solutions	http://www.signal9.com/
COOL-FIRE	Symbolic	http://www.symbolic.it/
Cowboyz Firewall	Cowboyz.com	http://www.cowboyz.com/
cIPro-FW	Radguard	http://www.radguard.com/
CSM Proxy Plus	CSM USA	http://www.csm-usa.com/
CyberGuard	CyberGuard Corporation	http://www.cyberguardcorp.com/
Cybershield	Data General	http://www.dg.com/
CyberwallPLUS	Network-1	http://www.network-1.com/products/products.htm
CYCON Labyrinth	CYCON Technologies	http://www.cycon.com/firewall/firewall.html
Digital Firewall Service	Digital	http://www.digital.com/
Elron Firewall	Elron Software	http://www.elronsoftware.com/fwindex.html/
NetSeer Light	enterWorks	http://www.enterworks.com/
F100	Netasq Secure	http://www.netasq.com/
FLUX EF Enhanced Firewall	Inter Networking Systems	http://www.ins.de/ins/engl/eflux.htm
Firewall-1	Checkpoint Software Technologies	http://www.checkpoint.com/
Fort Knox Firewall Device	Internet Devices	http://www.internetdevices.com/FKHOME.htm
FreeGate	FreeGate Corporation	http://www.freegate.com/
Fuego Firewall	Cendio Systems	http://www.cendio.com/fuego/
Gauntlet	Trusted Information Systems	http://www.tis.com/
GEMINI Trusted Security Firewall System	Gemini Computers	http://www.geminisecure.com/
GFX Internet Firewall System	Global Technology Associates	http://www.gta.com/
GNAT Box	Global Technology Associates	http://www.gnatbox.com/
HSC GateKeeper	Herve Schauer Consultants	http://www.hsc.fr/
IBM Firewall for AIX	IBM	http://www.ics.raleigh.ibm.com/firewall/
ICE Block	J. River, Inc.	http://www.jriver.com/
Instant Internet	Deerfield Communications	http://www.deerfield.com/instanti/
Interceptor	Technologic	http://www.tlogic.com/
InstaGate	Technologic	http://www.tlogic.com/

InterLock	MCI Worldcom Advanced Networks	http://www.ans.net/InterLock/
IPAD 1200	Netmatrix Internet Co.	http://www.ipad-canada.com/
IRX Firewall	Livingston Enterprises	http://www.livingston.com/Marketing/Products/
iWay-One	BateTech Software	http://www.batetech.com/
Juniper	Obtuse Systems	http://www.obtuse.com/juniper/
KarlBridge/KarlBrouter	KarlNet Inc.	http://www.gbnet.net/kbridge/
Kwall Firewall	BSJ Enterprises	http://www.kwall.com/
LANguard	GFI FAX & VOICE Ltd.	http://www.languard.com/
Lucent Managed Firewall	Lucent Technologies	http://www.lucent.com/security/
LuciGate	LUCIDATA	http://www.lucidata.com/
MIMEsweeper	Integralis	http://www.integralis.com/
M>Wall	MATRAnet	http://www.matranet.com/
NetCS	NetCS Informationstechnik GmbH	http://www.netcs.com/
NetGate	SmallWorks	http://www.smallworks.com/
NetGuard Control Center	LanOptics	http://www.ntfirewall.com/
NetRoad/FireWARE/FireWALL	Ukiah Software	http://www.ukiahsoft.com/
NetSafe	Siemens Nixdorf	http://www.swn.sni.be/Products/Internet/Netsafe/
Netscreen-100	Netscreen Technologies	http://www.netscreen.com/
Net SecurityMaster	SOLsoft SA	http://www.solsoft.com/
Netra Server	Sun Microsystems	http://www.sun.com/
Nokia IP & VPN Series	Nokia Telecommunications	http://www.iprg.nokia.com/
Normal Firewall	Norman Data Defense Systems	http://www.norman.com/
Novix	FireFox	http://www.firefox.com/
Orion	Zebu Systems	http://www.zebu.com/
Phoenix Adaptive Firewall	Progressive Systems	http://www.progressive-systems.com/
PORTUS	Livermore Software Laboratories	http://www.lsli.com/
PrivateNet	NEC Technologies	http://www.privatenet.nec.com/
Pyramid Firewall	DataTec	http://www.datatec.co.uk/
Raptor	Axent Technologies	http://www.axent.com/
SecurIT Firewall	Milkyway Networks	http://www.milkyway.com/
Sidewinder, SecureZone, Secure Computing Firewall for NT	Secure Computing	http://www.securecomputing.com/
Site Patrol	BBN Planet Corp	http://www.bbnplanet.com/
SonicWall	SonicWALL	http://www.sonicwall.com/
SPF-100, SPF-200	Sun Microsystems	http://www.sun.com/security
Sygate	SyberGen	http://www.sygate.com/sygate.html
T.REX: Open Source Firewall	Freemont Avenue Software	http://www.opensourcefirewall.com/
TUNIX Firewall	TUNIX Open Systems Consultants	http://www.tunix.nl/
TurnStyle Firewall System	Atlantic Systems Group	http://www.asg.unb.ca/
VCS Firewall	The Knowledge Group	http://www.ktgroup.co.uk/
Watchguard	WatchGuard Technologies	http://www.watchguard.com/
WebSENSE	NetPartners Internet Solutions	http://www.netpart.com/
WinGate	Deerfield Communications	http://www.deerfield.com/wingate/
ZapNet!	IPRoute/Secure	http://www.iproute.com/