CS42600 Spring 2019 Midterm 2, April 2, 2019
Prof. Chris Clifton

Turn Off Your Cell Phone. Use of any electronic device during the test is prohibited. As previously noted, you are allowed notes: Up to four sheets of 8.5x11 or A4 paper, single-sided (or two sheets double-sided).

Time will be tight. If you spend more than the recommended time on any question, **go on to the next one**. If you can't answer it in the recommended time, you are either giving too much detail or the question is material you don't know well. You can skip one or two parts and still demonstrate what I believe to be an A-level understanding of the material.

Note: It is okay to abbreviate in your answers, as long as the abbreviations are unambiguous and reasonably obvious.

In all cases, it is important that you give some idea of how you derived the answer, not simply give an answer. Setting up the solution correctly, even if you don't carry out the calculations to get the final answer, is good for nearly full credit.

1 Identity (5 minutes, 3 points)

Discuss the tradeoffs (from a security point of view) of these email address assignment schemes employed by a company.

A. Firstname.lastname@company.com

B. Randomstring@company.com

1 increases vulnerability to targeted attacks, such as spearfishing or trying to break into an account with access to specific data. 2 provides "security through obscurity" to such attacks - not the best, but something. However, the reduction in ease-of-use may prompt workarounds (such as posting email addresses) that would not only defeat such attacks, but also increase other vulnerabilities such as scraping websites to find email addresses for spam.

Other good examples include likelihood of misdirected email in either scheme (typos, name collisions), misidentification, etc.

Scoring: 1 usability distinction, 1 for security issue (there were many good examples).

2 Risk Analysis (5 minutes, 3 points)

Referring to the 6 basic steps of risk analysis from lecture, what step(s) are attack graphs used in? Briefly describe how. You can get partial credit by naming (as many as you can) of the six steps, or by describing what an attack graph is used for.

Attack graphs model the steps required for an attack to succeed, and can be annotated with the probability of each step succeeding to estimate the overall likelihood of an attack on a particular path succeeding.

- **Identify Assets**
- **Determine vulnerabilities - attack graphs.**
- **Estimate likelihood of exploitation - attack graphs.**
- **Compute expected annual loss**
- **Survey applicable controls and their costs**
- **Project annual savings of controls**

Scoring: 2 for good explanation of attack graphs, 1 for mapping appropriately to steps (good explanations existed to map this to several toher steps as well.) 2 for 4+ steps, 1 for 2-3 steps.

3 Authentication (3 minutes, 2 points)

What is authentication and how does it differ from Identification?

Authentication validates a purported identity, binding the principle to the identity. Identification is simply asserting the subject.

Don't confuse these with authorization, which is determining if an identity is allowed to perform an action. Authorization without identification and authentication rarely makes sense (particularly if we think of identity generally, such as "an authorized user" without knowing who specifically.)

Scoring: 1 for showing understanding of authentication, 1 for identity.

4 Anonymity (5 minutes, 2 points)

Why would an e-commerce company like Amazon resist anonymity?

Amazon wants to recommend products to people likely to buy them, either through product placement or advertising. True anonymity breaks the link between an individual's prior history and current activity, thus reducing the effectiveness of such recommendations.

There were other good examples - fraud prevention, product returns/credits. Simply "payment" wasn't a good enough answer (anonymous payment at Amazon is pretty easy - I'll let you figure it out.)

Scoring: 1 for some issue that requires identifying specific individual, 1 for the idea that anonymity prevents this.

5 Malware (5 minutes, 4 points)

What is the difference between a Virus, a Trojan Horse, and a Worm? Define all 3 and give an example for one of them.

Virus: Code that resides in a host code of some sort and spreads when that code is executed. Example: Stuxnet

Trojan Horse: Code that actually does something the user wants it to, but has (malicious) functionality hidden from the user.

Worm: A program that spreads copies of itself without requiring some other program to be executed. Example: Morris Worm

The key is that both viruses and worms replicate, but that a virus needs some host program to be executed.

Scoring: 1 point for each definition, 1 point for a correct example (either reasonably close name or description of a specific instance).

6 Audit (15 minutes, 6 points)

Assume we have a security policy where only authorized people are allowed to enter the building after hours. The building administrator maintains a list of who is authorized, adding and removing people as needed.

To enter the building, people have to enter the outer door, then once it is closed and locked and the sytem verifies that only one person is in the "airlock", the users swipes their ID and uses a fingerprint reader, the fingerprint is checked against one stored on the ID. If the fingerprint matches the one on the ID, and the ID is on the authorized person list, the inner door opens. People exit the building through a one-way turnstile, so there is no way into the building other than through the "airlock" door.

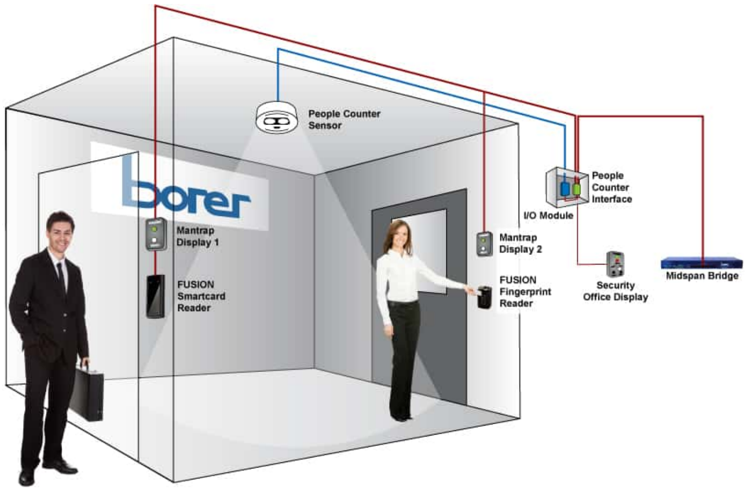

The system logs each card swipe, giving a log that looks like the following:

Time	*Door*	*ID*
4/1 21:30	North	5399205
4/1 22:40	South	9394203
4/1 22:44	South	5399205

6.1 Audit (5 minutes, 2 points)

Assuming the system works as designed (there is no way to enter without swiping an ID, people can only use their own ID, and every entry is logged), is this sufficient to audit to determine if the security policy “only authorized users can enter after hours” has been violated? If no, explain what is missing, if yes, explain how an audit would work.

No. While this system allows us to know who entered, it doesn’t capture if they were authorized *at the time of entry.* We would also need to log changes to the authorized list.

Scoring: 1 for ”no”, 1 for missing changes to authorized persons. 1 for ”yes” with good explanation assuming no changes to authorized users.

6.2 Authentication (3 minutes, 2 points)

Does this system use two-factor authentication? Explain.

Yes: The ID (something you have) and Fingerprint (something you are).

Scoring: 1 for ”yes”, 1 for ID+Fingerprint. 1 point for ”no” with good discussion that shows clear understanding of what two-factor authentication is.

6.3 Identity (3 minutes, 2 points)

Suppose we don’t want to have a log that can be used to determine who was in the building, but we still want to be able to audit if only authorized people entered the building. Explain what you would log to accomplish this.

The key is to eliminate the link to the individual, but still log that an access was made, and that it was authorized. Recording, for example, every attempt at access, but only logging the timestamp of the authorization list against which the ID was checked and if it matched or not, would enable auditing that policy was followed without logging which ID was used.

Why would we want to do this? For us old folks who remember “the McCarthy era” (look it up), a campus might not want to have a list of students who entered a building where a communist party meeting was held, as it could be subpoenaed by Congress and used to get the students in trouble. This isn’t part of the expected answer...

Scoring: 1 for eliminating link to individual, 1 for maintaining authorized log

7 Two-Factor Authentication (6 minutes, 6 points)

Which of the following types of attacks are prevented by two-factor authentication (e.g., BoilerKey). Your answers should be based purely on this being two factors, not other protections (e.g., encryption) that are used as part of this. For each, give a yes or no and briefly explain your answer.

A. Man-in-the-middle

No. It would be possible for a man-in-the-middle to capture and forward both types of authentication information. While some implementations (such as boilerkey) avoid this, it is not because of the second factor, but another part of the overall system (in the case of boilerkey using your phone, the fact that the second type uses a different communication channel.)

B. Replay

No. Again, the information about both types of authentication could be captured and replayed. Most two-factor authentication systems use a challenge-response or non-repeating token to avoid replay attacks, but this is separate from having two factors.

C. Brute force password cracking

Yes. While brute force password cracking may be effective against a password (something you know), it doesn't make sense against other types of authentication (e.g., a fingerprint or token).

8 Risk Analysis (7 minutes, 4 points)

Suppose we estimate that a phishing attack against the Chief Financial Officer of a company has a 25% chance of succeeding, and could result in early disclosure of financial results and a $100,000 fine from the Securities and Exchange Commission. A Denial-of-Service attack has a 5% chance of succeeding, and could shut the company sales site down for a day, resulting in the loss of $1,000,000 in revenue. Explain how you would use this analysis to determine where to invest your security budget, and other things you might want to know to make this decision.

The expected loss from a DoS attack is .05*$1,000,000; from the Phishing attack is .25*$100,000. So the expected loss from DoS is higher, but we'd still need to know the cost of mitigation to decide where to invest. Further relevant information could include the reduced risk as a result of the mitigation, and the expected frequency of attacks.

Scoring: 2 for proper estimate of expected cost (1 for at least having the idea), 1 for cost / benefit of security investments, 1 for tradeoff based on relative value of different investments

9 Mandatory Access Control (10 minutes, 5 points)

Assume an intelligence agency has multiple security levels (**H**igh, **M**edium, and **L**ow) and uses the Bell-LaPadula mandatory access control model, which prevents information from flowing from higher security domains to lower security domains.

9.1 Write Down (5 minutes, 3 points)

Assume a user with Medium clearance is editing two documents, one at the Medium level, and one at the Low level. The user tries to cut text from the Medium document and paste into the Low document. What should happen? Explain your answer.

Hint: There are three valid answers, you get 2 points for giving and explaining one correct answer, and a third point for giving and explaining a second correct answer.

Answer 1: L document contaminated with M information, so now at M level.

Answer 2: Paste to L document blocked.

Answer 3: User is trusted to only paste L information from M document to L document. This is a bit subtle - and generally would not be the way MAC would be implemented. Essentially, you are allowing the user to declassify information themselves.

Scoring: 1 for showing some understanding that you can't move information from high to low, 1 each for answers above or formal explanation.

9.2 Write Up (3 minutes, 2 points)

Would Bell-LaPadula allow a user at the Medium level to write to a document at the high level? Briefly explain why or why not.

Yes. They can't see the file, but can write to it. Note that it may not even be allowed to know if there is a file. But think of the Posix "open" - if you have access to the directory and you open in write mode, if the file doesn't exist it creates it. So you can write to a file without even knowing if it exists, much less reading from it.

Scoring: 1 for Yes, 1 for explanation. No okay if the explanation makes it clear that you assume that writing implies reading..

10 Formal access control models (12 minutes, 3 points)

We can model Bell-LaPadula Mandatory Access Control, which prevents information flowing from higher security levels to lower security levels, using the Harrison-Rizo-Ullman Access control matrix. One approach would be to have a Level subject that captures the security level of everything in the system:

	Level	Daniels	Clifton	Student	Evaluations	Answers	Exam
Level		3	2	1	3	2	1
Daniels							
Clifton						reading	
Student						writing	reading

For example, we could model a file opened for reading as:

```
open_read(Subject, Object) -- S opens O for reading, if authorized
  if  A[Level, Subject] >= A[Level, Object] then -- Subject at high enough level, authorized
    A[Subject,Object] = A[Subject,Object] + reading
```

10.1 HRU commands (5 minutes, 2 points)

Using the same basic style, model the open_write command so that it satisfies the Bell-Lapadula mandatory access control policy.

```
open_write(Subject, Object)
  if  A[Level, Subject] <= A[Level, Object] then
    A[Subject,Object] = A[Subject,Object] + writing
```

Scoring: One for check subject level $\leq$ object level, 1 for adding writing to matrix.

10.2 HRU usage (5 minutes, 1 point)

We use HRU to determine if a set of commands is "safe", that is, no right is leaked. Describe briefly what we would mean by "safe" when using HRU to evaluate if a system satisfies Bell-LaPadula.

That no execution of the commands can lead to a subject having can have reading on an object a higher level, or writing on an object at a lower level.

A point requires having both ideas: we need to ensure the constraint isn't violated, and that this is following any execution of the set of commands. I didn't expect many people to get this...