

# Computer Security

## Draft Exam with Answers. 2009.

Please note that the questions written here are a draft of the final exam. There may be typos in the questions that were corrected in the final version (available on the web page).

### Question 1. Access control

(a) In access control systems, what is a *capability*? [10%]

A capability is a piece of data possession of which proves authorisation to access a resource.

(b) (i) Explain an advantage of access control lists over capability lists. [5%]

Access control lists make it easy for an administrator to see who has access to a given resource. (This is harder with capabilities.)

(ii) Explain an advantage of capability lists over access control lists. [5%]

Capabilities may be transferred offline between users. This is generally not possible with access control lists.

(c) A hospital patient record system provides login accounts for nurses. It is desired to implement the following policy:

- (i) When a nurse registers a new patient, the nurse is granted access to the patient's records for a period of 90 days.
- (ii) A nurse possessing the right to access a patient record can give that right to another user (this facilitates staff shift changes). This may be done offline.

To implement this policy, the system works as follows. When a nurse registers a new patient, a capability to access the patient record for the following 90 days is generated. The nurse stores it on a USB stick, and may copy it onto other USB sticks to give to other users. When a user attempts to access patient records, she is prompted to upload the relevant capability. The capability has the following format:

*patient-id, issue-date, hmac(K, (patient-id, issue-date))*

where  $\text{hmac}(K, \dots)$  denotes a suitable keyed hash function with key  $K$ . The key  $K$  is a secret key known only to the patient record system. Any user in possession of this capability is able to access the records of the patient with *patient-id*, provided the date is within 90 days after *issue-date*.

- (i) Suppose nurse A registers a patient and receives such a capability. A passes it to B, B passes it to C, and C passes it to D. Is D able to use the capability? [4%]

Yes, provided the capability is still within its validity period.

- (ii) In order to stop long-lived capabilities being distributed widely, the hospital decides to adopt the policy that the nurse that initially registers the patient will have access to the records for 90 days, as before, but if she passes the capability to any other user, the validity should be 10 days from the issue date. Explain a new format for capabilities which would support this new policy. [10%]

If we include the user name in the capability, then the system can be programmed to grant 90 days to the named user, and 10 days to any other user. Thus, the new format is:

*patient-id, issue-date, user-id, hmac(K, (patient-id, issue-date, user-id))*

## Question 2. Trusted computing

- (a) What is the purpose of a *platform configuration register* (PCR) in the TPM? [10%]

The purpose of a PCR is to allow the TPM to record the platform state and configuration in a secure way.

- (b) Explain how a PCR  $p$  is updated with new data  $d$ . [10%]

PCRs may not be written directly; the process of incorporating data  $d$  into a PCR  $p$  is known as extending  $p$ , and corresponds to the assignment  $p := \text{SHA1}(p || d)$ .

- (c) The command `TPM_Seal` is used to encrypt data and bind it to some PCR values. Suppose `TPM_Seal` is called with arguments including the usual arguments concerning keys and authorisation data, and also the PCR  $p$  and associated value  $v$ . The command runs successfully and returns a blob.

- (i) Explain the conditions under which the blob can later be decrypted. [7%]

The blob can be decrypted only if the PCR  $p$  has value  $v$ , and the relevant authorisation data is given as part of the command.

- (ii) Explain a scenario in which `TPM_Seal` might be useful. [7%]

For example, in Microsoft BitLocker, the volume encryption key is sealed by the TPM, and can be decrypted only if the relevant PCRs have the correct values, thus ensuring that the system is in the correct configuration state.

## Question 3. Electronic voting

- (a) Explain the meaning of the following properties as used in the context of electronic voting systems.

- (i) *individual voter verifiability* [5%]

- (ii) *vote-privacy* [5%]

*Individual voter verifiability* asserts that, after the election, a voter may verify that her vote has been included in the declared result. *Vote privacy* means that a voter cannot be linked

with the vote that she cast.

(b) For each of the following voting systems, does it satisfy individual voter verifiability? Does it satisfy vote privacy? Explain your answers.

- (i) Ballot papers such as the ones shown below are distributed, and voters mark with an X their chosen candidate. (A ballot marked in any other way will be considered invalid.) Then an upturned hat is handed around the room. Each voter puts her ballot paper into the hat. At the end of this procedure, someone pins all the ballot papers on a publicly viewable noticeboard, and counts the votes.

<i>Candidate</i>	
C. Baldwin	
B. Barr	
J. McCain	
C. McKinney	
R. Nader	
B. Obama	

[8%]

Individual verifiability: no. There is no way for a voter to distinguish her ballot paper from that of another voter that voted in the same way. So she cannot be sure her ballot paper is included among those posted on the noticeboard.

Vote privacy: yes. A ballot paper on the noticeboard cannot be linked with any voter.

- (ii) PGP and email are used. An administrator creates a PGP key for the purpose of running the election. We may assume that each voter has an authentic copy of this key (it has been given to each voter in person). To vote, a voter sends an email to the administrator containing her name and her chosen candidate. The email is encrypted with the administrator's key. At the end of this procedure, the administrator publishes the list of voters and their selected candidates on a web page, like this:

<i>Voter</i>	<i>Candidate</i>
Mary Jones	Obama
Bob Smith	McCain
Anne Wood	Obama
Peter Green	McKinney

[8%]

Individual verifiability: yes. It is easy for a voter to verify that her vote is included in the published list.

Vote privacy: no. The published list links the voters and their votes.

(iii)

The protocol by Fujioka, Okamoto and Ohta (studied in lectures) is used. To help you remember the details, a figure illustrating the protocol is given below.

[Figure goes here]

[8%]

Individual verifiability: yes. A voter checks that the commitment to her vote is present on the bulletin board, and that the corresponding opened vote is also present.

Vote privacy: yes. Because the vote is blind-signed by the administrator, no-one (not even the administrator) can link a voter and her vote.

#### **Question 4.** Short notes.

Write short notes (maximum 100 words) on each of the following topics.

(a) Web application penetration testing [10%]

A web application penetration test is a method of evaluating the security of a web application by simulating an attack from a malicious source. It analyses the system for potential vulnerabilities that may result from poor or improper system configuration, known and/or unknown hardware or software flaws, or operational weaknesses in process or technical countermeasures. Web application penetration testing typically checks for vulnerability to URL manipulation, SQL injection, cross site scripting, back-end authentication, password in memory, session hijacking, buffer overflow, web server configuration, credential management, clickjacking, etc.

(b) ISO 27001 Information Security Management System [10%]

ISO/IEC 27001 is an ISO standard for evaluating security in information systems. ISO/IEC 27001 certification involves a review of the existence and completeness of key documentation such as the organization's security policy, Statement of Applicability (SoA) and Risk Treatment Plan (RTP). It also includes a detailed, in-depth audit involving testing the existence and effectiveness of the information security controls stated in the SoA and RTP, as well as their supporting documentation. Certification maintenance involves periodic reviews and re-assessments to confirm that the information system continues to operate as specified and intended.

(c) How could Alice surf the web without her activity being logged by her ISP? [7%]

Alice can use an anonymising proxy, such as anonymizer.com. The traffic between her and the proxy will be SSL encrypted, so that her ISP will see that she is communicating with the proxy but will not see the content of the communication (in either direction). In

particular, it will not see the URLs or contents of the web pages she is actually browsing. This method is insecure if the ISP is able to collude with the proxy. For added security, Alice can use The Onion Router (TOR), which hides her traffic in layers of encryption, each layer controlled by a different node. Compromise of all the nodes simultaneously would be required by the ISP to break Alice's privacy.

(d) In what way is the encryption scheme known as DES considered weak? [7%]

DES is considered insecure chiefly because its 56-bit key size is too small by the standards of today's computation power. In January 1999, distributed.net and the Electronic Frontier Foundation collaborated to publicly break a DES key in 22 hours and 15 minutes. There are also some analytical results which demonstrate theoretical weaknesses in the cipher, although they are unfeasible to mount in practice. The algorithm is believed to be practically secure in the form of Triple DES, although there are theoretical attacks. In recent years, the cipher has been superseded by the Advanced Encryption Standard (AES).