Security Status and Issues

&

Electronic Commerce on the Internet

Denis Arnaud

July 1995

Contents

List of Figures

Chapter 1

Introduction

1.1 The Internet

The Internet has exploded in the last two years thanks to the invention of the so-called "browser." A browser is a point-and-click software program that allows "surfers" to navigate around the Internet without knowing any Unix commands. The first browser was developed by the National Center for Supercomputing Applications (NCSA), a government agency. It is called Mosaic and can be downloaded for free off the Internet. Since then other companies have popped up to develop commercial browsers; Netscape Communications Corporation has snatched up the most market share so far (nearly 65%). With a browser and access to the Internet, you can order a pizza, listen to and purchase a CD, stroll through the Louvre, or view satellite photographs of Scotland. Search tools allow you to find all the florists on the Internet, and hypertext links enable hopping from one location to another. One second you are in Omaha, the next in Tokyo.

Browsers have introduced graphics to the Internet interface. Anyone who wants to share information about themselves over the Internet can create a personal "home page" with pictures of themselves and their pets, for instance, or biographical information, or a list of their favorite Internet sites. The World Wide Web is a body of documents connected through hypertext links, so that home pages can have references to other home pages, and a simple click of the mouse on a highlighted word or name takes you to another document dealing with that subject. Companies can use home pages as store fronts, displaying company information, descriptions of upcoming products, or catalogs with pictures of actual products, complete with order forms.

Although it may get congested from time to time, the Internet itself is extremely reliable. It was developed by the Department of Defense in the late 60s to establish a communication vehicle that would survive nuclear war. Internet transfer protocols route information around breakdowns in the network. There isn't actually any one network that is the Internet; it's made up of thousands of networks that connect to each other through common routers, and they all agree to carry each others' traffic. A number of different access providers and phone companies have set up backbone service to carry Internet traffic between the major routers. There is a lot of money flowing up from local access providers to these national players, guaranteeing that the infrastructure will continue to expand to meet demand.

Because so many resources are shared, the Internet is also very efficient. It costs a lot less to connect a business to the Internet than to lease telephone lines that customers dial into with their modems. Long-distance transactions take place with no long-distance charges. It will probably never be cost-effective to measure an Internet user's activity based on volume and distance of transactions, so most likely Internet users will continue to be charged for the size of the "pipe" connecting them to the Internet.

Businesses of all types are setting up shop on the Internet to take advantage of this new distribution channel and customer base. Whereas the Internet had been reserved for the research and education communities until recently, the arrival of commercially operated backbone networks has opened the way for commercial use. The number of commercial entities with an Internet presence has doubled in the last three months. Currently, twenty-five million Americans have Internet access. Two million of these have browsers; Microsoft's Windows 95 will include a built-in browser. Prodigy and CompuServe have recently added browsers to their list of services, and America Online will soon follow. Internet usage is growing at over fifteen percent a month.

The amount of information that can be exchanged and the speed with which it appears on the computer screen depend on the speed of the connection to the Internet, whether that is a modem or an ISDN line. Computer stores are selling 28.8K modems for $225 today, which are 15-20 times faster (with compression) than the standard 2400-baud modem on the market five years ago. Modems will keep getting faster, allowing more information, better graphics, and full-motion video to be downloaded more easily. However, in five years most households will probably buy their Internet access from their cable company, who will provide them with a 10 megabit-per-second connection through their cable wire. A 10-MB connection would download in one second a file that takes a 28.8K modem five and a half minutes to download.

1.2 Why is Security on the Internet an Issue?

In the "real" world, where information is exchanged in physical form, there is a great deal of security infrastructure, built up over literally thousands of years, that we take for granted.

The postal service is a good example. For a nominal fee, you are able to send a message anywhere in the world. If you want that message to be private, you enclose it within an envelope. Tampering or changing the message would require breaking the seal on the envelope and committing Federal crime. If you wish to make sure that your message was received, you ask for a return receipt. To make sure that the person who received the message was the one you intended him to be, you might check his or her signature. He or she might check yours.

Now imagine that someone has just constructed a brand-new postal service from scratch. The service operates at blinding speed, sorting and forwarding messages through various postal clearinghouses in seconds instead of days. But the new service has some drawbacks. Anyone can be a postmaster, and there is nothing to prevent a postmaster from making copies of the mail as it passes through his clearinghouse. Messages are written on postcards, and instead of signatures and return receipts, all message identification is written in identical block letters: "This message is from John at ZDS." So the postcard-sorters can't help but see entire messages; there is little to prevent a phony message from being swapped out for a genuine one; there is no guarantee that the persons on either end of the message exchange are the people on the address; and there is no means of verifying that a message was received.

Such a postal service exists today, and it is known as the Internet. By design, the Internet allows wiretapping, and it is estimated that 20of the message traffic sent via the Internet is copied and stored somewhere (by someone other than sender and receiver) for later reference. Most messages are sent as plain block text... an intercepted message can be read by different software platforms.

Businesses are now viewing the Internet as another possible distribution channel that can inexpensively reach a vast, international marketplace. Products can be attractively presented with rich graphics and detailed descriptions, and customers can have a convenient, fast communications channel in which to obtain a quotation and place an order.

Security, however, continues to be a point of concern and a barrier for commercial transactions over the Internet. The building block architecture of the Internet dictates that, for the most part, communications must pass though many computer systems to link the parties in a transaction. Each interconnecting computer system represents a potential security vulnerability. A security breach at any point potentially has a wide reaching effect on a large number of users and institutions. In addition, telecommunications facilities that transmit price quotations, product plans, or other guarded information are subject to penetration from unauthorized parties.

Several break-ins on government computers and those managed by Internet service providers have been well publicized. Most of these break-ins resulted from attackers exploiting vulnerabilities in host and network security that allow capture of a user's password. Such incidents raise concern among those using the Internet for non-commercial purposes, and they discourage businesses considering use of the Internet as a commercial channel. Unless security problems are fixed, companies whose survival may depend on proprietary trade secrets and closely guarded cost and price information, and financial institutions required to provide a payments infrastructure will not participate. Without the support of financial institutions, it will not be possible to gather the critical density of businesses necessary for a healthy marketplace to develop.

Chapter 2

What is Security Needed for?

2.1 Introduction

The security services that will enable widespread use and acceptance of commerce over the Internet include the following:

Authentication: whereby an individual, an organization, or a computer can prove its identity.

Authorization: the ability of a system, once identity is verified, to control access to specific resources.

Confidentiality: the ability to maintain the secrecy of the contents of a transmission between authorized parties.

Integrity: the capability of ensuring that a transmission arrives at its destination in exactly the same form as it was sent.

Non-repudiation of origin: the ability to ensure that when an entity sends an authenticated electronic communication, it cannot later deny the origin, or contents of the communication.

In applications of electronic commerce, encryption is typically used to provide security. However, the industry seems to recognize that cryptography, while a critical component of secure networking, is not enough to protect a business from fraud. Cryptography provides for transaction security, but does not do much to prevent unauthorized access to information and accounts. A good analogy would be using an armored van to transport cash from one branch to another, and then leaving it in the middle of the lobby once it reaches its destination. Information needs to be "put in the vault" once it reaches its destination, and another component is needed to do this: a firewall to secure this local network from Internet access.

Security methods, such as Internet firewalls, are very popular now, but many organizations may believe, or be led to believe, that an Internet firewall alone is sufficient for securing their network. It is like getting the most secure front door money can buy for your house but leaving the garage door unlocked, or the same, weak sliding door entrance from your back deck. Only if everyone plays by the same rules is it effective. Thus, we may expect to see smart cards and their related technology flourishing, as a mean to provide businesses with reliable access control and authentication processes: smart cards may be used to carry and to prove the user's identity, as well as other relevant personal information, and all this in a secure manner.

However, firewalls and smart cards are not enough. A risk and business analysis is almost always required, leading into the development of a security policy and the prescription of security mechanisms and methods for implementation within the local computing environment. Besides, doing this once is not enough. Threats change, vulnerabilities change, business requirements change, and the available counter-measures change: all of these must be periodically and routinely reevaluated.

2.2 Local Computing Environment

2.2.1 Introduction

Adequately secure systems deter, prevent, or detect unauthorized disclosure, modification, or use of information. Company information requires protection from intruders, as well as from employees with authorized computer access privileges who attempt to perform unauthorized actions. Protection is achieved not only by technical, physical and personnel safeguards, but also by clearly articulating and implementing company policy regarding authorized system use to information users and processing personnel at all levels.

2.2.2 The risks

The spread of computing power to individual employees via personal computers, local-area networks (LAN), and distributed processing has drastically changed the way we manage and control information resources. Internal controls and control points that were present in the past when we were dealing with manual or batch processes have not been established in many of today's automated systems. Reliance upon inadequately controlled computer systems can have serious consequences, including:

- Inability or impairment of the company's ability to perform its mission
- Inability to provide needed services to the public
- Waste, loss, misuse, or misappropriation of funds
- Loss of credibility or embarrassment to the others

To avoid these consequences, a broad set of information security issues must be effectively and comprehensively addressed.

2.2.3 The goals

Achieve cost-effective security

The dollars spent for security measures to control or contain losses should never be more than the projected dollar loss if something adverse happened to the information resource. Cost-effective security results when reduction in risk through implementation of safeguards is balanced with costs. The greater the value of information processed, or the more severe the consequences if something happens to it, the greater the need for control measures to protect it. The person who can best determine the value or importance of data is the functional manager who is responsible for the data. For example, the manager responsible for the company's budget program is the one who should establish requirements for the protection of the automated data which supports the program. This manager knows better than anyone else in the organization what the impact will be if the data is inaccurate or unavailable. Additionally, this manager usually is the supervisor of most of the users of the data.

It is important that these trade-offs of cost versus risk reduction be explicitly considered, and that management understand the degree of risk remaining after selected controls are implemented.

Assure operational continuity

With ever-increasing demands for timely information and greater volumes of information being processed, the threat of information system disruption is a very serious one. In some cases, interruptions of only a few hours are unacceptable. The impact due to inability to process data should be assessed, and actions should be taken to assure availability of those systems considered essential to agency operation. Functional management must identify critical computer applications and develop contingency plans so that the probability of loss of data processing and telecommunications support is minimized.

Maintain integrity

Integrity of information means you can trust the data and the processes that manipulate it. Not only does this mean that errors and omissions are minimized, but also that the information system is protected from deliberate actions to wrongfully change the data. Information can be said to have integrity when it corresponds to the expectations and assumptions of the users.

Assure confidentiality

Confidentiality of sensitive data is often, but not always, a requirement of company systems. Privacy requirements for personal information is dictated by statute, while confidentiality of other company information is determined by the nature of that information, e.g., information submitted by the company to its resellers. The impact of wrongful disclosure must be considered in understanding confidentiality requirements.

2.2.4 Control decisions

The official responsible for the company function served by the automated information system has a critical role in making decisions regarding security and control. In the past, risk was often unconsciously accepted when such individuals assumed the computer facility operators were taking care of security. In fact, there are decisions to be made and security elements to be provided that cannot be delegated to the operator of the system. In many cases, the user or manager develops the application and operates solely.

The cost of control must be balanced with system efficiency and usability issues. Risk must be evaluated and cost-effective controls selected to provide a prudent level of control while maximizing productivity. Controls are often closely connected with the system function, and cannot be effectively designed without significant understanding of the process being automated.

2.3 What Does the Use of the Web Imply About Security?

As is often the case with systems that provide a high degree of data sharing, there are a number of security issues for sites that want to take advantage of the web. They can be broken down into three major categories that will be exposed below.

2.3.1 Client threats

The major security issue with running client software comes about given the nature of the web: client programs interpret data that is downloaded from arbitrary servers on the Internet. If there are no checks on the contents of this imported data, there exists the potential for this data to subvert programs running on the client systems. These "Trojan horses" may take several forms, from malicious URLs to rogue code that is run through interpreters (such as PostScript) on the client system.

Another significant security issue applies mostly to sites that are protected by Internet firewalls. These sites have implemented security policies in their firewalls that describe the network services they want to permit within their organizations. A feature of URLs is that they support several resource types: file, Ftp, Http, Gopher, Wais, Nntp, Mailto, Prospero, Telnet, and Rlogin (other resources are in the works). It is likely that the firewall implementation permits only a subset of these resources. Therefore, since web servers provide these services independently of their normal mechanisms (independent of the normal Ftp channels, for example), it is possible to bypass a firewall security policy by using web services.

2.3.2 Server threats

Threats to servers consist broadly of unauthorized modification of server data, authentication and access control to server data, and compromising a server system by exploiting bugs in the web server software.

The HTTP protocol provides several methods (for example, PUT) for writing data on a server. It is possible that some of these methods may be used for unauthorized modification of server data. Work is ongoing to address security concerns in HTTP servers.

2.3.3 Point to point security issues

There are a number of security issues that arise between clients and servers. These security requirements depend on each other in many cases:

Authentication: For some potential uses of the web (e.g., electronic commerce), it is important that clients authenticate themselves to servers, that servers authenticate to clients, and that both authenticate to each other. It is important that a form of authentication be used that is not easily spoofed, such as cryptographically signed certificates.

Confidentiality: Confidentiality is important for web uses that involve sensitive data. Data exported by web clients or servers may need to be protected from eavesdropping. For example, services are available (on-line florist) that require the passing of credit card information. This requirement will be amplified when other kinds of data, such as employee records, government files, etc., begin traversing the web.

Integrity: Certain transactions have the requirement that transaction requests and/or contents remain unmodified during delivery. The electronic commerce example is a case where data integrity is critical to proper functioning of the system.

Availibility: This means all the systems are up and running when they are needed.

Chapter 3

How Can Security Issue Be Addressed?

3.1 Local Network: Security Policy

3.1.1 Principles

There are some common security attributes that should be present in any system that processes valuable personal or sensitive information. System designs should include mechanisms to enforce the following security attributes.

Identification and authentication of users

Each user of a computer system should have a unique identification on the system, such as an account number or other user identification code. There must also be a means of verifying that the individual claiming that identity (e.g., by typing in that identifying code at a terminal) is really the authorized individual and not an impostor. The most common means of authentication is by a secret password, known only to the authorized user. But one of the most reliable means of authentication is by a smart card, owned only by the authorized user.

Authorization capability enforcing the principle of least possible privilege

Beyond ensuring that only authorized individuals can access the system, it is also necessary to limit the users access to information and transaction capabilities. Each person should be limited to only the information and transaction authority that is required by their job responsibilities. This concept, known as the principle of least possible privilege, is a long-standing control practice. There should be a way to easily assign each user just the specific access authorities needed, and I believe that smart card technology could do this job well.

Individual accountability

From both a control and legal point of view, it is necessary to maintain records of the activities performed by each computer user. The requirements for automated audit trails should be developed when a system is designed. The information to be recorded depends on what is significant about each particular system. To be able to hold individuals accountable for their actions, there must be a positive means of uniquely identifying each computer user and a routinely maintained record of each user's activities.

Audit mechanisms

Audit mechanisms detect unusual events and bring them to the attention of management. This commonly occurs by violation reporting or by an immediate warning to the computer system operator. The type of alarm generated depends on the seriousness of the event.

A common technique to detect access attempts by unauthorized individuals is to count attempts. The security monitoring functions of the system can automatically keep track of unsuccessful attempts to gain access and generate an alarm if the attempts reach an unacceptable number.

Performance assurance

A basic design consideration for any information system should be the ability to verify that the system is functioning as intended. Systems that are developed without such design considerations are often very difficult to independently audit or review, leading to the possibility of unintended results or inaccurate processing.

Recoverability

Because companies can potentially be heavily dependent on a computer system, an important design consideration is the ability to easily recover from troublesome events, whether minor problems or major disruptions of the system. From a design point of view, systems should be designed to easily recover from minor problems, and to be either transportable to another backup computer system or replaced by manual processes in case of major disruption or loss of computer facility.

3.1.2 Access decisions

Once the automated system is ready to use, decisions must be made regarding access to the system and the information it contains. For example, many individuals require the ability to access and view data, but not the ability to change or delete data. Even when computer systems have been designed to provide the ability to narrowly designate access authorities, a knowledgeable and responsible official must actually make those access decisions. The care that is taken in this process is a major determining factor of the level of security and control present in the system. If sensitive data is being transmitted over unprotected lines, it can be intercepted or passive eavesdropping can occur. Encrypting the files will make the data unintelligible and port protection devices will protect the files from unauthorized access, if warranted.

3.1.3 Systems development process

All information systems software should be developed in a controlled and systematic manner according to company standards. The quality and efficiency of the data processed, and the possible reconfiguration of the system can all be affected by an inadequate development process. The risk of security exposures and vulnerabilities is greatly reduced when the systems development process is itself controlled.

3.1.4 Physical security

Information cannot be appropriately protected unless the facilities that house the equipment are properly protected from physical threats and hazards. The major areas of concern are described below.

Environmental conditions

For many types of computer equipment, strict environmental conditions must be maintained. Manufacturer's specifications should be observed for temperature, humidity, and electrical power requirements.

Control of media

The media upon which information is stored should be carefully controlled. Transportable media such as tapes and cartridges should be kept in secure locations, and accurate records kept of the location and disposition of each. In addition, media from an external source should be subject to a check-in process to ensure it is from an authorized source.

Control of physical hazards

Each area should be surveyed for potential physical hazards. Fire and water are two of the most damaging forces with regard to computer systems. Opportunities for loss should be minimized by an effective fire detection and suppression mechanism, and planning reduces the danger of leaks or flooding. Other physical controls include reducing the visibility of the equipment and strictly limiting access to the area or equipment.

Contingency planning

Although risks can be minimized, they cannot be eliminated. When reliance upon a computer facility or application is substantial, some type of contingency plan should be devised to allow critical systems to be recovered following a major disaster, such as a fire. There are a number of alternative approaches that should be evaluated to most cost-effectively meet the company's need for continuity of service.

Configuration management

Risk can be introduced through unofficial and unauthorized hardware or software. Another key component of information resource management is ensuring only authorized hardware and software are being utilized. There are several control issues to be addressed.

Records of hardware/software inventories, configurations, and locations should be maintained and kept up-to-date.

Complying with terms of software licenses

Especially with microcomputer software, illegal copying and other uses in conflict with licensing agreements are concerns. The use of software subject to licensing agreements must be monitored to ensure it is used according to the terms of the agreement.

Protecting against malicious software and hardware

The recent occurrences of destructive computer "viruses" point to the need to ensure that companies do not allow unauthorized software to be introduced to their computer environments. Unauthorized hardware can also contain hidden vulnerabilities. Management should adopt a strong policy against unauthorized hardware/software, inform personnel about the risks and consequences of unauthorized additions to computer systems, and develop a monitoring process to detect violations of the policy.

3.1.5 Data security

Appropriate security mechanisms have to be in place in order to allow responsible officials to designate access to data according to individual computer users' specific needs. Security mechanisms should be sufficient to implement individual authentication of system users, allow authorization to specific information and transaction authorities, maintain audit trails as specified by the responsible official, and encrypt sensitive files if required by user management.

3.1.6 Monitoring and review

A final aspect of information resource protection to be considered is the need for ongoing management monitoring and reviewing. To be effective, a security program must be a continuous effort. Ideally, ongoing processes should be adapted to include information protection checkpoints and reviews. Information resource protection should be a key consideration in all major computer system initiatives.

Earlier, the need for system audit trails was discussed. Those audit trails are useful only if management regularly reviews exception items or unusual activities. Irregularities should be researched and action taken when merited. Similarly, all information-related losses and incidents should be investigated.

A positive benefit of an effective monitoring process is an increased understanding of the degree of information-related risk in agency operations. Without an ongoing feedback process, management may unknowingly accept too much risk. Prudent decisions about trade-off between efficiency and control can only be made with a clear understanding of the degree of inherent risk. Every manager should ask questions and periodically review operations to judge whether changes in the environment have introduced new risk, and to ensure that controls are working effectively.

3.1.7 Personnel management

Information security is more a people issue than a technical issue. Personnel are a vital link in the protection of information resources, as information is gathered by people, entered into information resource systems by people, and ultimately used by people. Security issues should be addressed with regard to:

- People who use computer systems and store information in the course of their normal job responsibilities.
- People who design, program, test, and implement critical or sensitive systems.
- People who operate computer facilities that process critical or sensitive data.

3.1.8 Personnel security

From the point of hire, individuals who will have routine access to sensitive information resources should be subject to special security procedures. More extensive background or reference checks may be appropriate for such positions, and security responsibilities should be explicitly covered in employee orientations. Position descriptions and performance evaluations should also explicitly reference unusual responsibilities affecting the security of information resources.

Individuals in sensitive positions should be subject to job rotation, and work flow should be designed in such a way as to provide as much separation of sensitive functions as possible. Upon decision to terminate or notice of resignation, expedited termination or rotation to less sensitive duties for the remainder of employment is a reasonable precaution.

3.1.9 Training

Most information resource security problems involve people. Problems can usually be identified in their earliest stages by people who are attuned to the importance of information protection issues. A strong training program will yield large benefits in prevention and early detection of problems and losses. To be most effective, training should be tailored to the particular audience being addressed.

Most employees want to do the right thing, if company expectations are clearly communicated. Internal policies can be enforced only if staff have been made aware of their individual responsibilities. All personnel who access company computer systems should be aware of their responsibilities under company policy, as well as obligations under the law. Disciplinary actions and legal penalties should be communicated.

3.1.10 User-level security policy

Ultimately, computer security is the user's responsibility. Each user must be alert to possible breaches in security and adhere to the security regulations that have been established within their company. The security practices listed in

Appendix 1 are not inclusive, but rather designed to remind them and raise their awareness towards securing their information resources.

3.2 Access Security: Firewalls

3.2.1 Introduction

Many companies connect to the Internet, guarded by "firewalls" designed to prevent unauthorized access to their private networks. Despite this general goal, firewalls span a continuum between ease of use and security. This section describes some of the considerations and tradeoffs in designing firewalls. A vocabulary for firewalls and their components is given in Appendix to provide a common ground for discussion.

The rationale for installing a firewall is almost always to protect a private network against intrusion. In most cases, the purpose of the firewall is to prevent unauthorized users from accessing computing resources on a private network, and often to prevent unnoticed and unauthorized export of proprietary information. In some cases export of information is not considered important, but for many corporations that are connecting this is a major though possibly unreasoning concern. Many organizations will want simply to address the problem by not connecting to the Internet at all. This solution can be difficult to implement. If the private network is loosely administered or decentralized, a single enterprising individual with a high-speed dialup modem can quickly effect an Internet SLIP (Serial Line Internet Protocol) connection which can compromise the security of an entire network.

Often it is safe to say that a firewall needs to be put in place for the "CYA" (Cover Your Assets, a family publication) factor. Even though an employee could compromise proprietary information by carrying it offsite on a DAT (Digital Audio Tape) or floppy disk, the Internet represents a tangible threat, populated with dangerous "vandals." (The Vandals were a collection of tribes of roughneck barbarians who sacked Rome in 455 and looted it of all its portable wealth. Some use the term "hackers" to describe Internet snoopers, but "vandals," "crackers," or "jerks" is more appropriate) It could very easily cost a network manager his job if a break-in occurs via this route, even if the damage is no more extensive than could have been inflicted over a dialup line or by a disgruntled employee. Generally, for a would-be Internet site, the technical difficulties of implementing a firewall are greatly outweighed by the public relations problems of "selling" upper management on the idea. In summary, because Internet services are so highly visible, they are much more likely to require official oversight and justification.

3.2.2 Design decisions

In configuring a firewall, the major design decisions with respect to security are often already dictated by corporate or organizational policy; specifically, a decision must be made as to whether security is more important than ease-of-use, or vice versa. There are two basic approaches that summarize the conflict:

- That which is not expressly permitted is prohibited.
- That which is not expressly prohibited is permitted.

The importance of this distinction cannot be overemphasized. In the former case, the firewall must be designed to block everything, and services must be enabled on a case-by-case basis only after a careful assessment of need and risk. This tends to impact users directly, and they may see the firewall as a hindrance. In the second case, the systems administrator is placed in a reactive mode, having to predict what kinds of actions the user population might take that would weaken the security of the firewall, and preparing defenses against them. This essentially pits the firewall administrator against the users in an endless arms race that can become quite fierce. Users can generally compromise the security of their login if they try or are not aware of reasonable security precautions. If the user has an open access login on the firewall system itself, a serious security breach can result. The presence of user logins on the firewall system tends to magnify the problem of maintaining the system's integrity. A second important statement of policy is implicit in the "that which is not expressly permitted is prohibited" stance. This stance is more "fail safe," since it accepts that the administrator is ignorant of what TCP ports are safe, or what holes may exist in the manufacturer's kernel or applications. Since many vendors are slow to publicise security holes, this is clearly a more conservative approach. It is an admission of the fact that what you do not know can hurt you.

3.2.3 Levels of threat

There are several ways in which a firewall can fail or be compromised. While none of them are good, some are decidedly worse than others. Since the purpose of many firewalls is to block access, it is a clear failure if someone finds a loophole through it which permits them to probe systems in the private network. An even more severe situation would result if someone managed to break into the firewall and reconfigure it such that the entire private network is reachable by anyone. For the sake of terminology, this type of attack will be referred to as "destroying" a firewall, as opposed to a mere "break-in." It is extremely difficult to quantify the damage that might result from a firewall's destruction. An important measure of how well a firewall resists threat is the information it gathers to help determine the course of an attack. The absolute worst thing that could happen is for a firewall to be completely compromised without any trace of how the attack took place. The best thing that can happen is for a firewall to detect an attack, and inform the administrator politely that it is undergoing attack, but that the attack is going to fail.

One way to view the result of a firewall being compromised is to look at things in terms of what can be roughly termed as "zones of risk." In the case of a network that is directly connected to the Internet without any firewall, the entire network is subject to attack. This does not imply that the network is vulnerable to attack, but in a situation where an entire network is within reach of an untrusted network, it is necessary to ensure the security of every single host on that network. Practical experience shows that this is difficult, since tools like rlogin that permit user-customizable access control are often exploited by vandals to gain access to multiple hosts, in a form of “island hopping” attack. In the case of any typical firewall, the zone of risk is often reduced to the firewall itself, or a selected subset of hosts on the network, significantly reducing the network manager's concerns with respect to direct attack. If a firewall is broken into, the zone of risk often expands again, to include the entire protected network. A vandal gaining access to a login on the firewall can begin an island hopping attack into the private network, using it as a base. In this situation, there is still some hope, since the vandal may leave traces on the firewall, and may be detected. If the firewall is completely destroyed the private network can undergo attack from any external system and reconstructing the course of an attack becomes nearly impossible.

In general, firewalls can be viewed in terms of reducing the zone of risk to a single point of failure. In a sense, this seems like a bad idea, since it amounts to putting all of one's eggs in a single basket, but practical experience implies that at any given time, for a network of non-trivial size, there are at least a few hosts that are vulnerable to break-in by even an unskilled attacker. Many corporations have formal host security policies that are designed to address these weaknesses, but it is sheer foolishness to assume that publishing policies will suffice. A firewall enhances host security by funneling attackers through a narrow gap where there is a chance of catching or detecting them first. The well-constructed medieval castle had multiple walls and interlocking defense points for exactly the same reason.

3.2.4 Firewalls and their components

Overview

In discussing firewalls there is often confusion of terminology since firewalls all differ slightly in implementation if not in purpose. Various discussions on Usenet indicate that the term “firewall” is used to describe just about any inter-network security scheme. For the sake of simplifying discussion, some terminology is proposed, to provide a common ground:

Screening Router

A screening router is a basic component of most firewalls. A screening router can be a commercial router or a host-based router with some kind of packet filtering capability. Typical screening routers have the ability to block traffic between networks or specific hosts, on an IP port level. Some firewalls consist of nothing more than a screening router between a private network and the Internet.

Bastion Host

Bastions are the highly fortified parts of a medieval castle; points that overlook critical areas of defense, usually having stronger walls, room for extra troops, and the occasional useful tub of boiling hot oil for discouraging attackers. A bastion host is a system identified by the firewall administrator as a critical strong point in the network's security. Generally, bastion hosts will have some degree of extra attention paid to their security, may undergo regular audits, and may have modified software.

Dual Homed Gateway

Some firewalls are implemented without a screening router, by placing a system on both the private network and the Internet, and disabling TCP/IP forwarding. Hosts on the private network can communicate with the gateway, as can hosts on the Internet, but direct traffic between the networks is blocked. A dual homed gateway is, by definition, a bastion host.

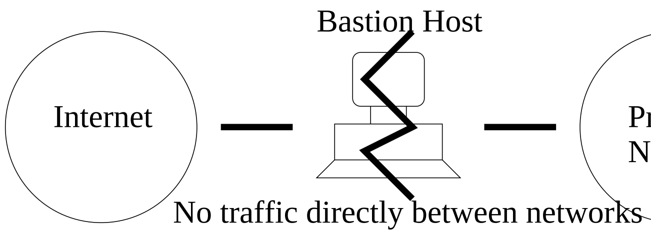

Figure 3.1: A typical Dual Homed Gateway

Screened Host Gateway

Possibly the most common firewall configuration is a screened host gateway. This is implemented using a screening router and a bastion host. Usually, the bastion host is on the private network, and the screening router is configured such that the bastion host is the only system on the private network that is reachable from the Internet. Often the screening router is configured to block traffic to the bastion host on specific ports, permitting only a small number of services to communicate with it.

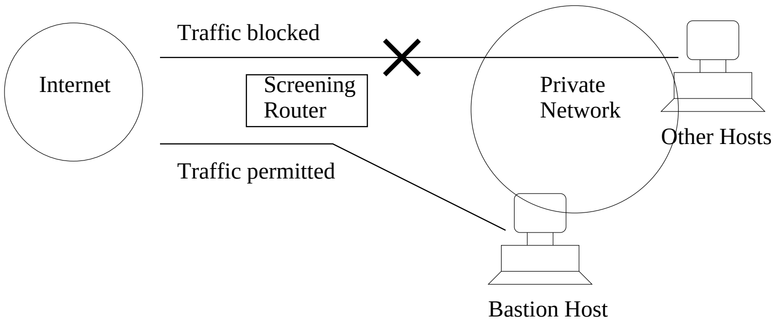

Figure 3.2: A typical Screened Host Gateway

Screened Subnet

In some firewall configurations, an isolated subnet is created, situated between the Internet and the private network. Typically, this network is isolated using screening routers, which may implement varying levels of filtering. Generally, a screened subnet is configured such that both the Internet and the private network have access to hosts on the screened subnet, but traffic across the screened subnet is blocked. Some configurations of screened subnets will have a bastion host on the screened network, either to support interactive terminal sessions or application level gateways.

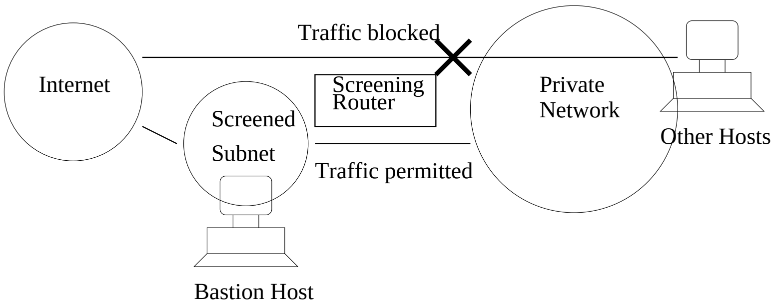

Figure 3.3: A typical Screened Subnet

Application Level Gateway (or "Proxy Gateway")

Much of the software on the Internet works in a store-and-forward mode; mailers and Usenet news collect input, examine it, and forward it. Application level gateways are service-specific forwarders or reflectors, which usually operate in user mode rather than at a protocol level. Generally, these forwarding services, when running on a firewall, are important to the security of the whole. The famous sendmail hole that was exploited by the Morris Internet worm is one example of the kinds of security problems an application level gateway can present. Other application level gateways are interactive, such as the FTP and telnet gateways run on the Digital Equipment Corporation firewalls. In general, the term "application level gateway" will be used to describe some kind of forwarding service that runs across a firewall, and is a potential security concern. In general, crucial application level gateways are run on some kind of bastion host.

Hybrid Gateways

Hybrid gateways are the "something else" category in this list. Examples of such systems might be hosts connected to the Internet, but accessible only through serial lines connected to an ethernet terminal server on the private network. Such gateways might take advantage of multiple protocols, or tunneling one protocol over another. Routers might maintain and monitor the complete state of all TCP/IP connections, or somehow examine traffic to try to detect and prevent an attack. The AT&T corporate firewall is a hybrid gateway combined with a bastion host.

Taking the components described above, we can accurately describe most of the forms that firewalls take, and can make some general statements about the kinds of security problems each approach presents. Assuming that a firewall fulfills its basic purpose of helping protect the network, it is still important to examine each type of firewall with respect to:

Damage control — If the firewall is compromised or destroyed, to what kinds of threats does it leave the private network open?

Zones of risk — How large is the zone of risk during normal operation? A measure of this is the number of hosts or routers that can be probed from the outside network.

Failure mode — If the firewall is broken into or destroyed, how easy is this to detect? In a post mortem, how much information is retained that can be used to diagnose the attack?

Ease of use — How much of an inconvenience is the firewall?

Stance — Is the basic design philosophy of the firewall "That which is not expressly permitted is prohibited" or is it "That which is not expressly prohibited is permitted"?

Firewalls using screening routers

Many networks are firewalled using only a screening router between the private network and the Internet. This type of firewall is different from a screened host gateway in that usually there is direct communication permitted between multiple hosts on the private network, and multiple hosts on the Internet. The zone of risk is equal to the number of hosts on the private networks, and the number and type of services to which the screening router permits traffic. For each service provided via peer-to-peer connection the size of the zone of risk increases sharply. Eventually it is impossible to quantify. Damage control is difficult as well since the network administrator would need to regularly examine every host for traces of a break-in. If there is no regular audit one must hope to stumble on a clue such as a mismatched system accounting record.

In the case of total destruction of the firewall, it tends to be very hard to trace or even to discover. If a commercial router (which does not maintain logging records) is used, and the router's administrative password is compromised, the entire private network can be laid open to attack very easily. Cases are known where commercial routers have been configured with erroneous screening rules, or have come up in some pass-through mode because of hardware or operator error. Generally, this configuration is a case of "That which is not expressly prohibited is permitted" as the ingenious user can fairly easily piggyback protocols to achieve a higher level of access than the administrator expects or wants.

Screening routers are not the most secure solution, but they are popular since they permit fairly free Internet access from any point within the private network. Many consultants and network service providers offer screening routers in a "firewall" configuration. It is uncertain if the various trade-offs involved are clear to the customer; thus the use of a screening router to protect sensitive information or trade secrets would not be recommended, since screening routers are very permeable from the inside.

Dual homed gateways

An often used and easy to implement firewall is the dual homed gateway. Since it doesn't forward TCP/IP traffic, it acts as a complete block between the Internet and the private network. Its ease of use is determined by how the systems manager chooses to set up access; either by providing application gateways such as telnet forwarders or by giving users logins on the gateway host. If the former approach is taken, the stance of the firewall is clearly "That which is not expressly permitted is prohibited"; users can only access Internet services for which there is an application gateway. If users are permitted logins, the firewall's security is then seriously weakened. During normal operation, the only zone of risk is the gateway host itself, since it is the only host that is reachable from the Internet. If there are user logins on the gateway host, and one of the users chooses a weak password or has their account otherwise compromised, the zone of risk expands to encompass the entire private network. From a standpoint of damage control, the administrator may be able to track the progress of an intruder, based on the access patterns of the compromised login, but a skillful vandal can make this quite difficult. If a dual homed gateway is configured without direct user access, damage control can be somewhat easier, since the very fact that someone has logged in to the gateway host becomes a noteworthy security event. Dual homed gateways have an advantage over screening routers from the standpoint that their system software is often easier to adapt to maintain system logs, hard copy logs, or remote logs. This can make a post-mortem easier for the gateway host itself, but may or may not help the network administrator identify what other hosts on the private network may have been compromised in an island-hopping attack.

Attacking a dual homed gateway leaves the attacker a fairly large array of options. Since the attacker has what amounts to local network access if a login can be obtained, all the usual attacks that can be made over a local network are available. NFS-mounted file systems, weaknesses in .rhosts files, automatic software distribution systems, network backup programs and administrative shell scripts – all may provide a toehold on systems on the internal network. Once a toehold is secured, it then provides a base from which to launch attacks back at the gateway itself. The weakest aspect of the dual homed gateway is its failure mode. If the firewalll is destroyed it is possible that a skillful attacker might re-enable routing and throw the entire private network open to attack. In the usual Unix-based dual homed gateway, TCP/IP routing is often disabled by modifying a kernel variable named ipforwarding; if systems privileges can be obtained or stolen on the gateway, this variable can be changed. Perhaps this seems far-fetched, but unless great care is paid to monitoring the software revision levels and configuration on the gateway host, it is not improbable that a vandal with a copy of the release notes for the operating system version and a login can compromise the system.

Screened host gateways

Several articles have described screened host gateways, and how to construct them. Generally, the screened host gateway is very secure, while remaining fairly easy to implement. Typically, a bastion host is configured on the private network, with a screening router between the Internet and the private network, which only permits Internet access to the bastion host. Since the bastion host is on the private network, connectivity for local users is very good, and problems presented by exotic routing configurations do not present themselves. If the private network is, as many are, a virtual extended local area network (e.g.: no subnets or routing) the screened host gateway will work without requiring any changes to the local network, as long as the local network is using a legitimately assigned set of network addresses. The zone of risk of a screened host gateway is restricted to the bastion host, and the screening router, and the security stance of the screened host gateway is determined by the software running on that system. If an attacker gains login access to the bastion host, there is a fairly wide range of options for attacking the rest of the private network. In many ways, this approach is similar to the dual homed gateway, sharing similar failure modes and design considerations with respect to the software running on the bastion host.

Screened subnets

A screened subnet is usually configured with a bastion host as the sole point of access on the subnet. The zone of risk is small, consisting of that bastion host or hosts, and any screening routers that make up the connections between the screened subnet, the Internet, and the private network. The ease of use and basic stance of the screened subnet will vary, but generally a screened subnet is appealing only for firewalls that are taking advantage of routing to reinforce the existing screening. This approach forces all services through the firewall to be provided by application gateways, and places the stance strongly in the "That which is not expressly permitted is prohibited" category.

If a screened subnet based firewall with inter-network routing blocked is attacked with an intent to destroy it, the attacker must reconfigure the routing on three networks, without disconnecting or locking himself out, and without the routing changes being noticed. No doubt this is possible, but it can be made very difficult by disabling network access to the screening routers, or by configuring the screening routers to only permit access from specific hosts on the private network. In this case, an attacker would need to break into the bastion host, then into one of the hosts on the private network, and then back out to the screening router — and would have to do it without setting off any alarms.

Another advantage of screened subnets is that they can be put in place in such a way that they hide any accidents of history that may linger on the private network. Many sites that would like to connect to the Internet are daunted by the prospect of re-addressing and re-subnetting existing networks. With a screened subnet with blocked inter-network routing, a private network can be connected to the Internet and changed gradually to new subnet and network addresses. In fact, this approach has been observed to significantly accelerate the adoption of new network addresses on loosely controlled private networks. Users will be more receptive to changing their host addresses if they can realize the benefits of Internet connectivity thereby, since hosts that are not correctly addressed cannot use the firewall properly. In most other respects, the screened subnet is very much dependent on the suite of software running on the bastion host. Screening a whole subnet provides functionality similar to the dual homed gateway or screened host gateway; it differs primarily in the extra level of complexity in routing and configuration of the screening routers.

Hybrid gateways

"Security through obscurity" is not sufficient in and of itself, but there is no question that an unusual configuration, or one that is hard to understand, is likely to give an attacker pause, or to make them more likely to reveal themselves in the process of trying to figure out what they are facing. On the other hand there is a real advantage to having a security configuration that is easy to understand, and therefore easier to evaluate and maintain. Since the hybrid gateway is mentioned here in the category of "something else" no attempt will be made to describe the indescribable. Some hypothetical hybrids may serve to show how hybrid gateways might differ from and be similar to the other types.

Let us postulate a hybrid gateway that consists of a box sitting on the Internet, which is capable of routing traffic, but also maintains a complete notion of the state of every TCP connection, how much data has gone across it, where it originated, and its destination. Presumably, connections can be filtered based on arbitrarily precise rules, such as: "permit traffic between host A on the private network and all hosts on network B on the Internet via the telnet service if and only if the connection originated from host A between the hours of 9:00 am and 5:00 pm and log the traffic." This sounds terrific, providing arbitrary control with great ease of use, but some problems simply refuse to go away. Consider that someone wishing to circumvent the firewall, who broke into the private network via an unguarded modem, might very easily set up a service engine that was piggybacked over the telnet port. This is actually a fairly easy firewall to destroy.

Another hybrid gateway might take advantage of various forms of protocol tunneling. Suppose the requirement is to connect to the Internet with very tight restrictions, but that a high degree of connectivity is required between the private network and an external network that is somewhat trusted (for example, a corporate R&D department needs to be able to run X-windows applications on a supercomputer at another facility). The usual archetypal gateways discussed here could provide general purpose e-mail connectivity, but for secure point-to-point communications, an encrypted point-to-point virtual TCP/IP connection might be set up with the remote system, after users had authenticated themselves with a cryptographic smart card. This would be extremely secure, and might be made fairly easy to use, but has the disadvantage that the protocol driver needs to be added to every system that wants to share communication. It is hard to make any guesses about the failure mode of such a system, but the zone of risk is neatly limited to all the hosts which are running the tunneling protocol driver, and to which the individual user has smart card access. Some of this might be implemented in hardware or in the routers themselves. In the future, it is likely that the rapid growth of the Internet will fuel more development in this area, and we will see various hybrid gateways arise. The basic issues surrounding configuring a firewall will probably remain the same as the ones discussed here.

Other firewall-related tools

There is active research and development on tools to aggressively seek out and identify weaknesses in an entire network, or to detect the patterns that might indicate when an attack is in progress. These tools range from the simple checklist to complex "expert systems" with inference engines and elaborate rule bases. Many firewalls today run software that is designed to go forth and gather information relating to possible attacks and their origins, often using and abusing tools like finger and SNMP. Unless true artificial intelligence is developed, however, these tools cannot guard against an unknown form of attack, since they cannot possibly match the creativity of a network vandal. While often billed as being "proactive" they are in fact reactive, and generally will serve only to catch systems crackers armed with last year's bag of tricks. Catching the small fry is still worth doing, but it is likely that they are less of a threat than the fellow who is so eager to break into your network that he is doing research and development in new system cracking techniques.

3.2.5 Comments

In dealing with firewalls, it is simply not reasonable to say that any particular approach is best, since there are so many factors that determine what the best firewall for a given situation may be. Cost, corporate policy, existing network technology, staffing, and intra-organizational politics may all easily outweigh the technical considerations presented here.

There are a few observations worth making about firewalls at a very general level. Firstly, a firewall is a leverage-increasing device from a network management point of view. Rather than looking at it as "all eggs in one basket," it can also be viewed as a trustworthy basket, and a single point from which a very important security system can be controlled. The size of the zone of risk is crucial to the design; if it is small, security can be maintained and controlled easily but if security is compromised, the damage can be more severe. The ideal would be to have such strong host-based security that a firewall would be redundant. Systems administration costs and a hard dose of reality prevents this ideal from being obtainable.

A second important aspect of firewall building is that it is not something to undertake in a vacuum. Many sites are connected with a simple firewall consisting of a screening router and nothing more because someone told them that it was "secure enough." In setting up a firewall one must trade off time and money, security, and risk. "Secure enough" depends largely on what you risk losing. Finally, it is important when approaching implementing a firewall to avoid the urge to start from scratch. System security is a lot like pregnancy; one is seldom only broken into a little bit, and it only takes a little mistake or a moment of inattention to find oneself in a delicate position. Leaning on the experiences of others, and learning from their mistakes and successes is very important. Setting up a firewall is definitely an area where having a wide background in experience to draw upon is important. The vandals on the network have a wide background in experience to draw upon as well, and a firewall administrator must communicate with others, and must keep up to date on other firewall-related happenings on the network. Static defenses do not work unless they keep up with emerging tricks of the trade, or one's firewall may be the next Maginot Line.

3.3 Data Security: Cryptography

3.3.1 What is cryptography?

According to a dictionary it is "The art or process of writing in or deciphering secret code." In terms of the topic at hand, it is the transforming of information into a form unreadable by anyone without a secret decryption key.

3.3.2 What is a key?

A key is a piece of data that can be used to encrypt or decrypt a given message.

In a common form of cryptography, secret key cryptography, the sender uses a secret key to encrypt some information. The sender must then get both the secret key and the encrypted information to the receiver securely. The problem with secret key cryptography, from the standpoint of transactions over the Internet, is that anyone who can get both the key and the encrypted information can decrypt the information. As we noted above, it is fairly easy for an unauthorized person to intercept my information on the Internet. If I send both an encrypted message and the secret key that decrypts the message over the Internet, then I risk having both intercepted and having the message decrypted by the interceptor.

3.3.3 What is public key cryptography (PKC)?

Public key cryptography (PKC) is a system based on pairs of keys, a public key and a private key, rather than on a single secret key. The keys are related in such a way that information that is encrypted with the public key can only be decrypted with the private key. This means that, if I have a pair of keys, I can make my public key freely available to anyone who wants it without worrying about who else might get it. If anyone wants to send me a secret message, they can use my public key to encrypt the message. Only I can decrypt the message, using my private key, which I keep safe.

The action of encrypting information with PKC keys is significantly slower than encrypting with a secret key, such as a DES (Data Encryption Standard) key. So, in practice, a system that uses PKC first generates a secret key according to a method such as DES. It uses the secret key to encrypt the message, and it uses the PKC key to encrypt the secret key. Then, it attaches the PKC-encrypted secret key to the secret key-encrypted message.

3.3.4 How does someone obtain a pair of keys?

A specialized piece of computer hardware and/or software generates key pairs on demand according to the method defined by the PKC standards.

3.3.5 How is information integrity assured?

When you put together your message to me and encrypt it with my public key, you also have your computer calculate something called a "message digest" and attach it to the message. The message digest is a small piece of data that results from performing a particular mathematical calculation on your message. It acts as the digital fingerprint of your message — the message digest of every message is different.

When I receive your message and decrypt it with my private key, I also calculate the message digest of your message. Then, I compare the digest I generated to the one you sent with your message. If they match, then I know that I received exactly what you sent.

3.3.6 What about authenticity?

When you obtain my public key, from me or my company or from some other organization, it comes with a "certificate". The certificate is a piece of cryptographically secure information attached to my public key that guarantees that the key you obtained is mine.

3.3.7 Is the certificate trustworthy?

Because it comes from an organization that you trust. If I hand you my driver's license to identify myself, you accept it as valid identification because you trust the organization that issued me the license, namely the government of my state.

If I am a catalog sales merchant and you obtain my public key from some place on the Internet, it might have both a certificate from my company and a certificate from my company's bank attached to it. This is a way of saying, "Look, it's not just Acme Sales Corp. that says this public key is from our designated representative, but First National Bank says so, too." Presumably, for First National Bank to issue such a certificate, I (or someone from my company) had to prove myself to be a representative of Acme to the appropriate person at First National.

As you can see, there is a notion of a "hierarchy of trust" associated with public keys and certificates. At the top of the hierarchy is some organization that we trust implicitly. With checks and credit cards, we have to trust the banks that issue them. With currency we have to trust the government that issued it.

A trusted organization that issues certificates on behalf of others is called a "certifying authority" (CA).

3.3.8 What about non-repudiatability of origin?

Using my private key and the message I am sending, I have my computer make a particular calculation that results in a piece of information called a "digital signature". Since it involves my private key, the signature can only have come from me. And since the calculation uses the message I'm sending, the signature is only for this message. I then have my computer attach the digital signature to the message, and I send to you.

When you receive the message, your computer does a calculation that uses the message, the digital signature and my public key. Since my public and private keys are related, the calculation can verify 1) that my private key was used to create the signature and 2) that the message you received was used to create the signature.

3.3.9 What about non-repudiatability of receipt?

If you and I are representatives of companies conducting business with each other, then we would both have PKC pairs that identify us. You would have my public key and I would have yours. We might agree that, once you receive a digitally signed message from me, you send me a receipt. The receipt would be my digital signature from the message I sent to you, with your digital signature attached to it. Using your public key, I would verify your digital signature. Using my private key I would verify my digital signature that you had sent back to me.

How can I tell when a message was originated?

I can attach a "digital time-stamp" to my message. First, I would have my computer calculate my digital signature and attach it to the original message. Next, I would have it calculate a message digest of the message plus the signature. Next, I would send the message digest (presumably over the Internet) to an external, trusted digital time-stamping service (DTS). (Bellcore has created a subsidiary to provide such a service.) The DTS would return a digital time-stamp consisting of:

- the message digest that I sent
- the date and time it was received by the DTS, and

- the digital signature of the DTS encrypted with the DTS's private key.

The "message" that the DTS signs would be the combination of my message digest (the digital fingerprint of my original message) and the time-stamp. When my computer received the time-stamp from DTS it would attach it to my message, add my digital signature and send you the whole thing.

When you received my message, your computer would verify the time-stamp using the public key of the DTS, and it would verify my signature using my public key.

3.4 Transaction Security

3.4.1 What are Kerberos, PGP, PEM, SSL, S-HTTP and IPSEC?

Kerberos

What is Kerberos? What is it good for?

Kerberos is a network authentication system for use on physically insecure networks, based on the key distribution model presented by Needham and Schroeder (in "Using Encryption for Authentication in Large Networks of Computers", Communications of the ACM, Vol. 21(12), pp. 993-999, December, 1978). It allows entities communicating over networks to prove their identity to each other while preventing eavesdropping or replay attacks. It also provides for data stream integrity (detection of modification) and secrecy (preventing unauthorized reading) using cryptography systems such as DES.

Kerberos works by providing principals (users or services) with tickets that they can use to identify themselves to other principals and secret cryptographic keys for secure communication with other principals. A ticket is a sequence of a few hundred bytes. These tickets can then be embedded in virtually any other network protocol, thereby allowing the process implementing that protocol to be sure about the identity of the principals involved.

Practically speaking, Kerberos is mostly used in application-level protocols (ISO model level 7), such as Telnet or FTP, to provide user to host security. It is also used, though less frequently, as the implicit authentication system of data stream (such as `SOCK_STREAM`) or RPC mechanisms (ISO model level 6). It could also be used at a lower level for host to host security, in protocols like IP, UDP, or TCP (ISO model levels 3 and 4), although such implementations are currently rare, if they exist at all.

There are several different versions and distributions of Kerberos. Most of them are based on an MIT distributions in one form or another, but the lineage is not always simple. Some of the distributions are freely available, some are stand-alone commercial products, and others are part of a larger free or commercial systems.

Several vendors currently sell relevant Kerberos material, and to name a few: CyberSAFE, Cygnus Support, Digital Equipment Corporation, Emulex Network Systems, OpenVision Technologies, TGV.

What is Bones? What is it for?

Bones is a system that provides the Kerberos API without using encryption and without providing any form of security whatsoever. It is a fake that allows the use of software that expects Kerberos to be present when it cannot be.

Why does it exist? Kerberos is a network security system which relies on cryptographic methods for its security. Since Kerberos' encryption system, DES, is not exportable, Kerberos itself cannot be exported or used outside of the United States in its original form.

As a partial solution to this problem, the Kerberos source code was modified by the addition of `#ifdef NOENCRYPTION` around all calls to DES functions. Compiling this version with the symbol NOENCRYPTION defined results in a system that looks like Kerberos from an application's point of view but that does not require DES libraries (and, as a result, does not speak the real Kerberos protocol and does not provide any security).

The final piece in this puzzle is a program called "piranha" which takes the Kerberos sources and removes all of the calls to the encryption routines, replacing it with the code which was under `#ifdef NOENCRYPTION`, producing the system known as Bones. Bones has the property that there is absolutely no question about whether or not it is legal to transport its sources across national boundaries, since it neither has any encryption routines nor any calls to encryption routines.

`#ifdef NOENCRYPTION` was not documented, and it was only intended to be used in the above manner. Someone who tries compiling Kerberos with that `#define` has in some sense "voided the warranty", and will get something which is both not secure and not Kerberos.

Pretty Good Privacy (PGP)

A little history

PGP is a remarkable phenomenon. Largely the effort of a single person, Philip Zimmermann, it provides a confidentiality and authentication service that can be used for e-mail and file-storage applications. In essence, what Zimmermann has done is the following:

- Selected the best available cryptographic algorithms as building blocks.

- Integrated the algorithms into a general-purpose application that is independent of operating systems and processors and that is based on a small set of easy-to-use commands.
- Made the package and its documentation, including the sourcecode, freely available via the Internet, bulletin boards, and commercial networks such as CompuServe.
- Entered into an agreement with a company (ViaCrypt) to provide a fully compatible, low-cost commercial version of PGP.

From its beginnings just a few years ago, PGP has grown explosively and is now widely used. A number of reasons can be cited for this growth:

- It is available for free worldwide in versions that run on a variety of platforms, including DOS, Windows, Unix, and Macintosh. And the commercial version satisfies users who want a product that comes with vendor support.
- It is based on algorithms that have survived extensive public review and are considered extremely secure (RSA for public-key encryption, IDEA for conventional encryption, and MD5 for hash, or summary, coding).
- It has a wide range of applicability and can be used by corporations that want to enforce a standardized scheme for encrypting files and messages as well as by individuals who wish to communicate securely with people worldwide over the Internet and other networks.
- It was not developed by, nor is it controlled by, any governmental or standards organization. For people with an instinctive distrust of “the establishment”, this makes PGP attractive.

Some people have hesitated to use PGP for several reasons. First, the developer lives and works in the United States, and because PGP is a cryptographic product, there is concern about export controls for organizations that wish to interoperate between U.S. and non-U.S. sites. However, some recent versions of PGP have initially seen the light of day outside the United States, which would seem to make U.S. export controls irrelevant.

Second, the freeware version of PGP was developed and disseminated without a license for the use of RSA, which is required within the United States but not outside. For people concerned with this issue, the commercial version of PGP has an RSA license. In addition, MIT has issued a new freeware version PGP with a licensed version of RSA.

What PGP does

PGP provides two services: encryption and digital signatures. Encryption allows a user to encode a file for storage locally or for transmission as an e-mail message. The local storage option is handy if you are worried about other people having access to files on your machine. The e-mail option enables PGP to be used for private exchanges over a network. PGP encrypts the entire contents of the message in such a way that only the intended recipient can decode and read the message. Anyone else who attempts to capture or copy the message en route will receive meaningless garble.

The digital signature service allows a user to 'sign' a document before transmission in such a way that anyone can verify that the signature is genuine and belongs with a particular document. If someone alters the message or substitutes a different message, the signature will no longer be valid. And any recipient can verify that the message has been signed by its true creator and not an imposter.

PGP's confidentiality and encoding services use the most popular public-key encryption scheme, known as RSA. All public-key encryption systems make use of an encoding and decoding algorithm and a related pair of keys. The input to the encryption algorithm is the text to be encrypted — known as plaintext — and a key. The algorithm takes the input and produces scrambled output known as ciphertext. To use the decryption algorithm, you input the ciphertext plus the key that matches the one used for encryption, and the original plaintext is produced as output.

The two keys used in any public-key encryption scheme, including RSA, are called the public key and private key. The public key, as the name suggests, is made public. The idea is to make your public key available to people with whom you correspond. You keep your private key secure, and it should be known only to you. These two keys can be used to provide confidentiality and encryption.

Here's how it works. Suppose Alice wants to correspond with Bob. If Alice prepares a message and encrypts it with Bob's public key, only Bob can decrypt the message using his private key. If Alice prepares a message and encrypts it with her private key, then anyone, including Bob, can decrypt the message. But only Alice could have encrypted the message, therefore the encrypted message is, in effect, signed by Alice.

It turns out that RSA, and all other known public-key algorithms, are time-consuming and inefficient. Therefore PGP, like most other encryption applications, does not use RSA directly to provide confidentiality and digital signatures.

For confidentiality, PGP encrypts messages with an efficient single-key or conventional encryption algorithm known as IDEA. It then uses RSA to encrypt, with the receiver's public key, the IDEA key used to encrypt the message. The receiver can use RSA to recover the IDEA key and use that key to recover the message.

For digital signatures, PGP uses an efficient algorithm known as MD5 to produce a summary code, or hash code, of the message that is, for all practical purposes, unique to that message. PGP then uses RSA to encrypt the hash code with the sender's private key. The receiver can use RSA to recover the hash code and verify that it is the correct hash code for the message. If it is correct, then only the alleged sender could have prepared the encrypted hash code.

Trusting public keys

The most difficult aspect of using PGP, or any public-key application, is getting your hands on the public keys of the people with whom you wish to correspond. You must make sure you have the true public key of each individual in your electronic Rolodex. Suppose I create a linked pair of public and private keys and send you the public key, declaring that I am Elvis. How do you know I am the real Elvis and not an impostor? If I am an impostor, I could send you signed messages and you would be sure they were from Elvis. If you send an encrypted message to Elvis, I can capture the message and recover the plain text.

PGP provides a number of tools and recommended procedures for obtaining public keys in trusted ways. One handy tool is the public-key fingerprint, which is nothing more than a string of printable characters based on the MD5 hash code of the key. For all practical purposes, the fingerprint of a key is unique. So, if Alice knows Bob's voice, Bob could send his public key to Alice via e-mail. Alice then could generate the fingerprint of that key, call Bob, and have Bob read the fingerprint over the phone to make sure there is a match.

Once you have a few trusted keys, you can make use of PGP's signature capability. If you have Bob's public key and you trust Bob to provide you with other public keys of other persons, Bob can send you John's key signed by Bob. That is, Bob takes John's public key and feeds it through the signature mechanism of PGP. Alice can use Bob's public key to ensure that John's key was provided by Bob and that the key has not been altered.

There also are a number of servers on the Internet that are public-key repositories. Most of keys are signed by one or more people. You can obtain someone's public key from the server and if you trust the signatories to the key, you can have faith that it is genuine. These public-key servers do not authenticate the keys; they merely serve as repositories.

One public-key directory that does attempt to provide authenticated PGP keys is SLED (Stable Large E-mail Database). The public keys in the directory are signed by SLED, indicating that the user's authenticity has been verified.

Privacy Enhanced Mail (PEM)

What is PEM?

PEM is an Internet standard for providing security services to electronic mail. It uses cryptographic techniques to provide message integrity checking, originator authentication, and confidentiality. It lets you know that a message hasn't been changed, who it's from, and, optionally, allows you to keep it secret from all but the intended recipients.

There is a set of Proposed Standard RFCs (Internet standards documents) that specify PEM. The four documents are RFCs 1421 (obsoletes 1113), 1422 (obsoletes 1114), 1423 (obsoletes 1115), and 1424 (new). The integration of MIME (Multipurpose Internet Mail Extensions) and PEM will either produce a new RFC for message format or cause RFC 1421 to be updated.

An overview of PEM was presented in the August 1993 issue (Volume 38, Number 8) of "Communications of the ACM" in an article entitled "Internet Privacy Enhanced Mail" by Stephen T. Kent.

What is TIS/PEM?

Trusted Information Systems (TIS), under ARPA sponsorship and in cooperation with RSA Data Security Incorporated (RSADSI), has released a reference implementation of Privacy Enhanced Mail (TIS/PEM) to the Internet community.

TIS/PEM is a Unix-based implementation (The authors are currently pursuing porting it to other operating systems) that has been integrated with Rand MH 6.8.3 and is easily integrated into other mail user agents. TIS/PEM is distributed in source form. It is openly available within the United States and Canada for non-commercial use (not for resale), but the export from the United States of the cryptography used in TIS/PEM is controlled by the United States government.

What is RIPEM?

Riordan's Internet Privacy Enhanced Mail (RIPEM) is an implementation of PEM. RIPEM allows electronic mail to have the four security facilities provided by PEM: disclosure protection (optional), originator authenticity, message integrity measures, and non-repudiation of origin (always). RIPEM is not really a complete implementation of PEM, because PEM specifies certificates for authenticating keys, which RIPEM does not handle at this time (however, their addition is planned). They have already been added to the Macintosh version.

RIPEM was written primarily by Mark Riordan ¡mrr@scss3.cl.msu.edu¿, but most of the code is now in the public domain, except for the RSA routines, which are a library called RSAREF licensed from RSA Data Security Inc. RIPEM has been ported to almost all platforms.

The certificates issue

PEM users have the choice of using bare public keys or public keys embodied in X.509 certificates. If we use X.509 certificates to bind distinguished names to RSA public keys, it is not necessary to join the Internet certification hierarchy or otherwise pay to use PEM.

TIS/PEM is capable of generating the bare keys and certificates that we may need. Joining the Internet certification hierarchy has the benefit of making it easier to verify others' mail and for them to verify ours. To join the Internet certification hierarchy, we must sign up our Certification Authority (CA) under a Policy-level Certification Authority (PCA).

A distinguished name is a hierarchical, globally unique name used to identify something or someone.

A Certification Authority (CA) vouches for the binding between users' distinguished names and RSA public keys within an organization or organizational unit. The CA's distinguished name is that of the organization or organizational unit and users' distinguished names are created by starting with the CA distinguished name and adding something to uniquely and unambiguously identify the user, like a common name.

Policy-level Certification Authorities (PCAs) vouch for the binding between a CA's distinguished name and RSA public key. By joining a PCA, others can verify our PEM messages by following the certification path to the Internet PCA certificate without having to have retrieved our RSA public key using secure, out of band means. PCAs may also make CA Certificate Revocation Lists (CRLs) and certificates available and provide other services for its members.

PCAs can be differentiated by the policy that they advertise. The policy includes the level of effort — and associated assurance — that a PCA uses to insure the correctness of the binding and the requirements they place on CAs which issue certificates under them. They can also be differentiated by the other services they offer and their price.

Several PCAs exist as part of the Internet certification hierarchy, including PCAs at RSADSI and TIS, and more may come online in the near future. Individual PCAs will have their own price schedules.

Secure Sockets Layer (SSL)

Introduction

Netscape Communications has designed and specified a protocol for providing data security layered between application protocols (such as HTTP, Telnet, NNTP, or FTP) and TCP/IP. This security protocol, called Secure Sockets Layer (SSL), provides data encryption, server authentication, message integrity, and optional client authentication for a TCP/IP connection.

SSL is an open, non-proprietary protocol, and has been submitted to the W3C working group on security for consideration as a standard security approach for World Wide Web browsers and servers on the Internet. Netscape Communications is working with the W3C and W3C members (including EIT) on developing and standardizing common, robust security mechanisms and protocols for the Internet. They will fully support such mechanisms and protocols as they become standardized.

The SSL protocol specification is freely available and Netscape Communications encourage everyone to examine the protocol and/or implement it themselves.

How SSL works

SSL provides a security "handshake" that is used to initiate the TCP/IP connection. This handshake results in the client and server agreeing on the level of security they will use, and fulfills any authentication requirements for the connection. Thereafter, SSL's only role is to encrypt and decrypt the bytestream of the application protocol being used (for example, HTTP, NNTP, or Telnet). This means that all the information in both the HTTP request and the HTTP response are fully encrypted, including the URL the client is requesting, any submitted form contents (including things like credit card numbers), any HTTP access authorization information (usernames and passwords), and all the data returned from the server to the client.

The Netscape Navigator includes embedded Certificate Authority (CA) keys for certain CAs, including Netscape Communications ' test CAs. As new CAs come online, they will embed their keys as well. These embedded keys allow the browser to verify the legitimacy of arbitrary servers. The Document Information dialog allow one to inspect both the identity of a given server as well as the identity of the CA that issued the server its certificate. SSL requires servers to have certificates issued by a Certificate Authority; and the Netscape Commerce Server includes a mechanism to easily acquire such a certificate.

Because HTTP+SSL (or "https") and HTTP are different protocols and typically reside on different ports (443 and 80, respectively), the same server system can run both secure and insecure HTTP servers simultaneously. This means a server can provide some information to all users using no security, and other information only securely, if the webmaster so choose. For instance, "storefront" and merchandise catalog could be insecure, while ordering and payment documents and forms could be secure.

Secure HTTP (S-HTTP)

Secure HTTP is an extension of HTTP that provides a variety of security enhancements for the Web:

- **Authentication** lets clients make sure they are communicating with the correct server, and servers make sure they are communicating with the correct client. Authentication is performed using digital certificates issued by Certifying Authorities such as VeriSign.
- **Encryption** makes data transferred over the network unintelligible to eavesdroppers.
- **Digital Signatures** provide two benefits. First, they verify that data transferred over the network was not changed en route. Second, they provide non-repudiation: The receiver of data can prove to a third party that the sender really sent the data.
- **Secure HTTP** is extremely flexible; it allows each application to configure the amount of security required. A transmission from client to server or server to client can be signed, encrypted, both, or neither.

Secure NCSA Mosaic and Spry SafetyWEB are two browsers that support the S-HTTP protocol.

Comparison between SSL and S-HTTP

Two comparisons of the S-HTTP and SSL protocols are possible: inherent capabilities and current implementations. Keep in mind that network security is a hot topic so developments continue at a rapid pace.

Because S-HTTP is an application-level protocol, it can provide non-repudiation of individual requests or responses through digital signatures. SSL is a lower-level protocol and does not have this capability.

S-HTTP, being an application-level protocol, is able to work with firewalls. SSL's transport-level encryption, in contrast, hides the application-level protocol from firewalls. A firewall that relays an SSL connection has no idea what data is being passed back and forth over the connection.

S-HTTP is more flexible than SSL, in that an application can configure the level of security it needs. Encryption and digital signatures can be expensive to compute, so in principle this flexibility can allow a server to handle more connections or respond more quickly. The other side of this trade-off is that SSL, a lower-level protocol, may be easier to optimize: It encrypts more bytes, but might balance this out with a lower cost per byte. Since SSL has fewer options it should be easier to set up and administer than S-HTTP.

Current SSL protocols and implementations are limited in a variety of ways relating to the handling of digital certificates; S-HTTP implementations are considerably more flexible. These SSL limitations are not inherent in the nature of SSL and will likely be removed as proves necessary.

Internet Protocol Security (IPSEC)

The IPSEC working group charter

IPSEC is not a general security architecture for the Internet, but is instead focused on the IP layer. IPSEC working group will develop mechanisms to protect client protocols of IP. A security protocol in the network layer will be developed to provide cryptographic security services that will flexibly support combinations of authentication, integrity, access control, and confidentiality. The protocol formats for the IP Security Protocol (IPSP) will be independent of the cryptographic algorithm. The preliminary goals will specifically pursue host-to-host security followed by subnet-to-subnet and host-to-subnet topologies.

Protocol and cryptographic techniques will also be developed to support the key management requirements of the network layer security. The key management will be specified as an application layer protocol that is independent of the lower layer security protocol. The protocol will initially support public key-based techniques. Flexibility in the protocol will allow eventual support of Key Distribution Center (KDC - such as Kerberos) and manual distribution approaches.

The IP authentication header (AH)

The IP Authentication Header is designed to provide integrity and authentication without confidentiality to IP datagrams. The lack of confidentiality ensures that implementations of the Authentication Header will be widely available on the Internet, even in locations where the export, import, or use of encryption to provide confidentiality is regulated. The Authentication Header supports security between two or more hosts implementing AH, between two or more gateways implementing AH, and between a host or gateway implementing AH and a set of hosts or gateways. A security gateway is a system which acts as the communications gateway between external untrusted systems and trusted hosts on their own subnetwork. It also provides security services for the trusted hosts when they communicate with the external untrusted systems. A trusted subnetwork contains hosts and routers that trust each other not to engage in active or passive attacks and trust that the underlying communications channel (e.g., an Ethernet) isn't being attacked.

The IP encapsulating security payload (ESP)

The IP Encapsulating Security Payload (ESP) is designed to provide integrity, authentication, and confidentiality to IP datagrams. The ESP supports security between two or more hosts implementing ESP, between two or more gateways implementing ESP, and between a host or gateway implementing ESP and a set of hosts and/or gateways. A security gateway is a system which acts as the communications gateway between external untrusted systems and trusted hosts on their own subnetwork and provides security services for the trusted hosts when they communicate with external untrusted systems. A trusted subnetwork contains hosts and routers that trust each other not to engage in active or passive attacks and trust that the underlying communications channel (e.g., an Ethernet) isn't being attacked. Trusted systems always should be trustworthy, but in practice they often are not trustworthy.

Gateway-to-gateway encryption is most valuable for building private virtual networks across an untrusted backbone such as the Internet. It does this by excluding outsiders. As such, it is often not a substitute for host-to-host encryption, and indeed the two can be and often should be used together.

For more information

You can see the IPSEC Working Group Charter at: ftp://ftp.ripe.net/ietf/ipsec/ipsec-charter.txt with all goals and milestones for this protocol, and/or see the Request For Comments (RFC) on the Security Architecture for the Internet Protocol at: http://www.cis.ohio-state.edu/htbin/rfc/rfc1825.

3.4.2 How can these protocols contribute to transaction security?

Kerberos

It allows entities communicating over networks to prove their identity to each other while preventing eavesdropping or replay attacks. It also provides for data stream integrity (detection of modification) and secrecy (preventing unauthorized reading) using cryptography systems such as DES.

It is important to realize that Kerberos is a one-trick pony. It provides for mutual authentication and secure communication between principals on an open network by manufacturing secret keys for any requestor and providing a mechanism for these secret keys to be safely propagated through the network. Kerberos does not, per se, provide for authorization or accounting, although applications that wish to can use their secret keys to perform those functions securely. Kerberos also does not provide password validation for individual workstations unless care is taken.

It is also important to understand that using Kerberos on time-sharing machines greatly weakens its protections, since user's tickets are then only as secure as the 'root' account (read: not very). Furthermore, dumb terminals and most X terminals do not understand the Kerberos protocol and as a result their cable connections remain insecure.

PEM

As far as I know, TIS/PEM and RIPEM are not used (or, at least, widely used) for securing electronic commerce transactions, but only for personal email messages. However, we may assume that some vendors will support PEM encrypted messages for their transaction processes if a significant number of customers desire so.

PGP

Some softwares developers have implemented a system by which NCSA Mosaic client and/or server call external PGP program which encrypt and decrypt their communications and thus provide secure communications between the server and the client, and ensure that a user is who he/she says he/she is.

However, the most "natural" mean to securely communicate with PGP is by exchanging encrypted email messages, since this can be performed an all platforms.

In both cases, the merchant can register his public key on a server such as SLED (http://www.four11.com) where the customer will be able to retrieve it. The merchant will have to provide the server with evidence (by another mean than the Internet) that he/she claims who he/she actually is, so that the customer will be able to trust his/her public key.

Several electronic storefronts, such as NetMarket, support purchase orders or customer sensitive information encrypted with PGP.

SSL

We can enter our credit card number on a secure (https) browser (we may assume that other browsers than Netscape Navigator will support SSL in a near future) form and transmit the form over the Internet to a secure Commerce Server without risk of an intermediary obtaining our credit card information. The security features offered by SSL compliant technology protect commercial transactions, as well as all other communications, from misappropriation and fraud that could otherwise occur as information passes through Internet computers.

Secure communications does not eliminate all of an Internet user's concerns. For example, the end-user must be willing to trust the server administrator with his or her credit card number before he or she enter into a commercial transaction. Security technology secures the routes of Internet communication; but does not protect from unreputable or careless people with whom we might choose to do business.

The situation is analogous to telling someone our credit card number over the telephone. We may be secure in knowing that no one has overheard our conversation (privacy) and that the person on the line works for the company we wish to buy from (authentication), but we must also be willing to trust the person and the company.

Server administrators must take additional precautions to prevent security breaches. To protect information, they must maintain physical security of their server computers and control access to software passwords and private keys.

A lot of electronic storefronts, such as MarketNet, now support SSL mechanism.

S-HTTP

S-HTTP specifies ways to add all the secure features needed in a transaction to the conversations that occur over the Internet between World Wide Web clients and servers. S-HTTP does not say how these elements will appear to users or what elements should be used with what transactions; that is up to the developers of the clients and servers and their customers.

However, a few companies like OpenMarket and NCSA have already released Web secure servers and clients that support S-HTTP, Terisa proposes toolkits to develop such systems, and CyberCash has developed a payment system based on this protocol.

IPSEC

The primary objective of the IPSEC work is to ensure that IPv4 and IPv6 will have solid cryptographic security mechanisms available to users who desire security. These mechanisms are designed to avoid adverse impacts on Internet users who do not employ these security mechanisms for their traffic. These mechanisms are intended to be algorithm-independent so that the cryptographic algorithms can be altered without affecting the other parts of the implementation. These security mechanisms should be useful in enforcing a variety of security policies.

As far as I know, the IPSEC is not yet implemented: this protocol is in a stage of "work in progress", but it may be worth to keep looking at what it will become.

3.5 What a Smart Card is Intended for – A Typical Example: the Authentication Issue

3.5.1 Smart card applications

Prepaid smart cards

Why should we use prepaid smart cards?

- The telephone card is just one example of the smart card being used as a prepaid card.
- Prepaid cards generate a considerable advantage in extra cash flow.
- Prepaid cards are good advertising supports because users of prepaid cards for a particular service are a well defined population. A card user will usually carry the card on him for weeks or months and use it several times before the total value is used.
- Prepaid cards can be used in vending machines, car washes, gaming machines, automatic laundries, photocopiers and other areas where coins are used.

- Card operated payphones reduce theft, and increase convenience for users.
- Of the different card technologies such as smart cards, magnetic cards and optical cards, smart cards have been selected by the telecommunication industry because they give the best global solution in terms of equipment and card cost, security and reliability.
- The prepaid card is the simplest form of smart card.
- Telephones which are able to use smart cards open up a totally new range of possibilities, using more sophisticated cards for services like home billing from public payphones and memorized phone numbers.

One example: making public telephones more convenient

In 1984 all French payphones were coin operated. In 1985, after having rejected alternative technologies such as magnetic or optical cards on the grounds of security and global cost, France Telecom decided to start installing smart card payphones. In 1994 nearly all public payphones use prepaid smart cards. 100 million phone cards were sold in 1993 and 115 million are expected to be sold in 1994. This is currently the largest smart card project in the world.

Any person wishing to use a payphone can buy phone cards from tobacco stores, post offices and some newspaper kiosks. The cards hold either 50 units (40 Francs = $6.7) or 120 units (96 Francs = $16). The duration of one unit depends on the distance and the time of day. A display on the telephone indicates how many units are left on the card. When no units are left a new card must be used.

The use of smart cards for payphones has had dramatic effects on the image of the payphone. In 1985, 44,000 payphones were broken into. As these phones were systematically replaced by the new card payphones, the motive was removed and the number of phones broken into fell to a mere 900 in 1989 and about zero in 1994.

Freed from the obligation of having the right small change at hand, people telephoned 50% longer from card phones. As card operated public phones ceased to be a target for theft, they were less often out of order. A survey of payphone users showed that the satisfaction index for card phones is 20 points higher than for coin phones, despite having to pay in advance for calls from a card operated public phone.

Now customers will also be given the possibility to store pre-recorded phone numbers or abbreviated dialing numbers and use their card to pay other services.

Phonecards are used in two additional ways: as a novel advertising support and as a collectors item just like stamps.

At the beginning of 1994, 66 operators in 60 countries are already using smart cards for payphones and, looking at the uptake trend, it is expected that over 100 countries will have adopted the smart card for payphones before the end of the century. The major smart card users apart from France are Germany and the Latin American countries. All systems are similar to the French system using a prepaid card. The world market for smart payphone cards is expected to leap from an estimated 250 million cards in 1993 to an estimated 1 billion cards in the year 2000.

For payphones, smart cards have proved to be cost effective: they can bring new services and are readily accepted by users. The same benefits can be found in many other applications where coin operated machines have suffered from the cost and loss of service through vandalism, coin collection, and low reliability. Smart cards are already being used for gas metering, car parking, car washing, arcade games, and vending machines.

If you put a coin into a machine today, you will most probably use a smart card for the same operation in five years time!

Why are smart cards replacing coins in the world payphone market?

There are several reasons:

- Reduced theft and maintenance cost.
- The electronics necessary to read a smart card are lower cost and more reliable than for any other card technology.
- Smart card payphones designed to accept simple prepaid smart cards allow the use of more sophisticated smart cards.
- Smart cards in payphones benefit users and operators.

Electronic money

Replacing coins and bank notes

In Singapore the CashCard, is starting to replace coins and bank notes for everyday purchases. The CashCard is a smart card which acts as an Electronic Purse, that is to say that it holds electronic money. As the national electronic purse program unfolds, it is predicted that most Singapore citizens will use the CashCard by 1997. They will then be able to pay with electronic cash for most everyday services, such as cinemas, car parks, tourist attractions, petrol, telephones, vending machines, sports, retail outlets, fast food, in fact virtually everywhere where cash is used today. The CashCard is available in several different initial values, S$20, S$50 and S$100 (the equivalent of $12, $31 or $62 in $US). The card is not linked to a bank account, however the possibility of combining the electronic purse card with an ATM cash withdrawal card or credit card is quite possible and will most probably be used.

Why use electronic money?

The main reasons for using electronic money are consumer convenience, reducing the cost and work associated with handling cash, and avoiding fraud. It becomes quite clear when examining these reasons that they are in fact quite closely linked.

Bringing more convenience for the consumer

Using a smart card to pay with electronic cash is faster and easier than using coins and bank notes; no more fumbling for that last little coin that you are certain that you saw in the bottom of your pocket (while 14 frustrated shoppers wait in a line behind you). Paying with electronic cash in most systems simply means checking that the amount on the till is the price you should be paying, and inserting the electronic purse in the terminal. Some systems may also have a button to push. Consumer research conducted in electronic purse pilot schemes show that consumer acceptance is very high.

Reducing the cost of handling cash

The cash that has been counted and recounted every time that it enters or leaves the till for each sale during the day, is then counted and recounted and recorded to be taken to the bank, where it is counted again. With electronic money, there is no change to be given, the sales terminal automatically counts the taking and wires it directly to the bank.

Fighting fraud

The smart card, in association with a well designed security system, has been demonstrated to be an effective weapon against payment fraud. The electronic purse is also an effective protection against counterfeit coins and bank notes.

The electronic purse: a worldwide phenomenon

Singapore is not alone in moving towards an electronic purse system. At the start of 1994 over 20 projects had been identified spread over the 5 continents, and to name a few, we can mention Mondex (which is still at an experimental stage in England and in the US, but which aims to become a worldwide electronic cash and payment systems standard), or Banksys –in Belgium– and Interpay –in Holland– Cash Cards (which claims to be the world's largest live electronic purse system). The technology is available, consumer acceptance has been demonstrated for card payment in the limited area of public payphones, and the benefits for consumers, merchants and banks are beginning to be well understood. By the year 2000, many working systems will be in place and perhaps in countries like Singapore, where an early adoption of the electronic purse has been made, electronic cash could well be more frequently used than today's coins and bank notes.

3.5.2 What is a smart card?

Presentation

A smart card is a credit card sized plastic card with a special type of integrated circuit embedded in it. The integrated circuit holds information in electronic form and controls who uses this information and how.

The type of information held, and the way this information is controlled, gives rise to the infinite variety of smart card uses. The plastic card has words and pictures printed on it which can be understood by the card holder and people he meets. The integrated circuit holds information in electronic form that can be easily, securely and accurately accessed by all sorts of electronic data processing equipment. This unique combination means that the smart card is a bridge between man and machine.

The plastic card

The size of the card is determined by an international standard (ISO 7816). This standard also defines the physical characteristics of the plastic, such as temperature tolerance and flexibility, the position of the electrical contacts and their functions, and how the integrated circuit talks to the outside world.

There are several types of plastic used for smart cards. The main types are PVC and ABS. PVC can be embossed but is not recyclable. ABS cannot be embossed but is recyclable. Within 5 years there are likely to be 1 billion smart cards used per year. So recycling cards will be important.

Access control to the information

Smart cards are smart because they control access to the information they hold. Nobody can add units to a telephone card (not even the telephone company). Only a doctor equipped with the correct card can read a patient's medical record.

The basic mechanisms of access control are described below, and illustrated using the examples mentioned in the first part of this document.

There are two types of access control:

Who can access the information?
Everybody, the card holder or a specific third party?

How information can be accessed?
Can it be read, added to or even modified or erased?

Who can access the information?
Examples:

- I should be the only person able to use my credit card.
- Only retailers who are affiliated to the credit card scheme should be able to accept it.

- My regular doctor should be the only one who can add information to my medical record.
- Other doctors should be able to consult it.
- Emergency services should be able to know my blood group and address.

Some smart cards will only allow access to information if they are given the correct password. If you possess the password you can have access to the information. A password may be a number that the card holder has to remember, it may be hidden in another smart card that only authorized people possess, or it may just be inside the specific telephone or sales terminal used with the smart card. Most sophisticated smart cards have the ability to recognize several passwords.

This means that in the case of a health card, the patient, the doctor, the hospital administration, and the health insurance organization could each have their own password giving them different access rights to sections of information held in a patient's health card.

Everybody
In some cases no password is used. So everybody can access the information. For example, in the health card, emergency medical information such as blood group and the patient's name are not protected by a password.

Card Holder Only
The most common form of password for card holders is a PIN (Personal Identification Number), a 4 or 5 digit number to be typed in at a key pad. The meaning of typing in a PIN is "This card is mine, I give my permission to debit my account or access my medical file." In theory, if a smart card is lost or stolen and it is protected by a PIN, it cannot be used. If somebody tries to guess the PIN, most smart cards will just lock-up after 3 wrong attempts. PINs have two disadvantages, people can forget them, and it is possible to 'steal ' a PIN by watching carefully when a card holder uses it. In the next few years better passwords for card holders which do not have these disadvantages are likely to appear.

Third Party Only
For example, when a doctor needs to consult a patient's health card, he needs a password, this password is hidden in his doctor's card. As only correctly trained doctors are issued with these cards, only they are able to consult patients' cards.

How can the information be accessed?

Read only
Some information in the smart card can only be read. The information is fixed like a book, where information cannot be added, modified or erased.

Some of the information in a prepaid telephone card is like this. This information contains a unique number for each card, the number of units in the new card, and an indication of who manufactured it. This information can be easily read, but cannot be modified in any way.

Add only
Some information in the smart card can only be added to but no erased, like engraving information in a stone. Information cannot be modified or erased, and information can only be added while there is room. If information can be added it can also be read. Units are used in a telephone card by adding information, like ticking off boxes. When all the boxes are ticked, no more units are left. Ticks can only be added and never erased, so nobody can refill the card.

Modify or erase
Information in the smart cards can be modified or erased like writing on a blackboard. This also implies that information can be added while there is room.

No access
Some information in a smart card can never be accessed. While this might appear rather useless, it is necessary for smart cards to be able to carry secrets the can never be revealed. The use of these secrets will be explained in "When passwords are not enough." With a careful use of passwords, it is possible to avoid incorrect use of information, either by error or by intent. It is possible to summarize the information in a smart card, and the access control of this information (who and how) in a table. Take the example of a medical card in the table above. Information can be freely distributed in a smart card, and yet only accessed by the people who are authorized.

When passwords are not enough

A smart card can restrict the use of information to an authorized person with a password. However if this information is then transmitted by radio or telephone additional protection is necessary.

One form of protection is ciphering, which is like translating the information into some unknown foreign language. Imagine two Martians talking loudly in a London bar, everyone can hear, but only the Martians can understand and use the information.

Some smart cards are capable of ciphering and deciphering (translating back to a form that can be easily understood), so that the information stored in them can be transmitted without compromising confidentiality.

The ciphering scheme usually used is actually far more complex than the situation of our two Martian friends in London. Smart cards can cipher into billions and billions of "foreign languages", and choose a different language at random every time they start a conversation. This makes eaves-dropping virtually impossible.

When a smart card and a computer system need to exchange confidential information over the telephone, they start by playing a complicated number game. The rules for this game are only known to genuine cards and genuine computers. It is impossible to play the game if the rules are not known, so any attempts to break into the system are easily detected and eliminated.

If the game is correctly played, the genuine card knows that it is talking to a genuine computer and visa versa. This is known as authentication and some serious information exchange can now take place.

During a successful authentication game one of the billions of possible "foreign languages" is selected, and only the players can know what it is. Information can then be transmitted in a language which is unknown to anyone else, even if they listened to the whole game.

The ability of smart cards to enable totally secure communication is used in the GSM radio phone network. Even when the network covers the whole of Europe, only radio phones with a valid subscription card will be able to use it.

How smart is a smart card?

Some smart cards are smarter than others. The simplest cards, like the payphone cards, are smart enough to be virtually impossible to copy or falsify, but offer no protection in case of loss. If you lose your payphone card, anyone can use it. Other smart cards have one password to restrict its use to one person or machine, and the most sophisticated cards manage several passwords and can use authentication and ciphering techniques to combine total freedom with total security.

The simplest smart cards can cost less than $1 in large quantities, whereas the most sophisticated can cost in the range of $6-10 according to how much information they can hold.

3.5.3 How can your company use smart card systems?

Two-factor authentication process

The major weakness in Kerberos based security is the reliance on passwords for the authentication of the end user. Passwords need to be reasonably easy to remember, or users will write them down, introducing yet another security vulnerability. Yet easy-to-remember passwords use only a small fraction of the key space, and are subject to dictionary attacks.

Smart cards can dramatically reduce password exposure by supporting a 2-factor, 2-step authentication mechanism (See, for example, J. Kohl and C. Neumann, "The Kerberos Network Authentication Service", Internet Engineering Task Force, Networking Group, Internet Draft RFC 1510, Sep 1993 or Marjan Krajewski, Jr., "Smart Card Augmentation of Kerberos", Proceedings of the Privacy and Security Research Group Workshop on Network and Distributed System Security, February 11-12, 1993, San Diego, CA, pp.119-123.).

With passwords only, a user must know one piece of secret information — the valid password for a principal name — in order to log in to a company cell. There is one step involved in authentication — supply a login name and its password to the system.

With smart cards, a user must possess a unique physical token — the smart card — as well as knowing one piece of secret information — a password associated with the card. Knowledge of the card password is of no use without possession of the card itself. Possession of the card is of no use without knowledge of the associated password.

With passwords only, there is one step to the authentication process — the user principal is authenticated to the computer. With smart cards, there is an additional step involved — before being authenticated to the computer, the end user must be authenticated to the smart card. The first step depends on knowing the correct password. The second step depends on possessing the correct smart card — the one containing the user's long-term key.

Therefore it is hereby proposed using the smart card's secure file system to store the user's long-term key. To access the key and log in to the computer, a user must possess the card and must also know the password required to access the key file on the card.

The details of this solution are contained in the following sections.

Card requirements

To achieve the additional security provided by this solution, cards must meet certain minimum physical and functional requirements.

Physical access requirements

The card must conform to the accepted International Standards ISO 7816-1, ISO 7816-2, and ISO 7816-3.

The card must be tamper-resistant. Common tamper-resistance features include: chip enclosed in a hard, opaque, tamper-evident sealing coating over a passivation layer; circuitry to disable card output if the passivation layer is removed or if temperature, voltage, or frequency fluctuates outside the specified operating range.

Cards with higher levels of tamper-resistance generally have a higher cost. The card issuer must evaluate with the card vendor the appropriate level of tamper-resistance for the issuer's application.

Failure to provide adequate tamper-resistance makes the long-term key stored on the card vulnerable to physical attacks, reducing security when a card is lost or stolen.

Access control requirements

The card file system must provide access controls that support the use of a password to protect access to the file containing the key. Passwords of up to 8 bytes in length must be supported. Without presentation to the card operating system of the password associated with the key file, there should be no access to the key stored in the card file system. Failure to conform to this specification negates the additional security that smart cards offer.

The card must support detection of sequential invalid access attempts (a request for a protected card operation without supplying the correct password for the operation). Upon detecting N sequential invalid access attempts on any operation (where N may be configurable), the card must deny all further access to all operations except for possibly one "unlock/erase" operation. This one operation, if supported, must erase the key stored in the card file system before restoring normal access to the card. Failure to conform to this specification leaves a lost or stolen card vulnerable to off-line dictionary password-guessing attacks. Accounts associated with lost or stolen cards are subject to compromise until an administrator invalidates the account in the Security Registry.

The card file system must support access controls that allow only the user to modify the password associated with the key file (with the possible exception described in the next paragraph).

An interesting implementation would have a way of providing administrative access to unlock/erase a locked card (see above), set a new key on the card and in the Security Registry, set a new card access password, etc., but that is a vendor-issue not required for the basic functionality specified here.

3.5.4 The future

Smart cards are just one particular way of using these special integrated circuits. Historically it came about because credit cards already existed, and the first smart card developments were for banking.

What is important about smart cards is that they are everyday objects that people can carry in their wallets, purses or pockets, and that they have the capacity to retain and protect information stored in electronic form.

There are already other smart objects in use. Smart keys are being used for pay-TV subscriptions. The smartness of smart cards comes from the integrated circuit embedded in the plastic card. The same electronic function could be performed by similar circuits embedded in other everyday objects, key rings, watches, glasses, rings, or earrings.

The way that people give a password with a PIN is not yet ideal. Biometric techniques will soon mean that a person can be reliably identified by his hand, finger print, the color of his eyes or the sound of his voice.

Soon it will be possible to hide electronic information in smart cards and give authorization for this information to be used with a spoken word or the touch of a hand.

The computer revolution has generated many advantages, but also introduced much form filling, errors and impersonal service. Smart cards along with other technologies such as biometrics, radio telephones and home computing, will give us the advantages without the disadvantages.

Smart cards are a relatively new technology that already effects the everyday lives of millions of people. All the signs are that this is just the beginning of something much larger that will influence the way that we shop, see the doctor, use the telephone and enjoy leisure.

Modern society needs an enormous amount of information. Computers give us the means to process this information. Smart cards give us a way of humanizing the way we handle this information. Smart cards give every individual the possibility to control how his personal information is used.

Society will have to create laws which limit the uses of smart card technology to those which benefit the individual, and increase the quality of life.

Chapter 4

Payment and Transaction Issues

4.1 The Rush Is on to Make Net Commerce Happen

4.1.1 Introduction

The Internet is capable of becoming one of the most powerful business and financial tools of our time–if the process becomes simple and complete enough. Some businesses launched on the Net sell goods and services, but the buyer must send a check via regular mail to complete transactions. Although this is a viable method for some applications, a richer suite of payment mechanisms is needed for the broader range of business that will be conducted on the Internet.

Like the traditional off-line market, buyers and sellers want a variety of options such as credit cards, debit cards, and cash–or in this case, digital cash. To meet this demand, a host of new companies are vying to establish their technologies as the means for conducting secure Internet commerce.

Online commerce has existed in the form of EDI (electronic document interchange) since the mid-1980s. EDI applications are transaction-oriented, ranging from government RFPs (Requests for Proposals) to purchase orders. EDI is in use throughout several major industries to automate and simplify the otherwise paper-intensive process of doing business. While performed today via secure Value-Added Networks, it is questionable to what extent EDI applications will translate to the Internet, especially since many of the transactions involve large sums of money. A convergence is on the horizon nevertheless.

4.1.2 What is available? What is coming?

CommerceNet, introduced in March 1994, is an industry consortium organized into working groups. Its members comprise an intriguing mix of Silicon Valley semiconductor manufacturers, software companies, financial institutions, information providers, and businesses of varying size. The consortium is a hothouse for electronic commerce and a great source of information, especially for businesses. As of October 1994, CommerceNet had more than 50 members and 15 subscribers (small to midsize companies).

Cisco Systems and Premenos Corp. recently announced a secure electronic commerce experiment that should be readyin early 1995. The system uses RSA Data Security Inc.'s public-key cryptography.

CyberCash Inc. is a well-funded start-up combining the talents of Dan Lynch and William. Lynch converted the ARPAnet–the predecessor to the Internet–to TCP/IP back in 1983 and has continually demonstrated his abilities as an Internet visionary. Melton is the founder of Verifone, a company that made real-time credit-card verification a reality for retailers. Together they are working on a model for Internet commerce that will accept digital cash, and that already deals with credit cards, or debit cards. Expect to see CyberCash make its official debut in the spring of 1995.

DigiCash was founded by Dr. David Chaum. The company pioneered E-Cash, a software-only product that allows digital cash transfer over the Internet. E-Cash was developed for smart cards and electronic wallets–used in tollbridges, for example. DigiCash will supply the technology through other organizations, including Encyclopedia Britannica and NCSA (National Center for Supercomputing Applications), which will release the products under various cooperation and trial programs. The user software, which allows both paying and receiving payment, will be distributed free of charge. DigiCash will encrypt transactions and ensure anonymity.

First Virtual Holdings Inc. is attempting to become the first Internet merchant banking system. The company brings together the Hollywood business savvy of Lee Stein, First Virtual's president and CEO, and three of the leading Internet gurus: Dr. Marshall T. Rose, Dr. Nathaniel S. Borenstein, and Einar Stefferud.

Stein, an attorney and accountant, provides financial advisory services to the music and entertainment industries. First Virtual principal Rose, one of the Internet wunderkinds, has written the book (actually several of them) on SNMP (Simple Network Management Protocol) and network management. Borenstein, the chief scientist, is the primary author of the MIME (Multipurpose Internet Mail Extensions) protocol that allows transfer of anything digital over the Internet, including multimedia documents. Stefferud, probably the world's foremost expert on e-mail, has the role of chief visionary.

Introduced and online in October 1994, the First Virtual payment system allows any Internet user to initiate retail transactions via e-mail. The system aims to help maintain Internet culture by allowing buyers to evaluate information prior to committing to a purchase, which is a concept similar to shareware. No sensitive data is exchanged online in First Virtual's system. Users set up accounts by telephone or fax, and credit-card payments are processed off-line.

Registration fees are $2 for consumers and $10 for sellers. For each transaction, sellers pay a 29-cent fee plus two percent of the transaction price and a $1 processing fee each time a payment is made to their account. Initial First Virtual merchants include Internet Resources Group, Internet Multicasting Service, and an Internet newsgroup filtering service founded by David Farber, an Internet luminary. The Department of Commerce is also considering adopting the system to offset distribution costs of electronic documents.

The Internet Shopping Network was founded by its president, Randy Adams, who formerly was CEO and chairman of IIAT Inc., a publicly held international education and training organization. The company, which was recently acquired by the Home Shopping Network, is an online microcomputer software and hardware superstore available to anyone with direct access to the Internet. Internet Shopping Network distributes more than 20,000 computer software and hardware products, plus InfoWorld, a computer news weekly. Membership is free, and prospective members need only pre-register with an approved Visa or MasterCard.

Microsoft is also getting into the act. It recently announced a venture with Visa International that will enable Windows users to make secure credit-card purchases over the Internet using software co-developed by the two companies. The software will be based on RSA Data Security's encryption technology and Visa's VisaNet system will be used to authenticate payments. According to Chairman Bill Gates 'Fall Comdex '94 speech, a 'Wallet PC' that handles financial transactions and other consumer retail interactions may also be part of our future.

Netscape Communications Corp., which provides an enhanced commercial World-Wide Web browser, announced that it will support credit-card purchases via the Internet through credit-card processor First Data Corp. The Netscape system uses RSA's encryption technology to scramble sensitive data. Merchants can purchase Netsite server software for $5,000, and buyers can access the system through free client software.

NetCash is another secure Internet project underway. Developed at the Information Sciences Institute (ISI) of the University of Southern California, it is a model for electronic currency that can be anonymously transferred over an unsecure network without the use of firewalls. NetCash allows service providers and users to select payment mechanisms with varying degrees of anonymity.

Unlike most electronic cash systems in which only one financial institution generates and accepts its own digital cash, NetCash is designed to support the transfer of electronic cash between currency servers. At this time, NetCash has not yet been made commercially available.

Open Market is another secure Internet transactions venture. The system supports secure payment, real-time credit-card authorization, account statements, administrative interfaces for storefront management, and a customer feedback mechanism. The company offers 'document fingerprinting, ' which automatically and uniquely numbers documents so that fraudulent distribution can be tracked.

Open Market's StoreBuilder kit ranges in price from $300 to $1,500 for setup and $50 to $300 in monthly fees, in addition to transaction fees and additional storage fees. Buyers can set up an account by sending credit-card info over the Net or by phone and fax. Open Market uses password and encryption schemes to secure its accounts.

Finally, Terisa Systems is a joint venture of EIT (Enterprise Integration Technologies) and RSA Data Security that is dedicated to producing a secure Mosaic product suite for Internet commercial transactions (see Section 4.2).

4.1.3 Connecting financial systems

Electronic commerce players are teaming up with banks, credit-card, EDI, and software companies, and more to create an infrastructure that can support business over the Internet. EDS, Visa, and First USA — which have teamed up with First Virtual — intend to stake a claim in the online transaction business. In fact, participation from this type of company is necessary so that people can buy goods and services in real time.

To make Internet commerce viable on a mass scale, it is necessary to connect the world's financial systems. Given that security is imperative to these organizations, why are they willing to risk doing business on the Internet?

CyberCash and First Virtual are not directly connecting the Internet to the global financial networks. CyberCash will use an intermediate highly secure server that is connected to both networks. First Virtual's system maintains security without having to use the encryption or signature technologies that the other systems employ. The system separates business on the Internet, which is insecure, from the bank processing system through firewalls and off-line batch processing handled by EDS, one of the world's leaders in banking transactions.

4.1.4 No numbers, no crime

Credit cards are currently the most popular way to buy something on the Internet. This method is too dangerous to be practical, according to First Virtual's Rose and Borenstein.

"It's very easy to write a program that looks for a sequence of numbers beginning with a known credit-card prefix," Rose said. That is why First Virtual never requires a potential information merchant or consumer to send credit-card or checking-account numbers over the Internet. Instead, First Virtual customers enter these numbers using an automated phone system to set up a First Virtual account.

The Internet Shopping Network also keeps credit-card information off the Net. Members enter their credit-card numbers using phone or fax when they first establish an account. Open Market's system also includes this option.

CyberCash deals with the problem using encryption. Its customers will have a public key and a private key and must use both to complete a transaction. The company will also use authentication and digital signatures for added security.

Debit cards can be used for Internet-based financial transactions today because the technology is in wide use among banks and retail points of purchase. It is likely that the percentage of businesses accepting debit cards for Internet purchases will be higher than the existing number in the off-line marketplace because no special hardware is required.

At the October World-Wide Web conference in Geneva, DigiCash CEO Chaum demonstrated how anyone with a personal computer could transfer DigiCash's E-Cash to any other workstation via e-mail on the Internet. The E-Cash system was shown integrated with Mosaic.

Digital cash is the binary equivalent to currency. It differs from credit and debit cards because it enables anonymous transactions.

"How much money are you willing to lose?" asked CyberCash's Lynch. "You don't, and I don't, walk around with $10,000 in your wallet. Usually people carry around what they are willing to lose. The same will be true for digital cash."

According to Chaum, DigiCash has a solution to the problem. "E-Cash is as secure as any government network because it employs the same underlying security mechanisms," he said. E-Cash ensures the high security required for electronic networks exclusively using 'blind signature' technology, an innovation in public-key cryptography.

4.1.5 Internet-savvy business

CyberCash is targeting banks and businesses, the Internet Shopping Network and Open Market are targeting businesses and consumers, and First Virtual is targeting all these and the government. In any case, it appears that there will be a solution for just about everyone.

Over the past 18 months or so, organizations have been setting up WWW and Gopher servers to distribute product and company information. Businesses clever enough to have pointers strategically located to and from other points on the Internet have taken the first steps in doing real business on the Internet, but that is just the beginning.

For businesses, the Internet represents the cheapest new business opportunity in history because many of the high costs of doing business — shipping, warehousing, duplication, and distribution — are handled at a fraction of the cost using servers and the network.

That said, businesses will not drop existing products and channels of distribution just to go online. It is very conceivable, however, that many businesses will use the Internet as an alternative form of distribution because of its large reach and relatively low cost. Many more small entrepreneurial businesses also will likely emerge if the Internet proves to be a viable business alternative.

4.2 A Tool to Make Web Commerce Secure

4.2.1 Introduction

With the explosion of interest in the World-Wide Web, hundreds, if not thousands, of companies have rushed to get their home pages online to provide general information, advertise, and offer catalogs. Many of them would like to accept orders from customers online.

However, these firms generally have found the level of security in Internet applications woefully inadequate. Companies want tools that allow them to know with some certainty that customers are who they claim to be so that billing can be charged against the proper accounts.

Similarly, consumers are concerned about security and privacy of online transactions. They want to know that personal account information or credit-card numbers will be seen only by authorized eyes. And as vendors want to verify that a credit-card number is indeed a customer 's, consumers want assurance that a person who purports to represent an online service really belongs to that service.

4.2.2 A head for business

Several dozen companies are working on extensions to Internet protocols to provide these types of security features (see Section 4.1). One of the first firms to announce such an effort was Silicon Valley's Enterprise Integration Technologies (EIT). EIT is working to provide Internet tools that will allow retail firms and shoppers to conduct business online as well as let businesses handle online transactions with their suppliers.

One of EIT's endeavors is support of CommerceNet, a consortium of California's Silicon Valley companies. The purpose of CommerceNet is to develop Internet technologies that will secure commercial interaction among firms in the Valley. A specific focus is to help foster their use of the Internet for commerce among these firms and customers across the globe.

Marty Tenenbaum, EIT's president, offered an example of an individual customer who wants to buy a personal computer from a retailer in the Valley. The customer would scan the company's online catalog using a secure version of a World-Wide Web browser such as Mosaic; no special security features would be required while the customer was browsing the catalog. Should the customer decide to order a machine, he or she would fill out a form with the desired options for the PC and submit the purchase order to the Web server at the retailer.

This is where security features become important. Although many consumers have sent credit-card numbers over the Internet via e-mail and World-Wide Web forms, those who are knowledgeable about Internet security are concerned about the information being intercepted, either on local networks or on any of the intervening networks between the consumer's desktop and the vendor's site. Tenenbaum's vision calls for a secure Mosaic, which would use public-key cryptography to encrypt the consumer's order to keep sensitive credit-card information safe from prying eyes.

4.2.3 The safe sell

The use of a credit card is analogous to a consumer filling out an order form and sending it in a paper envelope via postal mail. The customer is reasonably certain that the order and credit-card number will arrive safely. The vendor relies on a separate step — authorization and credit balance checking via the traditional credit-card networks — before fulfilling the order.

Once the PC vendor has received the order and verified the validity of the customer's account, the vendor must ship the product. Many personal computer vendors custom assemble each PC as orders arrive. In Tenenbaum's plan, the PC vendor would also rely on the Internet as a way for retail vendors to send orders to its suppliers.

Orders for everything from the chassis to the motherboard to the monitor would be sent over the Internet. With computerized inventory control, such orders could be placed online automatically, with the retailer able to operate using just-in-time inventory methods.

In fact, Tenenbaum sees virtual corporations springing up as a result of these technologies. Retail firms might become so intimately linked to their suppliers–with order processing occurring so fluidly and rapidly–that the nature of the relationship between suppliers and retailers would be changedaltogether.

Encrypted credit-card transactions are only one type of business that companies would like to conduct on the Internet. In the everyday business world, signatures are an essential element–for signing purchase orders, contracts, and the like. Most companies will accept facsimile transmissions of some documents, but Internet e-mail is easily enough 'spoofed ' that its use is discouraged for these applications. Some sort of digital signature is required to verify the integrity of such documents.

The major elements of security required for commerce on the Internet, therefore, are the following:

- Authentication of individuals and companies so both parties can trust that the parties they are dealing with are who they claim to be.
- Encryption of digital documents so that sensitive information cannot be decoded even if the messages are intercepted.

- Digital signatures that verify the integrity of documents, so senders and recipients can trust that documents have not been altered in transit.

4.2.4 Terisa is born

To turn these desires into reality, EIT joined with RSA Data Security Inc., a leading firm in digital-security technology, to form a company called Terisa Systems. The venture's first goal is to create a suite of client and server tools called the SecureWeb Toolkit. The tools will be used by World-Wide Web client and server vendors to build products for sale to end users. Terisa is working with the National Center for Supercomputing Applications (NCSA) — the organization that brought Mosaic to the Net — to build Secure NCSA Mosaic, which is expected to be freely available on the Internet.

EIT and its partners have defined a standard called Secure Hypertext Transfer Protocol, or Secure HTTP, for use in tools like Secure NCSA Mosaic. The standard calls for extensions to the Web's basic HTTP protocol to support needed security features Secure HTTP relies on public-key encryption technology patented by RSA Data Security (see Section 3.3).

How can users expect to cope with the encryption and authentication process, which at first glance seems to be enormously complicated? In theory, most of the details of the process are hidden to the end user.

A customer buys some sort of secure Mosaic package from a seller of Internet products — perhaps from his or her Internet service provider. The package is preconfigured to work with a given public-key encryption certifying authority (which verifies that the alleged owner of a public key is genuine). The user installs the package on his or her computer, and begins using it to browse the Web. When the user encounters a product catalog, and begins taking advantage of the security features, most of the details are handled by software.

By analogy, consider how you use an automated teller machine. A complicated protocol has been defined by the banking industry to handle transactions, but the mechanics of that protocol are hidden from you. All you need to know is how to insert your card, type your personal identification number, and request transactions. Ideally, a Secure Mosaic tool will offer the same level of simplicity.

The Terisa Systems approach is by no means the only mechanism for securing Mosaic or for conducting commerce on the Internet. A number of firms have announced plans for Internet secure-transaction schemes.

4.2.5 The clang of the cash register

A New Hampshire startup, NetMarket, began offering online services in the summer of 1994, selling products such as compact discs and flowers over the Internet. It claims to have conducted the first secured Mosaic transaction on August 10, 1994, using a secure version of X Mosaic that allows users to place orders. NetMarket's secure Mosaic uses PGP as the encryption code, which the company licensed from ViaCrypt.

Of course, relatively few end users are running X/Windows on Unix workstations. "PGP is the solution we're doing now," said Dan Kohn, president of Net Market. "It's technology we happen to like because it's cheap for providers and free for consumers." Kohn is open to adopting other standards more likely to be deployed for the masses as they evolve. "The best model would be for firms offering online services such as ourselves to pay reasonable fees for server technology and for secure clients to be free."

Another company working on secure Web client technology is Netscape Communications Corp., formerly known as Mosaic Communications Corp. The company, whose principals include Jim Clark (the founder of Silicon Graphics) and Marc Andreessen (the co-author of the original X Mosaic from NCSA), is developing commercial Web client and server products.

Netscape's approach to securing Mosaic also relies on RSA Data Security's public-key scheme, but its protocol enhancements differ from those proposed by Terisa Systems. It proposes to implement a Secure Socket Library to build a system it claims will be bulletproof. Netscape has announced a partnership with First Data Card Services, which is the world's largest credit-card authorization firm. Other partners include three banks: First Interstate, Old Kent, and Norwest Card Services.

4.2.6 Money makes the world go round

Even when new Web protocols are standardized and software is available, there will be lag time before security-enhanced client software is installed on millions of desktops. Not all companies are waiting for these developments. Some are devising ways to provide a measure of security today. One such firm is Open Market, which is building on technology devised at the Massachusetts Institute of Technology.

Open Market provides a StoreBuilder toolkit for firms that want to conduct business on the Internet. Its scheme works with existing Web clients and uses a payment URL (Uniform Resource Locator) that encodes price and other data in the information sent to the client program.

When a consumer selects an item, the URL is sent to a payment server, which authenticates the user and processes the payment. The payment server then redirects the client program to use an access URL. The model supports ordering of information goods such as delivery of magazine or database excerpts on demand as well as traditional hard goods. Open Market's partners include the Lexis/Nexis service and Digital Equipment Corp.

New security technologies and alliances are being announced at a breathtaking pace. For example, Microsoft has announced an alliance with Visa International. Microsoft recently announced its own online information service (called Marvel) and purchased Intuit, the leading maker of PC personal finance software.

Other new ventures such as First Virtual aim to provide transaction services without a need for security-enhanced software (see Section 4.1). Purchases are authenticated off-line via a relatively simple e-mail confirmation instead of more complex schemes.

During 1995, we probably will see many more alliances out to provide secure Internet transaction technologies. A shakeout also is likely as the market settles on technologies and players within and outside the Internet context.

Advocates of commerce on the Internet maintain that the more complicated schemes are essential for the Net to replace conventional means of conducting business. To prevent a Tower of Babel of complex new schemes, the newly formed W3 Organization is trying to foster cooperation among technology vendors. Tim Berners-Lee, the chief architect of the World-Wide Web, has convened meetings of major players on behalf of the consortium.

Berners-Lee is optimistic that there will be some convergence. He also points out that HTTP was designed under the assumption that multiple schemes would coexist. According to Berners-Lee, the top layer of Secure HTTP transactions will negotiate the level of security and the exact scheme to be used. This will allow users and vendors to take advantage of future encryption schemes that have not yet been invented.

Chapter 5

The legal issue: Clipper's Still Kicking

5.1 Introduction

Although the Clipper chip controversy seems to have died down, the Clipper initiative is not dead. The U.S. government's attempt to legislate an encryption technology that gives it a backdoor to monitor all our transmissions still lives. That backdoor is Clipper.

The argument is as old as encryption. On one side, you have those in authority who, for the good of the state, want to be able to read or listen to anyone's communications. On the other, you have civil libertarians who insist that everyone has a right to privacy against unreasonable searches.

This conflict exploded into the Internet's consciousness in 1994 when Vice President Al Gore and FBI director Louis Freeh pushed for the use of the Clipper chip and the passage of the Digital Telephony bill. Most of the public furor was directed at the chip, but Clipper and the Digital Telephony bill work hand in glove with each other.

5.2 The Clipper Controversy

Before examining the current state of Clipper, let's re-examine the root of the controversy. Clipper is a computer chip that encodes digital voice and data communications using an encryption algorithm called Skipjack, which is a product of the National Security Agency (NSA). The chip is meant for use in telephones, fax machines, and modems.

What makes Clipper different from other encryption devices, and what has drawn the ire of many citizens, is that the chip has a trap door that government investigators can open to wiretap Clipper-equipped devices. The door's key, which is unique to each chip, has two parts. In the government's plan, the two halves of the key would be held separately in escrow by the Treasury and

Commerce Departments. To use the keys, federal agents would need to obtain a warrant.

It sounds safe enough, but more than 45,000 Internet users disagreed and signed an electronic petition against Clipper. Many people fear the government could misuse Clipper to harm innocent people and businesses.

The FBI and NSA argued that the government was not about to become the Peeping Tom of the online world and that the security holes were needed for criminal investigations. Civil libertarians on the Internet replied that Clipper and the concepts behind it were unacceptable. They argued that even if the security holes were used for acceptable purposes, there was no guarantee that future administrations might not use this against them or their children. In particular, they objected to the keys being held by the government. But the federal government, in essence, shrugged and continued to push for Clipper acceptance in telecom and datacom businesses.

Clipper was not home free yet. Against a mounting storm of protest, the executive branch found Clipper under siege from all sides. This led to such odd alliances as the Electronic Freedom Foundation (EFF) working with conservative Republican senators against the Clipper initiative.

All this protest might have gone for naught until AT&T Bell Labs announced that one of its engineers had shown that Skipjack could be broken. Against a rising storm of protest, Vice President Gore wrote a letter on July 20, 1994 to House of Representatives member Maria Cantwell (Wash.) in which he said he believed there must be further public comment on Clipper.

Cantwell seized on this as a face-saving maneuver signaling the end of the executive branch's support for Clipper and announced victory for the anti-Clipper forces. She soon was seconded in this interpretation by the EFF, and the Clipper controversy sank from sight.

5.3 A Chip with Nine Lives

But the Clipper issue has not died. The NSA continues to work on an encryption algorithm that will be proof against all comers–except federal agents armed with the decryption key. Clipper itself, flaws and all, is still on the books as a voluntary Federal Information Processing Standard (FIPS) for government telecommunications.

If you want to export a telephone device today, you still have no choice but to ship Clipper-equipped products. Of course, finding a foreign customer that would buy such a machine might prove difficult.

Clipper, it should also be noted, was never designed to work with local area networks or the Internet. The NSA and National Institute of Science and Technology (NIST) are working to develop software encryption systems with key-escrow backdoors that would work on networks.

In short, has anything really changed? It doesn't look like it. No one (except employees in federal agencies that have adopted Clipper equipment) is required to use a key-escrow security system. Nevertheless, looks, as we shall see, can be deceiving.

But, you might ask, does it really matter? After all, we have public-key encryption systems using the Rivest-Shamir-Adleman (RSA) encryption algorithm. If you don't trust RSA–and there's reason to believe it might have a hole in it–there are other ways to prevent someone from reading your e-mail. For instance, Philip Zimmermann's popular Pretty Good Privacy (PGP) program uses RSA, International Data Encryption Algorithm (IDEA), and Message Digest Algorithm #5 (MD5) to triply ensure that your messages remain private. Zimmermann also is working on a voice PGP version to give telephone customers an alternative to Clipper.

PGP, however, like other encryption schemes, cannot be exported legally from the United States. As a result, Zimmermann is under investigation by the U.S. Customs Service for possible violations of the International Traffic in Arms Regulations (ITAR). It is conceivable that because it's clearly impossible to keep information from flowing freely on the Internet, there will be an attempt to regulate PGP. A requirement to register keys with the government seemingly would fit perfectly into Clipper's key-escrow framework.

Does this sound unlikely? Consider then, that Congress passed the FBI-backed Digital Telephony bill in late October last year. This legislation requires that common carriers, local telephone companies, and long-distance services must add openings for federal wiretaps to their phone systems. The bill also authorizes government agencies to access billing records. At least police and federal agencies will need search warrants to tap your digital lines and dig through your bills.

5.4 Uncle Sam is Listening

The telephone industry, after fighting this act tooth and nail, is slowly getting used to the idea. Its problem with the bill is not so much privacy considerations as a well-grounded fear that retrofitting older digital equipment and new digital phone systems will cost it a fortune in unnecessary expenses. The government will be supplying $500 million for retrofitting existing systems, but the telephone companies fear this will only be the first down payment on a bill that could swell to the billions of dollars.

The telephone companies have some small protection from a never-ending spiral of wiretapping bills. The legislation only requires that they pay 'reasonably achievable ' costs. The Federal Communications Commission (FCC) will decide, on a case by case basis, what costs fall into this category. If the FCC judges the costs unreasonable, a telephone company could petition the attorney general for funds to accomplish the operation. If the attorney general doesn't reimburse the company, the phone company could refuse to comply, remaining on the right side of the law.

5.5 Getting Clipped

The telephone industry isn't the only potential target of the Digital Telephony bill. Any public communications network that carries both voice and data also is affected. For example, if TV coaxial networks became avenues for two-way communications, they would be subject to the bill's regulations. Whether the Internet, with its relatively primitive voice-communications facilities, should be immune to this bill's regulations hasn't been decided.

There is a catch, however. If you 're calling your Internet provider via a long-distance line or you 're using an Integrated Services Digital Network (ISDN), your Internet communications would be susceptible to Digital Telephony-authorized taps.

Under the bill, telecom carriers not only have to ensure that the government can tap their communication lines, they have to aid the government in its snooping. Telephone companies are required to be able to find a suspect's numbers and report where his or her phone calls are originating and where they're going.

The bill could have been worse. One version would have given the attorney general the right to stop any technological advance that could interfere with the government's ability to make wiretaps. Another early model of the bill would not have required search warrants to engage in wiretapping.

5.6 An Orwellian Future?

Even from a civil libertarian point of view, there are a few things right about the Digital Telephony bill. Most notable is that a police officer or federal agent must obtain a court order to access your e-mail. In the past, all that was required was a subpoena, and generally speaking, subpoenas are easier and faster for an officer to obtain.

What does it all mean for you as an Internet user? In the short run, probably nothing. Look down the road a few years and it's another story. In the future, all telecom systems will be digital and, thus, tappable. It's doubtful that federal law enforcement agencies will put up with being able to easily tap communications without having the ability to decode the messages.

Chapter 6

Conclusion

6.1 General

The Internet is no longer just an academic and scientific research tool, it is rapidly being transformed into a model of a thriving new world economic paradigm. Many people have heard the terms "Information Highway" or "National Information Infrastructure", though few, including network experts, have a firm grasp or agree on what these terms mean. But, while the politicians and giant companies are fighting over how to build this infrastructure, the Internet is already transforming itself into the international "dirt roads" of network commerce.

For over a year, more and more new businesses have begun offering information, services, and even products on the Internet. The majority of these new commercial services are being offered through the use of the multi-media information distribution system World Wide Web (WWW), which allow exchanging documents, files, pictures, videos, sounds, and more across a network.

Mosaic, the WWW Navigator, has excited people all over the Internet because of its ability to empower them with the ability to explore the wealth of information which has been lurking throughout the Internet. What is perhaps more significant is the constant barrage of new businesses and services being offered every day through this medium. It is as if this one application has unleashed the hidden potential of computer networking for empowering all kinds of people with the ability to do business on and use the Internet. Many companies are already using Mosaic, and other tools on the Internet, as a means of performing market research, analyzing their competitors, conducting customer support, providing on-line brochures and catalogs, and even taking orders.

There are hundreds of companies which offer services through Mosaic, newsgroups, e-mail, ftp, gopher, WAIS, and many other Internet information navigation applications. Companies using the Internet are finding they can gain instant international market exposure, provide 24-hour support to their worldwide customers at lower costs than traditional methods, offer up-to-the-minute advertising, provide the latest product information or catalogs, and can even provide customers with the ability to take orders without a phone call. Furthermore, companies can use the Internet to collaborate with other companies (large and small) to work together more effectively. Transmitting information over the Internet is much faster in most cases than sending by fax or by traditional delivery services. Documents and data can be easily transferred saving countless hours and delivery costs.

Some larger companies have been connected to the Internet for years. But in many cases only a few of their employees have truly realized its benefits. The rapid growth and availability of new navigation tools have brought new business value to using the Internet. It may take time before these companies realize what they are missing.

The business opportunities on the Internet are developing too rapidly for a few scattered specialists in a company to follow. Using the Internet to a company's best advantage requires new training of a broad selection of their employees on the tools and informational sources available and on how to capitalize on the new opportunities. In addition, many of the most Internet-capable companies today have special internal teams, or use outside consultants, which are devoted to watching how the company can capitalize on this rapidly changing technology. Thankfully, new low-cost applications are making it easier to track and follow these rapid changes in the technologies.

The Internet has become an invaluable tool in the workplace. A number of experts from all types of fields are connected to the Internet, and many are outspoken and willing to help others with problems in their field. Getting answers to questions quickly can save lots of time and money. Furthermore, communication with peers can greatly enhance the level of expertise one gains while on the job. The Internet makes this communication easier every day.

While the number of businesses on the Internet has soared this year, many of these companies have not yet begun selling products over the Internet, and instead are simply providing information about themselves and upcoming products and services. Why? Because there is a lot of hesitation about security, and there have been many articles in the press highlighting recent break-ins. Unless security problems are fixed, companies whose survival may depend on proprietary trade secrets and closely guarded cost and price information, and financial institutions required to provide a payments infrastructure will not participate. Without the support of financial institutions, it will not be possible to gather the critical density of businesses necessary for a healthy marketplace to develop.

Cryptography and firewalls, based on company's security policy, can address most of the issues to protect a business from fraud. Cryptography provides for transaction security, while firewalls prevent unauthorized access to information and accounts.

All of this functionality can be implemented in different ways using different protocols. Secure e-mail software uses protocols such as PEM and PGP; secure web servers use either SSL or S-HTTP protocols. The SSL and S-HTTP web servers on the market today provide for confidentiality, integrity, and server authentication, but none yet offer digital signatures (Digital signatures are needed to perform customer authentication: to make sure that the company doing business on the Internet knows for sure who is at the other end). Many Internet shopping ventures have proceeded without digital signatures, because if a customer orders a product and provides a credit card number, it is not always necessary to know who has placed the order: the credit card number is enough for some companies. Banks, however, will probably want to wait for cryptographic software that incorporates digital signatures.

An electronic mail system using public key cryptography issues its own public-private keys, and a company desiring to do business over the Internet issues key pairs to customers. Currently, there are a few directory systems so that every person's public key can be available from several locations (this is the case for PGP public keys), but there is no central authority issuing key pairs so that there still remains to establish some certifying authority systems that will be unanimously trusted by merchants, customers and businesses.

Besides, commerce needs reliable and secure payment means in order to become in electronic version what it is in the "real" world. It is now clear that the technology necessary for secure electronic Internet payment systems already exists. Thus, achieving security for all parties, inclusive perfect untraceability for the buyer, is possible.

Currently, no proposal or system is dominant, but with high probability this will change within the next two years at most. However, the question "Which payment system will be used on the Internet?" will not have a single answer. Several payment systems will coexist: Micro-payments (less than $1), low-value payments ($1-$100) and high-value payments have significantly different security and cost requirements.

Possibly, high values will be transferred using non-anonymous, on-line payment systems based on asymmetric cryptography, implementing a cheque-like or credit-card-like payment model. As soon as smartcard readers are available at PCs and workstations, small amounts might be paid using pre-paid off-line payment systems that provide a certain degree of untraceability (like real cash).

Payment systems with and without tamper-resistant hardware at the buyer will coexist for some time. Ultimately, payment systems based on smartcards and electronic wallets (having their own display and keyboard, and communicating with the buyer's terminal via an infrared interface) will become dominant, since they clearly provide better security and enable the buyer to use untrusted terminals without endangering security. Probably, a few almost equivalent payment systems will coexist for the same areas of application (i.e., payment model and maximum amounts). The reasons are various "cultural" differences in the business and payment processes (e.g., between the U.S. and Europe), national security considerations that might disable some solutions in some countries, and competition between payment system providers.

Cryptography so far has been implemented application by application; a person running secure e-mail and a secure browser has to have two passwords, one for each application. It makes more sense long-term, however, for the cryptographic software to reside below the application level (like CP8 Software could intend to do), down at the operating system level, so that all applications on the computer, rather than just the specially modified application, can be made secure. By running cryptography down at the operating system level, businesses can deal with customers running insecure browsers such as Mosaic rather than having to give them or sell them a secure browser.

Applications based on encryption that can provide strong, reliable and robust security services exist in the marketplace today. Yet forces continue to prevent their widescale deployment within the Internet. Specifically, government officials have argued that cryptography is a threat to their missions to preserve national security and prevent crimes. They fear that encryption will be used to hide illegal activities, prevent authorized tapping of communications lines, and otherwise mask terrorist acts. Many nations have implemented laws that prevent or restrict the use, trade, export and import of algorithms, applications, and devices that use especially strong forms of cryptography. However, we can assume that a reasonable balance can be found between the groups that need strong and inexpensive cryptography for international electronic commerce and governmental agencies responsible for national security and crime prevention.

The implementation of an internationally agreed-upon cryptography policy and adherence to certain standard implementations of cryptographic mechanisms are necessary conditions for a strong, reliable electronic commerce infrastructure. Users must be able to purchase off-the-shelf, interoperable products that will easily plug into the infrastructure. If they cannot, the Internet will be perceived as a niche channel, not a robust commercial sales channel, exploited only by those with enough knowledge and resources to navigate the technical maze.

We can assume that the key to a robust, ubiquitous international electronic marketplace will be interoperability among many different hardware and software systems. Standard implementations of protocols along with a variety of user-driven options will be critical to user acceptance. The majority of users will prefer "plug-and-play" products and will not want to bother with the technical issues of encryption method, key length, challenge response protocols, and export status (Thus, providing a significant, seamless and reliable mean to partially secure electronic commerce, CP8 could allure many users).

A buyer on the Internet will need to interact with many servers from many organizations. In some cases, confidentiality and authentication may be required, in others only authentication or only confidentiality may be chosen. For certain transactions, the user may need to encrypt using a mechanism approved for export, such as an exportable version of RC2 from RSA Data Security. For other transactions, the user may wish to use triple DES, a strong form of DES encryption. These kinds of pick-and-choose, reliably implemented options will only be available if suppliers of products agree to standard protocols and technologies, that then become pervasive.

Eventually, we can assume that cryptographic mechanisms will be deployed at virtually every level of the global information infrastructure. Routers will use cryptography to authenticate each other, as will users to other users, programs to users, users to services, programs to hardware, and so on. Data links will be encrypted from a hand-held device to the computer on your desk, from host to network, and from router to router. Files on a local disk will be routinely encrypted, as are applications such as those that move sums of money from one bank to another or that transmit secret information such as a consumer's credit card number or a price quotation for goods or services.. This sort of universally suspicious action between interfaces at every layer of abstraction is the appropriate policy when the means to secure communications is inexpensive, the price of compromise so potentially high, and the path between end points crosses many boundaries of trust.

6.2 Prospectives: Opportunities for Your company

With Web browsers, products can now be attractively presented with rich graphics and detailed descriptions, and customers can have a convenient, fast communications channel in which to obtain a quotation and place an order. Security, however, continues to be a point of concern and a barrier for commercial transactions over the Internet.

Smart card technology, in allowing to prevent unauthorized access to information and accounts, can address most of the issues to protect a business from fraud. If smart card is used as a way to unambiguously identify a user or a computer, then there are at least two applications that Your company could exploit in order to allure new users and to gain advantage over its competitors.

Now that Web-related products have proved to be reliable, scaleable, efficient, user-friendly and cost-effective tools to convey and to present information all around the world, we can assume that these tools will be brought inside companies, as a way for a firm to broadcast its own internal information to each of its employees. Netscape Communications already proposes a version of its well-known browser that is specifically dedicated to internal use (See Appendix 8 : Putting the Web to Work Inside Your Business). For example, many companies have been on the Internet for several years and now develop internal information systems called Intranet. We may think that several firms will follow the same path.

The only security issue that remains when using internal information systems like Huevo, is to give the right person the right piece of information and to be sure that the price quotation, for example, will not leak outside the internal network. If each employee is issued a smart card, then the internal Web server (the Intranet) can be embedded with a piece of software allowing to personalize information access for each of these smart cards (i.e., for each of the employees): each user (employee) would see only the information pertinent to him/her, and access to information which he/she is not authorized to see would be denied to him/her; this can be done seamlessly for the user, without requesting him/her to enter anything like passwords but his/her smart card into the reader at the beginning of the session.

Thus, Your company has to work with Netscape Communications (or other firms distributing Web servers) in order to embed smart card capabilities inside the server software. It is important that security not be a feature added on top of other applications: it must be integrated so that the user will not be bored with it.

If we go further, we can even imagine that the smart card can be used not only as a mean to secure access to information, but also as a way to personalize the user environment and to provide him/her with settings he/she is familiar with. For example, when a user login into a computer by inserting his/her smart card, he/she will be recognized and retrieve all his/her bookmarks for his/her preferred browser.

Another advantage of the smart card is that it would allow employees connecting from outside the firewall to have access to their company's internal information. Internal information has to be encrypted for everybody else outside the firewalls. Web servers and clients like Netscape with SSL already allow the encryption feature, so that the same system needed for internal applications (like a smart card compliant Intranet) could be used for giving selected external access to the company's internal information.

There is another application where Your company could find a good way to gain advantage over its competitors.

After months of negative headlines about the raunchy sites kids are able to visit online and on the Internet, the people who bring cyberspace are getting ready to turn parents into police force (See Appendix 9 : Parents Cast in Cybercop Role). America Online, CompuServe, Prodigy, AT&T, MCI, Netscape, Microsoft and other heavyweights are banding together to create tools that could help parents begin to set limits on the places their children can visit online. The goal is not to replicate the type of ratings system used to judge the content of movies. Rather, the group wants to create tools that companies can use to build and market their own rating or filtering systems. Then, it will be up to parents to select the kind of family standards they want to set for their home computers. Software could be available by March.

Called the Platform for Internet Content Selection (PICS), the standards would allow outside companies to create their own ratings systems that could be marketed to families. Parents then could block their children's access to any site that does not carry the equivalent of a "Good Housekeeping" seal for family-suitable electronic content. Parents would have the option of buying ratings filter software for their home computers or relying on access providers to offer the filter options online. It is very clear that smart card technology could be the "natural" end-part of the access control by the parents. Browsers (only clients and not servers) should be altered in order to give the smart-card-equipped children restricted access to the only sites they are allowed to by their parents, according to the ratings broadcasted by the server.

PICS is still in a definition stage so that it may not be too difficult for Your company to participate to it and to urge for smart card utilization: membership for the alliance is open (See Appendix 9 : Parents Cast in Cybercop Role) to companies that want to participate in the standards-setting process, and these companies are expected to contribute input and resources that will help speed development of the technical standards.