Q.2 a. Differentiate between active and passive attacks. List some passive attacks and some active attacks.

Answer:

Passive attacks are in the nature of eavesdropping on, or monitoring of, transmissions. The goal of the opponent is to obtain information that is being transmitted.

Two types of passive attacks are the release of message contents and traffic analysis.
Passive attacks are very difficult to detect, because they do not involve any alteration of the data. Typically, the message traffic is sent and received in an apparently normal fashion, and neither the sender nor receiver is aware that a third party has read the messages or observed the traffic pattern. However, it is feasible to prevent the success of these attacks, usually by means of encryption. Thus, the emphasis in dealing with passive attacks is on prevention rather than detection.

Active attacks involve some modification of the data stream or the creation of a false stream and can be subdivided into four categories: masquerade, replay, modification of messages, and denial of service.
Active attacks present the opposite characteristics of passive attacks. Whereas passive attacks are difficult to detect, measures are available to prevent their success.
On the other hand, it is quite difficult to prevent active attacks absolutely because of the wide variety of potential physical, software, and network vulnerabilities. Instead, the goal is to detect active attacks and to recover from any disruption or delays caused by them. If the detection has a deterrent effect, it may also contribute to prevention.
Passive attacks: release of message contents and traffic analysis.
Active attacks: masquerade, replay, modification of messages, and denial of service.

b.Determine gcd (24140, 16762)

Answer:

gcd(24140, 16762) = gcd(16762, 7378) = gcd(7378, 2006) = gcd(2006, 1360) = gcd(1360, 646) = gcd (646, 68) = gcd(68, 34) = gcd(34, 0) = 34

c. The example used by Sun-Tsu to illustrate the CRT was
$x \equiv 2 \pmod 3$; $x \equiv 3 \pmod 5$; $x \equiv 2 \pmod 7$
Solve for x.

Answer:

M=3*5*7=105; M/3=35; M/5=21;M/7=15
The set of linear congruences
$35b1 \equiv 1 \pmod 3$; $21b2 \equiv 1 \pmod 5$; $15b3 \equiv 1 \pmod 7$

lead to solutions
b1=2 ; b2=1 ;b3=1
then
x= 2*2*35+1*3*21+1*2*15=140+63+30=233 (mod105)=23

Q.3 a. What are the essential ingredients of a symmetric cipher? Explain briefly

Answer:
A symmetric encryption scheme has five ingredients:

• Plaintext: This is the original intelligible message or data that is fed into the algorithm as input.

• Encryption algorithm: The encryption algorithm performs various substitutions and transformations on the plaintext.

• Secret key: The secret key is also input to the encryption algorithm. The key is a value independent of the plaintext and of the algorithm. The algorithm will produce a different output depending on the specific key being used at the time. The exact substitutions and transformations performed by the algorithm depend on the key.

• Ciphertext: This is the scrambled message produced as output. It depends on the plaintext and the secret key. For a given message, two different keys will produce two different ciphertexts. The ciphertext is an apparently random stream of data and, as it stands, is unintelligible.

• Decryption algorithm: This is essentially the encryption algorithm run in reverse. It takes the ciphertext and the secret key and produces the original plaintext

b. Define a P-box and list its three variations. Which variation is invertible?

Answer:
A P-box (permutation box) transposes bits.
We have three types of P-boxes in modern
block ciphers: straight P-boxes, expansion P-boxes, and compression P-boxes.
A straight P-box is invertible; the other two are not.

c. Explain synchronous stream cipher and one-time pad. Why one-time pad cipher is not practical?

Answer:
theoretically unbreakable (Claude Shannon)

☐the plaintext is combined with a random "pad" the same length as the plaintext.

☐Patent by Gilbert Vernam (AT&T) and Joseph Mauborgne

☐Encryption C=P ☐K

☐Decryption P=C ☐K

☐Claude Shannon's work can be interpreted as that any information-theoretically secure cipher will be effectively equivalent to the one-time pad algorithm. Hence one-time pads offer the best possible mathematical security of any encryption scheme, anywhere and anytime. one-time pad cipher is not practical because the keys need to be changed for each communication.

Drawbacks

it requires secure exchange of the one-time pad material, which must be as long as the message pad disposed of correctly and never reused

In practice Generate a large number of random bits,Exchange the key material securely between the users before sending an one-time enciphered message, Keep both copies of the key material for each message securely until they are used, and Securely dispose of the key material after use, thereby ensuring the key material is never reused.

Q.4 a. The criteria used in the design of DES focused on the design of the S-boxes and on the P function that takes the output of the S-boxes. List and briefly explain these criterions

Answer:

The criteria for the S-boxes are as follows.

1. No output bit of any S-box should be too close a linear function of the input bits. Specifically, if we select any output bit and any subset of the six input bits, the fraction of inputs for which this output bit equals the XOR of these input bits should not be close to 0 or 1, but rather should be near 1/2.

2. Each row of an S-box (determined by a fixed value of the leftmost and right- most input bits) should include all 16 possible output bit combinations.

3. If two inputs to an S-box differ in exactly one bit, the outputs must differ in at least two bits.

4. If two inputs to an S-box differ in the two middle bits exactly, the outputs must differ in at least two bits.

5. If two inputs to an S-box differ in their first two bits and are identical in their last two bits, the two outputs must not be the same.

6. For any nonzero 6-bit difference between inputs, no more than eight of the 32 pairs of inputs exhibiting that difference may result in the same output difference.

7. This is a criterion similar to the previous one, but for the case of three S-boxes.

The criteria for the permutation P are as follows.

1. The four output bits from each S-box at round i are distributed so that two of them affect (provide input for) "middle bits" of round (i + 1) and the other two affect end bits. The two middle bits of input to an S-box are not shared with adjacent S-boxes. The end bits are the two left-hand bits and the two right-hand bits, which are shared with adjacent S-boxes.

2. The four output bits from each S-box affect six different S-boxes on the next round, and no two affect the same S-box.

3. For two S-boxes j, k, if an output bit from Sj affects a middle bit of Sk on the next round, then an output bit from Sk cannot affect a middle bit of Sj. This implies that, for j = k, an output bit from Sj must not affect a middle bit of Sj. These criteria are intended to increase the diffusion of the algorithm.

b. Explain the avalanche effect using a suitable example.

Answer:
The Avalanche Effect

A desirable property of any encryption algorithm is that a small change in either the plaintext or the key should produce a significant change in the ciphertext. In particular, a change in one bit of the plaintext or one bit of the key should produce a change in many bits of the ciphertext. This is referred to as the avalanche effect. If the change were small, this might provide a way to reduce the size of the plaintext or key space to be searched.
(Give one example)

Q.5 a. Describe CTR mode. Write the encryption algorithm for CTR. Also list its advantages.

Answer: Page no 238 of Text Book.

b. What is a one-way function? What is a trap-door one-way function? Give an example of each.

Answer:

A **one-way function** is one that maps a domain into a range such that every function value has a unique inverse, with the condition that the calculation of the function is easy whereas the calculation of the inverse is infeasible:
A **trap-door one-way function** is easy to calculate in one direction and infeasible to calculate in the other direction unless certain additional information is known. With the additional information the inverse can be calculated in polynomial time.

Q.6 a. What are the motivations behind developing MACs based on hash functions? Describe design objectives and overall operation of HMAC.

Answer:
The motivations for this interest are

1. Cryptographic hash functions such as MD5 and SHA generally execute faster in software than symmetric block ciphers such as DES.

2. Library code for cryptographic hash functions is widely available. Following are design objectives for HMAC.

• To use, without modifications, available hash functions. In particular, to use hash functions that perform well in software and for which code is freely and widely available.

• To allow for easy replaceability of the embedded hash function in case faster or more secure hash functions are found or required.

• To preserve the original performance of the hash function without incurring a significant degradation.

• To use and handle keys in a simple way.

• To have a well understood cryptographic analysis of the strength of the authentication mechanism based on reasonable assumptions about the embedded hash function
(Page 355 of reference book)

b. Explain procedure of Message Digest (MD) generation using SHA-512.

Answer:
The algorithm takes as input a message with a maximum length of less than 2128 bits and produces as output a 512-bit message digest. The input is processed in 1024-bit blocks.
Step 1 Append padding bits. The message is padded so that its length is congruent to 896 modulo 1024 [length K 896(mod 1024)]. Padding is always added, even if the message is already of the desired length. Thus, the number of padding bits is in the range of 1 to 1024. The padding consists of a single 1 bit followed by the necessary number of 0 bits.

Step 2 Append length. A block of 128 bits is appended to the message. This block is treated as an unsigned 128-bit integer (most significant byte first) and contains the length of the original message (before the padding).

The outcome of the first two steps yields a message that is an integer multiple of 1024 bits in length. In Figure 11.8, the expanded message is represented as the sequence of 1024-bit blocks M1, M2, Á , MN, so that the total length of the expanded message is N * 1024 bits.

Step 3 Initialize hash buffer. A 512-bit buffer is used to hold intermediate and final results of the hash function. The buffer can be represented as eight 64-bit registers (a, b, c, d, e, f, g, h). These registers are initialized to the following 64-bit integers (hexadecimal values):

a = 6A09E667F3BCC908 e = 510E527FADE682D1

b = BB67AE8584CAA73B f = 9B05688C2B3E6C1F

c = 3C6EF372FE94F82B g = 1F83D9ABFB41BD6B

d = A54FF53A5F1D36F1 h = 5BE0CD19137E2179

These values are stored in big-endian format, which is the most significant byte of a word in the low-address (leftmost) byte position. These words were obtained by taking the first sixty-four bits of the fractional parts of the square roots of the first eight prime numbers.

Step 4 Process message in 1024-bit (128-word) blocks. The heart of the algorithm is a module that consists of 80 rounds
Step 5 Output. After all N 1024-bit blocks have been processed, the output from the Nth stage is the 512-bit message digest.

Q7 a. Describe briefly the kind of attacks on digital signatures

Answer:
Here A denotes the user whose signature method is being attacked, and C denotes the attacker.

• Key-only attack: C only knows A's public key.

• Known message attack: C is given access to a set of messages and their signatures.

• Generic chosen message attack: C chooses a list of messages before attempting to breaks A's signature scheme, independent of A's public key. C then obtains from A valid signatures for the chosen messages. The attack is generic, because it does not depend on A's public key; the same attack is used against everyone.

• Directed chosen message attack: Similar to the generic attack, except that the list of messages to be signed is chosen after C knows A's public key but before any signatures are seen.

• Adaptive chosen message attack: C is allowed to use A as an "oracle." This means the A may request signatures of messages that depend on previously obtained message–signature pairs

b. What problem was Kerberos designed to address? In the context of Kerberos, what is a realm?

Answer:
The problem that Kerberos addresses is this: Assume an open distributed environment in which users at workstations wish to access services on servers distributed throughout the network. We would like for servers to be able to restrict access to authorized users and to be able to authenticate requests for service. In this environment, a workstation cannot be trusted to identify its users correctly to network services.

A realm is an environment in which: 1. The Kerberos server must have the user ID (UID) and hashed password of all participating users in its database. All users are registered with the Kerberos server. 2. The Kerberos server must share a secret key with each server. All servers are registered with the Kerberos server.

c. Describe man-in-the-middle attack. How can such vulnerabilities be overcome?

Answer:
Suppose Alice and Bob wish to exchange keys, and Darth is the adversary. The attack proceeds as follows.

1. Darth prepares for the attack by generating two random private keys XD1 and XD2 and then computing the corresponding public keys YD1 and YD2.

2. Alice transmits YA to Bob.

3. Darth intercepts YA and transmits YD1 to Bob. Darth also calculates K2 = (YA)XD2 mod q .

4. Bob receives YD1 and calculates K1 = (YD1)XB mod q .

5. Bob transmits YB to Alice.

6. Darth intercepts YB and transmits YD2 to Alice. Darth calculates K1 = (YB)XD1 mod q .

7. Alice receives YD2 and calculates K2 = (YD2)XA mod q .

At this point, Bob and Alice think that they share a secret key, but instead Bob and Darth share secret key K1 and Alice and Darth share secret key K2. All future communication between Bob and Alice is compromised in the following way.

1. Alice sends an encrypted message M: E(K2, M) .

2. Darth intercepts the encrypted message and decrypts it to recover M.

3. Darth sends Bob E(K1, M) or E(K1, M¿) , where M¿ is any message. In the first case, Darth simply wants to eavesdrop on the communication without altering it. In the second case, Darth wants to modify the message going to Bob.

The key exchange protocol is vulnerable to such an attack because it does not authenticate the participants. This vulnerability can be overcome with the use of digital signatures and public-key certificates;

Q8. a. Describe briefly the five header fields defined in MIME.

Answer:
The five header fields defined in MIME are

• MIME-Version: Must have the parameter value 1.0. This field indicates that the message conforms to RFCs 2045 and 2046.

• Content-Type: Describes the data contained in the body with sufficient detail that the receiving user agent can pick an appropriate agent or mechanism to represent the data to the user or otherwise deal with the data in an appropriate manner.

• Content-Transfer-Encoding: Indicates the type of transformation that has been used to represent the body of the message in a way that is acceptable for mail transport.

• Content-ID: Used to identify MIME entities uniquely in multiple contexts.

• Content-Description: A text description of the object with the body; this is useful when the object is not readable (e.g., audio data).

b. How does PGP use the concept of trust? Describe the operation of the trust processing.

Answer:
Although PGP does not include any specification for establishing certifying authorities or for establishing trust, it does provide a convenient means of using trust, associating trust with public keys, and exploiting trust information.
The basic structure is as follows. Each entry in the public-key ring is a public-key certificate, as described in the preceding subsection. Associated with each such entry is a key legitimacy field that indicates the extent to which PGP will trust that this is a valid public key for this user; the higher the level of trust, the stronger is the binding of this user ID to this key. This field is computed by PGP. Also associated with the entry are zero or more signatures that the key ring owner has collected that sign this certificate. In turn, each signature has associated with it a signature trust field that indicates the degree to which this PGP user trusts the signer to certify public keys. The key legitimacy field is derived from the collection of signature trust fields in the entry. Finally, each entry defines a public key associated with a particular owner, and an owner trust field is included that indicates the degree to which this public key is trusted to sign other public-

key certificates; this level of trust is assigned by the user. We can think of the signature trust fields as cached copies of the owner trust field from another entry.
We can describe the operation of the trust processing as follows.

1. When A inserts a new public key on the public-key ring, PGP must assign a value to the trust flag that is associated with the owner of this public key. If the owner is A, and therefore this public key also appears in the private-key ring, then a value of ultimate trust is automatically assigned to the trust field.

Otherwise, PGP asks A for his assessment of the trust to be assigned to the owner of this key, and A must enter the desired level. The user can specify that this owner is unknown, untrusted, marginally trusted, or completely trusted.

2. When the new public key is entered, one or more signatures may be attached to it. More signatures may be added later. When a signature is inserted into the entry, PGP searches the public-key ring to see if the author of this signature is among the known public-key owners. If so, the OWNERTRUST value for this owner is assigned to the SIGTRUST field for this signature. If not, an unknown user value is assigned.

3. The value of the key legitimacy field is calculated on the basis of the signature trust fields present in this entry. If at least one signature has a signature trust value of ultimate, then the key legitimacy value is set to complete. Otherwise,PGP computes a weighted sum of the trust values. A weight of 1/X is given to signatures that are always trusted and 1/Y to signatures that are usually trusted, where X and Y are user-configurable parameters. When the total of weights of the introducers of a Key/UserID combination reaches 1, the binding is considered to be trustworthy, and the key legitimacy value is set to complete. Thus, in the absence of ultimate trust, at least X signatures that are always trusted, Y signatures that are usually trusted, or some combination is needed.

c. Describe one-way e-mail exchange architecture.

Answer:

Figure 16.1 *E-mail architecture*

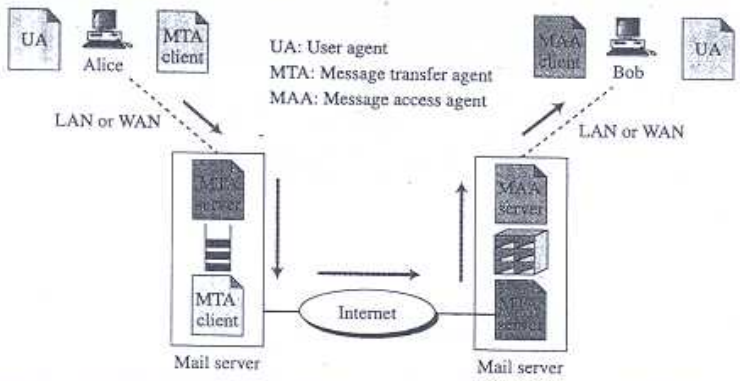

be connected to the e-mail server of an ISP through a WAN (telephone line or cable line). Bob is also in one of the above two situations.

The administrator of the e-mail server at Alice's site has created a queuing system that sends e-mail to the Internet one by one. The administrator of the e-mail server at Bob's site has created a mailbox for every user connected to the server; the mailbox holds the received messages until they are retrieved by the recipient.

When Alice needs to send a message to Bob, she invokes a **user agent (UA)** program to prepare the message. She then uses another program, a **message transfer agent (MTA),** to send the message to the mail server at her site. Note that the MTA is a client/server program with the client installed at Alice's computer and the server installed at the mail server.

The message received at the mail server at Alice's site is queued with all other messages; each goes to its corresponding destination. In Alice's case, her message goes to the mail server at Bob's site. A client/server MTA is responsible for the e-mail transfer between the two servers. When the message arrives at the destination mail server, it is stored in Bob's mailbox, a special file that holds the message until it is retrieved by Bob.

When Bob needs to retrieve his messages, including the one sent by Alice, he invokes another program, which we call a **message access agent (MAA).** The MAA is also designed as a client/server program with the client installed at Bob's computer and the server installed at the mail server.

There are several important points about the architecture of the e-mail system.

a. The sending of an e-mail from Alice to Bob is a store-retrieve activity. Alice can send an e-mail today; Bob, being busy, may check his e-mail three days later. During this time, the e-mail is stored in Bob's mailbox until it is retrieved.
b. The main communication between Alice and Bob is through two application programs: the MTA client at Alice's computer and the MAA client at Bob's computer.
c. The MTA client program is a *push* program; the client pushes the message when Alice needs to send it. The MAA client program is a *pull* program; the client pulls the messages when Bob is ready to retrieve his e-mail.

Q9. a. Briefly describe Data-expansion and Pseudorandom function in TLS.

Answer:
The data expansion function makes use of the HMAC algorithm with either MD5 or SHA-1 as the underlying hash function. As can be seen, P_hash can be iterated as many times as necessary to produce the required quantity of data. Forexample, if P_SHA-1 was used to generate 64 bytes of data, it would have to be iterated four times, producing 80 bytes of data of which the last 16 would be discarded. In this case, P_MD5 would also have to be iterated four times, producing exactly 64 bytes of data. Note that each iteration involves two executions of HMAC—each of which in turn involves two executions of the underlying hash algorithm.

TLS makes use of a pseudorandom function referred to as PRF to expand secrets into blocks of data for purposes of key generation or validation. The objective is to make use of a relatively small shared secret value but to generate longer blocks of data in a way that is secure from the kinds of attacks made on hash functions and MACs

b. Briefly describe the list of parameters for a session state in SSL.

Answer:
A session state is defined by the following parameters.
• Session identifier: An arbitrary byte sequence chosen by the server to identify an active o resumable session state.
• Peer certificate: An X509.v3 certificate of the peer. This element of the state may be null.
• Compression method: The algorithm used to compress data prior to encryption.
• Cipher spec: Specifies the bulk data encryption algorithm (such as null, AES, etc.) and a hash algorithm (such as MD5 or SHA-1) used for MAC calculation.
It also defines cryptographic attributes such as the hash_size.
• Master secret: 48-byte secret shared between the client and server.
• Is resumable: A flag indicating whether the session can be used to initiate new connections.
c. What steps are involved in the SSL Record Protocol transmission?

c. What steps are involved in the SSL Record Protocol transmission?

Answer:
The Record Protocol takes an application message to be transmitted, fragments the data into manageable blocks, optionally compresses the data, applies a MAC, encrypts, adds a header, and transmits the resulting unit in a TCP segment. Received data are decrypted, verified, decompressed, and reassembled before being delivered to higher-level users.

The first step is fragmentation. Each upper-layer message is fragmented into blocks of 2 14 bytes (16384 bytes) or less. Next, compression is optionally applied. Compression must be lossless and may not increase the content length by more than 1024 bytes.1In

SSLv3 (as well as the current version of TLS), no compression algorithm is specified, so the default compression algorithm is null.

The next step in processing is to compute a message authentication code over the compressed data. For this purpose, a shared secret key is used.

Next, the compressed message plus the MAC are encrypted using symmetric encryption. Encryption may not increase the content length by more than 1024 bytes, so that the total length may not exceed 214 + 2048. For stream encryption, the compressed message plus the MAC are encrypted. Note that the MAC is computed before encryption takes place and that the MAC is then encrypted along with the plaintext or compressed plaintext.
For block encryption, padding may be added after the MAC prior to encryption. The padding is in the form of a number of padding bytes followed by a one-byte indication of the length of the padding. The total amount of padding is the smallest amount such that the total size of the data to be encrypted (plaintext plus MAC plus padding) is a multiple of the cipher's block length. An example is a plaintext (or compressed text if compression is used) of 58 bytes, with a MAC of 20 bytes (using SHA-1), that is encrypted using a block length of 8 bytes (e.g., DES). With the padding-length byte, this yields a total of 79 bytes. To make the total an integer multiple of 8, one byte of padding is added.
The final step of SSL Record Protocol processing is to prepare a header consisting of the following fields:
• Content Type (8 bits): The higher-layer protocol used to process the enclosed fragment.
• Major Version (8 bits): Indicates major version of SSL in use. For SSLv3, the value is 3.
• Minor Version (8 bits): Indicates minor version in use. For SSLv3, the value is 0.
• Compressed Length (16 bits): The length in bytes of the plaintext fragment (or compressed fragment if compression is used). The maximum value is 214 +2048.

Text Book

Behrouz A. Forouzan, Cryptography & Network Security, Special India Edition.