The Future of Networks and Network Security

R. Nelson

Information System Security

48 Hardy Ave., Watertown, MA 02172, USA

(617) 924-9007 phone and fax; rnelson@cs.umb.edu

Abstract

Networks are evolving rapidly into huge, omnipresent, multiservice entities. They are connected worldwide into an Internet that has many different administrations, purposes, resource owners, and users. As the network grows, design parameters are exceeded and new vulnerabilities are introduced. Network security solutions must accommodate enormous changes in the network itself, in the network security requirements, and in the mechanisms and constraints that drive appropriate security mechanisms. As the network serves a larger and more diverse group of users, multiple, flexible security approaches will be necessary to meet their requirements.

Keywords

Internet, networks, security, denial of service, policies

1 INTRODUCTION

In its most general definition, security is the art of making a system behave according to its specification, even when there are forces trying to subvert its operation. Designing a secure system is never easy, but the problem is tractable for small, single purpose systems that are designed from the beginning with security in mind. The Internet is definitely **not** a system of this type. It is huge, constantly changing and under no single control. It was and is designed for flexible functionality rather than for security. This is our challenge as security professionals. How can we work effectively in such an environment?

We cannot secure the Internet. We cannot even secure an Intranet. We can certainly reduce the risk of putting our information resources on the network and of using the network for our various purposes. Network security must, however, be an ongoing process. We must understand the technology that we are trying to make safer, even as it grows, changes and becomes incompatible with our previous security approaches.

Security technology is also changing rapidly, so that we have new, more effective tools for our work. Cryptography is coming of age, and a public key infrastructure is being developed. Commercial tools can detect patterns of usage that may signal intrusions. Protocols are maturing and becoming more robust. Authentication mechanisms are more widely available, at affordable cost.

This paper looks at some of the rapidly changing network technology and at some recent network failures and attacks. It suggests some ways of restating and reinterpreting the Internet security problem so that we can apply security technology that is effective, flexible and adaptable to the changing environment.

2 THE CHANGING INTERNET

Today's Internet is almost unrecognizable to those who were involved in the early designs. It has evolved from a laboratory experiment to a business and household service, almost a necessity. All of the design parameters of the original network have changed by orders of magnitude. Growth has been explosive in capacity, services and number of users. The network is made up of many autonomous systems, with different administration, goals and services. No one, and no single country, is in control.

The size and scope of the current networks present challenging problems to those who need to manage and control network resources, including us in the security community. The network services and features vary from one provider to another and change from one day to the next. We cannot assume that a new behavior is a security problem; it may reflect a new, legitimate use.

The characteristics of the network affect not only the uses of the net and the expectations of the users, but also the vulnerabilities and the attacks. Many of the security approaches designed in the 1980s and implemented in the 1990s are independent enough of the new technologies to be effective. Some need updating and adaptation to the new environment; others have probably lost their usefulness.

In the complex Internet environment, users and providers have multiple, conflicting, changing goals and requirements. Security technology cannot possibly be effective without a clear specification of security policy and a clear model of anticipated threats and attacks. It is important that we security professionals understand the network technologies, how they are used, and how these uses can stress the system. This will allow us work with the network developers and providers on a realistic approach to security. We will be able to provide guidance about providing new services with a minimum of risk to the network, its resource owners and its users.

2.1 Integrated Services

The growth of the World Wide Web has brought integrated services into very widespread use. For some years, the major usage of the Internet was electronic mail, which is text data and typically small to moderate messages. This is still popular, but no longer dominant. Now the largest use may be transfer of web files using HTTP. Most of the information transferred is sound or image data; these are typically fairly large transfers. Another, growing, use is telephone conversations over the Internet; a popular toll-saving measure that can be done with cheap hardware and software. All of these forms of communications share common communications infrastructure, use IP, and traverse the heterogeneous Internet.

The integration of video and audio with the text and other data on the Internet has led to markedly different patterns of use. Transfers are more asymmetric than they were previously; requests to the servers are small and responses are large. It is interesting that packet voice was tried on the Internet in the 1980s and was basically a failure. The voice service worked on local networks, but the low bandwidth and high variability of delay on most paths made it impractical. The voice packets also required special handling by the routers: they were time sensitive and had to be delivered quickly or dropped. Voice and even video are now possible because of the vast increase in both communications bandwidth and network and end system buffer sizes.

A new initiative, called Internet II (Bradner, 1996), is addressing integrated services and new applications. This project is a collaborative effort among a number of universities, the US government, and private sector firms. Its goals include the extension of IP to support new applications and services while maintaining the concept of a single Internet bearer protocol. Significantly, this will be done by changing IP from a best efforts service to a differentiated service, with users able to specify guaranteed bounded delay, low data loss, and high capacity, for example. Current protocols, including IPv6, RSVP and RTP, will support this new structure.

From a security point of view, this user-selectable service presents both new vulnerabilities and new opportunities for control. Quality of service (QOS) parameters must be known to the routers on the route so that they can give the required service. Encryption protocols must permit this, requiring a larger bypass of data through the encryption process. QOS parameters must be authentic and protected from unauthorized change to protect against denial of service attacks on end users. Service providers must protect themselves against denial of service attacks by assuring authentication of QOS parameters and also by robust priority and throttling schemes.

2.2 High bandwidth and high usage

High bandwidth is a relative term. Rather recently, most home computers had modem connections to the local telephone company that ran at speeds from 2400 to 14400 bits per second. Current home computer modems double the previous top speed. Some subscribers are beginning to access the Internet through ISDN lines, at

T1 or higher speeds, and through cable television hookups at tens of megabits per second. Backbone speeds have also grown. When the ARPAnet, the early ancestor of the Internet, was begun, its backbone speed was 56 kbps. Internet access lines are often higher than this; backbone speeds may be 3 orders of magnitude greater.

It is estimated that by June 1996, there were 12 million Internet subscribers. Network users connected to their network service providers through telephone service have already overloaded local switches and caused some loss of telephone service. Telephone switches were designed for voice calls, lasting only a few minutes. Computer connections last much longer, using up the available connections.

Telephony communications models are used to predict load patterns and size systems. The applicability of these to the multimedia Internet world is now being questioned. William Stallings (Stallings, 1997) describes the phenomenon of self-similarity and how it affects buffers and links. The data traffic is not only burstier than voice traffic, but the bursts come in clusters. The effect is more overloading and more delays for the same network size and traffic. This fragility is a security vulnerability; denial of service attacks can be mounted with less effort if the network is overloaded.

2.3 Multiple services from multiple providers

The Internet is a large confederation of users and servers, with service providers of varying sizes and capabilities. Most data traversing the network crosses service boundaries. This means that servers must cooperate to provide end-to-end service. They also must protect themselves from unauthorized use of service and from denial of service attacks. The current protocol between domains is the Border Gateway Protocol (BGP), designed for limited cooperation and mutual protection of the domains. Quality of service negotiation will require a more complete protocol, and, probably more trust and cooperation between domains. Since not all providers can provide the same service, routing negotiations and computations will continue to increase in complexity, and complexity usually means less security.

2.4 Distributed Computing Power

One of the most significant changes in networks over the last 20 years is the distribution of computing power to millions of offices and homes. In the early days of networking, users had dumb terminals, network nodes were based on very small minicomputers, and processing was done by time-sharing mainframes. Threat and attack models for the net are still based on this model. The distributed computing power has blurred the distinctions between terminals, communications components and processing servers. It allows integrated services to be offered and managed. It has also given attackers both more capability to attack and more accessible targets.

Today, there are no dumb terminals; essentially every terminal is at least a personal computer with more memory and programming capability than many old mainframes. Capable users now connect to the Internet via direct connections, essentially becoming Internet nodes with individual IP addresses for the duration of the

connection. These user IP addresses are dynamically assigned, chosen either by the service provider or the user node. Assignment of IP addresses is a security concern, and there is no global policy or enforcement of this assignment. The user nodes should be network stubs, not carrying any through traffic, and should not participate as routers in network route computations. This restriction needs to be enforced by the service providers; there has been at least one incident where a major network was disrupted because a user node claimed its network number and this was not caught.

3 SECURITY ISSUES

Much of the security research and development over the last twenty years or so has been in the area of protecting confidentiality of information. This is clearly important, but it is not nearly broad enough. Network security must address the multiplicity and diversity of network resources, services, providers and users. It must allow multiple policies and even the interconnection of systems with conflicting policies.

Ideally, a secure system should behave according to its specification, even when there are forces trying to make it do otherwise. That is, it should enforce all the policies specified for it, and it should do so in a predictable, robust fashion, resisting all attacks and misuses. However, in the case of the Internet, there is no specification; in fact, there is no single authority able to specify system behavior, system architecture or system security policy. This has been true for some time and it will continue and increase as the Internet gets larger and even more diverse.

If the service provider sets the policies, users can be informed of the rules when they subscribe to the network service. The resources of the system include information, of course, but also computing and communications capabilities. The difficult part of this is to define policies that can be enforced with the available technology, that are consistent with the network services of the system, and that can be explained and sold to users. Then, as the system changes and new technologies and services are added, the security of the system must be analyzed periodically for *de facto* changes in policies and also for new vulnerabilities to new attacks. This dynamic and local nature of network security means that rigid definitions and models cannot be effective, and yet we need clear models and definitions to build secure systems.

3.1 Local security to enforce local policies

There is a practical solution to this problem. Each system must protect its own resources according to its own policies. It is impossible (or at least unlikely) that any widespread agreements will be made about policies or resource usage. It is also unrealistic to assume that all the millions of users of the Internet will abide by any rules. The diversity of interests has already spawned techniques for local protection, including firewalls and now Intranets. These will probably continue to be necessary, along with more specific, server- and application-oriented defenses.

Local security can apply to a single physical system, to a client-server community or to a distributed system sharing a common management organization. The policies, mechanisms and techniques applicable to each system depend on the kinds of resources to be protected and the kinds of services provided.

3.2 Security to suit the purpose

It is important that the purpose of the system be considered. Most system policies deal primarily with the confidentiality of information. It is also necessary to consider integrity of information and fair allocation of processing and communications capability. Identification and authentication policies and mechanisms are critical for all of these protections. Some recent examples illustrate how current networks are being successfully attacked. The nature of these attacks may give some insight into the future. Some attackers are technologically sophisticated, and they will continue to learn.

Modification

Recently, hackers infiltrated a US Air Force computer at the Defense Technical Information Center (Schiesel). They modified the Air Force home page so that it contained material critical of the Air Force and also links to other web sites with hacker information. The hackers claimed major penetration of Air Force and Department of Defense computers, but this is questionable. It is clear that the successful attack caused embarrassment and loss of service; multiple web sites were taken off line. The computers at this site did not contain confidential information; their purpose was distribution of generally available information. However, the web site was clearly not sufficiently protected from a modification attack. The World Wide Web was designed to be very open and to offer free access to information; sites that are accessed are intended to be modified, at least to record the access. Security of the web site as a whole is not adequate; the policies and mechanisms must be much more specific to avoid embarrassment or worse.

Impersonation or spoofing is another type of modification attack. In this mode, the attacker simply creates a web site with an attractive title, implying that it belongs to a government, a vendor, or whatever. Users lured to the site may not only get spurious information; they may also be induced to give away their private information to the attackers.

Denial of Service

Denial of service attacks are very difficult to prevent. Attackers can find the limited resources in the network by understanding its technology. This happened with the TCP SYN flooding attack last fall (Graff, 1996). The attack aims at those server system that accept TCP connections from other systems on the Internet, such as Web servers, FTP servers or mail servers. TCP connections begin when the initiator sends a SYN message to the responder. The responder acknowledges the connection request with a SYN-ACK message. The initiator then sends an ACK message and the connection is established. While this sequence is in progress, the system receiving the

SYN must maintain information about the connection while it sends the SYN-ACK and waits for the ACK. Timers are set to allow for the delay of the communication. The attack is to send numerous SYN messages to a target system, forging the source addresses to be those of a system that will not respond with the SYN-ACK. The target system's space for maintaining connection request information becomes full, denying service to legitimate users who cannot open connections. If the system is not robust enough to deal with this flooding, the attack can even cause a system crash.

The limited resource is the list of connection initiations in progress. Since connections are opened quickly in normal circumstances, most systems have been designed for only a short list. The attack can be countered to some extent by enlarging the list size, shortening the time a request is kept before it is discarded, and other parameter adjustments. It is important to understand that these measures do not guarantee protection against denial of service. They do reduce the risk. This kind of attack is made possible by the fact that users have access to computing power and fast network connections, as described in Section 2. These connections were exploited by some hackers in the large Internet community. Designs that were effective in their original environment can fail when conditions change.

4 CONCLUSIONS AND CAUTIONS

Security concerns will not stop network growth and change, even when new services and features bring added risks for both users and service providers. The challenge to the security community is to develop and apply our technology with respect for the usefulness of the network, and knowledge of the communications and processing systems we are trying to protect.

If we try to consider the Internet as a single entity to be secured, the task is overwhelming. The diversity means that there is no single definition of security for the network and certainly no single architecture or set of mechanisms. Even if we could determine a security approach for the network, we would soon be overtaken by the introduction of new services and features that did not fit our solution. Attackers have ever greater processing power and network access. Network and security components and protocols are being used outside the environments for which they were designed.

However, there is much that we can do to reduce network risk and make the resources on the network safer to use. By taking a more specific, local approach, we can provide good security to many network subscribers and service providers for current uses. We can also be alert to changes in network technology and reexamine our security solutions as the environment changes.

Network designers, operators and users all tend to want the newest, fastest, most feature-filled systems. Vendors will continue to develop and provide new products and services, and they will often do it with no consideration of security at all. Security concerns do not drive the network; they restrict and constrain it. We often try to keep change from happening because we cannot make the new version secure

(or even analyze it before it is fielded). This situation is likely to continue and to produce continued tension between the networkers and the security people.

The network providers and users are concerned about risk. And they need our help. We can be most effective if we understand that their primary interest is in the network functionality; we cannot stop the changes. We need to give them solutions for the changing network environment as it evolves. This means that our solutions must be flexible, modular, and functional. They must operate at the local system, where the service providers have some control. Most important, they must be consistent with the network services and architecture as they evolve.

We can develop useful network security strategies by working more closely with the system designers, understanding new network and computer technologies as they appear, and seeing the vulnerabilities of systems as they are revealed through attacks and misfunctions. Neither we nor the networkers have control of the Internet, but we do have some influence.

5 REFERENCES

- Bradner, S. et al (1996) Internet II Architecture, distributed on the Internet as a working document. Project Director is M. Roberts, roberts@educ.com.edu.
- Graff, M. (1997) Sun Microsystems Security Bulletin #00136, 9 October 1996.
- Schiesel, S. (1996) Air Force Computer Invaded as Hackers Forge Web Page. *Boston Globe*, December 1996.
- Stallings, W. (1997) Viewpoint: Self-similarity upsets data traffic assumptions. *IEEE Spectrum*, January 1997, 28-29.

6 BIOGRAPHY

Ruth Nelson is the president of Information System Security, a research and consulting company that she founded in 1993. She has been involved in network and computer security research since 1975, has participated in the design and security evaluation of many computer-based systems, has published numerous papers on security and is a frequent participant at security workshops and conferences. Ms. Nelson is a member of the Open Network Security Working Group, which meets to consider network security issues for US government and industry, and of a US National Research Council committee on Trustworthiness in Information Systems.