Information Security in Integrated Justice Applications: An Introductory Guide for the Practitioner

an IJIS Institute Monograph

Prepared by the Industry Working Group (IWG)
in cooperation with the Integrated Justice Information Systems (IJIS) Institute

Principal Author:
Alan Harbitter, Ph.D., Pec Solutions

Contributing Author:
Jeff Langford, Microsoft

[Inside front cover]

Information Security in Integrated Justice Applications: An Introductory Guide for the Practitioner

an IJIS Institute Monograph

IJIS Institute
720 7th Street, N.W., Third Floor
Washington, D.C. 20001
Tel.: (202) 628-8615
Fax: (202) 628-0080
URL: http://www.ijisinstitute.org

Principal Author: Alan Harbitter, Ph.D., PEC Solutions, Inc.
Contributing Author: Jeff Langford, Microsoft

Integrated Justice Information Systems Institute, Inc

720 7th St. NW, Third Floor, Washington, DC 20001-3716 • 202-628-8615

Dear Reader:

Information Security in Integrated Justice Applications: An Introductory Guide for the Practitioner provides an important overview for criminal justice practitioners on the security issues they face when planning for, procuring, and implementing integrated justice systems. It discusses important security technologies, such as encryption, public key infrastructure, biometrics, firewalls, and virtual private networks, among others, and discusses best practices in security for an integrated justice environment.

This paper was first developed by its authors – Dr. Alan Harbitter of PEC Solutions and Jeff Langford of Microsoft – as a product of the Security Subcommittee of the Industry Working Group (IWG). The IWG, which represents the information technology (IT) community, seeks to leverage the knowledge and experience of industry in furtherance of the integrated justice. Special thanks to IWG Chairman Paul Wormeli, Alan Harbitter, and Jeff Langford for their leadership on this important issue

On behalf of the Integrated Justice Information Systems (IJIS) Institute, I am pleased to publish and disseminate *Information Security in Integrated Justice Applications: An Introductory Guide for the Practitioner* as our first Institute Monograph. The Institute Monograph, which will be a quarterly publication, will leverage the best thinking of industry on issues salient to the implementation of integrated justice, ranging from emerging technologies to procurement and legal considerations in integrated justice.

Creating and disseminating products like the Institute Monograph goes to the core of the IJIS Institute's mission to foster public/private partnerships around the implementation of integrated justice. The IJIS Institute, which is a nonprofit organization and a sister agency to the IWG, achieves its mission by:

- delivering training and education to state and local governments on key technology issues and best practices in project management;
- providing technical assistance to state and local governments on technology and related issues;
- participating on boards and committees working to advance the field of justice system information integration;
- actively representing the industry perspective on information integration issues at key stakeholder conferences and meetings;
- developing relationships with key public sector and nonprofit associations; and
- undertaking demonstration projects that benefit the administration of justice.

We invite you to learn more about us and our products and services by visiting our site on the World Wide Web at www.ijisinstitute.org.

Sincerely,
Bob Shumate
President
IJIS Institute

Table of Contents

List of Illustrations

page

Figure

Tables

1. Introduction

It has become a national imperative that the justice community improves the way it uses information to enforce our laws. No longer can federal, state, and county agencies think and act in isolation. Information that is captured at every step in the justice process must be available for authorized access on a broad scale soon after it is created. This is the challenge of integrated justice: sharing investigative data; coordinating the events surrounding a court case; improving process efficiency; providing updates to the public; and building systems that link events, people, property, and evidence in a consistent and comprehensive way.

In the real world of implementing integrated justice information systems (IJIS), security concerns pose a significant challenge. The process of building IJIS systems increases the interface points among internal computer systems, trading partners, and with the public. Electronic interaction occurs more frequently and involves a broader range of information exchange and transactions. While the overall effect on productivity and service quality is extremely positive, valuable information resources will likely see increased exposure to the risks of privacy violation and security breach. The risk is real; the number and sophistication of cyber attacks is climbing. The Carnegie Mellon Computer Emergency Response Team (CERT) has documented that the number of incidents reported rose from 9,859 in 1999 to 21,756 in 2000, and the first half of 2001 has already seen 15,476 incidents.

The worst possible reaction would be to use security concerns as a reason to abandon efforts to improve the efficiency and effectiveness of the justice process. The good news is that the information security industry has developed very powerful technology in the last decade – technology that provides the tools necessary to build secure information systems and realize the objectives of safe information sharing. The bad news is that our enemies – those who seek to undermine privacy and confidentiality in the justice process – have become very intelligent. The information security practitioner must be equally educated in protective and preventive technology. The goal of this white paper is to provide the first step: to present an overview of information security technology and the best practices that can facilitate secure information sharing.

1.1 Scope

This white paper summarizes the technologies that are currently available and in common use to protect automated information systems and focuses on the application of these technologies to the justice environment. It does not cover privacy policy issues; the assumption is that any law enforcement, court, or corrections/custodial agency that is building or enhancing an automated system to provide integrated justice capabilities already has such policies in place and is now looking for technologies to implement those policies.

There are three fundamental security characteristics: confidentiality, integrity, and availability, sometimes represented by the mnemonic “CIA.” A complete security architecture provides services that address all three characteristics. Confidentiality services support the policies governing access to information and are designed to ensure that information is not exposed to unauthorized parties. Integrity services maintain the accuracy of information products to prevent unauthorized parties from modifying or compromising the integrity of information. Availability services provide confidence that information systems will be on the job when needed. (A significant threat to the availability of computer systems is a security attack called “distributed denial of service” – one of the most difficult attacks to prevent.) Each CIA characteristic is critical to the operation of integrated justice systems and, in many cases, to the safety of the criminal justice agency using those systems.

No technology can make an information system perfectly secure. Each system designer and owner must balance their own comfort with levels of assurance, risk, and investment in security products and procedures.

1.2 Objective

This report provides a practical overview of security technology for designers or owners of integrated justice information systems. The reader is assumed to have some knowledge of information technology, but little or no background in information security.

1.3 Unique Challenges

IJIS systems have unique requirements that complicate and work against efforts to secure information and control access to information, including the need to:

- Communicate with the public. Integrated justice brings with it an increased need to communicate electronically with the public. The communication can be one way, such as posting databases of sex crime offenders, or more sophisticated, such as an electronic case filing system that allows members of the legal community to file court documents over the Internet. In both cases, all three CIA services are required. Legal documents transmitted to the courts must be protected against exposure to unauthorized individuals, and both parties must have assurance that documents actually originated from the claimed source. The public must trust that criminal justice databases posted online are authentic. In many cases, the posting agency has a legal responsibility to guarantee the integrity of such postings. Finally, as justice organizations begin to replace paper-based process with electronic processes, system availability is mandatory to keep the wheels of justice in motion.

- Share information across organizational lines. The objective of integrated justice is to share information across organizational boundaries, starting at the law enforcement point of origin, potentially extending over several courts, and terminating perhaps in custodial or corrections organizations. Again, CIA is important. While information sharing is universally accepted as a mandatory requirement of integrated justice, from a practical perspective not all information can be made accessible to all parties – it would be counter-productive because it would impose the unnecessary additional risk of inappropriate exposure of sensitive information.

- Maintain practices that pass the test of use in the legal process. While integrity is important for all information systems, the justice process often imposes more rigorous requirements. In many cases, the integrity of electronic products must be guaranteed to an extent that can be proven in court procedures, which imposes unique requirements for technologies that establish electronic trust and ensure durable integrity.

1.4 Document Organization

Section 2 provides an overview of information security technologies so that readers have a basic understanding and familiarity with terms of the art. Section 3 presents best practices – examples of how technologies can be applied to build the elements of a secure integrated justice information system. Section 4 presents a brief guide to additional resources for building secure integrated justice systems.

2. An Overview of the Technologies

This section provides an overview of commonly used information security technologies: encryption, public key infrastructure, digital signature, biometrics, firewalls, intrusion detection, and virtual private networks. The area of information security is rich in both theory and products; a comprehensive treatment would likely fill a series of books. For this report, the topic is limited to a quick overview of only the most actively used technologies.

2.1 Encryption

Encryption is fundamental to all security technologies – most software and hardware security tools are based on one or more encryption algorithms. In its most basic form, encryption is used to scramble a message so that it can be unscrambled only by individuals who know the “key.” There are two broad categories of encryption algorithms: secret key and public key. In both cases, the keys are used to encrypt (scramble) or decrypt (unscramble) a digital message. The key generally consists of one or more very large numbers.

Secret Key Algorithms

Encryption dates back to the Roman Empire and the days of Julius Caesar. Caesar needed a way to transmit secret war plans to his commanders at distant points throughout the empire. He used what is now referred to as the “Caesar cipher.” This encryption algorithm replaces letters in a word by shifting a fixed number of letters to the left or right. For a Caesar cipher key of “3left” the encryption table would be:

Original Letter	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w	x	y	z
Encrypted Letter	x	y	z	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w

The message “Attack Greece” would be transmitted as “Xqqxzh Dobbzb.” This encryption can be unscrambled easily by trying all of the 26 possible shift keys. The Caesar cipher is a very simple example of a broad category of algorithms referred to as “secret key” encryption. Figure 2–1 shows how two parties, Alice and Bob, communicate confidentially using secret key encryption. (Tradition dictates that the two participants in an information security example are always called “Alice” and “Bob.”) Both Alice and Bob share a secret: the encryption key. Only Alice and Bob know this key. The key is used as input to the encryption algorithm to scramble the message in a certain way that can only be unscrambled with knowledge of the key. Secret key encryption is also called “symmetric key” because the same key is used both to encrypt and decrypt the message.

Current secret key algorithms are significantly more sophisticated than the Caesar cipher. Many are virtually unbreakable, even by the world’s fastest computers. Some popular, current secret key algorithms are DES (Digital Encryption Standard), Triple DES, AES (Advanced Encryption Standard), and RC4. All of them use a number as the key. One measure of the strength of a secret key algorithm is the length of its key (number of bits). A longer key is more secure because the way to unscramble a secret key encryption is to use brute force and try to guess the key.

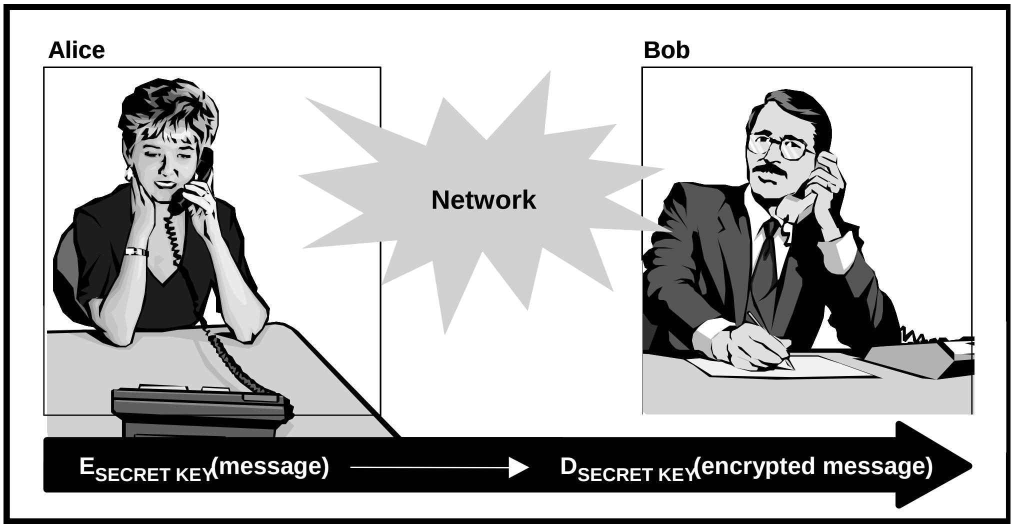
Alice
Bob
Network
$E_{\text{SECRET KEY}}$(message)
$D_{\text{SECRET KEY}}$(encrypted message)

Figure 2–1. Secure Communications Using Secret Key Encryption

DES, first introduced in the 1970s, is one of the oldest present-day encryption algorithms and is still the most widely used algorithm for secret key. It has a relatively short 56-bit key, which corresponds to a number between 0 and 72 billion trillion. Many cryptographers consider a 56-bit key to be completely inadequate. DES "cipher text" (that is, the encrypted message) has been decoded in as few as 22 hours without knowledge of the key. The first step toward improving DES was the introduction of Triple DES, which encrypts a message three times, each time with a different 56-bit key. While the reasons are beyond the scope of this discussion, cryptographers consider Triple DES to have a level of assurance equivalent to only a 112-bit key algorithm.

The National Institute of Standards and Technology, which sets standards for encryption in Federal Government applications, recently completed their selection of the successor to DES: the Advanced Encryption Standard. AES is faster and much more secure than DES. It provides the option of using 128-, 192-, or 256-bit key lengths. Because it is a new standard, however, there are fewer implementations.

It is easy to see how secret key encryption can be used in tools that provide confidentiality services. Secret key approaches can also be used as a basic integrity mechanism. If Alice is convinced that she shares a secret exclusively with Bob, Bob can attach the secret key to a message to prove to Alice that the message came from him. Of course he will still have to scramble the message before sending it to ensure that the shared secret is not revealed. This simple form of integrity is used in many security tools (such as the secure IP protocol, IPSEC), but it lacks an important characteristic called "non-repudiation." To be non-reputable, Alice must be able to prove to a third party that the message she got from Bob actually came from Bob. Non-repudiation is a characteristic that is required in order to implement digital signature. (Digital signature is covered later in this section.)

Public Key Encryption

Public key encryption is used for digital signature (Section 2.3) and rigorous authentication (Section 3.3). In public key algorithms, each key is actually a pair of keys. One portion of the key is private, known only to the individual key owner, and one portion is public and is posted to the world. The posting process is typically supported by a public key infrastructure (Section 2.2).

The public key encryption process is illustrated in Figure 2–2. If Alice wishes to transmit a message to Bob in confidence, she looks up Bob's public key and uses it to encrypt the message. Only Bob knows his private key, and the message can be decrypted only by using his private key. Public key is also referred to as "asymmetric key" because different keys are used for the encryption and decryption operations. Public key encryption is easier to manage than secret key encryption because it does not require a shared secret for secure communications to take place. For this reason, it can be applied when a large number of parties do not know each other prior to the first message exchange. This is certainly the situation that exists when a justice organization seeks to communicate securely with the general public.

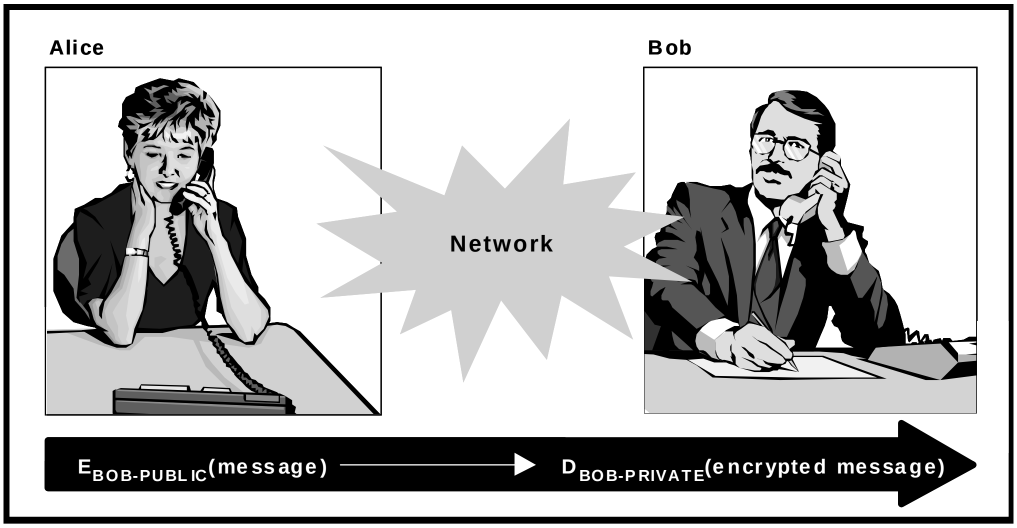
Alice
Bob
Network
$E_{BOB-PUBLIC}$(message)
$D_{BOB-PRIVATE}$(encrypted message)

Figure 2–2. Secure Communications Using Public Key Encryption

RSA, Elliptic Curve, and Diffie-Hellman are commonly used public key algorithms. Public key algorithms use much longer keys than secret key algorithms – often 1,024 bits or longer. A public key encryption is not unscrambled in the same way as a secret key encryption. An attacker will more likely try to reverse the key generation process and derive the private key by examining the encrypted message and the public key. It would be nearly impossible to use brute force to guess a key that is 1,024 bits long.

Another important difference between public and secret cryptography is operational efficiency. It takes considerably longer to perform the calculations necessary to encrypt a message using public key cryptography – sometimes 1,000 times longer. This can be a problem with e-commerce Web servers that use public key encryption to protect message exchanges containing sensitive information: the cryptography calculations place a large workload on Web site servers. For this reason, the most common application of public key cryptography is to securely pass a "session key" between Alice and Bob. The session key is a secret key that is used for symmetrical cryptographic operations to protect Alice-to-Bob message exchanges.

2.2 Public Key Infrastructure

Public key cryptography requires that half of the encryption key pair for each participating party is publicly posted and available from a trusted source. In addition to posting public keys, the public key pair must be managed. A facility must be in place to register keys to new owners, revoke expired or stolen keys, and possibly archive keys used for encryption. The computer systems and procedures that provide these key management services are called a public key infrastructure (PKI). A PKI generally includes the following components:

- Digital Certificate. A digital certificate is a computer file that binds a public key to an individual or organization. The primary mission of a PKI is to manage digital certificates. If Alice and Bob are sharing information and Alice needs to gain access to Bob's public key, she will request a copy of his digital certificate. The digital certificate is digitally signed by the Certification Authority to ensure its authenticity. The standard that governs the format of digital certificates is called "X.509."

- Certification Authority (CA). The certification authority is the central computer system in the PKI. It is responsible for maintaining valid digital certificates. It posts certificates to a directory for public access and revokes certificates that have expired or have otherwise become invalid. One technique for providing information about revoked certificates is to distribute a Certificate Revocation List (CRL). The level of assurance of a PKI is dependent upon effective certificate management.

- Registration Authority (RA). The registration authority is a system that consists of a computer and a person, perhaps a security officer, who is responsible for registering public key holders into the system. The RA enforces the portion of the security policy that governs how a key owner is identified. For example, Alice may be required to present a driver's license and a work ID before her public/private key pair will be registered into the system and a digital certificate posted.

- Directory: The directory is the database of digital certificates that can be accessed by users wishing to send an encrypted message or validate a digital signature.

Figure 2–3 shows how these components interact. Building a PKI involves defining the policies by which keys are granted, revoked, and archived; assembling the hardware and software required to configure the computers and networks that are used to implement the PKI; and establishing a staff to administer and manage the PKI system. Developing and operating a PKI from scratch can be an expensive undertaking. Although the cost is highly dependent on the scope of the application (that is, the number of participants and digital certificates managed), implementation costs can easily be in the millions of dollars.

Owing to the expertise and expense involved in building and operating a PKI, integrated justice program managers may choose to outsource the PKI function entirely or use one of many commercial certification authority service providers. As with any decision to outsource or use commercial services, the primary tradeoff is between the ability to tailor the PKI to meet unique needs (a benefit of a customized PKI) and the cost of doing so (use of commercial certification authorities can result in lower costs). Figure 2–4 lists a few of the trade-offs between building and operating a customized PKI versus using a commercial CA. For many justice applications, unique policy requirements mandate that the participating organizations build and operate their own customized PKI. It may be desirable for government agencies to band together and pool their resources to improve the cost-effectiveness of a PKI program.

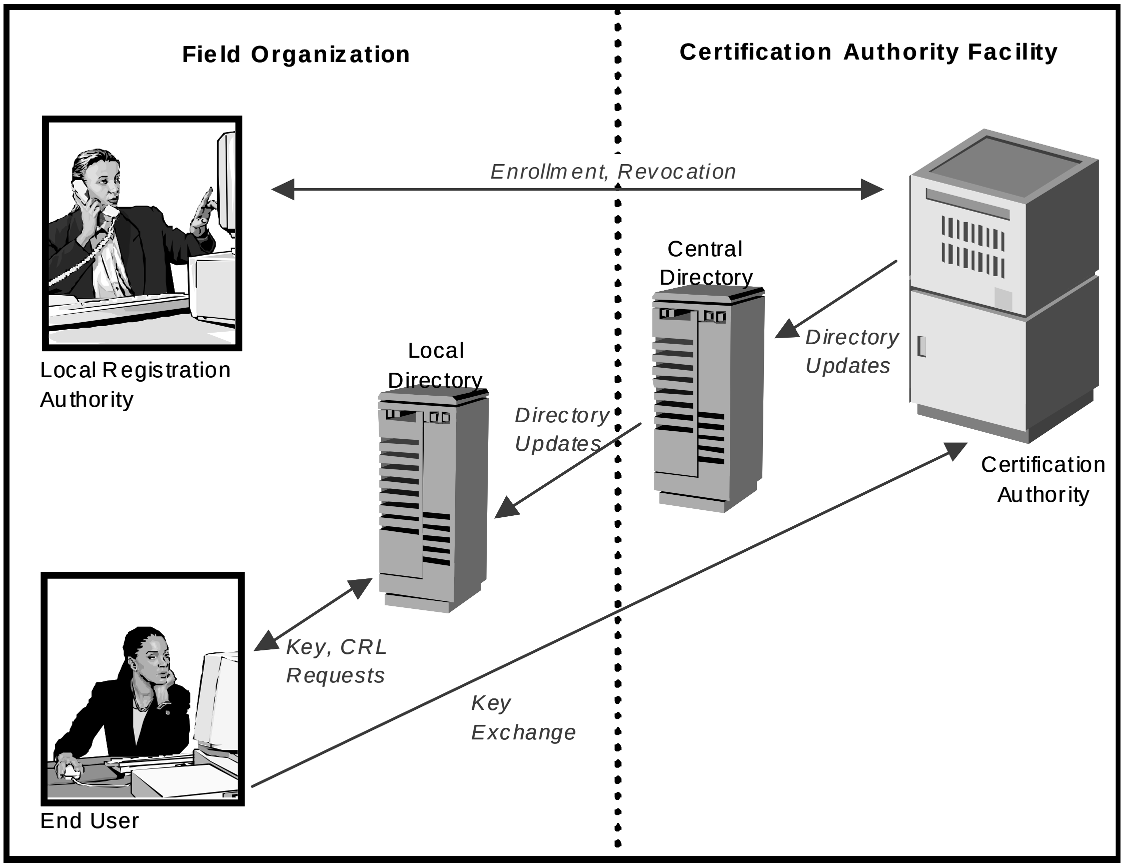

Figure 2–3. The Components of a Public Key Infrastructure

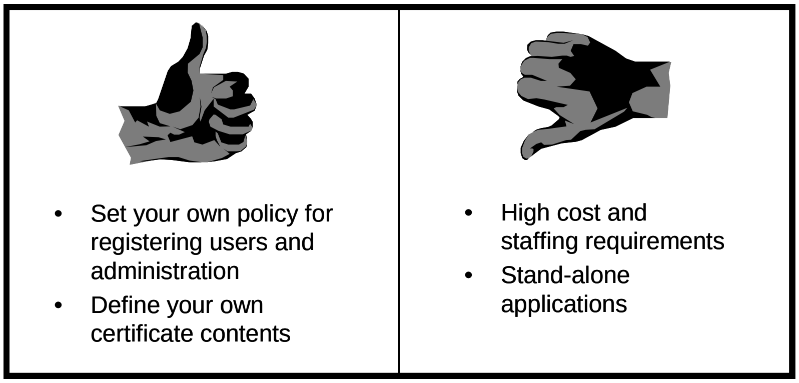

- Set your own policy for registering users and administration
- Define your own certificate contents

- High cost and staffing requirements
- Stand-alone applications

Figure 2–4. The Pros and Cons of a Customized PKI

2.3 Digital Signature

A digital signature accomplishes a function similar to a traditional pen-and-ink signature: it implements a legal commitment to a subject document and a signatory. The commitment might include a contractual consent to the terms specified in the document or a statement that the document is original and authentic.

Figure 2–5 describes how a digital signature can be generated. First, the document to which the signature will be affixed is digitally condensed. The result of this condensing function is often called a hash code or message digest. The message digest is much smaller in size (number of bytes) than the original document. The signatory, Alice, encrypts the digest with her private key. The encrypted digest is the digital signature. It can be transmitted independently from the document; it is unique to the document and the signatory.

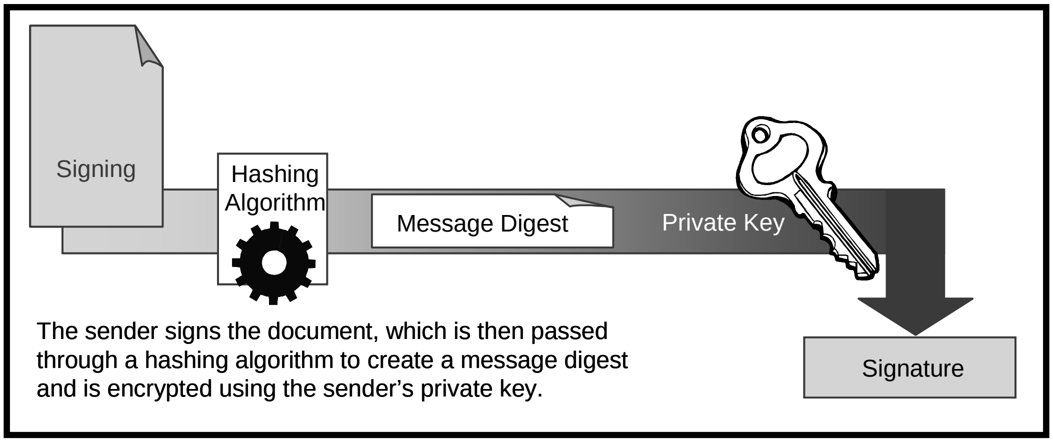

The sender signs the document, which is then passed through a hashing algorithm to create a message digest and is encrypted using the sender's private key.

Figure 2–5. Generation of a Digital Signature

To confirm that the signature is valid, Bob calculates the original document's hash, decrypts the signature with Alice's public key, and compares the two results. If the results match, the signature is valid. (If the original document is modified in any way, the hashes will not match.)

The digital signature accomplishes a function similar to a pen-and-ink signature but has a significantly different form. A more familiar automation of the pen-and-ink signature is the "electronic signature" – a digitization of the image of a pen-and-ink signature. This type of technology is often used in the retail industry to complete credit card transactions. Figure 2–6 compares electronic signatures to digital signatures. While an electronic signature is a more intuitive extension of the pen-and-ink signature, it lacks the rigor of a digital signature; nothing binds the electronic signature to a specific document. The image that is captured during a retail sale, for example, could easily be electronically "cut and pasted" onto any electronic file. The digital signature, on the other hand, is unique to the signed document and is much more difficult to forge or apply to other documents.

Digital signature technology has many potential applications to integrated justice. Court documents such as those exchanged in electronic case filing systems should be digitally signed by someone – a submitting attorney, clerk, or judge, for example. In addition, criminal justice agencies may want to digitally sign public postings, such as sex offender Web site information, to provide assurance that the information comes from an official source and has not been modified inappropriately.

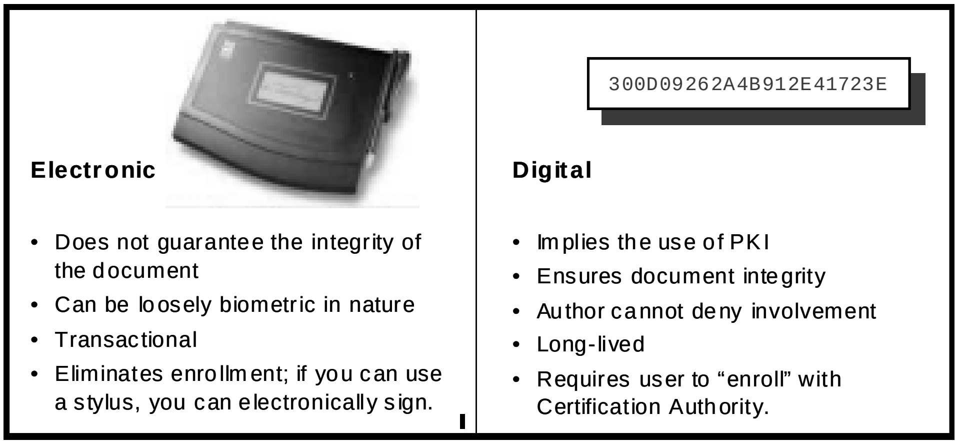

Electronic

- Does not guarantee the integrity of the document
- Can be loosely biometric in nature
- Transactional
- Eliminates enrollment; if you can use a stylus, you can electronically sign.

300D09262A4B912E41723E

Digital

- Implies the use of PKI
- Ensures document integrity
- Author cannot deny involvement
- Long-lived
- Requires user to "enroll" with Certification Authority.

Figure 2–6. Electronic vs. Digital Signature

If digital signatures are to be used in legal transactions such as electronic case filing, there should be laws to govern their use – and there are. One of the most well-known is the Electronic Signatures in Global and National Commerce Act (ESIGA), which was signed into law in June 2000. This federal law governs interstate and foreign commerce and endorses the use of electronic signatures for most of the situations in which pen-and-ink signatures are used. The law uses the term "electronic signatures," but uses the term generically so it allows other types of signatures. From a practical standpoint, digital signatures are the only technology that can be used for legal transactions. While the jurisdiction of that law is limited, it represents an important milestone because it supports the use of digital signature on a nationwide basis.

The laws that govern the use of digital signature at the state and local levels are unique to each state. Most states have digital signature legislation on their books. One of the first states to pass such a law was Utah; several states have based their laws on Utah's. The Utah law was established in 1996. Digital signature technology was just emerging at that time, so the law aimed at motivating the private sector to enter into the digital signature market. As such, it limits the liability of commercial certification authorities to ensure the integrity of electronic transactions. This is another factor to consider when weighing the decision of whether to outsource or develop and operate a PKI. Integrated justice system developers and managers should understand the implications of the law governing the use of digital signatures in their state.

2.4 Biometrics

Proper authentication of information system users is a basic security service. Section 3.2 describes the limitations associated with the most common form of authentication, password security. Biometrics offer more rigorous authentication by requiring physical identification in addition to passwords or PINs. Biometric methods take several different forms, and they result in varying levels of complexity depending on the type of information being accessed. Generally, one or more of three basic factors identifies individual information system users:

- Knowledge. Something you know; for example, a password, or PIN, or personal information such as a birth date or name.

- Possession. Something you physically have; for example, a key, ID card, or smart card.

- Biometric. A physical element of oneself; for example, a fingerprint, handprint, voice, or physical signature.

This section focuses on the third factor – biometrics. However, multiple factors can be combined to increase the level of assurance of the authentication process.

When evaluating different biometric devices and alternatives, it is important to consider the FRR (type I error, or “false rejection rate”) and the FAR (type II error, or “false acceptance rate”). The FRR measures the percent of rejections that should have been accepted (a valid user who used the device but was not properly identified); the FAR measures the percent of accepted or validated logins that should have been rejected (an invalid user who was improperly identified as a valid one). These two ratings are closely related. On average, today’s biometric devices typically have a 4 to 5 percent error rate. The correlation between the two rates can be expressed in the following manner: For a highly secure solution, the FAR would be 0 percent and the FRR would be 5 percent. If the FAR were to increase to 3 percent, the FRR would need to lower to 2 percent. All manufacturers provide their average FRR and FAR ratings. Other factors to consider are cost, environmental conditions (weather, dust, humidity), and intrusiveness to users.

The many different types of biometrics can be grouped into two categories: physical and behavioral. An example of a physical biometric is a fingerprint or iris pattern; an example of a behavioral biometric is a voice or keystroke pattern. The following paragraphs summarize physical and behavioral biometrics.

- Fingerprints. This is perhaps the most well-known and accepted form of physical biometric in use today. The uniqueness of fingerprints has been known for a long time, and fingerprints have been the de-facto standard identifier in the justice community for many years. It is not surprising that this is also the most common form of electronic biometric identification currently in use. The unique patterns of a given finger are analyzed and stored in a database and compared against a user attempting to gain entry into a system. If a matching pattern is found in the database, the user is granted access. The particular methods of validating a given pattern may differ (for example, minutiae or moiré fringe), but the end result is the same. Some newer scanners detect the temperature or electrical impulses of the digit being scanned, thereby confirming that the finger is attached to a living being at that moment. Fingerprints are very easy to obtain through scanning, and the technology is non-intrusive.

- Hand Geometry. This physical biometric method involves measuring and analyzing the shape of the hand. Different individual characteristics, such as length or width of a certain digit, are combined to ensure a unique pattern. This method can be quite accurate. It is relatively easy to implement and fairly non-intrusive.

- Retina Scanning. The retina of each eye is as unique as a fingerprint and relatively easy to scan. This works by mapping the layers of blood vessels on the retinal surface at the back of the eye. This physical biometric method requires that the person stand completely still for a period of time while focusing on a given object. While highly accurate, this method is not widely used due to its intrusive nature and the necessity to remove eyeglasses and, in some cases, contact lenses.

- Iris Scanning. Iris scanning is relatively new and very accurate. It works by comparing the color patterns in the iris with a sample or template stored in the database. This physical biometric method is somewhat intrusive, but not nearly as much as a retina scan. Although it is not necessary to remove eyeglasses, the method may not work on a person wearing colored contact lenses. This method is very easy and inexpensive to implement; a simple electronic camera device can be used to perform the scan.

- Facial Recognition: This area of physical biometrics has received much attention lately owing to the widespread appeal of its variety of methods. Facial recognition works by combining many different characteristics of the face such as size, shape, width, color, and even heat patterns. It is non-intrusive and fairly easy to implement, although its overall accuracy is not as good as fingerprints or eye scans.

- Voice Recognition: Voice recognition is not simply a matter of recognizing a person's voice, but rather an overall analysis of several different factors such as inflection, gait, and volume. Voice recognition is perhaps the cheapest method to implement because it does not require additional hardware. This behavioral biometric method is non-intrusive and easy to install but is not necessarily the most accurate.

- Signature Analysis: Signature analysis deals with capturing and monitoring several different aspects of a live signature. Users "sign" their name as usual on a device such as a touch screen or digitizing tablet while the system monitors the user. Characteristics such as velocity, pressure, and pattern are compared to a known sample. This behavioral biometric method is widely accepted as non-intrusive because all users already sign their name as a form of identification. The method is neither particularly expensive nor difficult to implement, but its overall accuracy has yet to be proven.

Perhaps more important than the type of biometric methods and devices is the overall strategy for deploying and implementing biometrics in an information system. Biometric methods are typically a very good way to identify an individual, but they should be used in conjunction with another method of verification. For example, if a fingerprint scanner is the sole method of verification, a user with an injured or bandaged hand may not be able to log on. This type of problem exists with all biometrics: a user with a cold sounds different; certain drugs affect the eyes; and heat, cold, dust and other environmental elements can affect the accuracy of some devices.

There are several factors to consider. It is important to know the operating location of the measuring device – whether it is a laptop installed in a police patrol cruiser or a desktop at the precinct. It may also be desirable to provide different authentication methods for different levels of information sensitivity.

2.5 Firewalls

Firewalls are the most common form of network protection. A wide range of firewall products are available, and all have the same basic mission: to examine all traffic transiting between the "private side" (internal local area networks) of an organization's network to the "public side" (connections to the Internet or other public networks) and to impose a policy regarding the traffic that may pass and the traffic that must be blocked. If a justice organization's private network is connected to a public network, a firewall is an absolute minimum requirement. To be effective, all information transfer between the public and private side should be monitored by the firewall. In other words, there should be no information flow that circumvents the firewall.

Figure 2–7 shows a firewall positioned to accomplish this function. The traffic filtering policy is most typically enforced through a “rules table” stored in the firewall. Table 2–1 is a simplified example of such a table. For every potential source, destination, and protocol triplet in the table, an action is specified. If the action is “accept,” the traffic flow is permitted; if the action is “deny,” it is blocked. A typical firewall rule might be to permit email traffic to be exchanged only with the email server. In Table 2–1, this rule is represented in row 5.

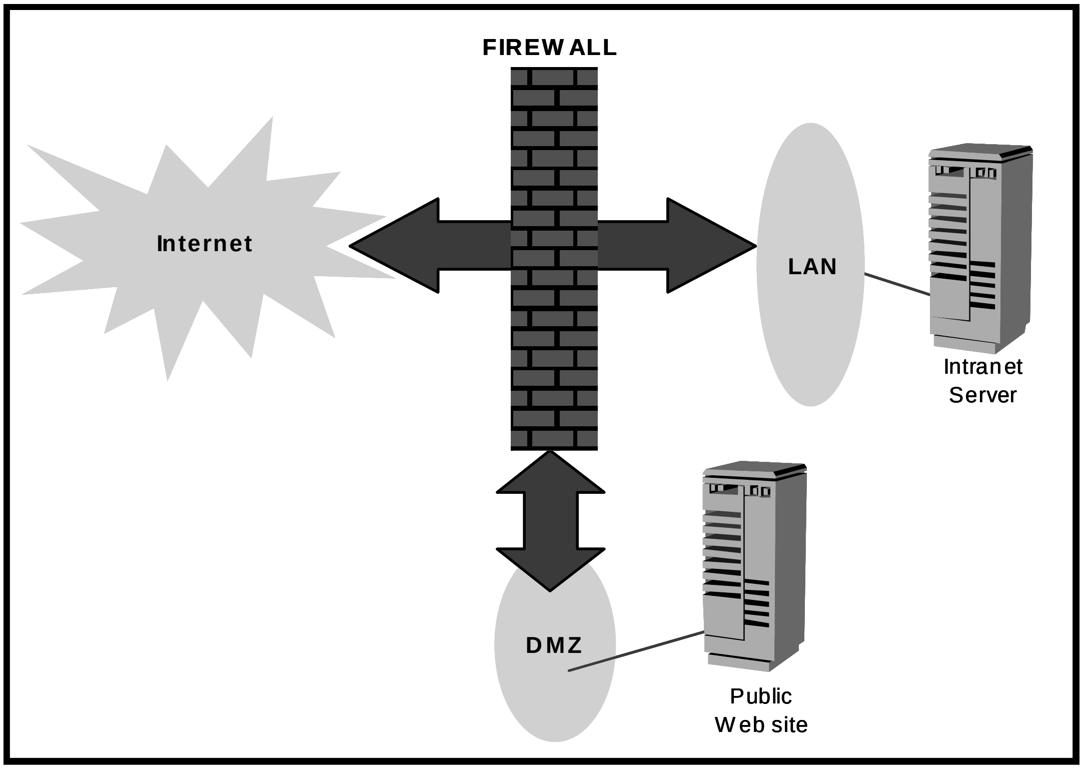

Figure 2–7. Typical Firewall Configuration

Table 2–1. Firewall Rules Tables – Simplified Example

Source	Destination	Protocol	Action
All external	All internal	Telnet	Deny
All external	Web site	HTTP	Accept
All external	All external	HTTP	Accept
All external	Web site	FTP	Accept
All external	Email	SMTP	Accept
Any	Any	Any	Deny

Notes

Telnet	Protocol for terminal users
HTTP	Protocol for Web page access
FTP	Protocol for file transfer
SMTP	Protocol for email

The firewall implements the rules from top to bottom. Table 2–1 has a catchall rule in the last row to deny any traffic that was not explicitly allowed by previous rules. This table represents a very conservative policy. In general, the level of protection that a firewall provides is dependent on the strength of the policy defined by system managers and represented in the rules table.

The firewall product market is very mature, and firewalls provide many functions beyond enforcing the rules table. The following paragraphs define and describe some of these functions:

- Application Proxy. An application proxy reduces the need for a message originating on one side of the firewall to cross over to the other side. Suppose an outside user wants to obtain a service from the private side of the firewall. For example, a user on the public side of the firewall might want to retrieve the contents of a Web page that is stored on the private side. If the firewall provides application proxy service, the request for the Web page will be processed on the firewall. The proxy will retrieve the page from the real Web server on the private side and deliver it in response to the external user's request; the external user's message never enters the network on the private side.

  Application proxy service provides a higher level of isolation between the public and private sides of the firewall and better security, but there are limitations to an application proxy firewall. Most private networks have a wide range of application services that are available to outside users, but it is difficult to find an application proxy that offers all of these services and matches the functional richness of the native servers. Further, the proxy can impose a performance bottleneck because it must provide application services in addition to filtering traffic.

- Authentication Server. A firewall may provide authentication service. Before a user's messages are passed from the public side to the private site, the user must prove identity by, for example, demonstrating knowledge of a password. This feature further protects access to private side information resources, but creates an additional administrative burden for users and network administrators.

- Caching. Firewalls that provide application proxy services can also provide caching. If there is a Web server on the private side of the firewall, for example, the proxy can store frequently requested pages so that external user requests to view those pages do not have to generate a request to the Web server. The same service can be provided for users on the private side of the firewall; frequently requested pages from public Web sites can be stored in the firewall, and users may view them without sending their requests out on the Internet. This feature has some security benefit, but it is generally provided to improve performance and reduce access time.

- DMZ. The "demilitarized zone" (DMZ) is a tongue-in-cheek term describing a third network that is logically cordoned off by the firewall. As Figure 2–7 indicates, the DMZ is often used to house information systems and stores, such as a Web server, that are frequently accessed by the public. The firewall rules table may enforce a more lenient policy for public access to the DMZ than it does for the private side of the network.

- Network Address Translation. Every device that is connected to a network – whether it is a personal computer, server, or printer – must have a network address. Many firewalls will change the network address of a source or destination device embedded in a message as that message crosses the firewall. The firewall keeps track of the real network address of the device and the translated address. Network address translation (NAT) serves several purposes. It shelters the real addresses of devices on the private side of the firewall from the public, and it allows private network operators the freedom to assign addresses on the private side of the network independently from their use on the public side.

- VPN. A virtual private network (VPN) is a technology that allows a public network, such as the Internet, to be used in a secure and private manner. Because the firewall guards the interface between the public and private worlds, it is frequently used to implement VPN services. A more thorough description of VPNs is provided in Section 2.7 of this report.

While firewalls are most frequently used as the barrier between the private and public sides of a network, they can be applied for other reasons. For example, within a private network there may be communities of interest that, for the purposes of enforcing a security or need-to-know policy, should be kept logically separate. For example, internal firewalls might be used in an IJIS to logically separate the networks that support law enforcement from courts or the district attorney's office from the public defender's office. An internal firewall allows information to pass between two justice organizations, but in a controlled manner.

2.6 Intrusion Detection

Intrusion detection (ID) is, perhaps, the most active area of current research in information security. ID products strive to identify, as quickly as possible, that an attack has been launched to gain inappropriate access or otherwise undermine the integrity of networks and information systems.

Two methods are typically used in ID. The first is to look for recognizable "signatures" of an attack – patterns in user sessions and programs that are followed when an attempt is being made to violate the security of an information system. This works well on known attack strategies but does not help identify new exploits.

The second method works by looking for anomalous behavior patterns that indicate an external user may be attempting to gain unauthorized entry and may be masquerading electronically as trusted staff. This method is more effective in identifying new threats. For example, most information system users log on during normal business hours. If a user who regularly logs on at 8:00 a.m. and logs off at 4:30 p.m. suddenly requests access at 3:00 in the morning, it is anomalous behavior and there may be reason for concern – a hacker may have stolen the user's electronic identity, intending to launch an attack.

An intrusion detection system monitors this level of activity and reports suspicious patterns to a system administrator or security officer. Some ID systems automatically implement simple protective measures in the event of a suspected attack, but it is usually up to the system administrator or security officer to determine whether the system is under attack and to take appropriate action. Current research is improving the detection capabilities and reducing the likelihood of false alarms.

Vulnerability analysis software complements intrusion detection by scanning the network configuration in search of holes through which attackers can gain entry – holes created by improper system configuration, out-of-date software, and weak password selections. The software then produces a report identifying the holes so they can be fixed.

2.7 Virtual Private Networks (VPNs)

VPNs allow the use of public networks for private communications. This is especially valuable to an integrated justice community that tends to be geographically dispersed. Compared to traditional private networks, VPNs can result in lower implementation and maintenance costs. Rather than use expensive point-to-point communications lines, VPNs allow agencies to establish private virtual connections by "tunneling" through public networks such as the Internet. One of several techniques encrypts the data stream before it enters the public network and decrypts it after it exits. The use of VPNs gives agencies access to the broad connectivity and ubiquity of the Internet for private communications.

There are generally three different scenarios for VPN connectivity: remote user access (also known as client-to-network), direct private connection (or server-to-server connection), and network bridge (or network-to-network). Figure 2–8 depicts all three of these scenarios.

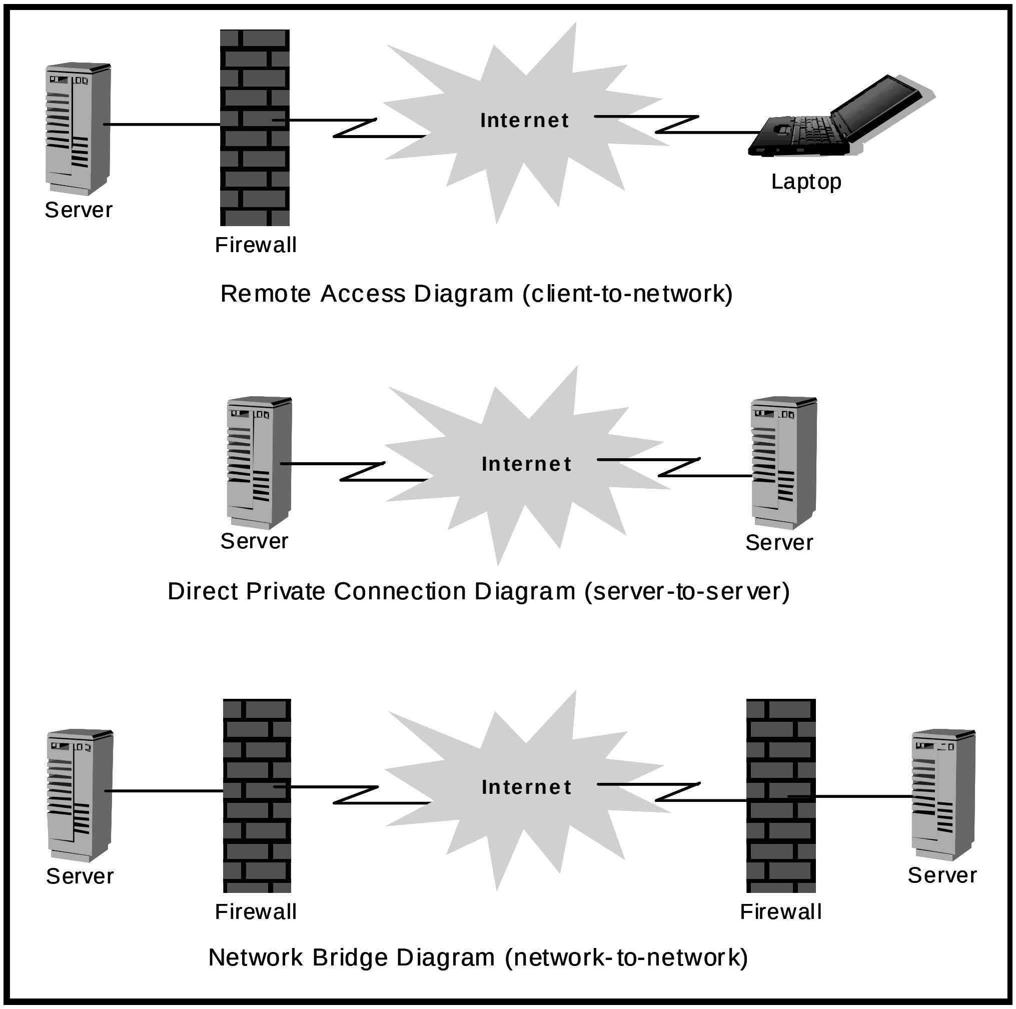
Server
Firewall
Internet
Laptop
Remote Access Diagram (client-to-network)
Server
Internet
Server
Direct Private Connection Diagram (server-to-server)
Server
Firewall
Internet
Firewall
Server
Network Bridge Diagram (network-to-network)

Figure 2-8. VPN Configuration Scenarios

The *remote access* user typically "dials in" to an agency or shared government LAN by going through an ISP (Internet Service Provider) to connect to the Internet. After the Internet connection is established, the user then connects to the agency's LAN through another login session. Once authenticated on the LAN through the VPN connection, the user (or client) may work just as if they are physically connected to their network. All data that travels through the VPN “tunnel” is encrypted and secure.

Direct private connection works similarly but is used to connect one server to another, rather than connect a client to a LAN. This method can be more secure, depending on the implementation. Although a firewall is not explicitly shown in Figure 2–8, one may be used. The primary difference here is that the VPN is not connecting a single computer to an entire LAN but rather to another single computer. It is important to note that it is the responsibility of each of the two servers to provide the VPN services required for connection in this instance. This method of connection is typically used to share only specific data contained on a single server while not opening up all the data (or computers) on the entire LAN.

A *network bridge* is a common method of connecting remote locations through the Internet. This scenario ties together two separate LANs through a VPN connection. Any computer can seamlessly send information to and from any computer located on the other side of the connection. It appears to the user (and any other service on the network or computer) that the two LANs form a single, logical network.

Some of the technologies used in VPNs are described in the following paragraphs:

- Point-To-Point Tunneling Protocol (PPTP) is the most common protocol used in VPN connections. By virtue of its design, PPTP can be passed through NAT firewalls to allow even unpublished network addresses to participate. PPTP is very easy to use, set up, and implement. It is included as a standard protocol in nearly every operating system available or in use today. PPTP can use a number of different strengths of encryption; 40-bit, 56-bit, and 128-bit are the most common.

- Internet Protocol Security (IPSec) is also a very common standard protocol for VPNs, although implementations vary somewhat between vendors. IPSec adds to the standard protocol used on the Internet (IP) features that encrypt the individual messages and provide message integrity. It does not work through NAT firewalls.

- Layer 2 Tunneling Protocol (L2TP) combines a number of existing protocols. L2TP does not offer any encryption itself; it must be used with another protocol if confidentiality is required. L2TP is typically used to package nonstandard protocols within IP packets in order to send data from disparate networks such as IPX or AppleTalk through the Internet.

VPNs are used for several reasons. The first one is cost: VPNs are inexpensive and relatively easy to set up and use. Second, VPNs do not require any physical changes to the network; if connectivity to the Internet already exists, it is easy to use VPN. Last, VPNs make it easy to join or separate any number of disparate, remote, or disconnected networks. VPN technology allows network managers to dynamically adjust to changing needs while ensuring that data is protected by encryption.

VPNs can also be used within an intranet for additional security. For example, if a county has just one LAN for all agencies and departments, VPN could be used to enable information sharing and lower maintenance costs. A single shared LAN may not work well for certain agencies such as the jail or the Department of Health because both have very sensitive data that requires protection. If the jail needs to send or receive sensitive health information about an inmate, the data should be encrypted or otherwise secured. A VPN makes it very easy to fulfill this requirement.

3. Best Practices

This section describes, in the context of an integrated justice system, typical applications of the technologies summarized in Section 2.

3.1 Network Security

Figure 3–1 presents a hypothetical and simplified integrated justice scenario that includes four participating entities – law enforcement, attorneys, courts, and the public – and assumes that these parties would like to share information in the following manner:

- The public would like to access information about trial dates on the court’s Web site.
- Police officers would like to check their court appearance schedule from their vehicles or home PCs.
- Attorneys would like to file motions through the Internet and electronically access judges’ rulings.

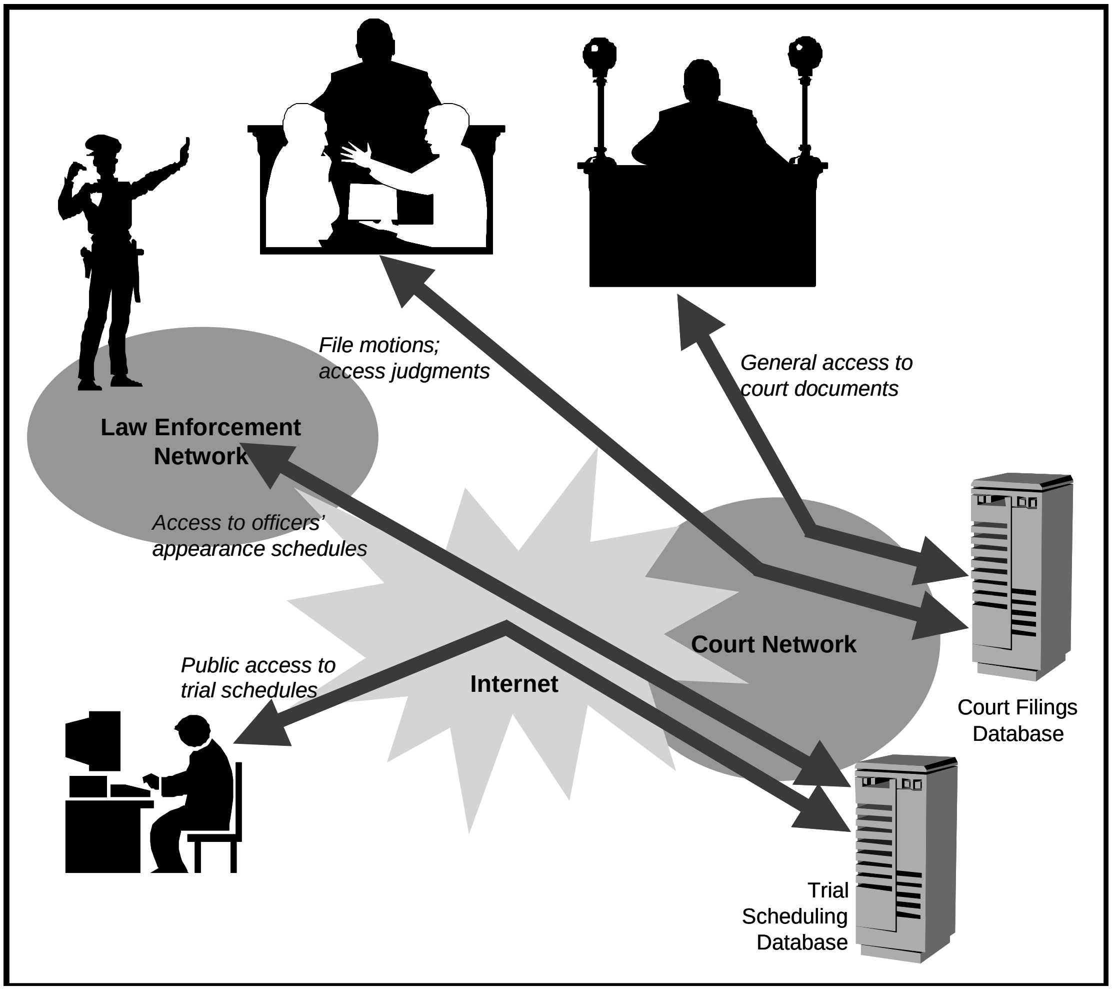

Figure 3–1. A Simplified Integrated Justice Scenario

In this scenario, confidentiality, integrity, and availability are important concerns. Assume that the court has had a bad experience with hacked Web sites and cannot allow the integrated justice automation project to proceed without assuring all stakeholders that security has been properly addressed. The IT staff implements the following security technology to help provide that assurance:

- The court places the server containing trial date information in its firewall DMZ. Because of previous problems with attacks on court Web sites, the court clerk digitally signs the Web pages that contain the dates. Periodically, a program that runs on a server on the private side of the firewall validates the signature. If the signature does not match, the system administrator is immediately notified.

- The court implements a VPN to allow police officers to check their court dates from a laptop in their patrol cars or home PCs. The VPN is implemented by the police and court firewalls. This configuration is shown in Figure 3–2.

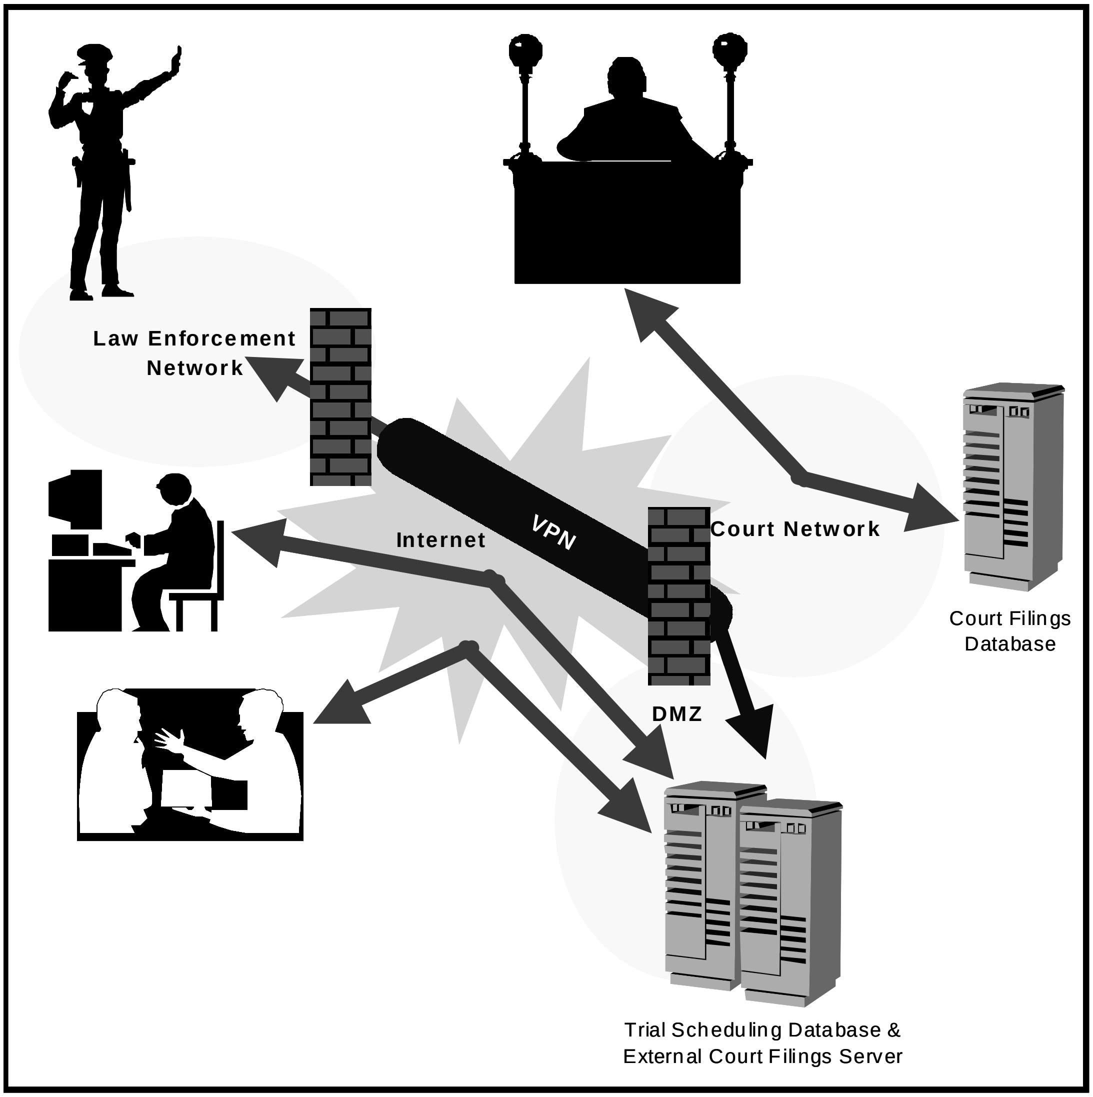

Figure 3–2. Security Technology to Enable Secure Information Exchange

- The requirement to be able to accept motions and post rulings online is a subset of an electronic case filing (ECF) application. This type of application requires both authentication and digital signature. In order to give judges and attorneys the ability to electronically sign documents, a public key infrastructure is required. In our example, we assume that the justice community participates in a PKI program run by a larger organization – perhaps the state government. Each judge is assigned a private key. The key may be stored, for example, in a smart card. The judge signs a document and then places the smart card into a reader attached to a PC and enters a PIN (Personal Identification Number). The court IT staff administers the distribution of private keys to participating attorneys; each key is distributed along with a PIN. The attorneys digitally sign their documents by running court software that requests their private key and PIN.

- Finally, the court IT staff installs intrusion detection software on a server to collect and analyze log information from the network and production servers. This software checks for known patterns of intrusion and alerts system administrators if there is suspicion that the system is under attack.

3.2 Fortifying Your System Against Hackers

Malicious hackers have been known to apply two criteria in selecting an information system to attack: (1) the system should be an easy target in which careless administration or lack of security precautions simplify gaining inappropriate access; and (2) either bragging rights, financial gain, or spite is associated with successfully compromising the system. Unfortunately, many government systems, and in some cases justice systems, meet one or both of these criteria. In this section, we will review some of the tactics employed to defend against attack.

Table 3–1 summarizes common attacks and the tactics that may be used to defend against them; the following paragraphs provide more detail.

Table 3–1. Defensive Tactics

Common Attack	Your Strategy (in brief)
Viruses and email denial of service	Anti-virus software on servers *and* workstation
Denial of service and distributed denial of service (IP flooding)	Good relationship with your ISP
Password attacks	Policy, education, automated enforcement
Product “holes,” worms, and Trojan Horses	Formal software patch/upgrade program
Sloppy configuration	Automated enterprise management
Social engineering	Policy and education

- Viruses and email denial of service. Probably the most common type of attack is an email-carried virus. Generally, the email will contain an executable attachment, entice the recipient to open the attachment, and then run the virus program on the recipient’s PC. The virus usually replicates itself by reading the address book on that host computer and sending copies of itself to every email address found. This action can result in a denial-of-service attack in which the associated email server, firewall, and potentially portions of the network are flooded with emails. The virus may also contain payloads that destroy or infect selected files on the host computer. Some sophisticated viruses launch from email without requiring the recipient to explicitly execute the attachment. For example, the recent NIMDA virus exploited a vulnerability in the Microsoft Outlook email client and executed if the user simply previewed email messages.

Most viruses can be prevented through the rigorous use of anti-virus software. Anti-virus software should run on all servers and clients in the enterprise. Virus control software on the email server can clear infected emails out of inboxes and prevent the virus from spreading further. Anti-virus software vendors provide updates for their files that identify known viruses. The system administrator should institute a regular program of installing these updates to all servers and clients in the enterprise.

Even with these precautions, the enterprise will remain vulnerable to new viruses that are not yet identified in the vendor's virus files. There are three ways to mitigate the risk from new viruses: (1) block all executable attachments in the email server before the corresponding messages are distributed to clients, (2) educate users on how to identify suspicious virus-carrying emails, and (3) train the system administrative staff to quickly respond to denial-of-service virus attacks by selectively shutting down sections of the network to isolate the virus and remove it from infected systems.

- Denial of service and distributed denial of service (IP flooding). IP flooding is a classic denial-of-service attack in which the attacker bombards a network with a high volume of data packets to clog network resources and slow or deny service to valid users. The source address may be real or "spoofed" (that is, the packets may appear to come from a network address that has been "stolen" for purposes of the attack). With the advent of intelligent routers and firewalls, a simple IP flooding attack is not very effective – system administrators can block traffic from the source address of the attack. A more sophisticated version is the distributed denial-of-service (DDOS) attack. In this exploit, the attacker sends a high volume of IP packets from many different IP addresses, making it hard to use address blocking to defeat the attack. The DDOS attack is difficult to launch, but it is also difficult to defeat. It usually requires that the organization under attack work with its ISP to cut the traffic flow off at its source and counter the DDOS.

- Password attacks. Passwords remain the most common form of identification and authentication. Unfortunately, passwords provide a very weak level of security. One reason is that users tend to pick simple passwords that are easy to remember. For example, there are approximately 50,000 words in the English dictionary. If a dictionary word is used as a password, it is a fairly quick and easy task for a computer program to try each one of the 50,000 and guess the password. When possible, system administrators should use software that enforces the selection of strong passwords (greater than 8 characters, with a mix of lowercase, uppercase, and special characters; and no simple words or names). Further, system administrators should periodically run security utilities that scan for weak passwords.

- Product "holes," worms, and Trojan horses: The complexity of system software has increased dramatically over the last decade. It is very difficult, almost humanly impossible, for vendors to produce software that is completely bug-free. Unfortunately, many of the bugs that exist in widely distributed products, such as those that provide Web, email, and file sharing services, can be exploited by hackers to launch attacks and gain unauthorized access. In some cases, a hacker will use a product hole to plant a worm or Trojan horse in a computer system. A worm is a program that replicates itself over a computer network and usually performs malicious actions, such as using up the computer's resources and possibly shutting the system down. A Trojan horse is a destructive program that masquerades as a benign application. Unlike viruses, Trojan horses do not replicate themselves, but they can be just as destructive. The defense against attacks that exploit product holes is to institute a program that regularly updates and installs security patches for key system software. This process becomes more complicated as the number and distribution of different software products in typical integrated justice systems and networks increase. Services are available that provide system administrators announcements of available updates and recommendations for installation.

- Sloppy configuration. A very common attack strategy is to probe for default system configurations and passwords that can provide a hacker unauthorized access. For example, some routers will use a default password such as "password" for the administrator's access. When a new router is installed, the default password may never get changed to a strong password. The defense against this type of attack is a good configuration management plan and execution. Automated enterprise management system products can be used to help ensure that the configurations of servers, routers, PCs, and other potential access points on the network are configured to prevent unauthorized access. Another strategy is to use network scanning software such as the shareware product "SATAN" to look for configuration vulnerabilities.

- Social engineering: Unfortunately the weakest link in any information security program is often the human link. In a simple example of a social engineering attack, a hacker looks up the name of the vice president's administrative assistant in the company phone book. The hacker then calls the network administrator claiming to be the administrative assistant and poses some fake emergency such as the vice president being out of town and forgetting the password to access email. The hacker may then ask that the vice president's password be reset to the default so the vice president can log in and check email and change the password. While it would take a fairly unsophisticated system administrator to be tricked by this type of ploy, there are others with clever twists that are very effective. The only defense against social engineering attacks is a good security training program that sensitizes users and administrators to potential attacks and makes them aware of proper security practices.

As final motivation for the importance of diligence in fortifying information systems, we present some of the observations drawn from studies conducted by a Carnegie Mellon Computer Emergency Response Team (CERT). The CERT research team analyzed reports of computer system intrusions over the time period from 1996 to 1999. Figure 3–3(a) shows what the CERT team *expected* to see in the pattern of intrusion reports. At the start of the analysis interval (month zero), a vulnerability in a popular computer system software is discovered. The number of reported intrusions that exploit this vulnerability gradually increase as the hacker community becomes aware of the potential attacks. When the software vendor issues a patch that corrects the vulnerability, the intrusions eventually drop off as system administrators install the patch or upgrade to a new version of the software. This was the expected pattern. It was not the pattern that the research team actually observed.

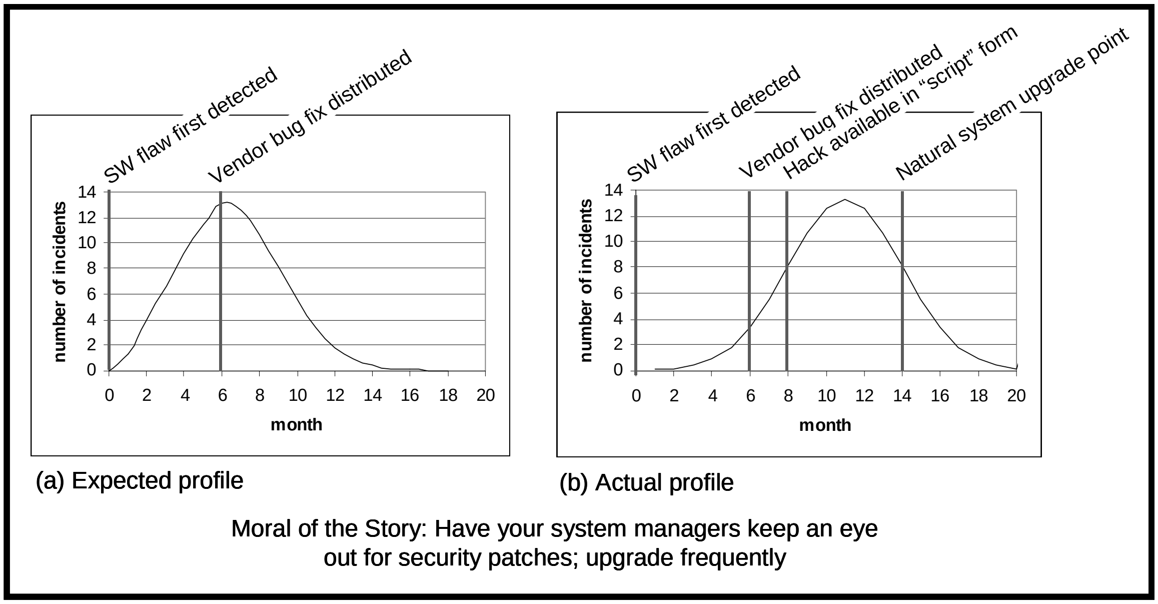

(a) Expected profile

(b) Actual profile

Moral of the Story: Have your system managers keep an eye out for security patches; upgrade frequently

Figure 3–3. Hacker Incident Reports

The *actual* observed pattern is represented in Figure 3–3(b). In this figure, the number of reported intrusions accelerates when the attacks that exploit this vulnerability are posted on the Internet in the form of automated scripts. These scripts make it easy for novice hackers to submit viruses and other programs that exploit the vulnerability. (The hacker community calls these less experienced hackers "script kiddies.") The researchers also found that the rate of intrusions reported does not start to diminish until many months after the software vendor distributes a corrective patch. The point at which there is a marked reduction in intrusion reports most likely corresponds to the point at which system software is upgraded – not for security purposes, but as a part of its natural upgrade cycle. This study shows that one way to reduce the frequency of intrusions is to be more rigorous about applying updates and security patches to server and system software. CERT estimates that 99 percent of all computer system attacks exploit vulnerabilities for which there is a known corrective measure.

3.3 Everything Else

Many technologies and technology applications that are important to information security in integrated justice cannot be addressed in this short overview because of space limitations. Three of these topics are described below; future updates to this document will include further detail on these and others:

- Security Policy, Architecture, and Training. Planning, design, and training are critical elements of maintaining information system security. Justice organizations should have a comprehensive security policy that includes descriptions of who is responsible for security, how information systems are protected, how individuals qualify for and acquire access privileges, and how information facilities are physically protected. Security technology is installed to implement the policy, so it is important that the policy is expressed clearly. The information security architecture provides the principles, strategy, and tactics applied to implement security policy. The architecture includes a technical architecture that describes how tools are implemented in the enterprise to provide confidentiality, integrity, and availability services. Finally, training is a critical element of information security. The human element is often the weakest link in the security chain. A good training program will familiarize system users and administrators with their responsibilities to protect valuable information assets.

- Database Security. Database security mechanisms are usually applied to provide "fine-grained" information control and implement need-to-know policies. In many situations, more than one organization will require access to subsets of information that are stored in the same physical database. This situation is becoming more frequent in integrated justice. For example, while all case-related data may be stored in the same database, it is very important to maintain separate data "views" for user groups such as those in the public defender's and district attorney's offices. While current database technology does not support rigorous multiple-level security models, most products allow programmers to restrict information access and modification privileges to user groups based on functional security classifications known as "roles." For justice organizations relying on application software that uses underlying databases, access rules may be implemented by that application. IJIS system designers and managers should be careful that all electronic paths to integrated databases implement their organization's data access and modification policies.

- Authentication and Access Control: Authentication and access control, in some ways, are the first line of defense in many information systems. These mechanisms provide a basic security function: they ensure that those wishing to gain access to information resources are indeed who they represent themselves to be and are accessing the system in ways consistent with their level of authority. There is increasing focus on authentication protocols and technology. Today, the most common form of authentication is password control, the limits of which have already been noted. In general, technologies for authenticating a potential user of an information system are organized into three identification factors: something you know, something you have, and something about you. An example of *something you know* is a password or a PIN. *Something you have* might be a smart card. *Something about you* is a biometric such as a fingerprint, iris pattern, facial pattern, or voice pattern. Highly secure systems generally use multiple factors. For example, a smart card is an excellent form of authentication. A smart card that contains a private key can be used to perform a public key encryption operation and conduct a very rigorous authentication. However, a smart card can be lost or stolen. As a result, engineers who incorporate smart cards into the security architecture may also specify that the user enter a PIN or submit to a fingerprint scan to unlock the key used on the card.

4. Summary and Resources

In all secure information system designs there must be a balance between the level of investment made in security and the risk of compromise. Maintaining that balance is ultimately the responsibility of the system owner. It requires that the owner understand costs associated with both the risk and the mitigation. This document has provided a basic introduction to security technology and a basis for understanding. However, information security is a broad and complex topic. This document concludes with a few information sources that augment this introduction.

4.1 Internet Sources

Much background and guidance information is available on the Web. The reader may be interested in exploring the following sites:

- NIST (csrc.nist.gov): The National Institute for Standards and Technology (NIST) Computer Security Resource Center (CSRC) sets standards for information security throughout the Federal government. The NIST Computer Security Resource Center (CSRC) distributes a wealth of information on security technology and practice. Their Web site contains much of this information.
- RSA (www.rsa.com): RSA is a public company that licenses one of the most commonly used public key encryption algorithms in the industry and markets many related products. Their Web site provides a frequently asked questions (FAQ) repository that is comprehensive and extremely informative.
- CERT (www.cert.org): CERT is associated with Carnegie Mellon University. It provides a nationwide service desk for reporting intrusions and finding information on preventive measures and related statistics.

4.2 Good Texts

The following text books provide detailed technical background on information security topics:

- Kaufman, C., R. Perlman, and M. Speciner. *Network Security, Private Communication in a Public World*. 1995, Englewood Cliffs, New Jersey: PTR Prentice Hall.
- Schneier, B., *Applied Cryptography*. Second edition, 1996, New York: John Wiley & Sons, Inc. 758.
- Howard, M. and D. LeBlanc, *Writing Secure Code.* 2001, Microsoft Press, ISBN 0–7356–1588–8.

4.3 The IJIS Institute

The IJIS Institute is a not-for-profit organization that fosters public/private partnerships in the implementation of integrated justice and provides services, expertise and consulting, either through grants or fee-for-service arrangements. Several industry participants in the Institute have significant experience in applying security technology to integrated justice applications. Please refer to their Web site at www.ijisinstitute.org to find more detail about working with the Institute.

[Inside back cover]

720 7th Street, N.W., Third Floor
Washington, D.C. 20001
Tel.: (202) 628-8615
Fax: (202) 628-0080
URL: http://www.ijisinstitute.org
http://www.ijis.org