Written exam in TSIT02 Computer security

08:00–12:00, 31st of March 2016

Jan-Åke Larsson
Institutionen för Systemteknik,
Linköpings Universitet

Permitted equipment: General dictionaries between English and another language without personal notes. (Thus not specific scientific dictionaries with or without formulæ.)

Solutions: Solutions will be posted on the course home page after the exam.

Grading: Grade 3 requires at least 29 points.

Other information: Answers can be in English or Swedish.

1. Describe with one sentence for each item what "CIA" stands for in computer security. Also, for each of these three, give an example of a threat that **only** applies to that category (as primary damage). Points will not be given for threats that apply to several categories, this is to make sure you understand the differences between the three. The examples should be specific, and include a description of the attack type and tools used for the threat you describe. You should describe each threat in at most three sentences and explain in at most two sentences why it is relevant for only one and not two or three of the CIA categories. (6p)

2. The following files are shown by an `ls -l` command on a typical Unix system:

```
-r-xr-sr-x 1 charlie  acct    70483   2008-01-04 22:53 accounting.so
-r--rw---- 1 alice    acct    139008  2008-05-13 14:53 accounts.db
-rwxr-xr-x 1 system   system  230482  1997-04-27 22:53 editor
-rw-r--r-- 1 alice    users   7072    2008-06-01 22:53 cv.txt
-r--r----- 1 bob      gurus   19341   2008-06-03 13:29 exam.pdf
-r--r----- 1 alice    gurus   6316    2008-06-03 16:25 solutions.tex
```

   Unix users `alice` and `bob` are both members of only the group users, while `charlie` is a member of only the group `gurus`. Application `editor` allows users to read and write files of arbitrary name and change their permissions, whereas application `accounting.so` only allows users to append data records to the file accounts. Draw up an access control matrix with subjects {`alice; bob; charlie`} and objects {`accounts.db; cv:txt; exam.pdf; solutions.tex`} that shows for each combination of subject and object whether the subject will, in principle, be able to read (R), (over)write (W), or at least append records (A) to the respective object. (9p)

3. There are three major categories of methods for user authentication. Which are the two categories (not specific examples of methods) most commonly used? Which is the third category for methods to authenticate users, the one that is more seldom used? Give an explicit example of a method that fits in each of the three categories. Examples from the third category have balance problems that do not exist for the other two categories. What two values must be adjusted to the requirements of the application, when you use this authentication category, and why does a change in one of these two values normally cause a change also in the other? (6p)

4. You own the store FooStore. The inventory can be updated by all employees through the cash registers. Deliveries are entered by Claire, Bob, and Elsie. Spoilt goods is deducted by these three and Dave. Price changes are made by Edward and Elsie. You handle new items and campaign prices. You now want to protect the data against unauthorized changes.

   (a) Can this be modelled with the Biba model? Why or why not? (2p)

   (b) Or is a Chinese Wall model description perhaps more appropriate? Why or why not? (2p)

5. On the 19^{th} of March at 19:30, several Swedish newspapers were attacked by a DDoS attack. Over 50 gigabit per second (100 million packets per second) of bogus traffic were sent to web pages, drowning out legitimate traffic. According to the Swedish Civil Contingencies Agency (*Myndigheten för samhällsskydd och beredskap, MSB*), 20% of the attack traffic consisted of SYN flooding.

   (a) What is a TCP SYN flood? (1p)

   (b) Would an IPSec deployment protect against DDoS attacks, or would it make it more vulnerable? Motivate your answer (2p)

   (c) Would a firewall protect against DDoS attacks, would it make it more vulnerable? Motivate your answer (2p)

   (d) Would an IDS protect against DDoS attacks, or would it make it more vulnerable? Motivate your answer (2p)

6. If Kerberos is used for single sign-on, the user password is never sent in any form, neither as plaintext nor encrypted, over the network at logon. How does a service test that a user knows the password of the stated identity (and thus indeed has that stated identity)? (4p)

7. How do you create a true digital signature with a fixed length? What is the difference, in creation and in what security it provides, between a true digital signature and a crypto-based checksum like a MAC? (4p)

8. You have a mobile phone and want to connect to a base station that does not belong to your home network. How does the base station agree on a shared secret key with your mobile device, in GSM? In your answer you can just describe the principles, if you do not remember exactly what each entity and variable is called in GSM.

   The base station never learns the basic device key, and thus cannot clone the device. But the base station can still be a threat in some scenarios. What is the weakness? What step/service is missing in GSM to form this weakness, but is added in UMTS? (4p)

9. Hash functions.

   (a) What does "collision resistance" mean in the context of hash functions? (1p)

   (b) What does it mean that it is hard to find a preimage in the context of hash functions? (1p)

   (c) When storing password lists, it is recommended to use a hash function. In what way? What type of attack does this protect against? (2p)

   (d) In addition to using a hash function when storing passwords, it is recommended to also use a salt. What is a salt? What type of attack does it protect against? (2p)

10. Why was the Clark-Wilson formal security model created? The lectures give two reasons, please mention both. Briefly describe a "Clark-Wilson triple". (3p)

11. What cryptographic method is used in DNSSec to protect DNS records? This type of protection needs a trust chain, where does the chain start in DNSSec and how does it reach the client node? What confidentiality issue is seen as a problem in DNSSec? (4p)

1 Solutions

1. Describe with one sentence for each item what "CIA" stands for in computer security. Also, for each of these three, give an example of a threat that **only** applies to that category (as primary damage). Points will not be given for threats that apply to several categories, this is to make sure you understand the differences between the three. The examples should be specific, and include a description of the attack type and tools used for the threat you describe. You should describe each threat in at most three sentences and explain in at most two sentences why it is relevant for only one and not two or three of the CIA categories. (6p)

   - Phishing is an attack that only reveals confidential data as primary damage (users' credentials like user ID and password). Phishing is when you send out a message, telling users of one protected site to go to a webpage of your design for some compelling reason and log in there. You, the attacker, collect all data delivered by the victims. Obviously just the collection of data does not influence data integrity or availability for the victim.
   - Web site defacements (altering contents of web pages) is a data integrity only offence. There are standard hacker tools using well-known and often not amended security holes to get editing access to the victim's web server. Just changing data on a web page will not reveal any new secrets to the attacker, and the victim can still access data. Virus infections are also only a data (program) integrity threat in its first step, the infection, since its only damage so far is an unsolicited piece of code on the machine. Both confidentiality and availability can be disturbed when the virus executes its full code, but just the infection disturbs neither.
   - Denial of Service attacks are directed only at availability. Denial of service can be achieved through flooding the victim with connection requests, then normally with the help of lots of infected computers. It can also be achieved by using known flaws in protocols and operating systems, which make the victim computer crash when you issue specific calls to it. These events do not spread secrets and do not change the victim's stored data.

2. The following files are shown by an `ls -l` command on a typical Unix system:

```
-r-xr-sr-x 1 charlie  acct    70483   2008-01-04 22:53 accounting.so
-r--rw---- 1 alice    acct    139008  2008-05-13 14:53 accounts.db
-rwxr-xr-x 1 system   system  230482  1997-04-27 22:53 editor
-rw-r--r-- 1 alice    users   7072    2008-06-01 22:53 cv.txt
-r--r----- 1 bob      gurus   19341   2008-06-03 13:29 exam.pdf
-r--r----- 1 alice    gurus   6316    2008-06-03 16:25 solutions.tex
```

   Unix users `alice` and `bob` are both members of only the group users, while `charlie` is a member of only the group `gurus`. Application `editor` allows users to read and write files of arbitrary name and change their permissions, whereas application `accounting.so` only allows users to append data records to the file accounts.

Draw up an access control matrix with subjects {`alice; bob; charlie`} and objects {`accounts.db; cv:txt; exam.pdf; solutions.tex`} that shows for each combination of subject and object whether the subject will, in principle, be able to read (R), (over)write (W), or at least append records (A) to the respective object. (9p)

- The solution is given in Table 1. The editor allows all users to write and read to the files with that permission. The accounting.so application has a setgid bit in addition to an execute bit set for all users. Therefore, all users executing accounting.so can append to files with the same permissions as a user in the acct group. Therefore, all users can append to accounts.db.

Table 1: Access control matrix.

	accounts.db	cv.txt	exam.pdf	solutions.tex
alice	RA	RW	-	R
bob	A	R	R	-
charlie	A	R	R	R

3. There are three major categories of methods for user authentication. Which are the two categories (not specific examples of methods) most commonly used? Which is the third category for methods to authenticate users, the one that is more seldom used? Give an explicit example of a method that fits in each of the three categories. Examples from the third category have balance problems that do not exist for the other two categories. What two values must be adjusted to the requirements of the application, when you use this authentication category, and why does a change in one of these two values normally cause a change also in the other? (6p)

   - Most used is "What you know", which can be passwords, PINs etc. Cards and tokens are examples of "What you hold/have/carry", the second category. Both are used in combination at ATMs.
   - The third category is "What you are" or biometrics. This can be fingerprints, iris patterns, handwritten signatures.
   - These suffer from the possibility of false acceptance or false rejection, since no two measurements are ever exactly the same.
   - If you have severe restrictions on the possible variation, you will reduce the risk for false acceptance, but raise the risk for false rejection, and vice versa if you allow for more variation between two measurements.

4. You own the store FooStore. The inventory can be updated by all employees through the cash registers. Deliveries are entered by Claire, Bob, and Elsie. Spoilt goods is deducted by these three and Dave. Price changes are made by Edward and Elsie. You handle new items and campaign prices. You now want to protect the data against unauthorized changes.

   (a) Can this be modelled with the Biba model? Why or why not? (2p)

- Biba concerns integrity, and has rules about writing, so Biba seems appropriate. The model is hierarchal, so the question really is if the access structure forms a hierarchy. Here, you have two kinds of data (number of items in store, and price) and a few user groups that can write to the two. The hierarchy is difficult to adapt, so both "yes" and "no" are accepted as long as the discussion contains the two mentioned items.

(b) Or is a Chinese Wall model description perhaps more appropriate? Why or why not? (2p)

- The Chinese wall cannot be used. It concerns confidentiality, and the relation between compartments of data. Here all data is available, since confidentiality is not the goal.

5. On the 19^{th} of March at 19:30, several Swedish newspapers were attacked by a DDoS attack. Over 50 gigabit per second (100 million packets per second) of bogus traffic were sent to web pages, drowning out legitimate traffic. According to the Swedish Civil Contingencies Agency (*Myndigheten för samhällsskydd och beredskap, MSB*), 20% of the attack traffic consisted of SYN flooding.

   (a) What is a TCP SYN flood? (1p)

   - A large number of TCP connection requests that intends to overwhelm the recipient.

   (b) Would an IPSec deployment protect against DDoS attacks, or would it make it more vulnerable? Motivate your answer (2p)

   - IPSec requires the recipient to perform calculations and store session data. Therefore, an IPSec deployment can make a DDoS attack worse. However, if the IPSec deployment includes a cookie mechanism this weakness can be mitigated and even protect against DDoS.

   (c) Would a firewall protect against DDoS attacks, or would it make it more vulnerable? Motivate your answer (2p)

   - A packet filter firewall in itself matches the source address, recipient address and port number of the packet with its internal set of rules. No provisions are made on the rate of packets. Therefore, the packet filter firewall makes no difference in case of DDoS attack.
     If the administrator adds a new rule to the firewall, blocking the originating addreses, then a firewall can protect against DDoS. This answer, however, must say that this can only be done manually by the administrator (or automatically by an IDS) to give full credit.

   (d) Would an IDS protect against DDoS attacks, or would it make it more vulnerable? Motivate your answer (2p)

   - An IDS can detect an attack but not take measures against it. It is up for a human operator to stop the attack. An IDS does therefore not make a difference in case of a DDoS attack.

6. If Kerberos is used for single sign-on, the user password is never sent in any form, neither as plaintext nor encrypted, over the network at logon. How does a service test that a user knows the password of the stated identity (and thus indeed has that stated identity)? (4p)

- Basic answer: Because the requesting user can present a ticket issued for the requesting client's identity, a ticket which contains the user ID and a session key encrypted and issued by the Kerberos system. In addition the requesting user can encrypt and decrypt messages sent with this session key, and the service knows that another Kerberos level has ensured that only the correct user knows this session key's value. A longer answer describes the whole chain: The user sends a login request encrypted with the user password to the Kerberos server. The Kerberos server uses the plaintext user ID to retrieve the password and decrypt. Thus only the correct user can get the Kerberos server to issue tickets in that user's name, since other users cannot create such requests that will be correctly decrypted. If it is a valid request, Kerberos creates a ticket server key for this user i, kit and sends this to the user encrypted with the user password together with a ticket server ticket, which contains the user ID and kit encrypted with the key shared by the ticket server and the Kerberos server. The user sends this ticket with a service request to the ticket server, which decrypts, checks the sender ID against the user ID in the encrypted ticket, which only the Kerberos server and this ticket server can encrypt/decrypt. If the identities match, the ticket server ticket was obviously issued for the requesting user, who is the only other entity who could have decrypted the key now shared between the user and the service. This is repeated on the next level with the ticket server creating a key for the user and the requested service and issuing it encrypted to the user and in an encrypted service ticket. Since only valid Kerberos entities can create tickets, which make sense when decrypted, and since these contain user keys, which only the correct user should know, false tickets cannot be created, and users trying to use intercepted tickets, will not have the required key to communicate with the server.

7. How do you create a true digital signature with a fixed length? What is the difference, in creation and in what security it provides, between a true digital signature and a crypto-based checksum like a MAC? (4p)

    - In order to have a fixed size you must first pass the message through a (collision resistant) hash function. The result is then transformed with asymmetric encryption using the private key, like for RSA, or is transformed in a specific asymmetric signature algorithm like DSA. For crypto-based checksums you use symmetric methods and normally chaining like CBC with DES or AES. Such checksums protect against interference from third parties, but they do not prove which version is correct, if sender and receiver have different claims on what was sent. Authorities cannot even check if any version has a correct checksum, if the sender or receiver does not give away the key and thus at the same time the possibility to create checksums for new messages. True digital signatures can only be created by the sending party, but can be checked by everyone.

8. You have a mobile phone and want to connect to a base station that does not belong to your home network. How does the base station agree on a shared secret key with your mobile device, in GSM? In your answer you can just describe the principles, if you do not remember exactly what each entity and variable is called in GSM.

The base station never learns the basic device key, and thus cannot clone the device. But the base station can still be a threat in some scenarios. What is the weakness? What step/service is missing in GSM to form this weakness, but is added in UMTS? (4p)

- Mobile phones share their secret customer key k_i with the home network. The phone sends its ID in clear to the home network. The home network creates a random number RAND, retrieves the customer key k_i, and creates with these two values a check value RES and a temporary key k_c. RES, RAND and k_c are sent to the visited network, which keeps RES and k_c and forwards RAND to the phone. The phone uses RAND and k_i to reconstruct RES and k_c and sends RES to the visited network. If the two versions of RES match, the visited network accepts the phone and uses k_c for encryption of the rest of the session.
- This opens up an opportunity for a "man-in-the-middle" attack, since the visited network does not authenticate itself to the home network or phone. Thus a false base station can ask a phone to switch off encryption entirely, or make the phone use a key known to the attacker. In UMTS the key is authenticated, so an eavesdropper on the radio part cannot make the phone accept a key not sent by the home network.

9. Hash functions.

   (a) What does "collision resistance" mean in the context of hash functions? (1p)

   - It should be difficult to find two inputs that produce the same hash value. An incorrect answer would be to say that it should be impossible.

   (b) What does it mean that it is hard to find a preimage in the context of hash functions? (1p)

   - Given y it is difficult to find x so that $h(x) = y$ for a hash function h. An incorrect answer would be to say that it should be impossible.

   (c) When storing password lists, it is recommended to use a hash function. In what way? What type of attack does this protect against? (2p)

   - It means hashing the password before storing it. If the password database leaks, the attacker won't immediately have the passwords of all users. Often, users use the same password in multiple sites, and the attacker will be prevented from easily using the same password on those sites.

   (d) In addition to using a hash function when storing passwords, it is recommended to also use a salt. What is a salt? What type of attack does it protect against? (2p)

   - A salt is a random but publicly known bit string that is added to the password before it is hashed. The password database will then contain the username, salt and salted+hashed password for each user. Salting prevents the attacker from computing rainbow tables for known hash functions, otherwise the same password will give the same hash value over many servers.

10. Why was the Clark-Wilson formal security model created? The lectures give two reasons, please mention both. Briefly describe a "Clark-Wilson triple". (3p)

    - The relative importance of Confidentiality versus Integrity differs between military and commercial applications — in commercial applications, Integrity is more important, to prevent fraud and accidental losses
    - Even trusted (trustworthy?) users make mistakes
    - A Clark-Wilson triple consists of (subject,object,tool): authenticated principals, data items (constrained and unconstrained), and permitted operations

11. What cryptographic method is used in DNSSec to protect DNS records? This type of protection needs a trust chain, where does the chain start in DNSSec and how does it reach the client node? What confidentiality issue is seen as a problem in DNSSec? (4p)

    - DNSSec uses digital signatures to protect DNS records. The trust chain starts at the DNS root and travels down the name server hierarchy to the client node. Many feel their network structure is confidential, and are worried that this can be easily retrieved in DNSSec.