Cryptography and the Internet: Lessons and Challenges

Kevin S. McCurley

Sandia National Laboratories*

Abstract. The popularization of the Internet has brought fundamental changes to the world, because it allows a universal method of communication between computers. This carries enormous benefits with it, but also raises many security considerations. Cryptography is a fundamental technology used to provide security of computer networks, and there is currently a widespread engineering effort to incorporate cryptography into various aspects of the Internet. The system-level engineering required to provide security services for the Internet carries some important lessons for researchers whose study is focused on narrowly defined problems. It also offers challenges to the cryptographic research community by raising new questions not adequately addressed by the existing body of knowledge. This paper attempts to summarize some of these lessons and challenges for the cryptographic research community.

1 Introduction

The Internet has been around for a long time, but the last year we have witnessed an explosion of interest and growth of the Internet. At the time of this writing, most of the interest surrounds the development of electronic commerce at the consumer level, but as a universal method of communication between computers we can expect many other interesting applications in the future, including such things as electronic stock markets and worldwide systems for retrieving computerized medical information.

Most of the current and future uses of the Internet have security considerations associated with them. Unfortunately, much of the Internet was designed without much attention to security. A large-scale engineering effort is currently underway to "bolt on some security" to many pieces of the existing Internet infrastructure, including the Domain Name Service (DNS), routing protocols (e.g., OSPF), and the hypertext transport protocol (HTTP). Moreover, as new capabilities are being developed, they are incorporating a variety of security mechanisms into them. For the most part they are using relatively unsophisticated cryptography (e.g., shared key MACS based on MD5).

The purpose of my lecture is to describe some lessons that this engineering effort provides to the research community, and describe some future challenges

* Author's address is Organization 9224, MS 1109, Sandia National Laboratories, Albuquerque, NM 87185, USA.

This work was supported by the United States Department of Energy under Contract DE-AC04-94AL85000

that can be expected to arise in securing the future global information infrastructure. Because of the short-term nature of this engineering work, the paper here will describe some of the lessons and challenges only in broad terms. It can be expected that after ten years, the value of some of the individual mechanisms that are developed by this engineering effort will be forgotten. At the same time, we are undergoing fundamental changes in the way we use and think about information, and the Internet drives this home.

At Crypto '96, Whitfield Diffie delivered a presentation in which he observed that widespread use of radio communication has been one of the biggest historical changes on the development of cryptography. Looking forward in time and predicting the next trend in cryptography is a risky business, but I believe that the Internet will also mark a sea change in the development and use of cryptography. The development of radio marked a tremendous change in the frequency and nature of communication, bringing with it new problems in securing communication due to the fundamental fact that radio is easily subject to eavesdropping. My reasons for believing that the Internet will bring a similar revolution in cryptography can be traced to three fundamental trends:

- The Internet is global. Communication across the Internet does not respect national boundaries and must conform to many different local cultures and standards of how information is handled. This has the potential to improve the level of understanding between different nations, but it also has the potential to highlight our differences and thereby spark conflicts. As computer networks become increasingly interlocked, parties communicating across the Internet will be increasingly distrustful, which makes cryptography all the more important.
- Communication between computers across the Internet involves some initiated directly by humans (current examples include email and web browsing), but will increasingly involve communication that follows automatic procedures for gathering and processing information. Cryptography has traditionally been a manual process between trusting parties, and new key management strategies will be required to address the increasing amount of automated communication.
- Communication across the Internet is not limited in scale by the size of the radio spectrum or physical limitations of distance, and can therefore scale to to an enormous volume of communication. This growth will not be easy, and serious problems of addressing and routing have yet to be experienced. This rapid growth will also strain our ability to devise effective key management mechanisms and cryptographic primitives.

For the most part, the most interesting problems in cryptography arise from how we *use* information, and not how we communicate it. The focus on the Internet is primarily justified by the fact that it marks a turning point in expanding the way we process information. New applications such as electronic commerce, data harvesting, and remote control of experiments bring with them a complicated set of requirements, and new cryptographic mechanisms will need be required to ensure their efficacy.

Because of the rapid chaotic development of the Internet, it would be inappropriate to concentrate too closely on details in this archival publication. Instead, this paper will flesh out some broad challenges and lessons that the Internet will force on cryptography, and leave the details of their current form to the oral presentation.

2 Lessons from the Current Engineering Effort

Much of the current work on cryptography and Internet security is engineering rather than science, and from a first glance would seem to be pretty routine deployment of existing techniques. Examples include various public key certification hierarchies and the IPSEC security enhancements to the underlying packet transmission protocol. From these engineering exercises there are valuable lessons to be learned for guiding future research directions, and in this section I will try to highlight just a few.

Lesson 1: What's in a name?

In designing cryptographic protocols that use public-key cryptography, theoreticians often ignore the difficulty of identifying public keys with the various parties of the protocol. Early engineering attempts to design a public-key certification hierarchy centered around the X509 standard, and was originally planned to have a single name space for all certificates. Unfortunately, for various reasons this has proved to be slow in coming, and there are now multiple emerging proposals for directory services and an associated public key certification hierarchy (e.g., [3]). Naming conventions are important not only to cryptographic key management, but are also important for authorizations within a larger context. It is common practice today to use address-based authorization because we are unskilled in recognizing the names of entities at a finer level of detail. In the future we should expect that associations of names for entities will be much more important, particularly in an environment where untrusting parties are introduced to each other and wish to carry out a mutually beneficial communication and/or computation. For example, a consumer who wishes to order food from a fast food restaurant might logically expect the domain kfc.com domain to be associated with the restaurant chain of the same name. Within a different context however, kfc.com might refer to another company, and kfc might refer to the initials of an individual. Such complications can be difficult to engineer around, but highlight the need to maintain a well-defined name space for entities in a cryptographic protocol.

Lesson 2: Firewalls are here to stay

The original design of the Internet was for a research environment in which little more was at stake than people's reputations, and it was assumed that nobody would take serious advantage of existing weaknesses. The commercialization of

the Internet infrastructure has evolved along roughly the same lines, where organizations whose members more or less trust each other will place themselves inside a "universal trust domain". When connecting these domains to the Internet as a whole, they concentrate their security at a single point of connection, known as a firewall. This eliminates the need to protect every machine and individual within the organization, and allows more freedom of sharing within the organizational boundary.

Firewall technology has proved to be very cost effective in minimizing human resources required to secure an organization, but are becoming unwieldy as the Internet embraces more and more protocols for carrying out communication and distributed protocols (e.g., video conferencing, active content in electronic mail, etc). Still, firewalls are likely to become more important in the future as organizations develop develop stronger internal information bonds.

The use of firewalls can greatly complicate cryptographic protocols however, since they are natural candidates for mounting "man in the middle" attacks. Cryptographic protocols will need to be developed to address the situation where an intermediary takes some of the responsibility for protecting parties against attacks, because the economic case for such engineering systems with firewalls is too strong to ignore.

Lesson 3: Implement security at the appropriate layer

Computer networks are traditionally described in terms of the seven-layer ISO model. In such a model, the highest level is where users interact with the network through applications, and the lowest level is a physical hardware level. The Internet protocol is more naturally thought of in terms of four layers, consisting of application (e.g., HTTP and SMTP mail), transport (e.g., UDP and TCP), network (e.g, IP, ICMP), and a physical link layer. The separation of a network design into layers allows for modular design, and separates the responsibility of the different layers. From a security perspective, it creates some confusion because the overall design of a network carries with it assumptions about how the different layers will interoperate. In order to build security into applications using this layered approach, we will need to choose the appropriate layer at which to apply cryptography.

In general it is axiomatic that the lower the layer, the higher the performance that can be achieved with cryptographic primitives. For example, at the physical link layer, encryption can easily be handled by hardware devices designed to handle data in appropriately sized chunks. Unfortunately the lower layers do not expose the security requirements of the underlying information, and for example the key management at the physical hardware layer might be problematic if there is a requirement to protect the confidentiality of information from multiple sources that share the physical layer. This is an example of the fact that the higher the layer, the easier it is to match cryptographic services with security requirements of the ultimate application.

The requirements of different applications are quite diverse, and cryptographic mechanisms need to match these requirements. For example, mail mes-

sages need only be checked after the entire message is received and read by the recipient, but TCP stream based applications like telnet need to be encrypted and authenticated in real time as bytes are received, before they are acted upon. It is unnatural to expect that the same algorithms and key management techniques would be used for each of these. Sometimes there is enough commonality between applications that a single mechanism can be applied to several. An example is the Secure Socket Layer (SSL), which provides security services for a range of applications that require a connection-oriented transmission service.

At a lower level, work is proceeding on providing basic security for the network layer, using independent IP packets. The IP security options (IPSEC) for the next generation of the IP protocol are intended to provide independent services of authentication and encryption, using a choice of several different algorithms. IPSEC includes provision for both authentication and confidentiality. There is currently no support for non-repudiation, which effectively limits the services that intermediaries such as firewalls can provide.

In addition, the IP layer relies upon various routing protocols to deliver packets. These routing protocols are of varying kinds, including both link-state and distance-vector approaches. Of the two, distance-vector is somewhat harder to protect with cryptography, since the information that is passed between routers is *derived* from information from other routers, but does not represent the original information supplied by those routers [1]. Hence digital signatures are of limited utility, since routers must still rely on their neighbors to validate information received from other routers before computing routes.

3 Challenges for the future

Challenge 1: The Definition of Alice

Most of the difficulty in engineering cryptographic systems has to do with understanding the trust relationships between different entities involved in a protocol. In theoretical work, we often speak of simple entities such as "Bob" and "Alice" as if they were themselves infinitely capable universal trust domains. In real systems, it is more natural to think of parties involved in the protocol as distributed systems in themselves. For example, when a person is using a web browser to investigate and purchase goods over the Internet, they are in fact acting as part of a system consisting of the person, their computer, their display system, their input/output devices, the operating system, the network, and possibly a cryptographic token. Exactly where Alice stops and the rest of the world begins is unclear for the purposes of analyzing cryptographic algorithms. In all likelihood, the infrastructure that she uses will also be used for other purposes, including possibly her employment and her personal life. In order for Alice to engage in a cryptographic protocol, she will need to store secret information, produce cryptographically secure random numbers, perform computations and communications, as well as deal with the goals of the protocol. In pre-electronic days, these were tasks that she was able to carry out with little more than paper

and pencil. The new electronic infrastructure offers to make her life "easier" by handling very complex data presentation and management tasks on her behalf. In order for Alice to have any trust in the system to act on her behalf, she may wish to understand these actions, but the complexity of modern information systems precludes this. This raises an (admittedly ill-defined) point regarding the analysis of cryptographic protocols: we should do as much as possible reduce the complexity of actions and information that Alice must place her trust in. As the complexity of information systems increases with time, it becomes increasingly important to narrowly define the complexity of the systems that Alice must trust.

Challenge 2: Flexible International Key Escrow

There is no doubt that there are instances when some parties will wish to have the encryption keys of other escrowed, whether for national security interests, political interests, or simply in an organization that wishes to protect it's information assets against the eventuality of an information custodian becoming unavailable. Putting aside the political issue of whether key escrow is desirable in a given situation, the problem of key escrow raises several interesting problems in the design of cryptographic protocols. First among these is the requirement to design a key management framework that will reflect the various access requirements that entities will place upon a key escrow service. Second is the need to minimize the overhead of such a system.

Challenge 3: Scalable cryptographic primitives

For many applications, our current state of knowledge concerning the amount of computation required to carry out various cryptographic primitives will severely limit the application of cryptography in the future. While computers continue to get faster at an astounding rate, there is still a continuing need to explore the boundaries on what the minimal amount of computation and communication required to perform specific cryptographic tasks. Examples include cryptographic hashing, key exchange, digital signature construction and verification, batch processing, and basic encryption. For example, IPv6 offers the option for all data across the Internet to be encrypted and authenticated on a packet-by-packet basis. While it is possible to encrypt and authenticate data streams at a very rapid rate already, there is continued pressure to use the computational capability at the endpoints to process the data stream content for the application rather than consuming resources for cryptographic protection.

Challenge 4: Electronic Commerce Issues

The primary reason for a rising interest by society in the Internet is directly derived from the perception that the Internet offers a promise of new ways to conduct commerce. A ubiquitous communication infrastructure provides a convenient way to offer information products and contact customers for electronic

commerce. Ideas for electronic cash that have originated in the cryptographic research community are now being seriously considered as a mechanism for supporting these new forms of electronic commerce. Delivery of information services begs for a lightweight payment protocol that supports a very low transaction cost. This in turn gave birth to the investigation of micropayment protocols such as Millicent [2], Payword, and Micromint [4]. As people figure out new ways to make money through a global communication and computing infrastructure, we can expect new requirements to come forth for electronic payment protocols.

Challenge 5: Denial of Service attacks

The ability to freely communicate with a vast number of parties leads to the need for parties to protect themselves against denial of service attacks. I have recently started receiving a tremendous amount of email whose purpose is to advertise a product. If such communication is not regulated in the future, then it will become an individual's responsibility to flexibly filter such nonsense. Cryptography offers a mechanism to address such concerns, in part based on the emergence of a cryptographically based electronic commerce system. Future attacks can be limited through the use of protocols that require payment in order to consume some resource. The scalability of payment systems is likely to be the deciding factor in their effectiveness.

In such a short paper, I cannot begin to describe the total range of problems and lessons that the Internet brings to cryptography. We should however expect major changes in direction to occur in years to come.

References

1. S. L. Murphy and M. R. Badger, "Digital Signature Protection of the OSPF Routing Protocol", Proceedings of the 1996 Symposium on Network and Distributed Systems Security, IEEE, 1996.
2. Mark S. Manasse, "The Millicent Protocols for Electronic Commerce, Procceedings of the 1st USENIX Workshop on Electronic Commerce, July, 1995.
3. Ronald L. Rivest, "SDSI - A Simple Distributed Security Infrastructure", preprint, 1996.
4. Ronald L. Rivest and Adi Shamir, Payword and MicroMint - Two Simple Micropayment Schemes, preprint, 1996. Available at `http://theory.lcs.mit.edu/~rivest/RivestShamir-mpay.ps`.

This article was processed using the LaTeX macro package with LLNCS style