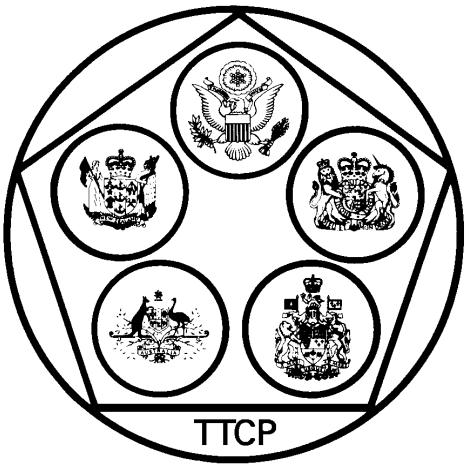

THE TECHNICAL COOPERATION PROGRAM

SUBCOMMITTEE ON NON-ATOMIC MILITARY RESEARCH AND DEVELOPMENT

TTCP DOCUMENT

Safe Use of the Internet for Defence Purposes

DOC - C3I - 1 1997

A Report from the Panel on Secure Information Systems, STP-11, C3I Group

PREFACE

This report was drafted at the first meeting of STP-11, Secure Information Systems, held at DERA Malvern, United Kingdom, April 21-25, 1997. The following STP-11 members contributed to the report:

Dr. Carl Landwehr, US, Naval Research Laboratory, Chairman and National Leader
Mr. Peter Drewer, AS, Defence Science and Technology Organisation, National Leader
Mr. Vincent Taylor, CAN, Department of National Defence, National Leader
Mr. Paul Devlin, CAN, Communications Security Establishment, Member
Mr. Alex Scott, UK, Defence Evaluation and Research Agency, National Leader
Mr. Dwayne Allain, US, USAF Rome Laboratory, Member
Mr. Michael Harrison, US, NCCOSC Research and Development, Member
Dr. John McLean, US, Naval Research Laboratory, Member

The Panel acknowledges the substantial and very helpful support of the DERA staff in preparing this report. The contributions of Simon Harding are particularly appreciated. Review comments by Mr. Michael Reed of NRL improved the report.

Safe Use of the Internet for Defence Purposes

A Report from the TTCP Panel on Secure Information Systems

June 2, 1997

Executive Summary

The choice for the Defence establishments of the TTCP member nations is not whether to use the Internet or not, but how to use the Internet safely and prudently.

There are clear risks in connecting systems to the Internet, particularly systems that process sensitive information. Prudent use of the Internet for defence purposes must assume, at present, that the Internet is a broadcast medium that is open to corruption and service denial. Traffic is world readable and world writable. The lack of authentication mechanisms in current Internet Protocols permits spoofing and connection highjacking. Relative to the public switched telephone networks, the Internet is unreliable. Although the Internet may prove a useful communication medium in emergency situations, it is risky to rely on its availability in times of heightened world tension.

Nevertheless, much routine Defence business is now appropriately conducted using the Internet. A significant portion of this traffic, although unclassified, is sensitive in one way or another, however, and although the Internet can safely carry this traffic, certain precautions should be taken:

- For traffic that is sensitive, even if unclassified, it is prudent to apply encryption within the management domain of the end system, for example, by using IPSEC mechanisms at the network layer or commercial encryption products such as Entrust at the application layer.
- For traffic whose authenticity or integrity is significant, apply digital signatures externally.
- For traffic whose timely delivery is critical, obtain positive acknowledgment of receipt, so that nondelivery is evident.
- Where measures such as digital signatures and encryption are used, pay close attention to how keys are generated, distributed, and stored. These factors limit the assurance these measures can provide.

These precautions will help ensure confidential, reliable, timely e-mail delivery. They do not, however, provide protection from another hazard of Internet connection: site intrusion. Such protection requires administrators of unclassified Defence systems connected to the Internet to:

- Carefully consider the configuration of internal networks and their vulnerability to attacks from the Internet. In most cases, it will be prudent to install a properly configured firewall between the internal networks and the Internet. In general, the administrator should be required to justify why a firewall is *not* needed, rather than the reverse.
- Maintain awareness of, and act upon, advice from CERT, AUSCERT, FIRST, and other relevant sources regarding vulnerabilities and attacks, and be sure available security patches are installed as they become available.

- Introduce measures that, so far as possible, prevent users from sending reusable passwords in the clear over Internet links and prevent internal systems from accepting cleartext reusable passwords from Internet links.
- Ensure that users receive appropriate security training, especially concerning the hazards of importing software (including plug-ins, Active-X controls, and Java applets, as well as more conventional forms) from Internet sources.

Administrators of publicly accessible Defence web servers, should in general

- Position the web server on the Internet side of the firewall on a separate processor dedicated to this purpose, configured with the minimum software required to provide the type of web support desired.
- If the server is to accept input via scripts (e.g., Common Gateway Interface (CGI) scripts), be aware of the vulnerabilities these mechanisms can contain and review new scripts for flaws before installing them

Administrators of classified networks should not connect their systems to the Internet without proper certification and accreditation. Current and emerging technology discussed in this report, including the Starlight Interactive Link for safe connection to the Internet from within enclaves, SINTRA/Pump technology for reliably moving low level information upward, and Onion Routing for impeding traffic analysis, can make such connections safer, cheaper, and more useful. Demonstrations of how these technologies and others, including Entrust and Purple Penelope, could be combined to support TTCP roster and calendar management, safe WWW browsing from within an enclave, importation of open source materials into a labelled domain, and coordination of intrusion detection systems should be pursued.

Technologies for Internet Protocol Security (IPSEC) and improved Domain Name System security now being developed will improve the security infrastructure of the Internet, but they will not soon obviate the measures we recommend. The open nature of the Internet means that its hardware and software will, for the foreseeable future, be open to compromise, so external mechanisms will need to be applied if Defence interests are to take advantage of its benefits without being subject to its weaknesses.

Finally, we recommend Defence R&D investments relevant to Internet security focus on development of high assurance network components, cryptography, and key management, and that Defence departments seek to influence commercial developments in Internet protocols, application program interfaces, commercial cryptographic protocols and public key infrastructures so that they will be useful in a defence, as well as commercial, environment.

CONTENTS

Safe Use of the Internet for Defence Purposes

A Report from the TTCP Panel on Secure Information Systems

1. Why Connect Defence Systems to the Internet?

The insecurities of the Internet are easy to document. Today, normal traffic on the Internet is transmitted in the clear, unencrypted. It is transmitted in standard formats using public protocols. It passes through switches and communication facilities that are publicly accessible and in many cases relatively uncontrolled. Under such conditions, almost any kind of attack is feasible. Eavesdropping unencrypted information is nearly trivial. Even if some level of encryption is applied, the source and destination of most traffic remains open. Identifying the participants in a routine conversation is not difficult even if the contents are somewhat hidden. On the other hand, it is also relatively easy to generate packets with bogus source address information, permitting spoofing attacks on many levels. Messages can be undetectably altered in transit unless cryptography is employed. Denial of service attacks can be mounted against specific portions of the network and may be very difficult to stop, so there may be unpredictable outages or delays in transmission.

Nevertheless, the defence bureaucracies, and, to an increasing extent, the defence forces of TTCP member nations, are finding the Internet a vital tool. Ships at sea, tanks on the battlefield, and aircraft on patrol have all used Internet links to retrieve information that they found essential. Governments who have attempted to implement systems not in the mainstream of technical development, which the Internet and internet technology have come to dominate, have encountered difficulties that provide a compelling argument for the use of public data networks. But such use must be supported by a solid security architecture, which is currently lacking.

More specifically, the Internet offers:

- *World-wide e-mail*: It is doubtful whether any other communication medium provides access to as many defence-related people and organizations while also buffering messages so that participants need not be simultaneously accessible in order to communicate. Telephone connectivity exceeds Internet connectivity, but buffering is provided only by answering machines or voicemail systems that are less widely deployed and lack the capability to transmit the scope and variety of information that can be sent digitally. On a mundane level, much of the coordination for meetings such as the one at which this report was produced now takes place through international e-mail. A few keystrokes can replace many telephone or telefax messages.
- *Access to open information sources via WWW*: Since the release of the NCSA Mosaic browser in 1993, use of the World-Wide-Web and its associated protocols and formats has mushroomed. Not only research publications and sources of entertainment can be found on the Web, but also current weather information and news that may affect defence operations are available. For many organizations, the Web is becoming a

primary way to provide and retrieve public information. All kinds of public software, including browsers, helper applications, other utilities for reading data in a variety of formats, games, and software patches, are routinely distributed in this fashion.

- *Real time communication*: E-mail can occasionally provide close to real time communication if both parties are connected and active, but voice and video communication via the Internet is also becoming available. At the present time, bandwidth limitations for voice and video are sometimes limiting, but there is a great deal of activity both in increasing the bandwidth available and in the development of efficient protocols to support video and voice transmission over the Internet. Use of the Internet for teleconferencing purposes is likely to grow substantially as these facilities become available because of the wide and relatively low cost access available to the Internet from many defence facilities and the scarcity and high cost of competing teleconferencing equipment. STP-11's predecessor, XTP-1, successfully presented its annual report to SGX via Internet/MBone links in the summer of 1995.
- *Electronic commerce*: Following the lead of industry, governments and defence departments are expected to adopt electronic commerce as their preferred means to conduct business with industry suppliers. Business transactions conducted over the Internet should be more efficient and less error-prone.
- *Resilience in non-critical situations*. While the Internet does not match the availability and reliability of the telephone and power grids in most parts of the world, it provides adequate reliability for most routine uses. Despite widely publicized forecasts of impending "Internet collapse" and substantial increases in its use, no major long term outages have been observed, and the outages that have occurred have generally been resolved fairly quickly. Increasing commercial use of the Internet should act both to stimulate the supply of facilities and to improve availability and reliability.
- *Easy accessibility and interoperability*. Access to the Internet in the TTCP countries is readily available from defence facilities. Increasingly, access from mobile platforms is commercially available. The Internet protocol suite has become a standard in the commercial marketplace, a goal many defence programs have sought but few have achieved. Commercial networking and operating systems are being converted to use the Internet protocols without government support.
- *Low cost*. Compared to alternative telecommunication facilities, Internet communication is cheap. In the U.S., the competitive rate for unlimited use of a 28.8KB Internet connection from a private residence is $20 per month at present. Other TTCP nations may face higher costs, and faster channels will be more expensive, but any reasonable assessment of the costs of alternative communications is likely to end heavily in favor of the Internet. For this reason, use of the Internet in place of dedicated leased lines is an attractive path for many organizations.

In the longer term, the development of the Global Information Infrastructure (GII) will blur the distinction between private and public data networks and the dominant data communication service will be based on a post-Internet public network. Defence Forces will be under increasing pressure to exploit advanced strategic communications across the GII in order to gain substantial cost reductions. This will require that a security architecture be provided for Defence use of the GII.

2. What Vulnerabilities Does an Internet Connection Pose?

Unfortunately, the benefits of the Internet are accompanied by some risks. We alluded to these briefly in Section 1; this section provides additional details on different classes of vulnerabilities. We first briefly summarize the architecture of the Internet in order to help the reader understand the points of vulnerability.

2.1 Internet Architecture

The Internet is a global, decentralized network of communication links, switching devices, and computers that communicates using the Internet protocol suite. The U.S. Federal Networking Council adopted the following definition of "Internet" on October 24, 1995 [FNC]:

> "Internet" refers to the global information system that (i) is logically linked together by a globally unique address space based on the Internet Protocol (IP) or its subsequent extensions/follow-ons; (ii) is able to support communications using the Transmission Control Protocol/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons, and/or other IP-compatible protocols; and (iii) provides, uses or makes accessible, either publicly or privately, high level services layered on the communications and related infrastructure described herein.

To understand where and how information flowing on the Internet can be disclosed, modified, or delayed, it is useful to understand something of its structure.

A. Links

The primary links of the Internet are provided by lines leased from common carriers in countries around the world. In other words, they are part of the public switched telephone networks (PSTNs), although lines used for Internet traffic are in general distinct from those used for voice traffic -- a line leased for Internet use will be dedicated to that purpose. The actual media that carry the traffic are thus as varied as the telephone networks themselves, ranging from copper wire to optical fiber and satellite links. Internet protocols have been used successfully over radio and cellular phone links as well.

B. Nodes

The primary nodes of the Internet are called routers. These are computers programmed to accept and forward packets of data (which today may represent real-time voice or video information as well as numbers or text) across the links. A router will be connected to two or more links. It examines each inbound packet for its destination address and, based on its routing tables, determines over which outbound link the packet should be forwarded.

In addition to the routers, there are host computers, which represent the sources and destinations of the packets that are routed through the network. Additional types of network nodes include gateways and firewalls; both of these can be considered types of routers. Gateways that connect networks using different protocols to the Internet will need to provide protocol translation services and addressing services. Firewalls can filter incoming and outgoing traffic, translate addresses and more.

Routing tables can be updated dynamically in order to permit the network to adapt both to outages and to new links. Although the original theory behind the Arpanet's switching scheme was to provide fully decentralized and dynamic routing, the tremendous size to which the Internet has grown has led to somewhat more hierarchical and static routing regimes in practice. Each packet in a series originating from a host attached to a local Internet Service Provider (ISP) and destined for a host attached to an ISP in another country is likely to traverse the same route, and that route is unlikely to include random ISPs in either country. Rather, the packets will be sent from a local ISP "upwards" in the source country to major routers and circuits controlled by the large common carriers, "across" to the major carriers in the destination country, and "down" to the local ISP that connects the destination host. The routing on the links controlled by the major common carriers can vary, but it is relatively static and is unlikely (though not impossible) for it to cause packets to traverse arbitrary routers in out-of-the-way corners of the Internet. There may, of course, be extensive local networks in place at both the host and destination locations, and the packets may be broadcast widely within those areas depending on the configurations and protocols in use.

C. Addressing

Internet Protocol Version 4 (IPv4) network layer addresses consist of four octets (32 bits, usually represented as four decimal numbers from zero to 255, separated by dots: 132.250.80.57), logically divided into a link number and a host number (originally, link number consisted of a net number and subnet number, but this distinction proved unuseful [Perl92]). A host may be connected to more than one link and may therefore be reached by more than one IP address. Packets contain both source and destination addresses, and software in the sending node is responsible for providing both of these. In addition to the source and destination, the IP header may contain a number of optional fields, including a method to specify a route through the network and to require the destination host to route packets back to the source over the same route (this is the loose source route option [ChBe94]).

The next generation Internet Protocol (Version 6, or IPv6 [SBAM96]) expands the IP address space substantially, allowing 128 bits for source and destination addresses. IPv6 addresses are assigned to interfaces, not nodes; any of the unicast addresses assigned to any of the node's interfaces can be used to identify that node. Both the structure and the intended use of these addresses is complex and details are still being decided. We do not address these issues here.

D. Domain Name System

Packets traverse the Internet using numerical addresses, but users and programs usually deal with more mnemonic addresses in the form of Domain Names. The Domain Name System (DNS) provides the infrastructure for translating domain names into IP addresses [RFC-1034,RFC-1035, RFC-2065].

Syntactically, a domain name is a sequence of character strings separated by dots ("."), such as "dsto.defence.gov.au". Although software that processes domain names preserves the case of the character strings, the case adds no semantics, so "DSTO.Defence.gov.AU" has the same meaning as the previous example. The domain name space is a tree structure. Each node and leaf has a label, from one to 63 octets in length and a (possibly empty) resource set. The root node (only) has a null (zero length) label. The domain name of a node is the list of labels from that node to the root. To simplify implementations, the total number of octets that represent a domain name is limited to 255.

In fact, DNS can provide more than a simple translation from an object name to an IP address. The resource records stored for a domain name that corresponds to a host, for example, may indicate what operating system and version number the host is running. Other resource records can designate a host that processes incoming mail for the specified domain, identify a name server for a domain, or map an alias to the real domain name for a host. DNS can also provide an inverse mapping from IP address to host name; there is no enforced relationship between the inverse mapping and the forward mapping, however [ChBe94]. Many host names can map to a single IP address, but given an IP address, DNS returns a single corresponding domain name.

The size of the Internet dictates that the DNS database is distributed among many servers (called Name Servers), none of which has a complete copy. A name server knows the parts of the domain tree for which it has complete information; it is said to be an Authority for those parts of the space. Authoritative information is organized into units called Zones, and these zones can be automatically distributed to the name servers that provide redundant service for the data in a zone. A name server must periodically refresh its zones from master copies in local files or foreign name servers. A name server may also cache information about other parts of the domain tree for which it is not an authority.

To use DNS, a program typically invokes a local procedure called a Resolver. The resolver contacts a name server (likely, but not necessarily, on a different machine) to satisfy the request. The name server either provides the information requested by the resolver or a pointer to another name server that the resolver can try.

The original DNS database was designed to support queries to a statically configured database. Changes to the database were expected to be infrequent and were to be made as external edits to a zone's Master File. Mechanisms for permitting dynamic updates to DNS records have recently been proposed [VTRB96].

Although DNS is a critical operational part of the Internet infrastructure, it has no strong security mechanisms to assure data integrity or authentication. However, extensions to provide these services to security aware resolvers or applications through the use of cryptographic digital signatures are under review [RFC-2065]. These mechanisms are also being extended to the proposed dynamic update mechanisms [East97].

E. Protocols: IP, ICMP, UDP, and TCP

The Internet Protocol (IP) provides an unreliable, connectionless "best effort" delivery service that routes datagrams (packets) towards a specified IP address. IP also includes a protocol for reporting errors, the Internet Control Message Protocol (ICMP); routers use ICMP messages to report delivery failures, misroutings, congestion, and related problems to each other. Both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are layered on top of IP. Application programs can use TCP or UDP to send messages to applications running on remote hosts [Come91].

TCP uses IP services to implement a reliable, connection-oriented transport service. TCP tries to guarantee that messages it receives for transmission are delivered to the correct address uncorrupted, without duplication, and in order. User Datagram Protocol (UDP) uses IP services to provide connectionless datagram service. UDP packets may be lost, duplicated, or delivered out of order.

Both UDP and TCP incorporate the notion of "ports" to distinguish traffic sent to the same IP address but for different recipients. A port number in both protocols is a 16-bit integer. An application on one machine can send UDP datagrams to different processes (listening to different UDP ports) on a remote machine by addressing the datagrams to different ports. UDP queues traffic for different ports independently. For TCP, the connection is the fundamental abstraction, and a connection is specified by its two endpoints. Each endpoint is a pair of integers (host, port), where "host" is the host's IP and "port" is a TCP port number on that host. This arrangement permits, for example, a program that accepts incoming mail to use only one local TCP port even though it may be communicating over many connections concurrently. TCP and UDP port numbers are independent, since each message specifies its protocol, as well as its destination IP address and port number, and in both protocols certain protocol numbers are used, by convention, as the addresses for particular services. Such a reserved, and advertised, port number is called a "well-known port." The well known port for e-mail delivery, for example, is TCP port 25; UDP port 53 provides access to the Domain Name Service (as does TCP port 53, but UDP is normally used for initial DNS queries).

A TCP connection is opened with a "three-way handshake": (1) the initiating host sends a "SYN" segment with its IP address and an arbitrary number sequence number N, (2) the destination host replies by sending an acknowledgment and another arbitrary sequence number M, and (3) the initiator completes the protocol by acknowledging the second message. Closing connections is slightly more complicated, as there are various contingencies to consider -- for example, there may be segments en route at the time when one of the participants requests the that connection be closed

In UDP, of course, there are no connections, so there is no protocol to set them up or tear them down.

F. Application Layer Protocols: Telnet, FTP, SMTP, HTTP, SNMP

The protocols that most Internet users see most directly are those at the application layer. These protocols support the transport of e-mail (SMTP) and files (FTP), the initiation of terminal sessions on remote hosts (Telnet), and the operations of World Wide Web Browsers (HTTP) and many other functions. These protocols may use TCP, UDP, or both to accomplish their functions. Typically, a host that supports a particular service, such as e-mail delivery, will have a process or processes that listen for requests phrased in the appropriate protocol (such as SMTP) over a particular port, as noted above.

2.2 Vulnerabilities

Any system that is connected to the Internet and uses the Internet to communicate with other systems inherits many well-known vulnerabilities. Some of these arise in the Public Switched Telephone Networks (PSTNs) that carry Internet traffic among routers and hosts. Others come from the routers and hosts themselves. Many have to do with the ways in which people use the Internet and come to rely on it as a means to access both information and processing resources. This section lists a number of such vulnerabilities according to the problems they can cause.

A. Network

Sniffing. Computer networks are shared communication channels. It is simply too expensive to dedicate local loops to the switch (hub) for each pair of communicating computers. Sharing means that computers can receive information that was intended for other machines. To capture the information going over the network is called sniffing.

The most popular way of connecting computers is through Ethernet. The Ethernet protocol works by sending packet information to all the hosts on the same circuit. The packet header contains the proper address of the destination machine. Only the machine with the matching address is supposed to accept the packet. A machine that is accepting all packets, no matter what the packet header says, is said to be in promiscuous mode.

Because, in a normal networking environment, account and password information is passed along Ethernet in clear-text, it is not hard for an intruder who gains control of a machine to put it into promiscuous mode and, by sniffing, compromise all the machines on the net.

Wardialing. Although the system security policies of most protected networks will limit uncontrolled connections to the Internet, many such connections exist illicitly. Such a connection, typically provided by a modem connected to the telephone system, can be exploited by any attacker who identifies it.

A common method of identifying such connections is wardialing, in which an attacker programs a computer to connect to a series of telephone numbers (e.g. all numbers starting 01684 89). Any answers which are either a modem or a fax machine are logged for further investigation.

As these connections are often installed without consent or knowledge of the higher levels of the organization, the organization rarely audits attempted connections, and so wardialing is difficult to detect.

B. Protocol

Data Link Layer Security. Address Resolution Protocol (ARP), which is used to translate Ethernet addresses on a LAN to IP addresses, is open to manipulation. For instance, Unix System V does not check whether received ARP packets are associated with an outstanding request. This could result in malicious responses to ARP requests and unsolicited updates to ARP tables. The most likely impact is a denial of service, though "man-in-the-middle" attacks, in which one address masquerades for another, are also possible.

Network Layer Security. The implementation of IP is generally robust but can be manipulated.

- Routing is fairly open. This could lead to data not conforming to configured routing.
- IP Packets can be injected directly onto the network.
- ICMP has no authentication, which could permit manipulation of routing
  - a malicious user could subvert local routing tables
  - ICMP could also permit unsolicited address mask reply packets

Network security can be further compromised by the protocols used to manage the network routers. The Simple Network Management Protocol (SNMP) has poor authentication, and unless the routers are correctly configured, they are vulnerable to malicious reconfiguration.

Transport Layer Security. TCP has weak port number mechanism, so that there is no guarantee that non-UNIX systems (e.g. PCs running Windows) conform. TCP checksumming of IP packets is not strong, leading to a potential for forgery, injection of

data and tailgating of packets. The randomness of TCP initial sequence numbers varies across UNIX systems, leading to a potential to inject packets into a connection between two users.

IP Origin Forgery. The origin of an IP message can be forged relatively easily. This in itself is not a serious vulnerability. However, as many higher level protocols use the IP origin as a form of identification, it becomes serious. For instance, the r-commands, which allow managers of one UNIX system to control another, use IP source address as a primary authentication method.

C. Application

Many of the higher level protocols can be exploited to attack systems connected to the Internet. Many vulnerabilities are well known, such as those in versions of the sendmail system that allow an attacker to gain root privileges quickly. At this point, the attacker can stop audit of his actions, delete any previous audits, install Trojan horse software, read, modify or delete user applications or data, and then use the current system as a platform for launching an attack on further systems.

Although sendmail may contain the most infamous bugs, many other protocols and software components contain similar bugs and vulnerabilities. Many of these bugs originate from simple errors in the program code, such as a failure to check array bounds. Indeed, with the applications originating in university environments, and being written primarily with functionality rather than robustness in mind, it is not surprising that so many protocols and servers are open to attack.

As an example of vulnerabilities in newer protocols, consider the World Wide Web.

World Wide Web. The risks associated with World Wide Web (WWW) security vulnerabilities affect confidentiality, integrity and availability of information. A WWW server provides access to local information that the whole Internet can potentially access. If a Web server can be subverted then any or all the following may occur:

1. Confidential documents held on the server may be accessed by unauthorised users.

2. Public documents, e.g. Web pages, may be changed, thus destroying the integrity of the information. This has been graphically illustrated by recent attacks on Web servers belonging to the CIA, DoJ and the British Labour Party.

3. Remote users may execute commands on the Web server's host machine providing access to operating system level functionality.

Recently identified Web server vulnerabilities include a bug which allows remote users to download and read the contents of executable scripts, thus potentially accessing sensitive information such as database passwords or network configurations.

Many Web servers provide additional functionality through server side processing. This includes the use of Common Gateway Interface (CGI) scripts normally written in a scripting language such as Perl. CGI scripts may unintentionally leak information about a system or may be tricked into executing commands. An example is a widely published CGI script which inadvertently allows remote users to write files into directories to which they don't have access.

Client side network browsers may also suffer from security vulnerabilities. Most browsers provide functionality to allow additional applications to be specified when accessing a particular document type. It would seem natural to specify a word processing application to view documents. This, however, potentially leaves the application vulnerable to macro viruses.

Increasingly, there has been a move towards providing rich networked functionality via technologies such as Java and ActiveX. Java programs are precompiled and stored on a Web server. These mini-applications are known as Java applets and are downloaded to a network browser and executed locally. ActiveX is a technology for distributing software over the Internet. The ActiveX analog of a Java applet is called a *control*. An ActiveX control may be embedded in a Web page, where it may be accessed via a network browser.

Java has been explicitly designed to address security issues through various mechanisms which essentially restrict the behaviour of applets. However, a number of implementation problems have been identified allowing Java applets to execute arbitrary machine instructions, interfere with other applets and bypass the Java Security Manager.

ActiveX controls are not restricted in any way and rely on digital certification. If a browser accesses an ActiveX control which is not signed or the certificate is not recognised then a dialog box is used to warn the user. This behaviour is controlled by properties held in the browser. ActiveX controls have been published on the Internet which close down machines, format hard disks and install viruses. Relying on users to determine which ActiveX controls are safe to use is a dangerous strategy. History teaches that users quickly become bored with "warning" dialog boxes that appear very frequently. Users either stop paying attention to them or find a way to prevent their appearance.

D. Above The Application Level

Finally, it is worth considering classes of vulnerability that occur above the application layer.

Trojan Horses. A Trojan horse is a malicious, security-breaking program that is disguised as something benign, such as a directory lister, archiver, game, or, in one notorious 1990 case on the Macintosh, a program to find and destroy viruses. More recent Trojan horses have used the macro facilities within MS Word, allowing them to be constructed using an intuitive programming language.

Viruses. A virus is a program that searches out other programs and infects them by embedding a copy of itself in them. When these programs are executed, the embedded virus is executed too, thus propagating the infection. This normally happens invisibly to the user. A virus cannot infect other computers without assistance. More recent viruses have been based on document macro capabilities. The virus may do nothing but propagate itself and then allow the program to run normally. Usually, however, after propagating silently for a while, it may write messages on the terminal or play tricks with the display (some viruses include sophisticated display hacks). The damage a virus does, once it gains control of a machine, is limited only by the forbearance and competence of its author. Nothing else prevents a virus from causing irreversible damage, such as destroying all of a user's files.

In the 1990s, viruses have become a serious problem, especially among PC and Macintosh users (the lack of security on these machines enables viruses to spread easily, even infecting the operating system). The production of special anti-virus software has become an industry, and a number of exaggerated media reports have caused outbreaks of near hysteria among users; some users blame everything that doesn't work as expected on virus attacks. In some cases, the fear of a virus infection can waste as many resources as an actual occurrence, since users may be driven to adopt extreme precautions and broadcast unnecessary warnings.

Documents. A final vulnerability exists when exporting complex documents. It is often impossible to fully review the contents of a computer file, so that classified text can escape without proper sanitisation. As an example, most users of MS Word use the *fastsave* option. When text is deleted from a document and the document is then saved again, Word does not delete the text: it merely inserts a note not to display the text. The text can be recovered using another program which cannot recognise Word's instruction to itself.

3. What Off-the-shelf Technology Can Reduce these Risks?

A number of off-the-shelf technology products are available to mitigate the vulnerabilities outlined in the preceding section. These products can be grouped under the following broad headings: authentication mechanisms, encryption, intrusion detection, security management tools, firewalls and guards. Businesses throughout the world are starting to use such methods - judicious combinations of firewalls, encryption and authentication techniques, for example - to create corporate "Intranets" that are both connected to the Internet and reasonably protected from it. They recognise that their networks are not completely protected from the threats posed by connection to the Internet, but believe on balance that the net benefits outweigh the risks.

3.1 Authentication Mechanisms

A vulnerability in systems connected to the Internet is the system password file. If a system password is compromised, then the system may be wide open to a variety of attacks. Thus it is imperative that an unguessable password is chosen, and also that that password never be passed over the Internet in the clear. Tools currently exist that permit the use of one-time passwords over the network; thus the disclosure of the password does not compromise the system. Shortly, tools will become common to allow a user of a system to authenticate itself to a server across the Internet via an encrypted session.

One such tool is Kerberos, a publicly available and widely used system that can provide reliable authentication over open networks such as the Internet. Kerberos is a secret key authentication system that involves a central database keeping copies of the secret keys of all users. It uses DES for encryption and authentication (though the latest version permits other algorithms), and allows entities to communicate over networks and to prove their identity to each other while preventing eavesdropping or reply attacks. It also provides for data stream integrity and secrecy. Its dependence on a central database limits the ability of Kerberos to scale to very large (e.g., 100,000 or more) user communities, but it has been used effectively in systems with thousands of users. This is not to say it is without vulnerabilities; its most current version, which was developed to remedy some shortcomings in earlier versions, is numbered 5.0.

Systems now exist which will allow encrypted and/or digitally signed information to flow across the Internet between cooperating sites. At present the sites must use proprietary

software and tokens and supporting infrastructure, such as those developed by Entrust Technologies. However, this confidentiality- and integrity-preserving technology is standards-based and should maintain compatibility with similar emerging technologies. The technology is being taken up by a number of software vendors as an integrated security solution to enhance usability.

The foregoing helps to protect and preserve data transmitted over the network. The use of one-time passwords helps to protect the site computer from packet sniffing attacks. However, as noted in the previous section there are many and varied attacks that can be mounted against systems connected to the Internet. By hiding the site computers behind carefully configured firewalls, some of the risks to the site of connecting to the Internet can be reduced to acceptable levels

3.2 Encryption

Encryption and digital signature (analogous to a handwritten signature) offer solutions to five important security requirements: confidentiality, access control, integrity, data origin authentication, and non-repudiation. Commercial encryption products (and some public domain implementations) are available that provide encryption services at different protocol layers, including the application layer, session layer, and the IP layer. There are also significant efforts underway to standardize access to encryption services by application programs (Crypto Application Program Interfaces, or CAPI), so that a variety of standards-based products can come to market (see [ICE 97] for further information).

Application layer encryption refers to encryption services applied on a per-application basis. An application layer encryption system based on public key certificates effectively involves three main components - a User Agent (UA) to which a user is authenticated; a repository of public keys for the potential decryptors of encrypted data; and an associated encryption/digital signature capability. The UA acts on behalf of the authenticated user, so any digital signature created by the UA on his behalf can only be as strong as the original authentication between the user and the UA. Encryption is performed by the system for specific recipients, which may include the originator. Encryption systems are typified by "Entrust" from Entrust Technologies, Inc. In its basic form, the system includes a Public Key Infrastructure and a tool for encrypting and signing or decrypting and reading signatures. Entrust also includes a toolkit for software vendors to allow them to integrate these capabilities into their products. At present there are already e-mail, e-form and word processing applications which are "Entrust aware". So far no applications which have been made or are evolving to be "Entrust aware" are real-time applications (e.g. video-conferencing), although there appears to be no fundamental restriction against this use.

Session encryption acts in a similar way to an application encryption system, except that the authenticated entities will be the agents participating in the session. These may be, for example, a web browser and server. All data transmitted between the agents during a session will be encrypted but will be in the clear when presented to the applications. Typical session encryption would not authenticate users but rather applications on hosts. Examples of current session-based systems are SSL (Secure Sockets Layer), developed by Netscape and widely deployed in Web browsers and servers and "Defensor" from Sagus Securities, Ltd., which uses the Entrust infrastructure for its security services. Defensor is designed to set up secure sessions between PC workstations or between workstations and servers or firewalls.

IP layer encryption is used by some commercially available encryptors that permit using the Internet as a substitute for a dedicated private line, although these tend to be relatively expensive and key management may be an issue. A pair of these encryptors, each acting as a gateway between a private network and the Internet, will route all traffic to each other via Internet routers, which will see cleartext IP headers but encrypted payloads. The address of the actual destination system can be hidden in the encrypted payload, so that the Internet routers only see the addresses of the encrypting gateways. Firewalls that provide Virtual Private Networks (VPNs) use essentially this approach. Although this technique has its uses, it effectively limits protected communications to those sites that are behind the encrypting firewalls; generalized web browser and e-mail delivery cannot be protected. The Internet Protocol is being extended with a security option (IPSEC) that will permit more flexible use of encryption at the IP layer.

3.3 Intrusion Detection

Despite the best efforts of the protocol designers, implementers, and system administrators, it is prudent to assume that attacks will occur and some, unfortunately, will succeed. Therefore, it is vitally important that a means to detect and respond to these attacks is installed to protect critical information services. Both Commercial and Government Off the Shelf intrusion detection products exist that detect and provide alerts to known attack methods.

The intrusion detection systems that currently exist are categorised into two types; network-based and host-based. Network-based intrusion detection systems pull every packet as they enter your network and examine a series of them for strings matching known attack methodologies. When a possible attack is discovered, the administrator is alerted and must take action to prevent further intrusions from occurring. These tools are similar to virus checkers, in that, as new attacks are discovered, the tool must be modified to allow their detection. Sample GOTS products available are Network Security Monitor (NSM) and Network Intrusion Detection System (NIDS). COTS products available are Net Ranger, Intruder Alert, CMDS and others. Tripwire [KiSp94] is a freely available and widely used Unix utility that can alert a system administrator to changes in file systems that may signal an intrusion.

Host-based intrusion detection systems are required to be run on each individual host within your network. These systems detect probes of host ports, password guessing and other known attack methods, which are captured as part of the auditing features (if turned on) of your host. These host-based tools run as a background process on your hosts and provide a warning as they are discovered.

There is some overlap between host-based and network-based string matching capabilities, and running the two in concert will in some cases provide duplicate alerts. These techniques cannot detect every potential attack, and they will provide some false alarms; however, they will significantly reduce your risk of undetected intrusions. As a minimum, a network-based intrusion detection system should be installed at the Internet connection to your intranet and behind a firewall, if installed.

3.4 Security Management Tools

A number of tools exist that can be used by both attackers and system managers to test the security of a system. Listed below are some of the more widely available ones.

SATAN
Security Analysis Tool for Auditing Networks (SATAN) is a network vulnerability toolkit, using a WWW front-end. In its simplest (and default) mode, it gathers as much information about remote hosts and networks as possible by examining such network services as finger, NFS, NIS, ftp and tftp, rexd, and other services. The information gathered includes the presence of various network information services as well as potential security flaws - usually in the form of incorrectly setup or configured network services, well-known bugs in system or network utilities, or poor or ignorant policy decisions. It can then either report on this data or use a simple rule-based system to investigate any potential security problems. Users can then examine, query, and analyse the output with an HTML browser, such as Mosaic, Netscape, or Lynx. While the program is primarily geared towards analysing the security implications of the results, a great deal of general network information can be gained when using the tool -- network topology, network services running, types of hardware and software being used on the network, and more.

However, the real power of SATAN comes into play when used in exploratory mode. Based on the initial data collection and a user configurable rule set, it will examine the avenues of trust and dependency and iterate further data collection runs over secondary hosts. This not only allows the users to analyse their own network or hosts, but also to examine the real implications inherent in network trust and services and help them make reasonably educated decisions about the security level of the systems involved. SATAN should prove to be most useful when used by the system or security administrators who own or are responsible for the security of the systems involved.

ISS
Internet Security Scanner (ISS) can perform a scan on a host or network, to test for a common set of security flaws and errors in configuration. It is designed to carry out a simple vulnerability test on a network of computers in order to highlight systems that are vulnerable and may be used to gain access to more secure and important systems.

ISS highlights the problem with having a mixture of secure and insecure machines on the same network. All hosts should be secured to the same level as the highest machine on the network. If a host is easy to penetrate then it will become the weak link in the security chain. An intruder can use the weak link to exploit the configurations of other hosts. The most sensitive hosts may be configured securely, but if they trust a host that isn't then they too are vulnerable.

tiger
The tiger software suite is produced by the Texas A & M University. It is a collection of Bourne shell scripts, C programs and data files that are used to perform a security audit of UNIX systems. It has pre-defined configuration databases for AIX 3.x as well as many other systems.

tiger has one primary goal: report ways in which root can be compromised. While checks are performed for other purposes, most of the checks are directed at this goal. The primary assumption made is that any UID other than 0 can be obtained and that any GID can be obtained by unauthorised persons.

The checks performed extend to cover other means of gaining root access, such as *cron*, *inetd* and *setuid* executables. Access through these methods is checked to see if any user, other than root, can alter any of the configuration files associated with these utilities. Specific checks are performed to see if anything in the root executable path can be modified

by a normal user. All user accounts will have these checks performed, but special attention is paid to the root account.

crack
crack is a dictionary based password guessing tool. To have a more complete understating of the implications of this kind of attack a brief description of the UNIX password encoding system is required.

The password typed in by the user is encrypted using the DES algorithm. The encryption is then modified by deleting selected sections and it is then encrypted again (for details on this, and other aspects of Unix security, see [GaSp 96]). The result is a 14 character string that cannot be reversed. To check the validity of a password typed by the user the computer simply encrypts the string using the same technique and then compares the two strings. If they are identical then the password is valid.

snoop
A promiscuous-mode IP packet sniffer. The present Sun operating system, Solaris, provides to administrators a command *snoop* for the capture and inspection of ethernet packets. By default it uses both the network interface and the streams buffer modules to capture packets, then displays a single line summary of each packet seen. It can also provide a far more detailed packet description, including information from the ethernet layer upwards to the top layer. This utility can also be used to monitor certain protocols and to sniff for authentication data, for example the rexec service.

3.5 Firewalls

A firewall is a description of a system (one or more pieces of hardware and software) that acts as a barrier between two network segments such as a protected enclave handling sensitive information and the Internet. A firewall can be considered the technical implementation of a security policy. It upholds the security policy of a network when connecting that network to a second network which has a less stringent security policy.

Firewalls come in three types: packet filtering, which is usually implemented with a screening router, circuit gateways, and application gateways, which are usually implemented with a dual-homed host and proxy servers. In the dual-homed implementation, the firewall uses two separate network connections and doesn't allow data to pass directly between the two.

When operating as an application gateway, the firewall will examine specific application protocols to decide whether connections are permissible. The range of supported application protocols varies but most firewalls examine such popular ones as Telnet, HTTP or FTP.

Configured as a proxy server, the firewall interacts over the Internet on behalf of internal users, making it appear that all outgoing traffic emanates from the firewall and shielding internal node addresses from the Internet.

The security provided by firewalls is sometimes likened to the crunchy outer shell of a candy that is soft and chewy on the inside, in that once the firewall is penetrated (or circumvented) the intruder may find the internal systems easy to manipulate. The implication is that firewalls must not be seen as a panacea that reduces the need to administor the internal systems so that they remain as secure as possible.

3.6 Guards

Guards provide a controlled connection between networks of differing sensitivity levels. For example, the US Standard Mail Guard (SMG) enables users on SECRET networks to send and receive unclassified e-mail without attachments. An SMG lies between a high-level network and a low-level network, such as the Internet, and ensures that no message traffic from the high-level network that is labeled SECRET passes through to the low-level network unencrypted. It also ensures that no attachments, which may contain Trojan Horses, can pass from the low-level network to the high-level network. Other forms of guards can automatically downgrade highly formatted data. Because it is so difficult to assure, based only on an examination of its content, that a message to be downgraded in fact conveys no high level information, a guard should best be considered a device to detect user mistakes, rather than as a protection against malicious software.

4. Emerging Technologies for Safer Defence Connections to the Internet

A great deal of commercial and government software and hardware development seeks to capitalize on the capabilities provided by Internet connectivity. Some of this development, such as the efforts devoted to firewalls and secure transaction technologies for electronic commerce, are relevant to safe use of the Internet. This section highlights particular emerging technologies relevant to safe use of the Internet for defence purposes that have been undertaken or facilitated by participants in STP-11. The following section proposes collaborative demonstrations that employ the technology described here.

This section is not intended to be exhaustive. There is other technology not described here that may also be used to reduce the risks of Internet connection. For example, the IPSEC extensions to IPv4 and IPv6, documented in [RFC-1825 - RFC-1827], (currently under revision, see [Bell96] for an account of some problems), are intended to provide end-to-end integrity, confidentiality and authentication services at the IP layer. Public Key Infrastructures (PKIs) under consideration and development in various nations could simplify authentication and signing procedures. CESG and NSA developments such as Kilgetty and FORTEZZA, respectively, may be used to improve security in various ways, and CESG's CASM programme has developed an architecture inteneded to counter a variety of vulnerabilities in e-mail systems, in part through the use of a public key infrastructure being developed under the CLOUD COVER project. In the U.S., the Defense Message System (DMS) project is intended to provide secure e-mail services.

Nevertheless, STP-11 finds the technologies described below to be of particular interest.

4.1 Entrust Overview

Entrust is a family of public-key cryptography software products for encryption and digital signature on computer networks with fully automated key management. Encryption provides for the confidentiality of information such as personnel data, business plans and design files. Digital signature provides strong authentication of the originator and the prompt detection of any data tampering. Examples of applications supported by Entrust include e-mail, e-forms, and file protection.

Entrust provides the network infrastructure required to meet the exploding demand for public-key cryptography technology. As such, Entrust is a unique product in a new

product category known as Public-Key Infrastructures (PKIs). There are a number of products on the market that use public-key cryptography and provide desktop encryption and digital signature capabilities, however they are geared to individuals or small groups. The Entrust/Client application provides similar capabilities for public-key cryptography to these products, but is supported by an automated key management infrastructure which scales to enterprise levels of tens of thousands of users and beyond.

Entrust provides a standards-based public-key cryptography key management solution that ensures key updates are automatic, transparent to users and have no additional cost attached. This ensures that, as network security is deployed on a large scale, network administration costs are minimized, yet critical centralized controls are maintained.

4.2 Starlight Family of Devices

Starlight is a family of limited functionality, high assurance add-on security devices currently under development by the Australian Defence Science and Technology Organisation (DSTO), and are intended to be retrofitted to untrusted COTS workstations so as to provide enhanced levels of security functionality. The initial two devices are the Interactive Link and the Trusted Path (or Export) Module. DSTO has produced prototypes of these devices; production versions of both are intended to be made available as commercial products.

Starlight Interactive Link

The Interactive Link [ANGMYY 96] is a trusted "information pull" device, and is designed to enable a user on a classified machine to interact with a machine on a lower classified network allowing him to bring up a window displaying low-side information on the high side without compromising the confidentiality of classified information on the high side. Typically, the low-side network would be at Sensitive But Unclassified level, and would have access to the Internet via a firewall, thereby supporting access to Internet email and Web browsing.

The key components of the Interactive Link are a trusted keyboard switch, a trusted data diode (a simple device that permits data to flow through it in one direction only) and the use of X-windows protocols. When the user's keyboard and mouse are switched to the low side, they appear to belong to an X-terminal on the low side network. In response to keystrokes/mouse movements, the low-side processor generates X-window display information. This display information is then copied up through the data diode to the high side network, where it is made to appear in a window on the user's classified workstation (note that buffering is required on the "high" side of the diode to ensure that there is a reliable connection to that network).

Interactive Link demonstrators have been distributed and they are currently undergoing user trials. It should perhaps be mentioned that there needs to be only one data diode per connection between networks. The IL hardware is therefore reduced, on a per-user basis, to a simple keyboard switch, plus an associated interface to the low-side network. In addition to the IL hardware, there is also associated application-level software on both the high and low sides in order to support the transfer of X display information between the networks.

It should be emphasised that the secure operation of the Interactive Link does not depend upon the user's workstation or software being trusted in any way; this also includes the

software supporting the operation of the IL itself. Naturally, security is also dependent on the user not typing in any classified information while his keyboard is connected to the low network: however, if desired, all keystrokes sent to the low-side server could be audited.

While the current demonstrators work with Sun and HP Unix workstations and/or X-terminals, access to a Windows NT server has also been demonstrated via third-party software products such as NTRIGUE, WinDD or WinCenter Pro (which convert WinNT display protocols to X-protocol data for remote X displays). Using this approach in addition to the Interactive Link, applications such as Microsoft Word which are running on a low-side NT server can be brought up in a high side window.

Starlight Trusted Path Module

The Trusted Path Module or TPM is a trusted "information push" device, and will allow a document which has been created or edited on the high side to be reviewed and exported to the low side in a highly trusted way (i.e. so that no additional unintended information is released). The TPM is accessible over the network and also attaches to the user's workstation as described in the following scenario. A user on the high network creates or edits a file via his workstation or X-term, and then wishes to export it to the low network. The file is first loaded over the network into the TPM, which then takes over the user's workstation monitor and allows the entire contents of the file to be reviewed in a highly trusted way. (This imposes some limitations on the complexity of documents which can be reviewed, but mostly there are possible work-arounds such as exporting sanitised text sections of the final document and then putting the final complex document together on the low-side system, via the Interactive Link. When the user is satisfied that the file to be exported contains no sensitive information, he indicates this via an accept button on the TPM, and a cryptographic seal is applied (via a device such as a Fortezza card). Provided that this seal remains intact (i.e. the file has not been modified in any way since leaving the TPM) it can be exported to the low side via a trusted gateway (or, alternatively, it could be exported directly from the TPM to the low-side network). Either way, all exports to the low side should be audited.

Starlight Summary

Through high integrity limited functionality devices such as the Interactive Link and Trusted Path Module, a form of multilevel secure information system can be assembled using several COTS computing networks each operating at a single, system high, security level, with the transfer of information between the various levels being under the control of Starlight devices.

4.3 Purple Penelope

Purple Penelope (PP) extends the functionality of Windows NT to support domain based working. The domain based approach considers security from the viewpoint of business requirements, not implementation mechanisms. This approach is intended to lead to security measures that are well suited to an organisation's working practices and modern computer systems. Specific features, described below, include:

a. Discretionary labelling;
b. Private and shared filestores;
c. Standard applications operate unaffected;
d. Applications may be customised to reflect the domain function;

e. Export sanction by user;
f. Role based access controls;
g. Accounting.

Discretionary labelling
Discretionary labelling offers user friendly, intuitive labelling which encourages users to correctly label information through its entire life-cycle. Such labelling helps to prevent the accidental release of inappropriate material.

The most noticeable feature of PP to the user is the presence of a screen stripe which conveys labelling information about the current application and the contents of the clipboard. The file manager has been extended to allow the labels of individual files or groups of files to be viewed or changed via the screen stripe.

Private and shared filestores
Individual users have their own private file storage area. The user has complete discretion over the labels given to information held in such private stores. The labels may be changed either via the application software that created the file or via the filemanager and screen stripe.

Shared storage areas are available on a read only basis to the entire domain. Access to the shared area is mediated on the basis of the users clearance.

Standard applications
Standard software applications operate in their usual manner on PP. By default, each invocation carries a label which is maintained by user intervention, either directly via the screen stripe or indirectly via cut and paste operations.

Customised applications
Many modern software applications offer the ability to customise their look and feel via macros. In PP, customisations have been performed on MS Word and Access to demonstrate close integration with PP. Macro additions to MS Word enable the introduction of markings at any point in a document which are subsequently maintained according to the label given to the document. Such customisations enable the business function of the domain to be supported via the introduction of templates reflecting business requirements. Thus, document templates with markings in the headers and footers may be constructed or, indeed, any template which carries the markings in appropriate places.

It is intended that other software products should be customised to allow them to be used in a PP friendly manner. A clearly defined API for PP has been published and a number of vendors are currently assessing PP to determine the appropriateness of augmenting their products in such a manner.

A further example of augmenting COTS products has been demonstrated in PP by the use of a labelled DBMS product for discretionary labelling. All labelled DBMS products were designed for use within a mandatory labelling environment with the labelling being applied at a relatively coarse granularity. To illustrate the flexibility of the approach, the chosen product, Trusted Oracle 7, was used to support a requirement for both discretionary and fine grained labelling. The underlying approach to achieving database sharing within a domain is, however, the subject of ongoing research.

Export Sanction
Whenever a user requires that information is copied or moved from the private file store, an export sanction must be given by that user. Such explicit sanctions by the user must occur for the movement of files to the shared filestore and to floppy discs. Similarly, the release of e-mail messages is subject to such sanction. The implementation of this sanction prevents Trojan Horse software from compromising the confidentiality of information.

Role based access controls
The shared filestore of PP is to be augmented with role based access controls within the next few weeks which should enable the business function of the domain to be more accurately reflected.

Accounting
Whenever a user accesses the shared filestore, or sends e-mail messages they should be held accountable for their actions. Although not currently performed in PP, work is underway to do this, by recording all export actions in an audit trail. Initial work has concentrated on e-mail, with a record being maintained of all outgoing e-mail traffic. This remains to be integrated into PP.

Future Developments
A further mechanism, called sub sessions, is planned for integration into PP. This mechanism will enable users to connect securely to applications running below their clearance level. An example where such a mechanism will be used is to allow connection to legacy systems.

In parallel with the main development of PP, the domain based approach underpinning PP is being researched to ensure that a PP product can be integrated into a complete system solution.

Exploitation of PP is ongoing. A non-exclusive license agreement has been signed by a small US company and talks are on-going with a larger US company. It is anticipated, however, that a PP like product will be brought to market by the end of 1997. In support of this activity, a number of beta versions of PP have been issued to UK MoD projects for assessment of the approach.

Much of the development work on PP has, so far, concentrated on the issues of working within a single domain. Future work will focus upon support to inter domain working. Thus, e-mail (based on X400 rather than the initial PP demonstration of SMTP); firewalls supporting shared databases and filestores; and mediated access to Web pages are being addressed.

Although the current implementation of PP is based on NT 3.51, work is underway to move to NT 4. This process should be completed by September 1997.

4.4 SINTRA/Pump Technology

If two systems or networks operate at different security levels it may be permissible for data to be sent from the lower level system to the high level system, but high level data must be prevented from flowing downwards. However, imposing a strict one-way flow between systems is likely to cause operational problems, because most communication protocols and distributed application systems require a two-way flow of information to achieve reliable communication.

The NRL Pump [KML 95], is a device that permits upward flow and restricts downward flow to acknowledgments only. Further, the timing of the acknowledgments is obscured in order to limit its exploitation as a covert timing channel. In the SINTRA project [FGKLMMP 95], prototype Pump implementations based on both the Honeywell XTS-300 and Windows NT 4.0 have been demonstrated in conjunction with commercial database management systems and replication servers to demonstrate how a Low database can be replicated at a higher security level and consistency between the two databases can be maintained by propagating updates of the low system to the high system via the Pump.

At present, a hardware prototype of the NRL Network Pump is under development; this effort is expected to yield a detailed specification that will permit operational Pumps to be ordered from commercial system builders. This technology might be applied to provide safe, reliable upward communication between systems connected to the Internet and systems connected to higher level networks.

4.5 Onion Routing

Onion Routing [SGR 97] provides a mechanism for hiding which hosts are communicating with each other across the Internet. The mechanism works as follows: an application, instead of making a (socket) connection directly to a destination machine, makes a connection to an Onion Routing Proxy on some remote machine (for example, a local firewall might run an Onion Routing Proxy) . That Onion Routing Proxy builds a route through several other Onion Routers to the destination. Each Onion Router can only identify adjacent Onion Routers along the route. When the connection is broken, all information about the connection is cleared at each Onion Router. Data passed along the anonymous connection appears different at each Onion Router, so data cannot be tracked en route, and compromised Onion Routers cannot cooperate. Onion Routing lives just beneath the application layer and is designed to interface with a wide variety of unmodified Internet services by means of proxies. Onion Routing has been implemented on Sun Solaris 2.4; proxies for WWW (HTTP) and TELNET have been implemented as well. Proxies for e-mail (SMTP) and FTP are forthcoming.

Onion Routing differs from other anonymity services in two ways: communication is real-time and bi-directional; and the anonymous connections are application independent. Onion Routing does not provide anonymity at the application layer (though measures can be added for this purpose if required). Applications may (and usually should) identify their users over the anonymous connection. However, the use of a packet switched public network should not automatically reveal who is talking to whom. This is the traffic analysis that Onion Routing complicates.

At present, NRL is running an Onion Routing prototype that is available for public use. To try it out, a user need only configure a browser to use as its HTTP proxy port 9000 on host *onion-router.nrl.navy.mil*. Binary and source code for the system (which is in a beta test mode at this writing) are being made available to parties interesting in running onion routers.

4.6 Coordinated Intrusion Detection

Current state-of-the-art intrusion detection devices/systems are brittle and can only detect most known types of attacks. Furthermore, these detection systems are scattered throughout the Defense Infrastructure, making correlation of the variety of sensors all but impossible. In addition, a mechanism for the early warning of possible attacks to other

country responsible agencies is lacking. Rome Laboratory is in the process of building an Information Protection Integration Infrastructure, which will use a COTS Real Time Process control environment to integrate and correlate Information Protection sensors within an enclave, such as a base network, and subsequently across networks. This environment will provide dynamic, near real time, interactive visualization of the security state of network circuits, components and nodes using existing GOTS/COTS products such as NSM/NIDS, Network Management Tools, and others. Tools for analysing audit and measurement data that were not designed for detection can nevertheless be used for this purpose if their results are appropriately displayed. These include logging tools, routers, firewalls, and audit reduction tools.

5. Useful Demonstrations of Existing and Emerging Capabilities

In many cases the current and emerging technologies hold the potential for providing access to desired Internet capabilities with substantially reduced risk. In this section we suggest some potential demonstrations that could form the basis of TTCP collaborations. Some of these could also be of substantial benefit for both improving and securing TTCP cooperation over the Internet.

5.1 Internet-based TTCP Roster and Calendar Services

Coordination of TTCP panels and groups might be significantly simplified and improved if newsletters, roster information and meeting dates and locations were readily available to qualified individuals on the Internet. The technology described in this report can be used to demonstrate such a capability. The key issues are to authenticate that requests to retrieve or update roster or calendar information come from valid TTCP participants and that the information returned via the Internet is not subject to eavesdropping *en route*.

Alternative demonstrations might be:

a. Organize a web server capable of two-way authentication via SSL. Provide TTCP members with certificates so that mutual authentication is possible. Configure a web server to accept update requests for panel rosters and calendar entries from designated Chairs and support personnel.
b. Use national Public Key Infrastructure as certificate source for authentication purposes.
c. Use Tradewaves (COTS product) to make web clients and server "Entrust aware." Issue TTCP members copies of Entrust to use in conjunction with this software. Demonstrate authentication of users, mediation of user requests by Tradewaves, encryption of responses sent over Internet.

5.2 Internet Access from within a Secure Enclave Without Downgrading

Authorized connections from workstations within secure enclaves to the Internet today typically require either an expensive multilevel secure workstation or a guard system that is relied on to permit only appropriate traffic to be sent from the enclave down to the Internet. Starlight Interactive Link permits the keyboard and mouse of an untrusted workstation attached to a High network within an enclave to be switched safely to a server on a Low network outside the enclave. Data returned from the Low side is transmitted to the workstation's display via a one-way communication link. Because only the user's keystrokes and mouseclicks pass from the enclave to the low network (the Internet, in this demonstration) there is no need to downgrade them as they cannot have been polluted with information from the untrusted high processor. This architecture holds the promise of both higher security and lower cost than current implementations.

5.3 Anonymous Internet Browsing from within an Enclave

The Onion Routing scheme permits anonymous Internet browsing by routing mulitply-encrypted "onions" through the Internet. Each onion router performs decryption and mixing operations that make it difficult for an outside observer to correlate its input and output messages, and all traffic over an onion-connection is encrypted between the first and last onion-router in a path. No single onion-router (except the first in a path) can determine more than its predecessor and successor router in a path. This demonstration would combine the use of Starlight IL to reach the Internet from an enclave with the use of Onion Routing to prevent the resulting Internet traffic from being associated with the High enclave from which it originated.

5.4 Updating a High Database from Low Sources Without Downgrading

A common problem in military systems is that information originating from Low sources and input to a Low database is difficult to incorporate in systems running at High. Often, updates performed at Low are batched and manually inserted in the High system. This demonstration would combine Starlight Interactive Link and SINTRA/Pump technology to show how a user in a High enclave can safely access an open source on the Internet and cause an update to a Low database. The SINTRA/Pump technology will cause the Low update to be propagated to High quickly and reliably without causing a downward flow of information. This represents a more secure and lower cost alternative to using either a multilevel database or using a guard to downgrade queries for Low information that originate within a High enclave.

5.5 Safe Internet Access from a Discretionary Labelled Environment

Purple Penelope demonstrates how discretionary labels can be added to a widely used COTS operating system (Windows NT) at relatively low cost. This effort would demonstrate the use of Starlight Interactive Link from a Purple Penelope system to gain access to the Internet for purposes of importing open source materials into the labelled environment.

5.6 Trusted Release Demonstration

A benefit of the discretionary labelling provided by Purple Penelope is better guidance to users who wish to sanitize a document for use at a lower security level. The Starlight Trusted Path Module is designed to permit secure review of documents intended for release by assuring that untrusted software is not involved in displaying the document to the reviewing user. This demonstration would combine the Purple Penelope and Starlight TPM to provide strong support for safe release of appropriate information to lower security levels.

5.7 Information Protection Integration Infrastructure

The Information Protection Integration Infrastructure described in Section 4 could demonstrate more effective detection of Internet-based attacks. The planned environment can be replicated at several laboratories to demonstrate and test its ability to detect correlated but distributed attacks, to view the state of different enclaves, and to warn other enclaves when an attack is detected. The correlation can determine whether the data support the "intrusion or misuse" theory; whether separate events are related to the same event, attempt or group, and whether distant events are related. It should also reduce the data required to perform intrusion detection.

6. Limitations

The products described in Sections 3 and 4 mitigate security problems but do not eliminate them. All of these products have limitations. These limitations are fairly well-understood for Defence-developed "high assurance" products such as Starlight, Pump, and products coming out of the MISSI development program. However, even Defence-developed products are not perfectly understood because of the recent move from "risk avoidance" to "risk management" strategies and because of new forms of attacks. There are also problems with implementations of theoretically sound security mechanisms For example, many standard implementations of secure protocols include hidden channels.

For COTS products the situation is even worse. These products are developed for commercial, rather than for Defence, grade security. There is also less understanding of what these products actually do, and any understanding of current security relevant features is soon made irrelevant by new versions. Lack of standardization (e.g. in Windows NT) exacerbates this problem. For both COTS and high-assurance products, things get worse when components are integrated into a functional system.

It should also be noted that export restrictions on cryptographic technology limit the availability of COTS cryptographic solutions in TTCP member countries. Such restrictions have also hindered international collaboration and interoperability and may prove to be an issue for IPv6.

Finally, it should be noted that none of these products addresses concerns outside of its domain. They do not replace the need for good security management and user security training. System and network configurations that do not address known attacks will not be solved by secure components, and systems will be subject to Trojan Horse attacks as long as users are permitted to introduce software from untrusted sources. These products also fail to protect systems from malicious insiders or mistakes. (The relevance of this to the Internet is that Internet connections provide a new way of transmitting sensitive information outside of a system.) Finally, it should be noted that these products require adequate care in initially establishing trust for use in later authentication procedures and certificate/key management.

7. Conclusions and Recommendations

The choice for the Defence establishments of the TTCP member nations is not whether to use the Internet or not, but how to use the Internet safely and prudently.

Prudent use of the Internet for Defence purposes must assume, at present, that the Internet is a broadcast medium that is open to corruption and service denial. Traffic is world readable and world writable. The lack of authentication mechanisms in current Internet Protocols permits spoofing and connection highjacking. Relative to the public switched telephone networks, the Internet is unreliable. Although the Internet may prove a useful communication medium in emergency situations, it is risky to rely on its availability in times of heightened world tension.

A great deal of routine Defence business is now conducted using the Internet, and appropriately so. A significant amount of this traffic, although unclassified, is sensitive in

one way or another, however, and although the Internet can be safely used for this traffic, certain precautions should be taken:

- For traffic that is sensitive, even if unclassified, it is prudent to apply encryption within the management domain of the end system, for example, by using IPSEC mechanisms at the network layer or Entrust at the application layer.
- For traffic whose authenticity or integrity is significant, it is prudent to apply digital signatures externally.
- For traffic whose timely delivery is critical, positive acknowledgment of receipt should be obtained, so that nondelivery is evident.
- Cryptographic measures such as digital signatures and encryption provide assurance only to the extent that the keys they employ are correctly generated, distributed, and protected, so externally applied cryptography must be accompanied by adequate key generation and distribution. The infrastructure to support these functions is nascent.

These precautions will help ensure confidential, reliable, timely e-mail delivery. They do not, however, provide protection from site intrusion. Such protection requires administrators of unclassified Defence systems connected to the Internet to follow some additional recommendations:

- Carefully consider the configuration of your internal networks and their vulnerability to attacks from the Internet. In most cases, it will be prudent to install a properly configured firewall between the internal networks and the Internet. In general, the administrator should be required to justify why a firewall is *not* needed, rather than the reverse.
- Maintain awareness of, and act upon, advice from CERT, AUSCERT, FIRST, and other relevant sources regarding vulnerabilities and attacks, and be sure available security patches are installed as they become available.
- Introduce measures that, so far as possible, prevent users from sending reusable passwords in the clear over Internet links and prevent internal systems from accepting cleartext reusable passwords from Internet links.
- Ensure that users receive appropriate security training, especially concerning the hazards of importing software (including plug-ins, Active-X controls, and Java applets, as well as more conventional forms) from Internet sources.

Administrators of publicly accessible Defence web servers, should in general

- Position the web server on the Internet side of the firewall on a separate processor dedicated to this purpose, and configured with the minimum software required to provide the type of web support desired.
- If the server is to accept input via CGI scripts, be aware of the vulnerabilities these mechanisms can contain and review new scripts for flaws before installing them

Administrators of classified networks should not, of course, connect their systems to the Internet without proper certification and accreditation. As discussed in sections 3-5, there is current and emerging technology that will soon make it both safer and cheaper to make such connections.

The technologies now being developed will improve the security infrastructure of the Internet, but they will not soon obviate the measures we recommend in this section. The open nature of the Internet means that its hardware and software will, for the foreseeable future, be open to compromise, so external mechanisms will need to be applied if Defence interests are to take advantage of its benefits without being subject to its weaknesses.

Finally, we make the following general observations about where Defence R&D investments are (and are not) required in relation to Internet security:

Areas in which Defence must develop its own products:

- High assurance network components
- High assurance cryptography
- High assurance key management

Areas in which Defence needs to influence commercial development:

- Internet protocols
- Application Program Interfaces
- Commercial cryptographic protocols
- Public Key Infrastructure

Areas in which Defence can rely on commercial development:

- Low assurance Internet component infrastructure
- Routers
- Protocol stacks
- Applications

Acknowledgments

The authors are grateful for the hospitality and support extended by DERA Malvern in the creation and production of this report. They particularly wish to thank Mr. Simon Harding in this respect. Review comments by Mr. Michael Reed of NRL improved the report.

8. References

[ANGMYY 96] Anderson, M., C. North, R. Griffin, R. Milner, J. Yesberg, K. Yiu, "Starlight: Interactive Link." *Proc. 13th Annual Computer Security Applications Conference*, San Diego, Dec., 1996, IEEE CS Press.

[Bell96] Bellovin, S.M. Problem Areas for the IP Security Protocols. Proc. of the Sixth Usenix UNIX Security Symposium, July, 1996. Available at URL: *ftp://ftp.research.att.com/dist/smb/badesp.ps*

[BrMa96] Bradner, Scott, and Allison Mankin, Eds. *IPng, Internet Protocol Next Generation.* Addison-Wesley, Massachusetts, 1996.

[ChBE94] Cheswick, William R. and Steven M. Bellovin. *Firewalls and Internet Security: Repelling the Wily Hacker*. Addison-Wesley, Massachusetts, 1994.

[Come91] Comer, Douglas E. *Internetworking with TCP/IP, Vol. I*. Prentice-Hall, New Jersey, 1991.

[East97] Eastlake, D., 3rd. Secure Domain Name System Dynamic Update. Internet-Draft <draft-ietf-dnssec-update-04.txt>, 4 Feb 1997.

[FGKLMMP 95] Froscher, J.N., D.M. Goldschlag, M.H. Kang, C.E. Landwehr, A.P. Moore, I. S. Moskowitz, C. N. Payne. "Improving Inter-Enclave Information Flow for a Secure Strike Planning Application," *Proc. 12th Annual Computer Security Applications Conference*, IEEE CS Press New Orleans, Dec., 1995.

[FNC] Federal Networking Council Resolution: Definition of "Internet" Oct. 24, 1995. Available at URL: *http://www.fnc.gov/Internet_res.html*

[GaSp 96] Garfinkel, S. and E. Spafford. *Practical UNIX and Internet Security, Second Edition.* O'Reilly & Associates, Inc., 1996, ISBN: 1-56592-148-8.

[KML 95] Kang, M.H., I.S. Moskowitz, and D. Lee, "A Network Pump," *IEEE Transaction on Software Engineering,Vol. 22*, No. 5, May., 1996. pp.119-129.

[KiSp 94] Kim, G.H. and E.H. Spafford, " Writing, Supporting, and Evaluating Tripwire: A publically Avaliable Security Tool," Purdue Technical Report CSD-TR-94-019, Computer Science Department, Purdue University, 12 March 1994. This report, and much relevant information is available at the COAST computer security archive, URL: *http://www.cs.purdue.edu/coast/*

[ICE97] URL: *http://www.tis.com/docsw/research/crypto/ice/index.html#WEBLINKS* This web page includes links to a wide variety of information sources on standards for cryptographic applications programming interfaces and standard interfaces for security services.

[Perl92] Perlman, Radia. *Interconnections.*Addison-Wesley,Massachusetts, 1992.

[RFC-1034] Mockapetris, P. Domain Names - Concepts and Facilities. RFC 1034, Nov. 1987.

[RFC-1035] Mockapetris, P. Domain Names - Implementation and Specification. RFC 1035, Nov. 1987.

[RFC-1825] Atkinson, R. Security Architecture for the Internet Protocol. RFC 1825, August, 1995.

[RFC-1826] Atkinson, R. IP Authentication Header. RFC-1826, August, 1995.

[RFC-1827] Atkinson, R. IP Encapsulating Security Payload. RFC-1827 August, 1995.

[RFC-2065] Eastlake, D., 3rd, C. Kaufman. Domain Name System Security Extensions. RFC 2065, Jan., 1997

[SGR 97] Syverson, P. F., D.M. Goldschlag, and M.G. Reed, "Anonymous Connections and Onion Routing," *Proc. IEEE Symp. on Security and Privacy*, Oakland, May, 1997. See also URL: *http://www.itd.nrl.navy.mil/ITD/5540/projects/onion-routing/*

[VTRB96] Vixie, P. (Ed.), S. Thomson, Y. Rekhter, J. Bound. Dynamic Updates in the Domain Name System. Internet Draft, <draft-ietf-dnsind-dynDNS-11.txt>