LECTURE 1

Information security introduction

Telecommunication systems department

Lecturer: assistant professor Persikov Anatoliy Valentinovich

INFORMATION SECURITY DEFINITION

Information security means protecting information and information systems from unauthorized access, use, disclosure, disruption, modification or destruction.

The terms **information security**, **computer security** and **information assurance** are frequently incorrectly used interchangeably. These fields are interrelated often and share the common goals of protecting the **confidentiality**, **integrity** and **availability** of information; however, there are some subtle differences between them.

Confidentiality ensuring that information is accessible only to those authorized to have access.

Data integrity is a term used in computer science and telecommunications that can mean ensuring data is "whole" or complete, the condition in which data is identically maintained **during any operation** (such as transfer, storage or retrieval), the preservation of data for their intended use, or, relative to specified operations, the a priori expectation of data quality. Put simply, **data integrity** is the assurance that data is consistent and correct.

Availability is possibility of system (or it components) to be accessible.

DIFFERENCES OF TERMS

Computer security can focus on ensuring the availability and correct operation of a computer system without concern for the information stored or processed by the computer.

Governments, military, corporations, financial institutions, hospitals, and private businesses amass a great deal of confidential information about their employees, customers, products, research, and financial status. Most of this information is now collected, processed and stored on electronic computers and transmitted across networks to other computers.

Should confidential information about a business' customers or finances or new product line fall into the hands of a competitor, such a breach of security could lead to lost business, law suits or even bankruptcy of the business. **Protecting confidential information** is a **business requirement**, and in many cases also an ethical and legal requirement.

For the **individual**, information security has a significant effect on privacy, which is viewed very differently in different cultures.

LEGISLATIVE REGULATION

We will concentrate our attention to USA legislative norms for information security:

- Clinger-Cohen Act of 1996.
- Federal Information Security Management Act (FISMA) of 2002.
- Office of Management and Budget (OMB) Circular A-130.

The Clinger-Cohen Act (CCA), formerly the **Information Technology Management Reform Act of 1996** (ITMRA), is a 1996 United States federal law, designed to improve the way the federal government acquires, uses and disposes information technology (IT).

The Clinger-Cohen Act supplements the information resources management policies by establishing a comprehensive approach for executive agencies to improve the acquisition and management of their information resources, by:

- focusing information resource planning to support their strategic missions;
- implementing a capital planning and investment control process that links to budget formulation and execution; and
- rethinking and restructuring the way they do their work before investing in information systems.

FISMA

The Federal Information Security Management Act of 2002 ("FISMA", 44 U.S.C. § 3541, et seq.) is a United States federal law enacted in 2002 as Title III of the E-Government Act of 2002 (Pub.L. 107-347, 116 Stat. 2899). The act recognized the importance of information security to the economic and national security interests of the United States. The act requires each federal agency to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including those provided or managed by another agency, contractor, or other source.

FISMA has brought attention within the federal government to **cybersecurity** and explicitly emphasized a "**risk-based policy for cost-effective security.**" FISMA requires agency program officials, chief information officers, and inspectors general (IGs) to conduct annual reviews of the agency's information security program and report the results to Office of Management and Budget (OMB). OMB uses this data to assist in its oversight responsibilities and to prepare this annual report to Congress on agency compliance with the act.

In FY 2008, federal agencies spent $6.2 billion securing the government's total information technology investment of approximately $68 billion or about 9.2 percent of the total IT portfolio.

FISMA FRAMEWORK

FISMA defines a **framework for managing information security** that must be followed for all information systems used or operated by a U.S. federal government agency in the executive or legislative branches, or by a contractor or other organization on behalf of a federal agency in those branches. This framework is further defined by the standards and guidelines developed by NIST.

Major parts of framework are:

- Inventory of information systems;
- Categorize information and information systems according to risk level;
- Security controls;
- Risk assessment;
- System security plan;
- Certification and accreditation;
- Continuous monitoring.

INVENTORY OF INFORMATION SYSTEMS

FISMA requires that organizations have in place an information systems inventory. According to FISMA, the head of each agency shall develop and maintain an inventory of major information systems (including major national security systems) operated by or under the control of such agency The identification of information systems in an inventory under this subsection shall include an identification of the interfaces between each such system and all other systems or networks, including those not operated by or under the control of the agency.

The first step is to determine what constitutes the "information system" in question. There is not a direct mapping of computers to information system; rather, an information system may be a collection of individual computers put to a common purpose and managed by the same system owner.

NIST SP 800-18, Revision 1, Guide for Developing Security Plans for Federal Information Systems provides guidance on determining system boundaries.

CATEGORIZE INFORMATION AND INFORMATION SYSTEMS ACCORDING TO RISK LEVEL

All information and information systems should be categorized based on the objectives of providing appropriate levels of information security according to a range of risk levels. The first mandatory security standard required by the FISMA legislation, FIPS 199 "Standards for Security Categorization of Federal Information and Information Systems" provides the definitions of security categories. The guidelines are provided by NIST SP 800-60 "Guide for Mapping Types of Information and Information Systems to Security Categories."

The overall FIPS 199 system categorization is the "**high water mark**" for the impact rating of any of the criteria for information types resident in a system. For example, if one information type in the system has a rating of "Low" for "confidentiality," "integrity," and "availability," and another type has a rating of "Low" for "confidentiality" and "availability" but a rating of "Moderate" for "integrity," then the entire system has a FIPS 199 categorization of "Moderate."

SECURITY CONTROLS

Information systems must meet the minimum security requirements.

These requirements are defined in the second mandatory security standard required by the FISMA legislation, FIPS 200 "Minimum Security Requirements for Federal Information and Information Systems". The minimum security requirements cover seventeen security-related areas with regard to protecting the confidentiality, integrity, and availability of information systems and the information processed, stored, and transmitted by those systems:

1) access control
2) awareness and training
3) audit and accountability
4) certification, accreditation, and security assessments
5) configuration management
6) contingency planning
7) identification and authentication
8) incident response
9) maintenance
10) media protection
11) physical and environmental protection
12) planning
13) personnel security
14) risk assessment
15) systems and services acquisition
16) system and communications protection
17)system and information integrity

RISK ASSESSMENT

The organization's risk assessment validates the security control set and determines if any additional controls are needed to protect agency operations (including mission, functions, image, or reputation), agency assets, individuals or other organizations.

The resulting set of security controls establishes a level of “security due diligence” for the federal agency and its contractors.

A **risk assessment** starts by identifying potential threats and vulnerabilities and mapping implemented controls to individual vulnerabilities. One then determines risk by calculating the likelihood and impact that any given vulnerability could be exploited, taking into account existing controls. The culmination of the risk assessment shows the calculated risk for all vulnerabilities and describes whether the risk should be accepted or mitigated.

SYSTEM SECURITY PLAN

Organizations should develop policy on the system security planning process.

NIST SP-800-18 introduces the concept of a **System Security Plan**.

System security plans are living documents that require periodic review, modification, and plans of action and milestones for implementing security controls. Procedures should be in place outlining who reviews the plans, keeps the plan current, and follows up on planned security controls.

The System security plan is the major input to the security certification and accreditation process for the system. During the security certification and accreditation process, the system security plan is analyzed, updated, and accepted. The certification agent confirms that the security controls described in the system security plan are consistent with the FIPS 199 security category determined for the information system, and that the threat and vulnerability identification and initial risk determination are identified and documented in the system security plan, risk assessment, or equivalent document.

CERTIFICATION AND ACCREDITATION

Once the system documentation and risk assessment has been completed, the system's controls must be **reviewed and certified** to be functioning appropriately. Based on the results of the review, the information system is **accredited**. The certification and accreditation process is defined in NIST SP 800-37 "Guide for the Security Certification and Accreditation of Federal Information Systems".

Security accreditation is the official management decision given by a senior agency official to authorize operation of an information system and to explicitly accept the risk to agency operations, agency assets, or individuals based on the implementation of an agreed-upon set of security controls.

Security accreditation provides a form of quality control and challenges managers and technical staffs at all levels to implement the most effective security controls possible in an information system, given mission requirements, technical constraints, operational constraints, and cost/schedule constraints. By accrediting an information system, an organization official accepts responsibility for the security of the system and is fully accountable for any adverse impacts to the agency if a breach of security occurs.

CERTIFICATION AND ACCREDITATION

The information and supporting evidence needed for security accreditation is developed during a detailed security review of an information system, typically referred to as **security certification**.

Security certification is a comprehensive assessment of the management, operational, and technical security controls in an information system, made in support of security accreditation, to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system.

The results of a security certification are used to reassess the risks and update the system security plan, thus providing the factual basis for an authorizing official to render a security accreditation decision.

CONTINUOUS MONITORING

All accredited systems are required to monitor a selected set of security controls and the system documentation is updated to reflect changes and modifications to the system. Large changes to the security profile of the system should trigger an updated risk assessment, and controls that are significantly modified may need to be re-certified.

Continuous monitoring activities include configuration management and control of information system components, security impact analyses of changes to the system, ongoing assessment of security controls, and status reporting. The organization establishes the selection criteria and subsequently selects a subset of the security controls employed within the information system for assessment. The organization also establishes the schedule for control monitoring to ensure adequate coverage is achieved.

SECURITY CLASSIFICATION FOR INFORMATION

An important aspect of information security and risk management is **recognizing the value of information** and **defining appropriate procedures and protection requirements for the information**. Not all information is equal and so not all information requires the same degree of protection. This requires information to be assigned a security classification.

The first step in information classification is to identify a member of senior management as the owner of the particular information to be classified. Next, develop a classification policy. The policy should describe the different classification labels, define the criteria for information to be assigned a particular label, and list the required security controls for each classification.

Some factors that influence which classification information should be assigned include how much value that information has to the organization, how old the information is and whether or not the information has become obsolete.

Laws and other regulatory requirements are also important considerations when classifying information.

SECURITY CLASSIFICATION FOR INFORMATION

The type of information security classification labels selected and used will depend on the nature of the organization, with examples being:

- In the **business sector**, labels such as: **Public**, **Sensitive**, **Private**, **Confidential**.
- In the **government sector**, labels such as: **Unclassified**, **Sensitive But Unclassified**, **Restricted**, **Confidential**, **Secret**, **Top Secret** and their non-English equivalents.
- In **cross-sectoral formations**, the Traffic Light Protocol, which consists of: **White**, **Green**, **Amber** and **Red**.

All employees in the organization, as well as business partners, must be trained on the classification schema and understand the required security controls and handling procedures for each classification. The classification a particular information asset has been assigned should be reviewed periodically to ensure the classification is still appropriate for the information and to ensure the security controls required by the classification are in place.

SECURITY CLASSIFICATION FOR INFORMATION

There are four colors (or traffic lights):

- **RED – personal for named recipients only**. In the context of a meeting, for example, RED information is limited to those present at the meeting. In most circumstances, RED information will be passed verbally or in person.
- **AMBER – limited distribution**. The recipient may share AMBER information with others within their organization, but only on a 'need-to-know' basis. The originator may be expected to specify the intended limits of that sharing.
- **GREEN – community wide**. Information in this category can be circulated widely within a particular community. However, the information may not be published or posted on the Internet, nor released outside of the community.
- **WHITE – unlimited**. Subject to standard copyright rules, WHITE information may be distributed freely, without restriction.

NETWORK SECURITY

Network security consists of the provisions and policies adopted by a network administrator to prevent and monitor unauthorized access, misuse, modification, or denial of a computer network and network-accessible resources.

Network security involves the authorization of access to data in a network, which is controlled by the network administrator. Users choose or are assigned an ID and password or other authenticating information that allows them access to information and programs within their authority.

Network security covers a variety of computer networks, both public and private, that are used in everyday jobs conducting transactions and communications among businesses, government agencies and individuals. Networks can be private, such as within a company, and others which might be open to public access. Network security is involved in organizations, enterprises, and other types of institutions.

Main goals of network security (according to RFC standard series) are:

Peer entity authentication
Non-Repudiation
Systems Security
Unauthorized Usage
Inappropriate Usage
Denial of Service

Peer entity authentication. What we mean by this is that we know that one of the endpoints in the communication is the one we intended. Without peer entity authentication, it's very difficult to provide either confidentiality or data integrity. For instance, if we receive a message from Alice, the property of data integrity doesn't do us much good unless we know that it was in fact sent by Alice and not the attacker. Similarly, if we want to send a confidential message to Bob, it's not of much value to us if we're actually sending a confidential message to the attacker.

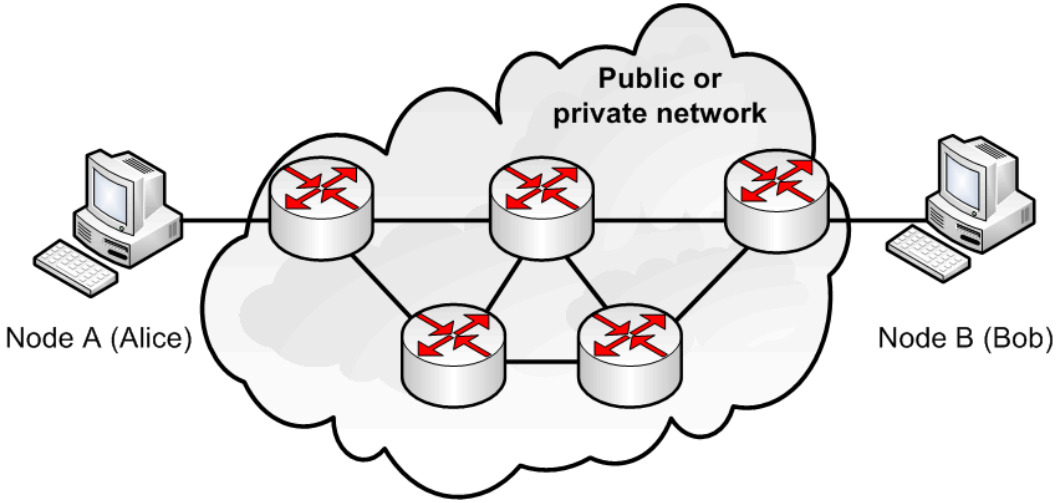

In messaging situations, you often wish to use peer entity authentication to establish the identity of the sender of a certain message. In such contexts, this property is called **data origin authentication**.

Non-Repudiation. A system that provides endpoint authentication allows one party to be certain of the identity of someone with whom he is communicating. When the system provides data integrity a receiver can be sure of both the sender's identity and that he is receiving the data that that sender meant to send. However, he cannot necessarily demonstrate this fact to a third party. The ability to make this demonstration is called **non-repudiation**.

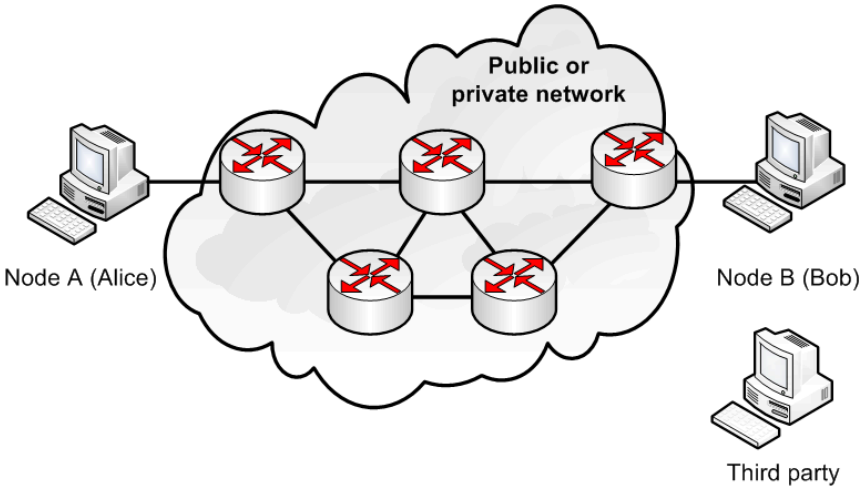

There are many situations in which non-repudiation is desirable. Consider the situation in which two parties have signed a contract which one party wishes to unilaterally abrogate. He might simply claim that he had never signed it in the first place. Non-repudiation prevents him from doing so, thus protecting the counterparty.

NETWORK SECURITY

Systems Security. In general, systems security is concerned with protecting one's machines and data. The intent is that machines should be used only by authorized users and for the purposes that the owners intend. Furthermore, they should be available for those purposes. Attackers should not be able to deprive legitimate users of resources.

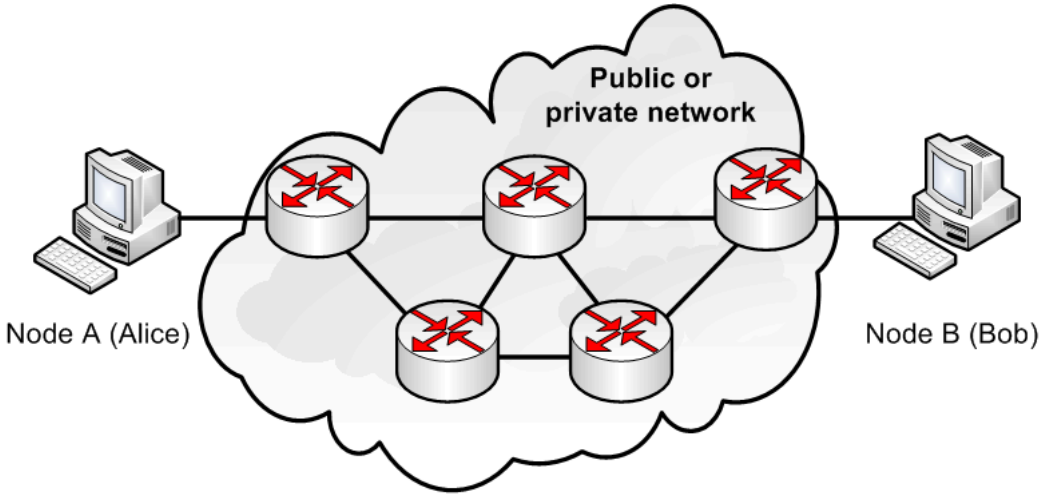

Web browser security	Network stack security	Protocol implementation security
Windows security	UNIX security	Cisco IOS security

NETWORK SECURITY

Unauthorized Usage. Most systems are not intended to be completely accessible to the public. Rather, they are intended to be used only by certain authorized individuals. Although many Internet services are available to all Internet users, even those servers generally offer a larger subset of services to specific users. For instance, Web Servers often will serve data to any user, but restrict the ability to modify pages to specific users. Such modifications by the general public would be UNAUTHORIZED USAGE.

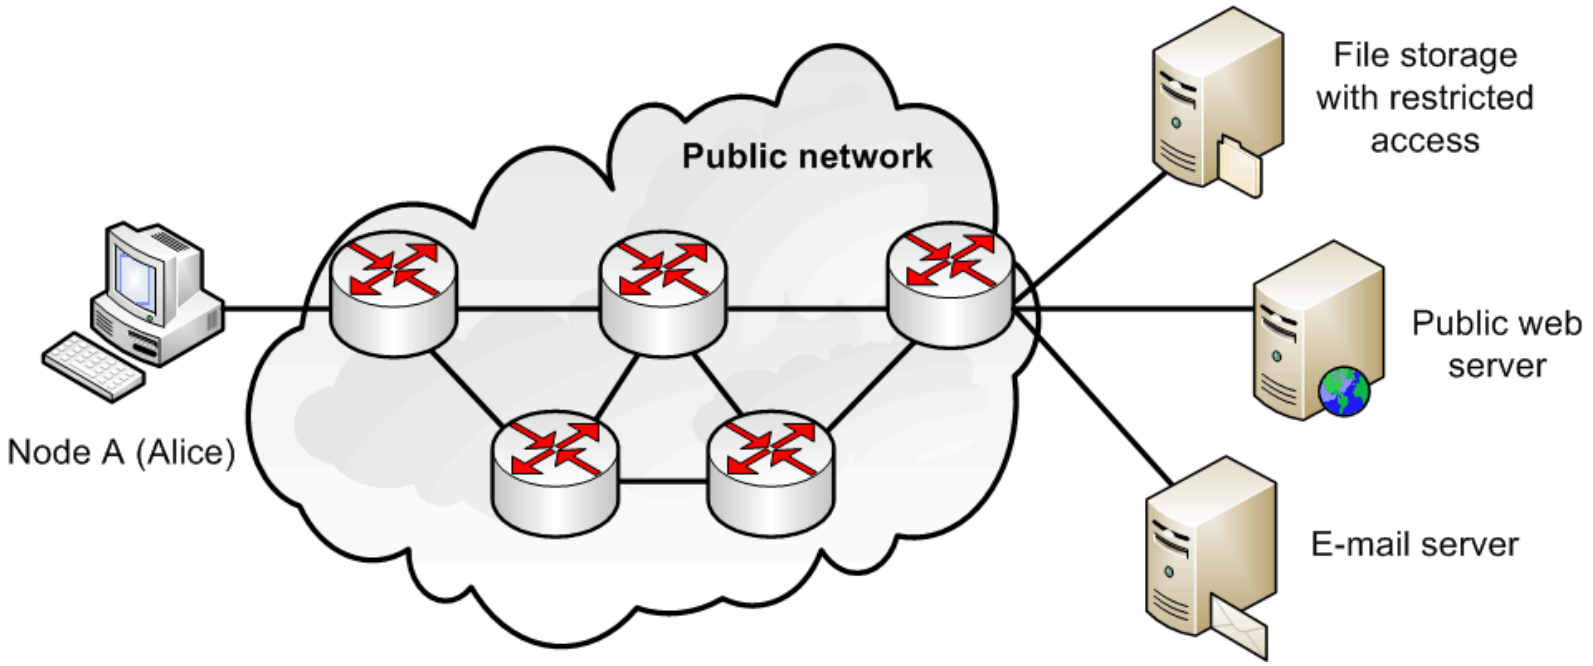

NETWORK SECURITY

Inappropriate Usage. Being an authorized user does not mean that you have free run of the system. As we said above, some activities are restricted to authorized users, some to specific users, and some activities are generally forbidden to all but administrators. Moreover, even activities which are in general permitted might be forbidden in some cases.

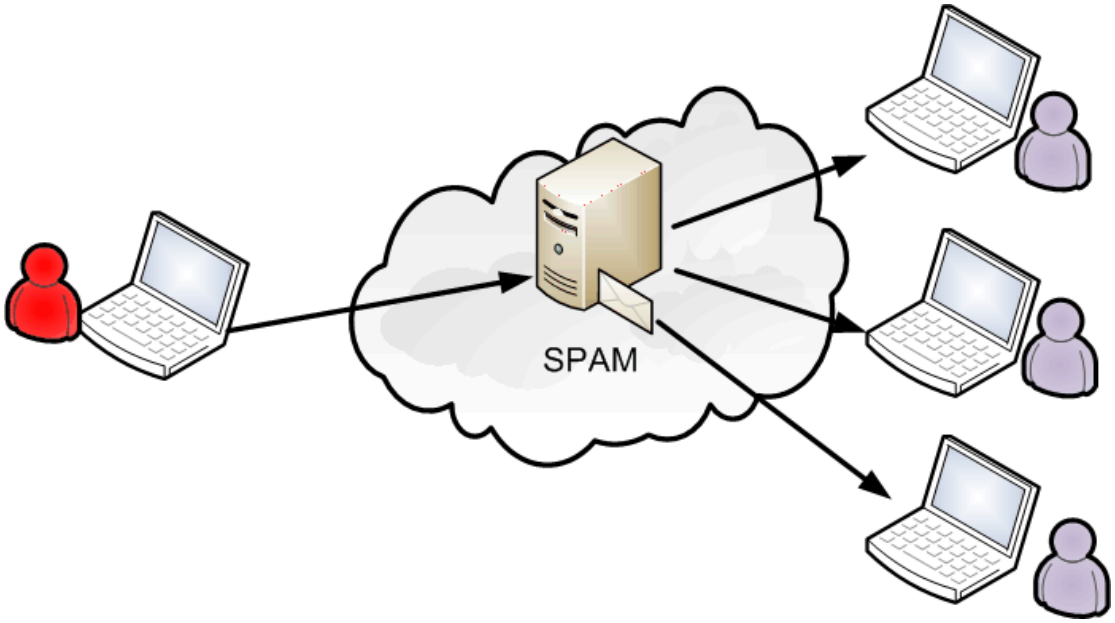

For instance, users may be permitted to send email but forbidden from sending files above a certain size, or files which contain viruses. These are examples of INAPPROPRIATE USAGE.

NETWORK SECURITY

Denial of Service. Recall that our third goal was that the system should be available to legitimate users. A broad variety of attacks are possible which threaten such usage. Such attacks are collectively referred to as **denial of service** attacks. Denial of service attacks are often very easy to mount and difficult to stop. Many such attacks are designed to consume machine resources, making it difficult or impossible to serve legitimate users. Other attacks cause the target machine to crash, completely denying service to users.

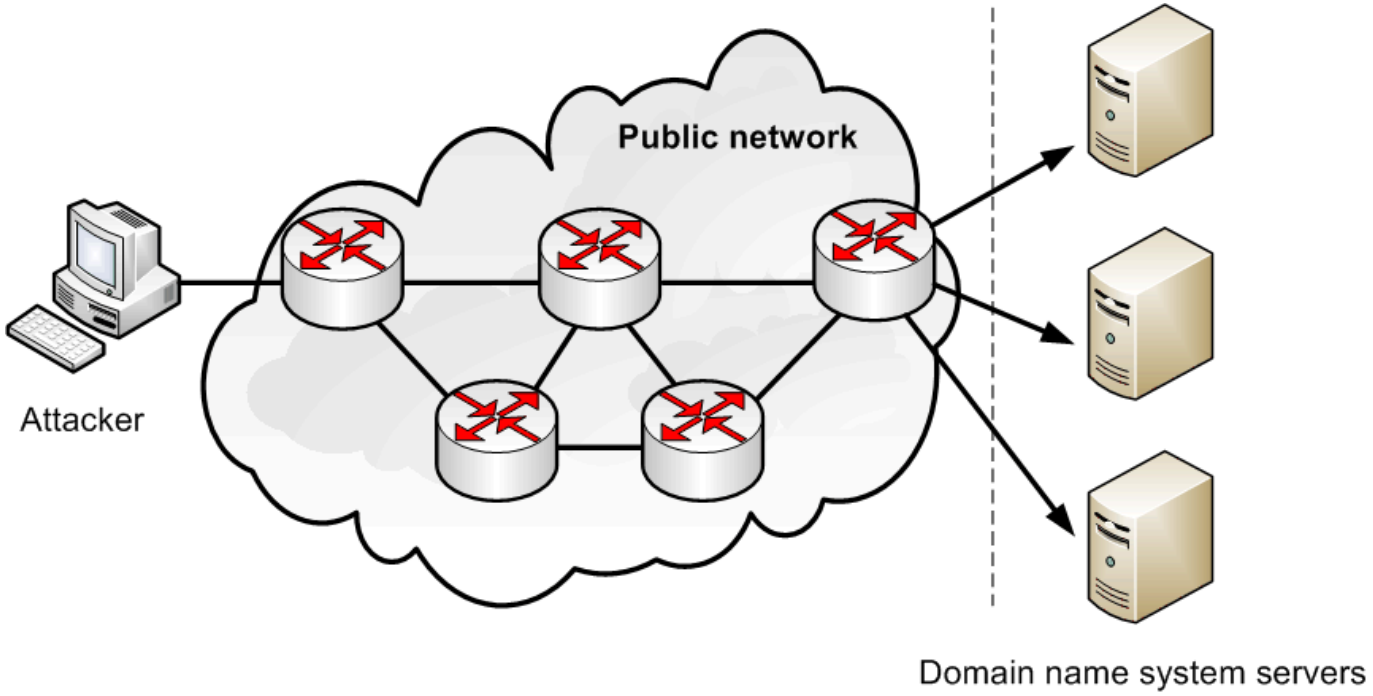

THE INTERNET THREAT MODEL

A **threat model** describes the capabilities that an attacker is assumed to be able to deploy against a resource. It should contain such information as the resources available to an attacker in terms of information, computing capability, and control of the system. The purpose of a threat model is twofold. First, we wish to identify the threats we are concerned with. Second, we wish to rule some threats explicitly out of scope. Nearly every security system is vulnerable to a sufficiently dedicated and resourceful attacker.

The Internet environment has a fairly well understood threat model. In general, we assume that the end-systems engaging in a protocol exchange have not themselves been compromised. Protecting against an attack when one of the end-systems has been compromised is extraordinarily difficult. It is, however, possible to design protocols which minimize the extent of the damage done under these circumstances.

A resourceful and dedicated attacker can control the entire communications channel.

However, a large number of attacks can be mounted by an attacker with fewer resources. A number of currently known attacks can be mounted by an attacker with limited control of the network.

THE INTERNET THREAT MODEL

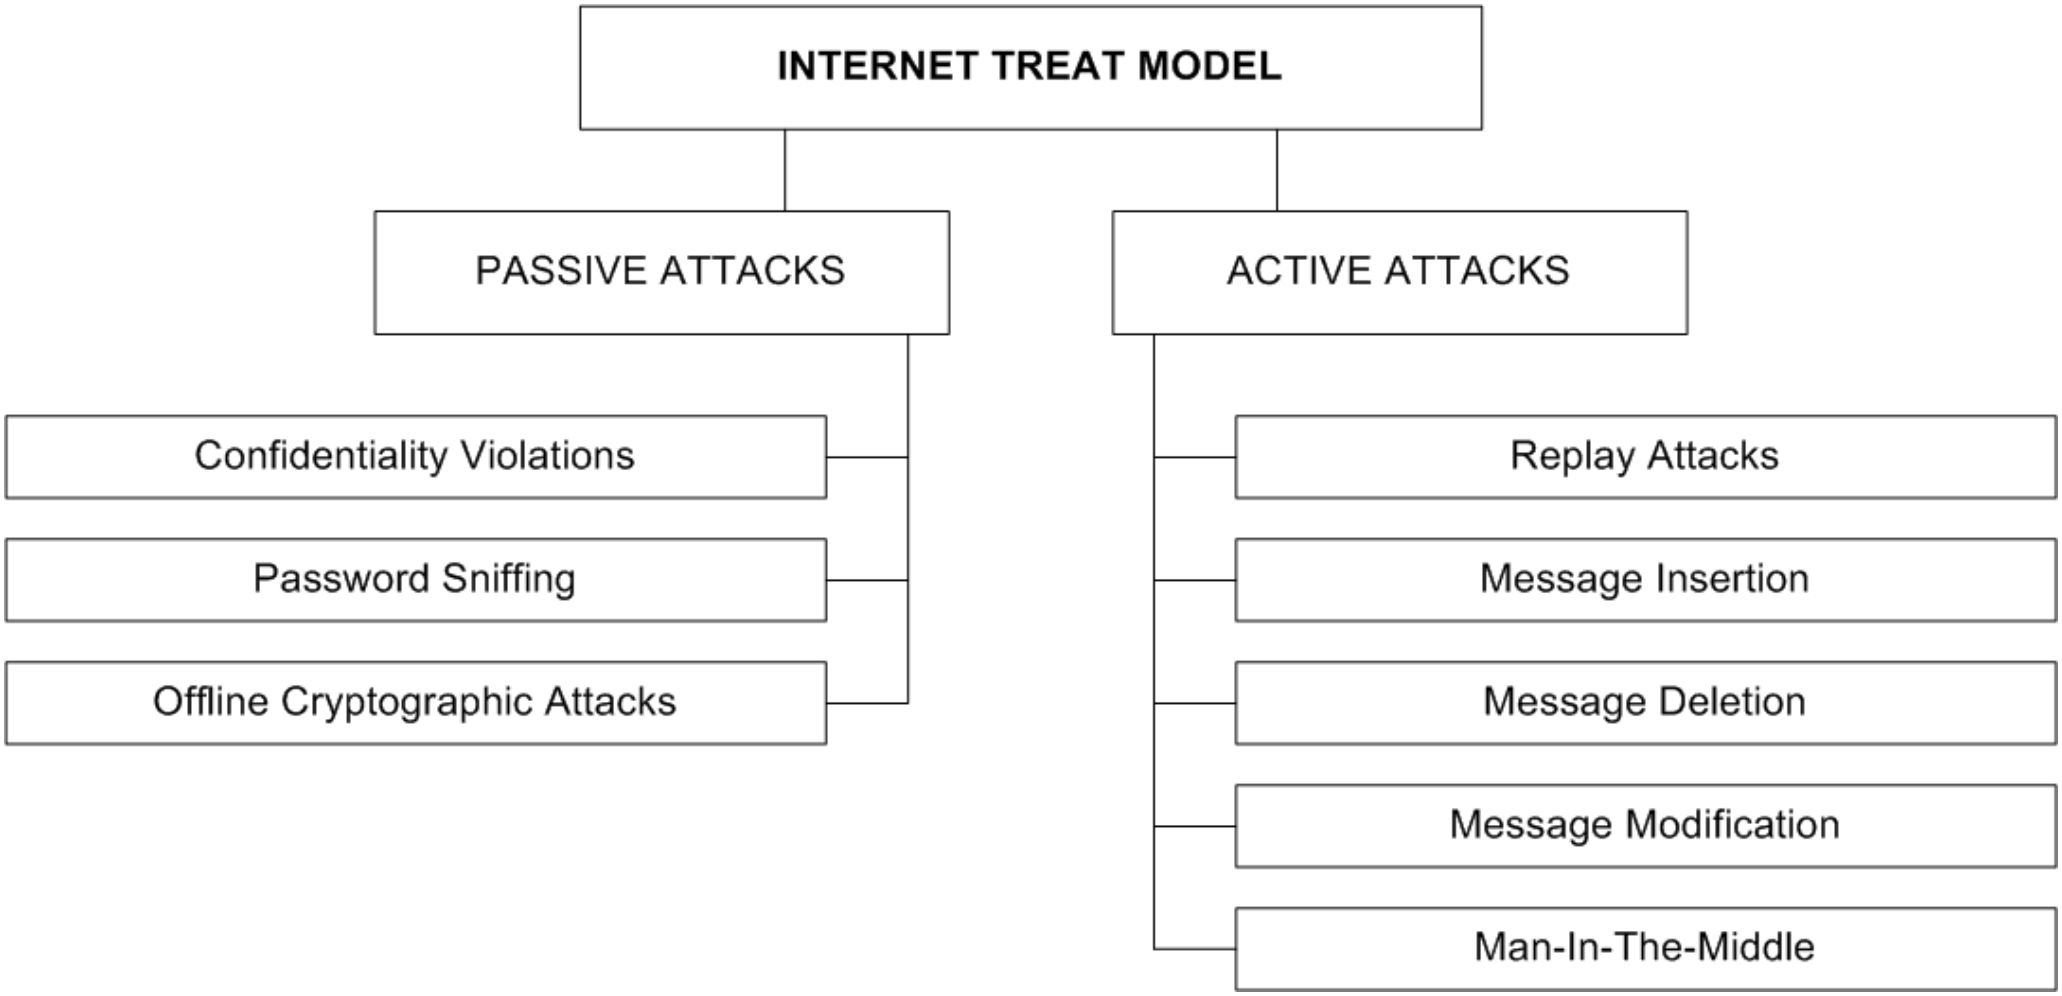

PASSIVE ATTACKS

In a passive attack, the attacker **reads packets** off the network but **does not write them**.

The simplest way to mount such an attack is to simply be on the same LAN as the victim. On most common LAN configurations, including Ethernet, 802.3, and FDDI, any machine on the wire can read all traffic destined for any other machine on the same LAN. Note that switching hubs make this sort of sniffing substantially more difficult, since traffic destined for a machine only goes to the network segment which that machine is on.

Similarly, an attacker who has control of a host in the communications path between two victim machines is able to mount a passive attack on their communications. It is also possible to compromise the routing infrastructure to specifically arrange that traffic passes through a compromised machine. This might involve an active attack on the routing infrastructure to facilitate a passive attack on a victim machine.

Wireless communications channels deserve special consideration, especially with the recent and growing popularity of wireless-based LANs, such as those using 802.11. Since the data is simply broadcast on well known radio frequencies, an attacker simply needs to be able to receive those transmissions. Such channels are especially vulnerable to passive attacks. Although many such channels include cryptographic protection, it is often of such poor quality as to be nearly useless.

PASSIVE ATTACKS

Confidentiality Violations. The classic example of passive attack is sniffing some inherently private data off of the wire. For instance, despite the wide availability of SSL, many credit card transactions still traverse the Internet in the clear. An attacker could sniff such a message and recover the credit card number, which can then be used to make fraudulent transactions. Moreover, confidential business information is routinely transmitted over the network in the clear in email.

Password Sniffing. Another example of a passive attack is PASSWORD SNIFFING. Password sniffing is directed towards obtaining unauthorized use of resources. Many protocols, including **telnet**, **pop**, and **nntp** use a shared password to authenticate the client to the server. Frequently, this password is transmitted from the client to the server in the clear over the communications channel. An attacker who can read this traffic can therefore capture the password and REPLAY it. In other words, the attacker can initiate a connection to the server and pose as the client and login using the captured password.

Note that although the login phase of the attack is active, the actual password capture phase is passive. Moreover, unless the server checks the originating address of connections, the login phase does not require any special control of the network.

PASSIVE ATTACKS

Offline Cryptographic Attacks. Many cryptographic protocols are subject to OFFLINE ATTACKS. In such a protocol, the attacker recovers data which has been processed using the victim's secret key and then mounts a cryptanalytic attack on that key. Passwords make a particularly vulnerable target because they are typically low entropy. A number of popular password-based challenge response protocols are vulnerable to **dictionary attack**. The attacker captures a challenge-response pair and then proceeds to try entries from a list of common words (such as a dictionary file) until he finds a password that produces the right response.

Historically, it has also been possible to exploit small operating system security holes to recover the password file using an active attack. These holes can then be bootstrapped into an actual account by using the aforementioned offline password recovery techniques. Thus we combine a low-level active attack with an offline passive attack.

ACTIVE ATTACKS

When an attack involves writing data to the network, we refer to this as an **active attack**. When IP is used without IPsec, there is no authentication for the sender address. As a consequence, it's straightforward for an attacker to create a packet with a source address of his choosing. We'll refer to this as a **spoofing attack**.

Under certain circumstances, such a packet may be screened out by the network. For instance, many packet filtering firewalls screen out all packets with source addresses on the **internal** network that arrive on the **external** interface. Note, however, that this provides no protection against an attacker who is inside the firewall. In general, designers should assume that attackers can forge packets.

However, the ability to forge packets does not go hand in hand with the ability to receive arbitrary packets. In fact, there are active attacks that involve being able to send forged packets but not receive the responses. We'll refer to these as **blind attacks**.

Note that not all active attacks require forging addresses. For instance, the TCP SYN denial of service attack can be mounted successfully without disguising the sender's address. However, it is common practice to disguise one's address in order to conceal one's identity if an attack is discovered.

Each protocol is susceptible to specific active attacks, but experience shows that a number of common patterns of attack can be adapted to any given protocol.

ACTIVE ATTACKS

Replay Attacks. In a replay attack, the attacker records a sequence of messages off of the wire and plays them back to the party which originally received them. Note that the attacker does not need to be able to understand the messages. He merely needs to capture and retransmit them.

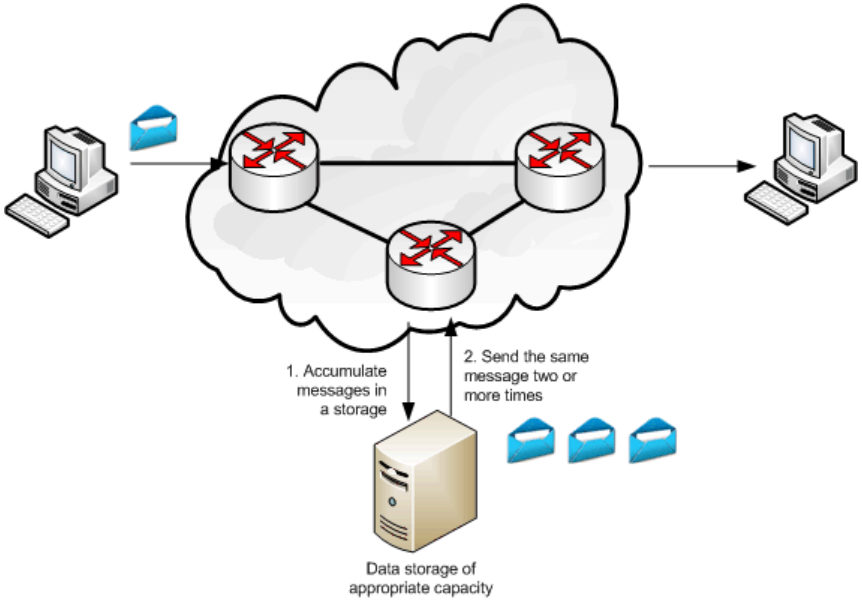

For example, consider the case where an S/MIME message is being used to request some service, such as a credit card purchase or a stock trade. An attacker might wish to have the service executed twice, if only to inconvenience the victim. He could capture the message and replay it, even though he can't read it, causing the transaction to be executed twice.

ACTIVE ATTACKS

Message Insertion. In a MESSAGE INSERTION attack, the attacker forges a message with some chosen set of properties and injects it into the network. Often this message will have a forged source address in order to disguise the identity of the attacker.

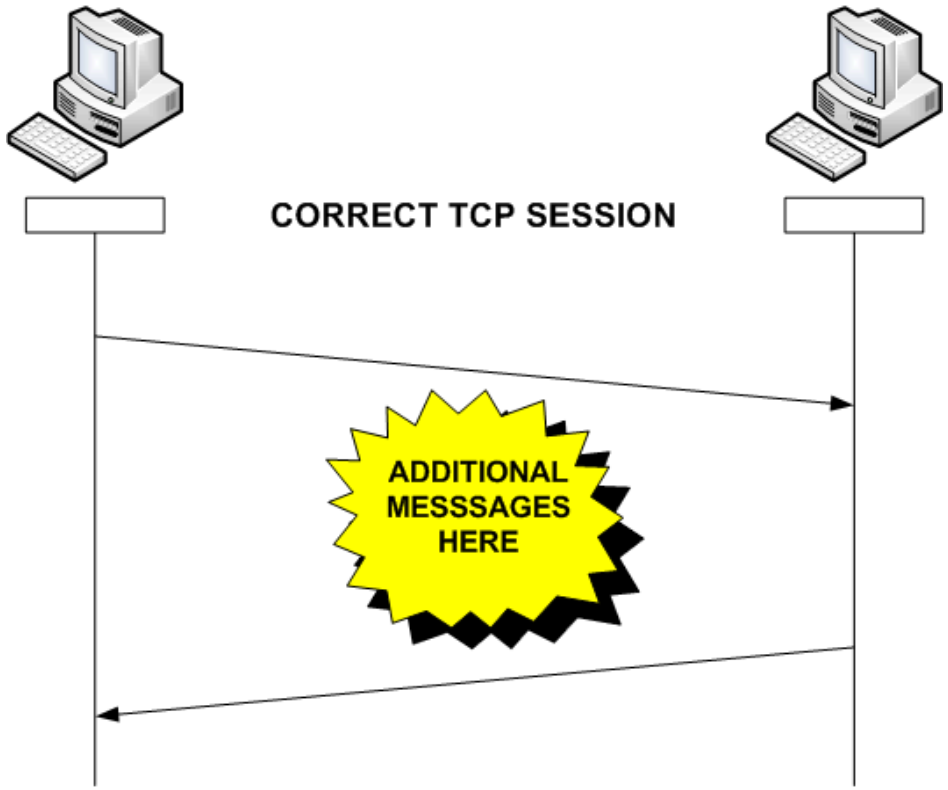

For example, a denial-of-service attack can be mounted by inserting a series of spurious TCP SYN packets directed towards the target host. The target host responds with its own SYN and allocates kernel data structures for the new connection. The attacker never completes the 3-way handshake, so the allocated connection endpoints just sit there taking up kernel memory.

Typical TCP stack implementations only allow some limited number of connections in this "half-open" state and when this limit is reached, no more connections can be initiated, even from legitimate hosts. Note that this attack is a blind attack, since the attacker does not need to process the victim's SYNs.

ACTIVE ATTACKS

Message Deletion. In a message deletion attack, the attacker removes a message from the wire. Morris' sequence number guessing attack often requires a message deletion attack to be performed successfully. In this blind attack, the host whose address is being forged will receive a spurious TCP SYN packet from the host being attacked. Receipt of this SYN packet generates a RST, which would tear the illegitimate connection down. In order to prevent this host from sending a RST so that the attack can be carried out successfully, Morris describes flooding this host to create queue overflows such that the SYN packet is lost and thus never responded to.

Message Modification. In a message modification attack, the attacker removes a message from the wire, modifies it, and reinjects it into the network. This sort of attack is particularly useful if the attacker wants to send some of the data in the message but also wants to change some of it.

Consider the case where the attacker wants to attack an order for goods placed over the Internet. He doesn't have the victim's credit card number so he waits for the victim to place the order and then replaces the delivery address (and possibly the goods description) with his own. Note that this particular attack is known as a CUT-AND-PASTE attack since the attacker cuts the credit card number out of the original message and pastes it into the new message.

ACTIVE ATTACKS

Man-In-The-Middle. A MAN-IN-THE-MIDDLE attack combines the above techniques in a special form: The attacker subverts the communication stream in order to pose as the sender to receiver and the receiver to the sender.

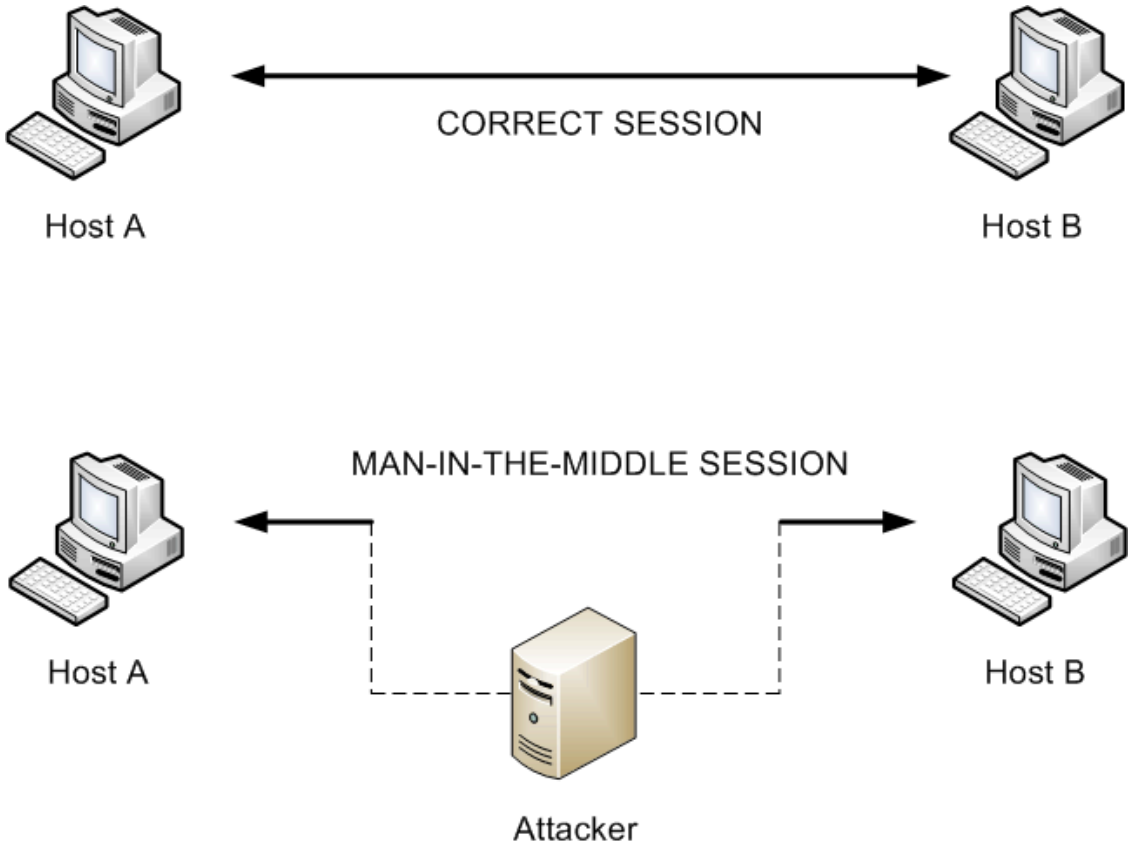

ACTIVE ATTACKS

Man-In-The-Middle. This differs fundamentally from the above forms of attack because it attacks the identity of the communicating parties, rather than the data stream itself. Consequently, many techniques which provide integrity of the communications stream are insufficient to protect against man-in-the-middle attacks.

Man-in-the-middle attacks are possible whenever a protocol lacks **peer entity authentication**. For instance, if an attacker can hijack the client TCP connection during the TCP handshake (perhaps by responding to the client's SYN before the server does), then the attacker can open another connection to the server and begin a man-in-the-middle attack. It is also trivial to mount man-in-the-middle attacks on local networks via ARP spoofing – the attacker forges an ARP with the victim's IP address and his own MAC address. Tools to mount this sort of attack are readily available.

Note that it is only necessary to authenticate one side of the transaction in order to prevent man-in-the-middle attacks. In such a situation the peers can establish an association in which only one peer is authenticated. In such a system, an attacker can initiate an association posing as the unauthenticated peer but cannot transmit or access data being sent on a legitimate connection. This is an acceptable situation in contexts such as Web e-commerce where only the server needs to be authenticated (or the client is independently authenticated via some non-cryptographic mechanism such as a credit card number).

THANKS FOR ATTENTION