ECE 750: Usable Security and Privacy - Final	Marks obtained ↓
Date: Aug 8, 2024, Total questions: **4** Total points: **50**	
ID: Name:	Time: 2.5 hrs

Page:	2	3	4	5	6	7	9	Total
Points:	8	6	10	10	6	5	5	50
Score:								

Instructions

No aids allowed. All you are allowed is a pen and pencil.

Use space provided. Answer the questions in the spaces provided. If you run out of room for an answer you can use the back of the page but clearly note in the original answer space that the back was used.

Point value in right-hand margin. The number of points each question is worth is listed in the right hand margin. The number of points and the space provided are roughly indicative of how long of an answer is expected.

Pencils and pens allowed. Both pens and pencils are allowed. Pencil marks need to be clearly present or erased. If it is unclear if a mark is or is not present then it is up to the discretion of the grader and this point cannot be contested later.

Scratch paper. At the end of the exam there are two blank pages marked "SCRATCH PAPER", you may tear these off and use them as scratch paper.

Passwords

1. Imagine an attacker was able to gain access to the shadow password file on a Linux server. The file contains usernames, hashed passwords, and their associated salts. The attacker wants to find the passwords of users like root.

   (a) We learned about online and offline password guessing attacks. Is the attacker most likely to use an online or an offline attack in the above situation? Briefly describe what about the situation makes your chosen attack the most appropriate for the attacker. (4)

   Solution: Offline. Because the attacker has the file content and can therefore directly make guesses against the hash+salts. In an offline attack they can create and test their guess locally which is the fastest option. In an online attack they would be limited by the network speeds, page reload time, and any guessing rate limits.

   Solution: POST MARKING COMMENTS:

   - Pre-computed hash tables (aka rainbow tables) will not work in this case because the passwords used a salt. The salt means that the attacker must compute hash(password,salt)==passwordHash for every possible password and every password entry in the file. There is no way to do a precomputation because the salt for every password is different.

   (b) Imagine a user is assigned a 4 character password randomly selected from a possible set of 56 characters. The password is used on a website that uses lockout. If more than 10 passwords are attempted against 1 user name, then the system locks the account for 24 hours. Under what assumptions might this password be considered sufficiently secure? (4)

   Solution: The "don't care" region of password strength. If lockout is correctly implemented then the attacker is limited in their guessing to a relatively small number of passwords. As long as the resource being protected isn't too sensitive. A short password can be ok if it is truely random.

2. Long random passwords are better for security because they are harder to guess. But long truly random passwords are bad for usability because they are challenging for people to remember.

   To better understand what length random password people can reliably remember Edith designs an online user study. She recruits people using an online survey platform like Amazon Mechanical Turk and requires them to have a high positive rating on the platform. She uses survey software to create an online experiment where users see information and answer questions via a webpage. Users in her study go through the following steps:

   1. Consents to be part of research.
   2. Software randomly decides if they will get a 6, 9 or 12 character password, and then creates and records a random password of that length containing at least one upper case, one lower case, one number, and one symbol. The user is shown the password and told they will need to enter this password later in the study.
   3. On the next page the user is asked to enter the password to make sure they remember it, if they enter the wrong password, they are shown the password again and reminded they need to remember it.
   4. The user then fills out the IUIPC scale.
   5. The user is asked to enter their password. No feedback is given about if it is correct or wrong.
   6. The user then fills out the SeBIS scale.
   7. The user is asked to enter their password. No feedback is given about if it is correct or wrong.
   8. The user fills out a set of demographics questions including age, gender, native language, and current ocupation.
   9. The user is asked to enter their password. No feedback is given about if it is correct or wrong.
   10. They are thanked for taking part in the survey.

   To analyze the results, Edith counts how many times participants enter their password correctly on steps 5, 7, and 9 producing a score between 0 and 3. She then uses t-tests to see if there is a statistically significantly difference in the mean number of errors between each pair of password lengths.

   Answer the questions below about the study. Feel free to use the step numbers above in your answers.

   (a) List the **independent** variable(s) Edith is collecting in her study. (3)

   > **Solution:** She collects and uses: the condition (step 2, password length). She also collects demographics (step 8). She also collects the IUIPC scale (step 4), and SeBIS (step 6) though it is unclear how they are used.

   (b) List the **dependent** variable(s) Edith is collecting in her study. (3)

   > **Solution:** Password entry accuracy which is a sum of steps 5, 7, and 9. The password entry in step 3 is ignored in the analysis and so is normally not a dependent variable.

(c) The study has several internal validity problems. Describe one of them. (6)

Solution: The largest issue is that users are being asked to remember the password over a very short time and being frequently asked to recall it. The situation is quite different from how normal password recall is meant to happen. Online users are also quite likely to write the password down rather than try and recall it from memory. These issues mean that Edith is unlikely to be able to accurately answer her question

(d) Like all studies, this study has some external validity limitations. Describe one external validity limitation. (4)

Solution: The most obvious external validity issue is the users chosen. Similar to Assignment 2, the people being studied are all from an online survey platform. They take surveys every day and are generally quite familiar with the Internet. The study results are unlikely to directly align with how people less skilled with internet technology might behave.

Solution: POST MARKING COMMENTS:

- Limiting participation to those that have high scores on Mechanical Turk is a normal thing to do in studies. It is done to avoid having participants who are known to cheat or that use robots to fill in answers. It is an external validity issue, but having lots of participants that are known to not pay attention when taking surveys is a larger issue. We did not learn this point in class, so no marks were taken in regards to it.

Public/Private Key Cryptography

3. (a) Fill in the blanks in the following text describing Alice and Bob correctly communicating securely. Each blank needs to be filled in using one of the following words: (6)

   - Encrypt, encrypted, encrypts
   - Decrypt, decrypted, decrypts
   - Sign, signed, signs, signature
   - Public
   - Private

   Alice wants to send an encrypted and signed message to Bob who she has met in the past. When they last met in person they verified and then ***signed*** each other's _____ keys using their respective _____ keys.

   To send an email to Bob, Alice first ***signed*** the message using her _____ key and then _____ the resulting message using Bob's _____ key. Alice then sends the message over an untrusted connection.

   Bob receives the message and first _____ it using his _____ key. He then verifies that the message really was from Alice by verifying the ***signed*** using Alice's _____ key.

   Solution: Alice wants to send an encrypted and signed message to Bob who she has met in the past. When they last met in person they verified and then **signed** each other's **public** keys using their respective **private** keys.
   To send an email to Bob, Alice first **signs** the message using her **private** key and then **encrypts** the resulting message using Bob's **public** key. Alice then sends the message over an untrusted connection.
   Bob receives the message and first **decrypts** it using his **private** key. He then verifies that the message really was from Alice by verifying the **signature** using Alice's **public** key.

   (b) Research, such as the *Why Johnny Can't Encrypt* paper, often mentions "confusing metaphores" in regards to encryption interfaces. Give a specific example of a metaphor related to encryption that is likely confusing for people. (4)

   Solution: The most obvious answer is a key or a lock. Both are commonly used in encryption technology to refer to the act of encryption using the metaphor of a key and lock. Unfortunately physical keys work rather differently than encryption keys, particularly in regards to public key cryptography where there are two keys that both interact with the same lock (encrypted text) in different ways.

(c) Would Alice and Bob's security in the interaction above be improved by using a Certificate Authority (CA)? State 'Yes' or 'No' and then briefly explain what a CA would do to help Alice and Bob OR explain why their current security practices are equal to or better than what a CA would provide. (6)

Solution: No. Alice and Bob directly verified their public keys and signed them in the past. A CA's role is to attest that a particular public key is associated with a specific person. Since Alice and Bob already directly verified and signed each other's keys, they do not need a third party to do the verification. Therefore a CA would not improve the security of their interaction. If anything it would add the risk that the CA was corrupt which would increase the risk. CAs are run by many entities, including entities tied to governments. So when two countries are in conflict, sometimes the CA for one country will sign bad certificates enabling a digital attack.

Solution: POST MARKING COMMENTS:

- This question and the points below assumes you answered Question a correctly. If you did not, that was taken into account.
- A Certificate Authority *only* provides verification of identity by using its private key to sign the server's public key. It puts the signature in a "certificate" that contains information beyond just the key like the domain the key is for. It is used during a TLS setup by the browser, app, or email provider. But a CA itself does not provide TLS or in any way perform the encryption.
- TLS itself happens over an untrusted connection. A trusted connection would be something like a dedicated wire used only for this purpose or a courier who is paid to hand deliver a message.
- TLS and PGP use the same level of encryption. The keys are different, so there is a tiny improvement in terms of someone being able to break the encryption. But it is rare now days that anyone outside of research tries to break encryption directly. It is much easier to do things like bribe a CA to sign the wrong key.
- Not covered in this course, so I do not expect you to know. But there are situations where applying encryption more than once can degrade the encryption and make the message less secure.

Design Layout

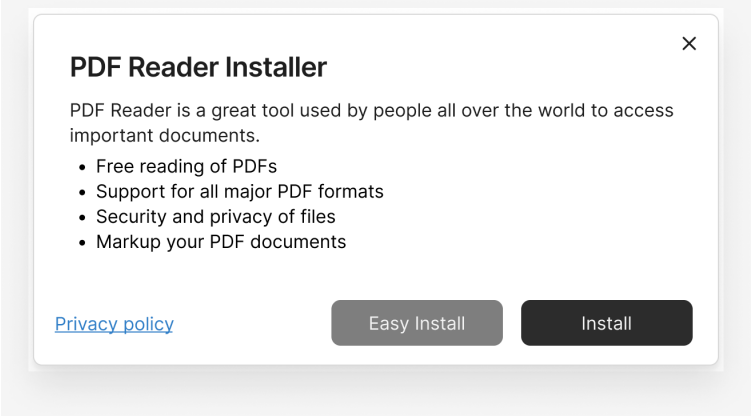

(a) PDF Reader main installation screen.

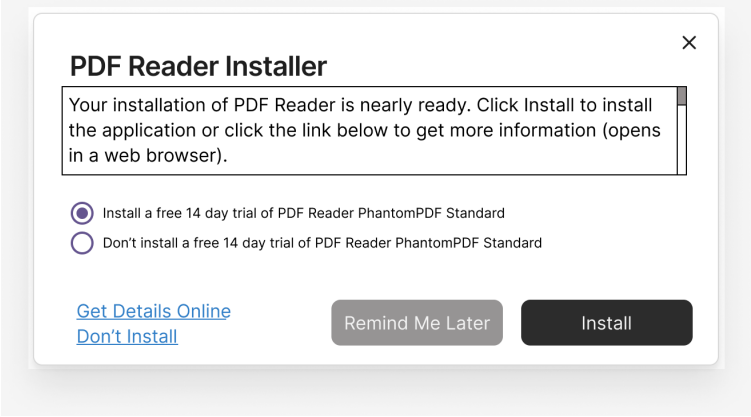

(b) Installation screen for PDF Reader addon PhantomPDF Standard.

Figure 1: Two screens the user encounters when trying to install PDF Reader.

4. Figure 1b is an example of deceptive design. The user is trying to install an application called "PDF Reader". The first screen they see is the main installation screen (a). They click 'Install' and screen (b) appears. Screen (b) is offering to install a free trial of "PDF Reader PhantomPDF Standard" in addition to the "PDF Reader" software the user is trying to install. Clicking "Install" on screen (b) will install both PDF Reader Installer *and* PDF Reader PhandomPDF Standard. Clicking "Don't Install" on screen (b) will install only PDF Reader Installer which is what the user wants to install.

   (a) Use the Gestalt Principles on visual grouping to explain why a user might have trouble realizing that the "Don't Install" option in Figure 1b is a button similar to "Install" rather than a web link similar to "Get Details Online". (5)

   Gestalt Principles: 1. Proximity 2. Similarity 3. Continuity 4. Closure 5. Symmetry 6. Figure/Ground 7. Common Fate

QUESTION 4 ANSWER SPACE

Solution: The most obvious are Similarity and Proximity. The two buttons for Install and Remind Me Later are designed to be similar in size, font, and they are both dark colors. They are also located right next to each other. The blue text on the left is not similar to the buttons, it is also located on the other side, so not in close proximity. Instead the two blue underlined lines look similar to each other and different from the buttons.

Considering the user saw Figure 1a first, they are also likely to assume that buttons look similar to the Install button in Figure 1a, and that web links look like the "Privacy policy" link in (a).

Solution: POST MARKING COMMENTS:

- I expect the Gestalt Principles to be used to *explain*. Simply stating the buttons are not similar is not enough. I state in the question that the user does not consider them to be similar. I expect to see information in the answer beyond what is in the question.
- Figure/Ground causes the person to perceive smaller objects as on top of larger objects. It is easily applied to buttons in that users will see the button text as "on top".

(b) Draw an improved sketch of Figure 1b. The design does not need to be super detailed, but all the buttons need to be clearly labeled and it should be clear where text from the original design was moved or adjusted. (5)

A copy of Figure 1b also appears on the scratch paper at the end of the exam.

Solution: There are many confusing elements to Figure 1b that should be improved.

- The radio buttons ave "Install" and "Don't install" as do the buttons. Its unclear what happens if the user selects "Don't install" and then click's "Install".
- The user is trying to install software, they may not want this free trial, but they are trying to install PDF Reader. So they want to "Install". Its unclear if they click "Don't Install" does that cause *both* software to not be installed or does it cause just the free trial to not be installed.
- From Question 4, there is also an issue around design itself. The user is expecting a button and probably not a link.

Solution: POST MARKING COMMENTS:

- Easiest solution is to change to three buttons:
  - Remind me later
  - Install free trial
  - Continue without trial
- Other easy approach is to let the user choose what is installed by providing multi-select boxes that let them select what is installed. This allows the user to do what they want to do – Install PDF Reader – without confusion.
  - ☐ PDF Reader
  - ☐ Free 14 day trial of PDF Reader PHantomPDF Standard
  - Buttons: "Install" and "Cancel"
- Common error was to create a "Don't Install" button where it was unclear if it was only the free trial not being installed or if both the free trial AND PDF Reader install was being canceled.
- Putting information about the "free trial" in a large block of text the user is unlikely to read increases the deceptive design by making it easier for the user to accept the free trial install without knowing they accepted it.

SCRATCH PAPER - ok to tear off

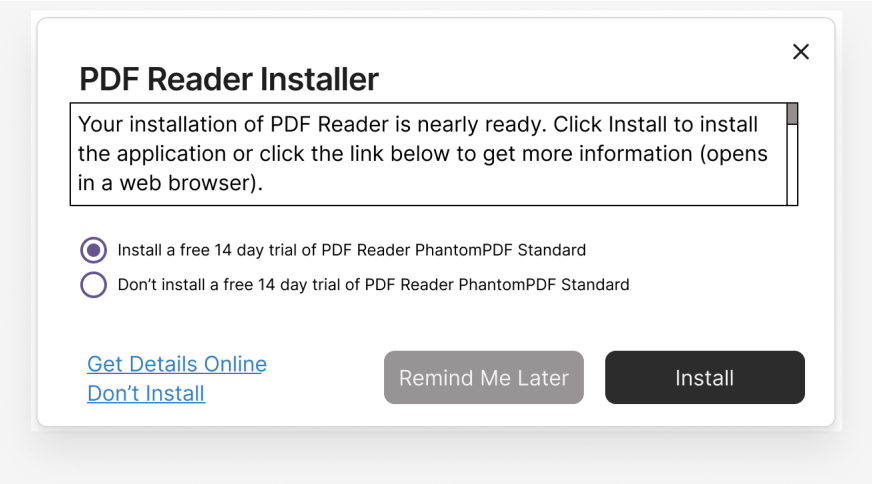

Figure 2: Copy of Figure 1b

SCRATCH PAPER - ok to tear off