Written exam in TSIT02 Computer security

14:00–18:00, 12th of January 2017

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 24 points

Other information: Answers can be in English or Swedish.

1. What three entities must exist for a risk to exist? (3p)
2. (a) What basic property is needed for a secure cryptographic hash function? (1p)
   (b) Even if the hash function has this property, checking the hash of some downloaded document is not enough to guarantee that it is not manipulated in transit. How is authenticity guaranteed? There are two types of systems, and each gives one point. (2p)
3. The Access Control Matrix (ACM) is a matrix which shows the full relation between all subjects and objects in the system.
   (a) Why does no real system use the ACM? (1p)
   (b) If we instead give a "per subject" list of which objects that that are permitted, what is that called? (1p)
   (c) In the course we have also talked about formal models for access control, two examples of which are Bell-LaPadula and Biba. These models have a drawback in that they are complex and expensive to implement. Why then, would one use such a formal model? Describe (don't just list) the advantage. (2p)
4. In both Windows and Unix systems, the system administrator account can modify a user's password. Can the administrator also read the user's passwors? Why/why not? Motivate your answer. (1p)
5. The most common method of defending against CSRF attacks is to use a token. Describe what properties we require of the token and the basic steps the server and client take to generate and validate the token. (6p)
6. Passwords for user logins should never be stored on the server in plain text form as we have seen in the lectures. Instead, they should be hashed and salted. However, many sites around the web use only hashing and don't do salting.

(a) An attacker with access to the unsalted password hashes can perform an attack that uses a space-time tradeoff to break many hashes. What is this method called? (1p)

(b) Describe the steps involved in performing this attack on unsalted, hashed passwords. Why is it called a time-space tradeoff? (3p)

(c) What is a salt? How is a salt generated? Describe how a salt defeats the type of attack asked for in question a). (2p)

(d) Should the salt be kept secret? Motivate. (1p)

7. Fingerprints, irises and handwritten signatures are properties used for biometric identification of users. What is meant with enrollment, false rejection rate and false acceptance rate? Describe three important properties of biometric authentication methods which are different from other ways of authenticating users. (6p)

8. Access control is often simplified by using groups. But groups can cause an ambiguity interpreting privileges, requiring new rules. The course mentions two different main principles for such interpretations. Describe them shortly and also explain how they relate to the use of the negative privilege "None"! (3p)

9. In a statistical database, a group of three is too small to conceal numerical values. This is true even if you only are allowed to query about MAX, MIN, and AVERAGE. Why? (2p)

10. (a) Is a firewall a useful tool in defence against DoS attacks? Describe shortly how it can be used or motivate why it is irrelevant for that specific threat. (2p)

    (b) Is an IDS a useful tool in defence against DoS attacks? Describe shortly how it can be used or motivate why it is irrelevant for that specific threat. (2p)

11. Draw a diagram of Kerberos and briefly describe each participant. What secrets are shared between the participants? What final data item does the client gain in the process? (5p)

12. In Swedish law, why can an e-mail without a digital signature not be considered an "urkund"? (2p)

Solutions

1. What three entities must exist for a risk to exist? (3p)

   - Threat: the cause of damage to asset(s)
   - Vulnerability: the unwanted system property that enables the threat
   - Damage: the adverse effect of an unwanted event

2. (a) What basic property is needed for a secure cryptographic hash function? (1p)

   - Collision resistance: it should be difficult to find another message (alternatively two messages) that hash to the same value.

   Even if the hash function has this property, checking the hash of some downloaded document is not enough to guarantee that it is not manipulated in transit. How is authenticity guaranteed? There are two types of systems, and each gives one point. (2p)

   (b)
   - A secret must be added, something known to only the sender or receipient. This woud give a MAC, which is a secret-key system.
   - Alternatively, a public key system can be used to sign the hash value. This is known as a digital signature.

3. The Access Control Matrix (ACM) is a matrix which shows the full relation between all subjects and objects in the system.

   (a) Why does no real system use the ACM? (1p)

   The ACM has one row for each subject and a column for each object. In any real system, such a matrix will become extremely large and will be very inefficient.

   (b) If we instead give a "per subject" list of which objects that that are permitted, what is that called? (1p)

   Capability list.

   (c) In the course we have also talked about formal models for access control, two examples of which are Bell-LaPadula and Biba. These models have a drawback in that they are complex and expensive to implement. Why then, would one use such a formal model? Describe (don't just list) the advantage. (2p)

   A formal model gives a testable approach and a security theorem. Therefore, as soon as the requirements of the theorem are fulfilled we can be sure the system will remain secure.
   Common mistakes: The fact that BLP enforces C and Biba enforces I is not an "advantage". The question relates to formal models *in general*. Just repeating the whole BLP and Biba model descriptions is an incorrect answer.

4. In both Windows and Unix systems, the system administrator account can modify a user's password. Can the administrator also read the user's passwors? Why/why not? Motivate your answer. (1p)

In both Unix and Windows, reading passwords has nothing to do with permissions. It is instead protected by being salted and hashed, and no administrator privileges can break a cryptographic one-way function. In other words, the protection is due to mathematics which not even the root/admin user can break.

Common mistakes: Even though Windows UAC is based on Biba it is an incorrect answer to refer to this as the reason. The admin/root accounts traverse normal authorization which mean they can read any file, including password lists. In addition, just listing reasons why it *should* be impossible for the admin to read passwords is also not a good enough answer. Often in computer security, things that *should* be stored safely are not.

5. The most common method of defending against CSRF attacks is to use a token. Describe what properties we require of the token and the basic steps the server and client take to generate and validate the token. (6p)

   When the server generates a page, it generates a token. This token is unique for the client's session and must be hard to predict. In addition, it must be short-lived. The server stores the token in the client's session and inserts the token as a parameter into the links in the page.

   When the client submits a request, the server verifies that the token is present, that it matches the client's session, and that is hasn't expired. Only then is the request granted and otherwise it is rejected.

6. Passwords for user logins should never be stored on the server in plain text form as we have seen in the lectures. Instead, they should be hashed and salted. However, many sites around the web use only hashing and don't do salting.

   (a) An attacker with access to the unsalted password hashes can perform an attack that uses a space-time tradeoff to break many hashes. What is this method called? (1p)

   Dictionary attack.
   Common mistake: While rainbow tables is similar in form to a dictionary attack it is a further refinement of it. For full credit in this case, your answer in b) should actually describe a rainbow table and not just a dictionary attack.

   (b) Describe the steps involved in performing this attack on unsalted, hashed passwords. Why is it called a time-space tradeoff? (3p)

   First, take a dictionary of common words. Then, take the hash value for each word in the dictionary and store it. Next, compare the hashed password to the hashed dictionary words. If there is a match, you found the password. It's called a time-space tradeoff because it reduces the time necessary for the attack at expense of needing more space.
   Common mistake: We generally assume that the hash function is known to the attacker. See Kerckhoff's principle.

(c) What is a salt? How is a salt generated? Describe how a salt defeats the type of attack asked for in question a). (2p)

A password salt is a randomly generated string that is unique for each user. It is appended to the password before hashing. In this way, even when using a common word for a password, the salt ensures that the resulting hash value can't be compared to a dictionary.

(d) Should the salt be kept secret? Motivate. (1p)

There is no need for the salt to be secret. It is usually stored in open, next to the salted hash. The point of the salt is not to be secret, but simply to defeat the dictionary attack by slowing it down. Remember that an attacker with access to hashed passwords probably also has access to the salts anyway!

7. Fingerprints, irises and handwritten signatures are properties used for biometric identification of users. What is meant with enrollment, false rejection rate and false acceptance rate? Describe three important properties of biometric authentication methods which are different from other ways of authenticating users. (6p)

(a) Enrollment is when a user first is registered as a user and the biometric features are capture and stored in the user database for the first time.

(b) False rejection rate is the probability that a legitimate user is declined access.

(c) False acceptance rate is the probability that an unauthorized user is granted access.

(d) A biometric feature is inseparable from the body of a person. It can not be removed or is at least very difficult to remove.

(e) Biometric properties are not constant and thus every capture of the property will be different. There will be scale to how close a capture feature is to the value captured during enrollment. In contrast with passwords which are either correct or incorrect. Some persons lack specific features e.g. not everyone has a fingerprint (may not even have a finger). Everyone can get a password.

Common mistake: It is important to note that we are asking about *differences* between biometrics and other types of authentication. Many have seen the word "property" and just repeated common *properties we require of biometry* which is something completely different.

8. Access control is often simplified by using groups. But groups can cause an ambiguity interpreting privileges, requiring new rules. The course mentions two different main principles for such interpretations. Describe them shortly and also explain how they relate to the use of the negative privilege "None"! (3p)

Since users can belong to several groups with different permissions, and can also have individual permissions, you must decide what then applies. One principle is "first relevant entry", which means that access is allowed or denied according to the first entry in an ACL, which applies to the user. The other main principle is "any permission", which means that if there is any record in the ACL, which allows access for the user, then access is allowed. If you use "first relevant entry", you can allow entry

for an already existing group, but exclude one (or more) single user by adding a "none" entry for that user before the group entry in the ACL.

9. In a statistical database, a group of three is too small to conceal numerical values. This is true even if you only are allowed to query about MAX, MIN, and AVERAGE. Why? (2p)

   One value is given by MAX, the other by MIN, and the third can be obtained from the formula 3×AVERAGE-MIN-MAX.

10. (a) Is a firewall a useful tool in defence against DoS attacks? Describe shortly how it can be used or motivate why it is irrelevant for that specific threat. (2p)

    - A firewall can block offending sender sites, once you have identified them, A firewall can block the port used by the DoS attack, disabling this service, but keeping the rest going, and allowing the internal network to operate without disturbance.

    (b) Is an IDS a useful tool in defence against DoS attacks? Describe shortly how it can be used or motivate why it is irrelevant for that specific threat. (2p)

    - An IDS can detect probing and identify probing sites, so that they can be blocked. It can detect attempts at finding and using a known vulnerability for an attempt to crash the system. If you realise that a bot infected computer taking part in a massive DoS attack is blocked for its user and thus itself a DoS victim, you can argue that finding bot software on your computer or traces of its activity is also a DoS defence.

11. Draw a diagram of Kerberos and briefly describe each participant. What secrets are shared between the participants? What final data item does the client gain in the process? (5p)

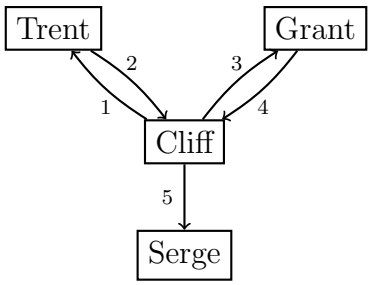

- The client Cliff wants to connect to a server Serge. He first contacts the authentication server (or trusted authority) Trent and then the access-granting (authorizing) server Grant, and then he can get access to Serge.
- There are shared secrets between Cliff and Trent, between Trent and Grant, and between Grant and Serge.
- The end result is that Cliff holds a ticket that he can present to Serge to get access to the service provided.

12. In Swedish law, why is an e-mail without a digital signature cannot be considered an "urkund"? (2p)

Among other requirements, an "urkund" must be *reliably verifiable*. A mere e-mail isn't verifiable by itself, however with an e-mail signature this becomes possible.