Name…………………………………..….ID………..………..Section……..……Seat No……..….

Sirindhorn International Institute of Technology
Thammasat University

Final Examination: Semester 2/2007

Course Title	:	CSS 322 – Security and Cryptography
Instructor	:	Dr Steven Gordon
Date/Time	:	Monday 3 March 2008, 13:30 – 16:30

Instructions:

③ This examination paper has __ pages (including this page).

③ Condition of Examination
Closed book (No dictionary, no calculator)

③ Students are not allowed to be out of the exam room during examination. Going to the restroom may result in score deduction.

③ Turn off all communication devices (mobile phone etc.) and leave them under your seat.

③ Write your name, student ID, section, and seat number clearly on the answer sheet.

③ The space on the back of each page can be used if necessary.

Part A - Multiple Choice Questions [22 marks]

Select the most accurate answer (only select one answer). Each correct answer is worth 2 marks. You receive 0 marks for an incorrect answer or no answer.

1. Stateful packet inspection:
   a) Allows a firewall to reject (drop) packets based on the content of emails (e.g. spam, viruses)
   b) Allows a firewall to reject (drop) packets that do not belong to an open TCP connection
   c) Allows a firewall to reject (drop) packets that contain malicious HTTP GET requests
   d) Requires a single homed bastion host
   e) Requires a dual homed bastion host
   f) Requires a screened subnet with demilitarized zone (DMZ)

2. When compared to a public key cryptosystem, a symmetric key cryptosystem is:
   a) Easy to distribute keys, but less computationally efficient
   b) More computationally efficient, but difficult to distribute keys
   c) More secure, but difficult to distribute keys
   d) Easy to distribute keys, but less secure
   e) More computationally efficient, but less secure
   f) More secure, but less computationally efficient

3. A computer virus:
   a) Is a program that searches for other systems to infect, connects to a remote system and copies itself to that remote system and executes.
   b) Goes through the phases of being dormant, propagating, triggering and execution.
   c) Is host independent
   d) Is non-replicating
   e) Is a program modification that contains unauthorized access to functionality

4. If you had access (e.g. login as an administrator) to the SIIT gateway router, for all messages passing through that router, you could:

   a) Read the contents of the messages if they were encrypted only with TLS
   b) Read the contents of the messages if they were encrypted only with IPsec (transport mode)
   c) Read the contents of the messages if they were encrypted with any encryption algorithm/protocol
   d) Not read the contents of the messages if they were all encrypted with IPsec (transport mode)
   e) None of the above.

5. If Transport Layer Security is used to secure data (e.g. web pages) between a client and server, TLS uses:
   a) Public key algorithms for data confidentiality and MD5 or SHA1 for data integrity
   b) Public key algorithms for key exchange and Diffie-Hellman for data integrity
   c) Message authentication codes for data integrity and symmetric key algorithms for data confidentiality
   d) Symmetric key algorithms for key exchange and message authentication codes for authentication
   e) None of the above.

6. John’s X.509 certificate (which is signed by the certificate authority Steve) contains:
   a) Only John’s private key (and no other keys)
   b) Only John’s public key (and no other keys)
   c) John’s private key and Steve’s public key
   d) John’s public key and Steve’s public key
   e) John’s private key and Steve’s private key
   f) John’s public key and Steve’s private key

7. When compared to using a single certificate authority, an advantage of using a hierarchy of certificate authorities (CAs) is that:
   a) The CA’s do not need to trust each other
   b) The lifetime (or validity) of the certificate can be made longer
   c) It is more efficient to validate certificates signed by different CAs
   d) Users only need to register and obtain certificates from their local CA, rather than the central CA
   e) The man-in-the-middle attack can be avoided

8. In the figure illustrating the use of a hash function for authentication, if the following property of hash functions *does not hold*, then an attacker could discover the secret S:

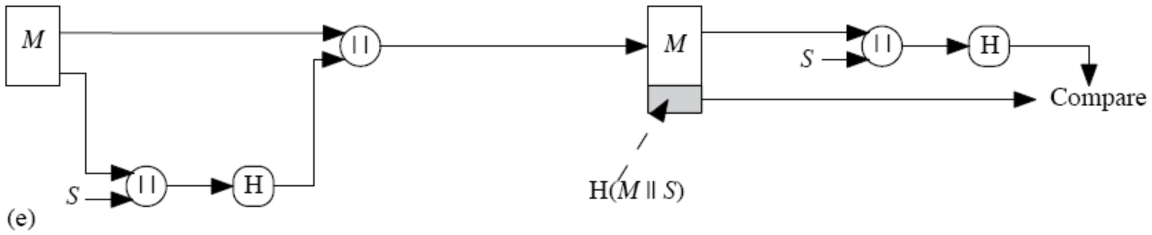

   a) Hash function produces a small hash value
   b) Efficient to calculate the hash function
   c) One-way property
   d) Weak collision resistant
   e) Strong collision resistant

9. The security of RSA:
   a) Cannot be attacked using timing attacks
   b) Cannot be attacked using a brute force approach
   c) Depends on the difficulty in multiplying two large prime numbers
   d) Depends on the difficulty in factoring large numbers into their primes
   e) Is much stronger than AES
   f) Is much stronger than 3DES

10. The Diffie-Hellman exchange in Transport Layer Security (TLS/SSL) is used for exchanging:
   a) Nonce values between two users
   b) Secret values between two users
   c) Sequence numbers between two users
   d) Certificates between two users
   e) Encrypted files between two users

11. A distinct feature of a metamorphic virus (compared to other types of virus):
   a) A normal application that contains unexpected additional functionality
   b) Stored in main memory, and hence infects other applications that execute
   c) Stored on the boot sector of a hard disk, and hence infects applications when the computer starts
   d) Changes its behavior and appearance when propagating (making copies of itself)
   e) Changes only its appearance when propagating (making copies of itself)

Part B – General Questions [78 marks]

Question 1 [10 marks]

Using the RSA algorithm, where the ciphertext C = 146 and public key PU = {e=7, n=187}, determine the plaintext P and private key PR. Show your calculations.

(Hint: if you do not have a calculator, the following table gives some selected calculations which may be useful)

a	**b**	**c**	**a^b mod c**	**A**	**b**	**c**	**a^b mod c**
146	3	187	82	146	15	187	12
146	4	187	4	146	16	187	69
146	5	187	23	146	17	187	163
146	6	187	179	146	18	187	49
146	7	187	141	146	19	187	48
146	8	187	16	146	20	187	89
146	9	187	92	146	21	187	91
146	10	187	155	146	22	187	9
146	11	187	3	146	23	187	5
146	12	187	64	146	24	187	169
146	13	187	181	146	25	187	177
146	14	187	59	146	26	187	36

Answer:

According to RSA, d and e must be multiplicative inverses in ∅(n). That is:

$$de \equiv 1(\bmod \phi(n))$$

where e = 7 and n = 187.

By trial and error we can find the prime factors of 187 to be 17 and 11. And hence:

$$\begin{aligned}\phi(n) &= (p-1)(q-1)\\ &= (17-1)(11-1)\\ &= 160\end{aligned}$$

Therefore:

$$de \bmod 160 = 1$$

And so de can take the values: 161, 321, 481, …

Since e = 7, if d = 23, then ed = 161. Therefore the value of d is 23 and the private key PR = {187, 23}.

To calculate the plaintext:

$$\begin{aligned}P &= C^d \bmod n\\ &= 146^{23} \bmod 187\\ &= 5\end{aligned}$$

Question 2 [16 marks]

The following C code shows an example DNS program (`dnsnametoip.c`).

```
#include <string.h>
#include <stdio.h>
char* getIpFromDns(char* strDnsName)
{
      if (strcasecmp(strDnsName, “www.raytheon.com”) == 0)
      {
            return “192.168.0.1”;
      }
      else if (strcasecmp(strDnsName, “www.w3.org”) == 0)
      {
            return “192.168.0.2”;
      }
      else if (strcasecmp(strDnsName, “www.slashdot.org”) == 0)
      {
            return “192.168.0.3”;
      }
      else if (strcasecmp(strDnsName, “www.kerneltrap.org”) == 0)
      {
            return “192.168.0.4”;
      }
      return “Unknown”;
}

int main(int argc, char* argv[])
{
      char strBuffer[255];
      strcpy(&strBuffer[0], argv[1]);
      printf(“%s\n”, getIpFromDns(strBuffer));
      return 0;
}
```

a) Explain what a user needs to do to cause a *buffer overflow* in the above program. Also explain why the buffer overflow occurs. [3 marks]

Answer:

The user of the program should provide as the first argument (argv[1]) a string which is longer than 255 characters. This is because strbuffer has memory allocated for 255 characters. The strcpy() function copies argv[1] into strbuffer – if argv[1] is longer than 255 characters then the remaining characters will be written into the memory adjacent strbuffer. That is, it will overflow the strbuffer memory/buffer.

b) Explain a change to the code that will avoid the buffer overflow. Also explain how the new code avoids the buffer overflow. [3 marks]

Answer:

Using strncpy() instead of strcpy() can avoid the buffer overflow. By using strncpy(&strbuffer[0], argv[1], 255), only the first 255 characters of argv[1] will be copied into strbuffer. Therefore, even if an argument longer than 255 characters is used, the buffer will not be overflowed.

The following C code shows an example program that starts a command line shell (`exploitcodeusingExecve.c`):

```
#include <unistd.h>

int main()
{
      char* argv[1];
      argv[0] = "/bin/sh";
      argv[1] = NULL;
      execve(argv[0], argv, 0);
      return 0;
}
```

c) An attacker wants to perform a buffer overflow attack. Explain the steps the attacker may take for this exploit code to be inserted into memory using the DNS program of part (a). [3 marks]

Answer:

First the code must be compiled, and then disassembled to obtain the executable machine instructions as a set of byte codes. The hexadecimal byte codes are then represented as a string (C provides an escape sequence to include hex numbers in a string). The string is passed as an input to the DNS program. The string is then written into memory from the start of strbuffer.

d) When the function `getIpFromDns()` in the DNS program is called, an instruction pointer is created that points to a position in memory. Before an attack is performed, what is stored in memory at the position the instruction pointer points to? [2 marks]

Answer:

The next line of code (or more precisely, machine instruction) that the program will execute when the function returns. In the DNS program that is, "return 0".

e) Explain the purpose of the attacker including a new instruction pointer after the exploit code. Include a discussion of the value of the instruction pointer and how it is used. [3 marks]

Answer:

The new instruction pointer should point to the position in memory where the exploit code is stored. The instruction pointer is written after the exploit code, with the intention of overwriting the old instruction pointer. As a result, the calling function will execute the code pointed to by the new instruction pointer (instead of the old instruction pointer). That is, the calling function will execute the exploit code.

f) Often the attacker cannot determine the exact position in memory where the exploit code is inserted. What does the attacker do to increase the chance that the exploit code will still be executed, even if the incorrect memory position is used [2 marks]

Answer:

The attacker inserts many No-operation (NOP) byte codes before the exploit code.

Question 3 [12 marks]

In some UNIX systems the password file contains a list of usernames for users on the computer system, along with a hash of the password for each user. Assume an attacker has access to the password file.

a) Explain an advantage the attacker has in discovering passwords, compared to an online attack. [2 marks]

Answer:

The attacker is not limited by time or number of guesses of passwords. The attacker can also execute the attack on their chosen machine (e.g. which may be very fast), rather than the computer system that the passwords are used for (which may be a normal PC).

As an additional security measure, a *salt* can be included when hashing the password. That is, a 12-bit random value (the salt) is appended to the plaintext password before the hash is applied. The password file then stores the username, the salt, and the hash value (that is Hash(password + salt)).

b) Explain how the addition of the salt helps prevent detection of duplicate passwords in the password file. (Hint: in your explanation, compare what happens if the salt is not used, and then if the salt is used) [3 marks]

Answer:

If there are two duplicate passwords, and no salt used, then the hash values would be identical. Then if a attacker knew one password (and corresponding hash value), they would then know every other user who has an identical password. If salt is used, even though the passwords are the same, the salt values should be different, and therefore the resulting hash is different. Therefore, an attacker cannot identify who has the same password.

c) In addition to preventing detection of duplicate passwords, and despite the salt being known to the attacker, how else does using the salt value increase security of the password file. (Hint: again consider what happens when the salt is and is not used – focusing on the chance of the attacker guessing a correct password). [3 marks]

Answer:

Without the salt, the attacker can guess a password and encrypt it. If ANY of the users on a system use that password, then there will be a match. With the salt, the attacker must guess a password and then encrypt it once for each user, using the particular salt for each user. Hence using a salt means there will be more guesses.

d) Explain why increasing the salt size from 12-bits to a much larger size (e.g. 24 or 36 bits) *does not* make password guessing (practically) impossible? (Hint: consider part (c)). [4 marks]

Answer:

The security depends on the number of users, not the size of the salt. From part (c), if the number of users (n) is large, then the attacker has to encrypt the guessed password with a larger number of salt values (n). However, just increasing the salt size still means the same number of encryptions (albeit with a larger plaintext). Example: if n = 10, the attacker must try 10 different salt values (no matter what size) with the guessed password. If n = 100, the attacker must try 100 different salt values with the guessed password.

Question 4 [8 marks]

Assume a digital signature is applied on the hash of a message (not the message itself), and A sends the message (M_A) and digital signature to B. If the hash function is not weak collision resistant, then explain how an attacker C can forge A's digital signature on a new message from C (called M_C). Make sure you explain why B cannot detect this forgery.

Answer:

A sends M_A and $E_{PRA}(H(M_A))$ to B, but it is intercepted by C. If the hash function is not weak collision resistant then C may find a message M_C with hash value of h_C such that:

$$H(M_C) = h_C = H(M_A)$$

That is, the hash of both messages are the same.
Then C sends M_C and $E_{PRA}(H(M_A))$ to B. Note that C does not modify $E_{PRA}(H(M_A))$ – it just resends that portion. Now if B checks the signature at a later stage then finds that:

$$H(M_C) = D_{PUA}(E_{PRA}(H(M_A)))$$

That is, the signature is correct (although M_C is different from the original message).

Question 5 [5 marks]

The figure below shows the steps applied to application data by the SSL Record Protocol.

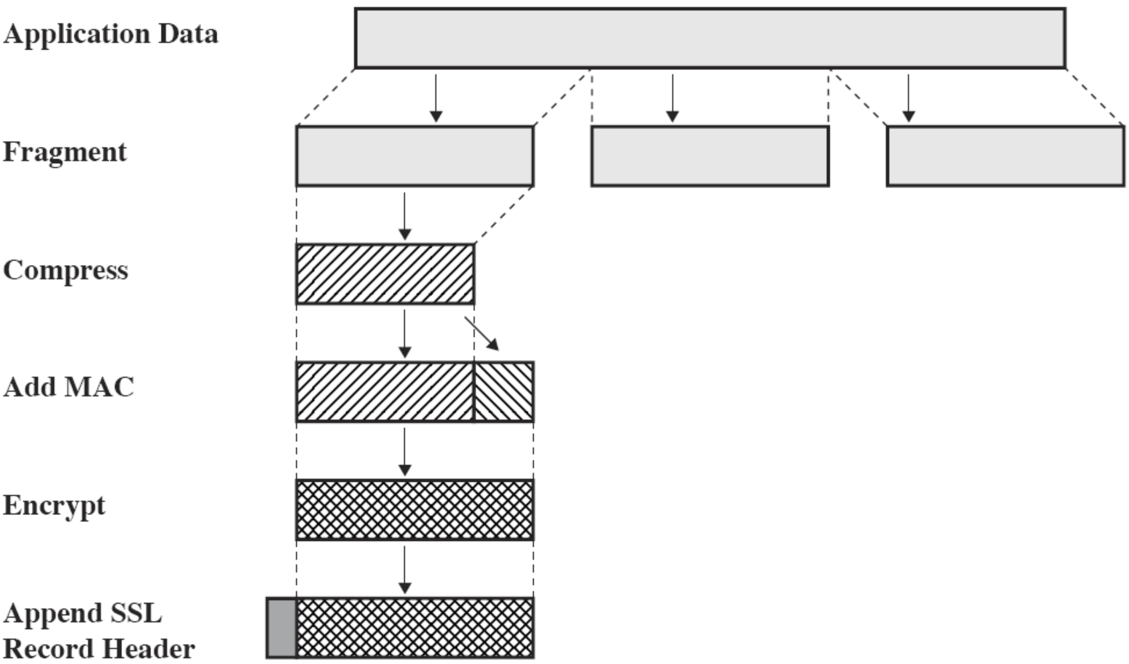

a) What two security services does the SSL Record Protocol provide? [2 marks]

Answer:

Confidentiality and integrity

b) List and explain an advantage and disadvantage of using SSL compared to using IPsec. [3 marks]

Answer:

Advantage of SSL: Implemented more widely; Often implemented in applications, and therefore can be easily distributed in applications (as opposed to operating systems)

Disadvantage of SSL: Only applicable for TCP applications (not UDP).

Question 6 [9 marks]

The figure below shows two networks, each with three PCs, connected across an internet. Refer to the IP addresses of nodes by their name, for example: the IP address of PC1 is PC1.

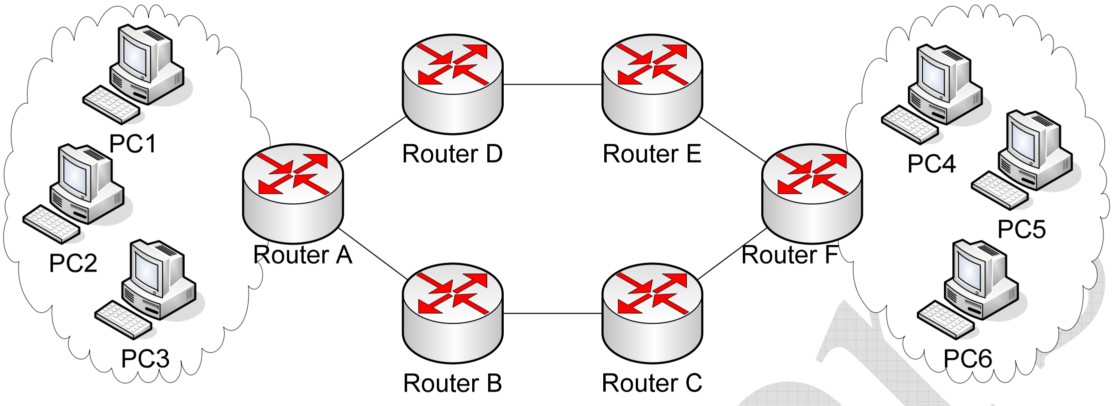

In parts (a), (b) and (c), assume IPsec with Encapsulating Security Payload and Transport Mode are used to secure data transfer between PC1 and PC4.

a) List three security services provided by the IPsec connection. [1.5 marks]

Answer:

Any of: confidentiality, authentication, integrity, access control, anti-replay.

b) If an IPsec packet is received by Router E, does Router E know that PC1 and PC4 are communicating with each other? Explain why. [1.5 marks]

Answer:

Yes, the IP addresses of PC1 and PC4 are included in the clear.

c) If an IPsec packet is received by Router E, does Router E know the application being used by PC1 and PC4? Explain why. [2 marks]

Answer:

No, the port numbers are in the TCP/UDP header, which is encrypted.

In parts (d) and (e), assume IPsec with Encapsulating Security Payload and Tunneling Mode are used to secure data transfer between the two networks of PCs, with the network gateways being tunnel end-points. PC1 and PC4 communicate via the tunnel.

d) If an IPsec packet is received by Router E, does Router E know that PC1 and PC4 are communicating with each other? Explain why. [2 marks]

Answer:

No, because in tunneling mode the original IP packet from PC1 to PC4 is encrypted.

e) List and describe an advantage and disadvantage of using tunneling mode (versus transport mode). [2 marks]

Answer:

Advantage of tunneling: IPsec only needs to be configured on routers, not on all PCs; hide the IP addresses of the originating and destination nodes;

Disadvantage of tunneling: cleartext sent between PCs and tunnel end-point (must trust local network); performance bottleneck at end-points;

Question 7 [10 marks]

A single homed bastion host is one way to configure a network using both a packet filter firewall/router, as well as an application level firewall (bastion host).

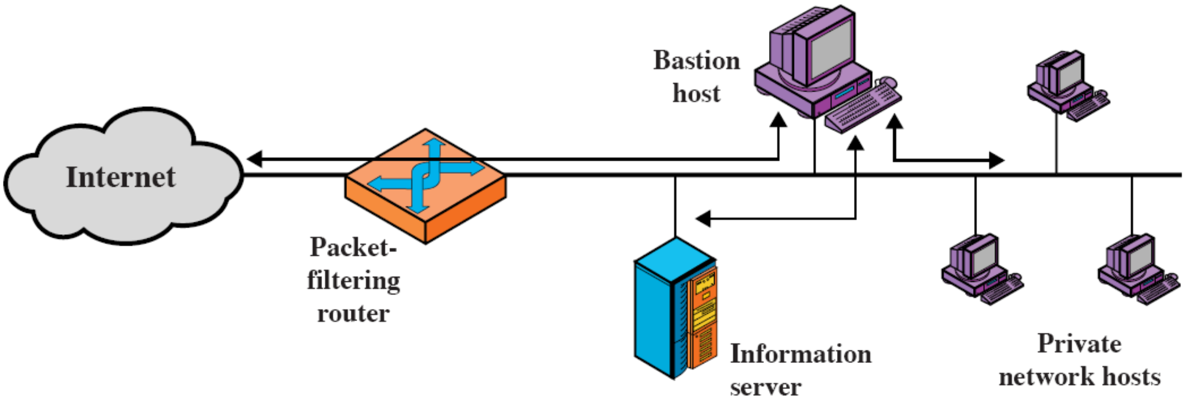

a) Explain the difference between a packet filter firewall/router and an application level firewall. Include an advantage of each type of firewall. [3 marks]

Answer:

A packet filter firewall inspects only TCP/IP packets, that is looks only at the values of the IP and TCP/UDP headers, and makes accept/drop decisions based on that information. Advantage is that this is very simple, and can be implemented efficiently.

An application level firewall also inspects application level content, like web page responses, domain names, email contents and so on. This is more secure because can look at more information than a packet firewall.

b) Describe two rules that must be configured on the packet filter router in this configuration. (You only need a explanation of what the rules will do – you do not need to give specific rules) [3 marks]

Answer:

Rule 1: Direct all traffic from outside to the bastion host
Rule 2: Only accept traffic from inside the network that has come from the bastion host

c) What is an advantage of using this configuration (as opposed to using only an application level firewall)? [2 marks]

Answer:

In most cases, attacker must break two lines of defence to get into network: break the packet filter and the application level firewall.

d) What is the difference between this configuration and a dual homed bastion host configuration? Explain how the dual homed configuration provides an advantage over the single homed configuration. [2 marks]

Answer:

A dual home bastion host configuration has two physical connections of the bastion host. The advantage of this is that even if the packet filtering router is compromised, all traffic must still go through the bastion host (in single homed, the compromised filter router may be configured to bypass the bastion host).

Question 8 [8 marks]

a) If A sends an unencrypted message M to B, and also sends the Message Authentication Code, what two security services does the system provide? [2 marks]

b) Explain what happens if an attacker C intercepts the message from A to B, modifies M (e.g. changes M to N) and then forwards the modified message, with the original MAC attached, to B. (Include an explanation of whether B detects an attack) [3 marks]

c) Explain what happens if C changes both the message M (to N), and calculates and sends the new MAC based on N? (Include an explanation of whether B detects an attack) [3 marks]

Answers:

a. Authentication, integrity

b. B will check the MAC by calculating the MAC for the received message N, and comparing it to the received MAC. They will be different and so B will assume there has been an attack (or error).

c. C does not know the key used by A to calculate the MAC. Hence C uses a different key. B will receive and check the MAC by calculating the MAC for the received message N using key K, and compare it to the received MAC (which was calculated as MAC(N,AnotherKey). Since the MACs were calculated with different keys, they will be different. B will assume there has been an attack (or error).