COM-301 2023
Final Exam MRE

General remarks

- It is important to answer the question you are asked. As it can be seen in the errors described below. There are a number of cases where answers just do not correspond to the question, deviate from the question assumptions, or forget to answer some part of the question. Only answers that actually correspond to the question received points.
- When you are asked to write a response in the given amount of lines it is important that you do not go beyond. We considered it to be a bit wider or have a couple of words in the next line due to some correction. In answers that blatantly ignored the instruction, we only considered the first lines when doing the correction.When questions ask for a justification, it is important to give the justification. Answers that only gave properties names, high-level attacks, or partial explanations did not receive full points.

Open questions

Crypto: Star Wars (Q9 and Q10)

Q9: Learning Yoda's message

Error: Recovering the message from the signature. It is not possible to recover a message from a signature. Signatures do not have a `decrypt' function that allows to recover the signed content. Signatures only allow for *verification* in which *given a message* the verifier can be assured that the signature has been computed on that message.
This error was common in mini-exam 3 2020, and mini-exam 2 2021. The MRE document of mini-exam 3 in 2020, explains this issue in detail.

Error: Encrypting the same string twice outputs the same ciphertext.
In order to not be vulnerable to frequency attacks (like classical ciphers, e.g. Caesar) secure encryption schemes do not output the same ciphertext even if the input plaintext is repeated. The question is explicit that the encryption scheme does not have vulnerabilities, thus any attack based on this kind of reasoning is incorrect. Yet, because the correct answer was in this direction, we did give partial points to these answers.

Error: The adversary performs a man-in-the-middle attack.
The question states "an adversary eavesdropping on the conversation". Any attack that

requires the adversary has further capabilities, like a man-in-the-middle-attack, is therefore not valid.

Q10: Extra information to perform attack

Error: Assuming bruteforce of the keys is feasible.
For secure schemes (as the one in the question that explicitly has no vulnerabilities), it is not possible to brute force the key. Otherwise they are not secure. Answers contradicting the scenario in the question are not valid.
This error was also common in the Midterm 2022.
(Besides, if the key was not strong enough, then you don't need the extra information. You could just brute-force the key before without knowing the plaintext (there is only yes and no). So it is not the fact that you have known plaintext that would give you an advantage).

Error: Key recovery is not possible because Obi Wan's secret key is not used during the exchange.
It is not correct to assume that key recovery attacks are not possible when the adversary "only" knows (public key, plaintext, ciphertext) tuples, where indeed the secret key is not used. There exist insecure crypto schemes for which the secret key can be recovered from plaintext/ciphertext pairs or from the public key itself (the public and private keys are mathematically related). To be complete, the answer would need to say why the key recovery given (public key, plaintext, ciphertext) is not possible *in the context of the question*, i.e., that the assumption that the schemes are secure removes this vulnerability.

Web Security: Quizzle

Q11: Vulnerability + Defence

Error: Mentioning attacks without describing the vulnerability.
Some answers mentioned the attack without stating the vulnerability in the code of save_answer, which the question explicitly asks for. Without a concrete link to the code in the question we cannot assess whether a student actually understands the question or is simply paraphrasing the slides.

Error: Mentioning possible vulnerabilities that are not in save_answer.
Many answers described vulnerabilities in different parts of the system like authentication and updating the time for the csrf_token although the question explicitly asked about *vulnerabilities in save_answer*. Even if correct, these answers have no points as they do not address the question.

Error: Using old csrf_tokens in an attack.
Some answers discussed problems with using previous csrf_tokens. In the question, it is explicitly mentioned that the token is "refreshed" i.e. updated in every access. Therefore, the use of previous tokens is not useful (they do not exist in the database) and these attacks are not valid answers.

Error: Talking about software-level attacks and vulnerabilities or cross-site scripting.
Some answers mentioned software-level attacks and vulnerabilities while others discussed cross-site scripting. The software-level attacks are out of scope since we don't give information about low-level interactions between save_answer and the OS to warrant a vulnerability. Cross-site scripting is about embedding a script in a page that someone else will run when opening it. It has nothing to do with the functionality of save_answer.

Q12: CSRF and Session Cookie

Error: Assuming insecure authentication when evaluating save_answer checks.
Many answers argue that the checks are not enough because a fake session_id created by the user could be used to pass the check. The question explicitly states that authentication is secure. Thus, users cannot create session_ids themselves. Other answers try a set of vague attacks at the authentication stage to argue why checks at save_answer are insufficient also disregarding the assumption that authentication is secure.

Error: Arguing that session cookie is not enough because it can be edited while the csrf_token cannot.
Some answers note that it is necessary to have both csrf_token and session_id because csrf_token cannot be modified by the user while session_id can. This is not true, any user with capabilities to edit cookies will easily be able to change the csrf_token value that is stored in their browser.

Error: Partial answers describing the purpose of only one token, that just answer the second part of the question about the suitability of save_answer.
Many answers describe the purpose of one token (session_cookie or csrf_token) and not the other; or, they answer the second part of the question directly speaking why one of the tokens is not sufficient as a check. The question asked for three items, and any answer that contains less than three resulted in point deduction.

Error: Confusing authorization with authentication.
Some answers confuse "authorized" with "authenticated". Authentication refers to verifying that the user is legitimate. Authorization refers to verifying that the user is allowed to use a specific resource. So, answers along the lines of "the checks are sufficient because only a logged in student can make requests" (authentication) is insufficient because it does not take into account that being authenticated does not mean that the "logged in student" can save a particular answer.

Error: Saying the purpose of csrf_token is to track time.
Many answers mention that the purpose of a csrf_token is to track exam time. While started_at is stored alongside csrf_token as a timestamp, the save_answer does not use started_at inside for any time-related task.

Q13: Attack

Error: Mentioning vague steps for the attack and/or stopping halfway.
Many answers simply state vaguely the steps that Mario must do to impersonate Luigi. Ex: "we get the session_id of luigi with a malicious script". We did not give full points to these

answers. Some answers describe only half of the attack by describing how to get session_id without explaining how this session_id can be used to update Luigi's answers.

Error: Attacks requiring Mario to do more than injecting a script at the start of the exam.
Many answers suggest attacks that do not respect the requirement in the question that Mario can only do the attack by preparing a script beforehand that they inject at the start of their exam. Examples of answers that violate this constraint are: any answer that make Mario do live interactions with the server during the exam, writing code in the "answer" field of questions, and answers that suggest communication between Mario and Luigi during the exam.

Error: Not following the code execution to the end and describing attacks that make Mario fail his own exam.
Many answers try to update the "answer" parameter without following the rest of the code and not noticing that line 29 fails. Since line 29 fails, Mario cannot complete their exam which is a requirement of the problem (first line of the problem: "Mario and Luigi need to take a web security exam on "Quizzle" to graduate."). Students that try to account for the problem at line 29 got partial credit.

Network Security: FELP

Q15: Tampering with the webpage Alice requests

Error: Stating that an attack, e.g., includes traffic redirecting or direct manipulation of the traffic, would permit serving a page different than that served on an HTTPS connection
Most of the answers included attacks which either contain redirecting to a fake server with DNS hijacking, or direct tampering of the HTML webpage. These attacks cannot be applied in this setting because the server uses **HTTPS** protocol (as indicated in the hyperlink) which means that the connection is protected using TLS.
The TLS protocol provides server authentication, thus one cannot serve content from any other server; as well as integrity, thus traffic cannot be tampered with in transit.
This error was also common in mini-exam 4, 2021.

Error: Stating that tampering is not possible due to the use of DNS-over-HTTPS.
The fact that a webpage is accessed via the HTTPS protocol (as indicated by the link "https://…"), only indicates that the connection *with the web server* is secured. It does not imply that the connection to the DNS server is also done over HTTPS. Thus, such an answer was not accepted as valid.
Even if the DNS requests were done via a secure channel, the explanation above about why an HTTPS-based connection cannot be tampered holds.

Q16: Block Alice's access

Error: Preventing Alice from submitting with a Denial of Service that would also deny service to Bob.
The question stated that Bob also wants to submit. Proposing an attack that would hurt Bob utility therefore does not comply with the scenario in the question. Since the attack would indeed prevent Alice from submitting, we applied a small point deduction to these answers.

Privacy: P2PM

Q17: ISP of the student dormitory

Error: Do not explain how a (global) adversary that observes both Franz and George's connection can learn that George and Franz are communicating. Some answers stated that the dormitory's ISP constitutes a global adversary that can observe traffic at both the entry and exit node but did not explain how this enabled the ISP to learn that George and Franz are communicating. These answers received only partial points. To justify that the question statement was wrong, it is not enough to state a correct threat model and that the system does not protect against this threat. A full security argument also needs to include why or how the threat applies.

Q18: Russian ISP

Error: Miss that Tor is an overlay network and that thus a Russian ISP can use BGP hijacking to reroute traffic. Most answers missed that, *as explained in detail in the solutions to the* Theory Exercise 12.1, the onion routers of the Tor network run at the application not at the network layer. Therefore, a Russian ISP could use BGP hijacking to reroute traffic in such a way that they can see both traffic coming in and leaving the anonymous communication network, i.e., a Russian ISP can become a global adversary.

Software Security: Token Server (Q19 and Q20)

Q19: Fuzzer reduced coverage

Error: Not giving an example that the question explicitly asks for.
In part 1, the problem description explicitly asks for an example from a certain part of the program. However, some answers only say that the fuzzer is not aware of the input structure but fail to provide a concrete example. This kind of answer only got partial points, since without a concrete link to the code in the question we cannot assess whether a student actually understands the question or is simply paraphrasing the slides.

Error: Mixing sanitization and fuzzing.
For both parts, some answers mentioned sanitization. Sanitization does not help in these questions.
For part 1, the problem is asking why the fuzzer has a limited coverage, a sanitizer is not for helping the fuzzer reach larger coverage. This can be a mixed concept between sanitization and guided fuzzing.
For part 2, it is very dangerous to assume that a sanitizer can help catch any logic error in the code. Sanitization can help enforce some policies on memory, undefined behaviour etc. However, logic errors can happen without any of these.

Q20: Leaking the Key

Error: Answering an improvement instead of why fuzzer does not help.
Some answers explained how the vulnerability in the code could be found. But this was not what the questions asked. The question asked why the fuzzer does not help. Correct answers that do not address this question have zero points. Answers that are wrong lead to point reduction.

Error: Saying that a fuzzer in part 2 does not help because it has limited coverage.
The vulnerability exploited in part 2 cannot be prevented by any fuzzer, regardless of the coverage. A fuzzer mutates inputs, and looks for crashes, but it does not check the semantics of the output. Thus, it cannot detect that the output contains the key.

Common Issue: Endianness does not make a difference on what the key is from looking at the token.
Some answers argue that what the key exactly is depends on the endianness of the device. We did not remove points as long as the student shows that they understand the program in other parts of their answer. However, we would like to point out that this argument is actually a misunderstanding of endianness and its implication on data representation. In part 2, the key and token are stored as char in bytes, they are always copied in bytes to and from a buffer. The endianness will have an impact on the order of stored bits of a byte, however, the device will always interpret the byte in the stored way.

StartStuff

Q21: Propose a covert channel

Error: Proposing a channel that is not covert. Adding new scripts/files that contain secrets in plain text is not a covert communication channel. Modifying access_script.sh to print the content of secret files on error is also not covert. Such methods are easily detectable by all employees rather than being hidden as a covert channel should be.

Error: Writing down to leak secrets. Proposing that TOP SECRET clearance employees write to PUBLIC level files is against the BLP policy by definition.

Error: Trick a TOP SECRET clearance employee to leak secrets. The question asks to provide a mechanism where a TOP SECRET clearance employee is willingly colluding with a PUBLIC employee to leak secrets via a covert channel. Tricking the TS employee is not needed here.

Q22: Attack by PUBLIC clearance subject

Error: Attacks involving TOP SECRET clearance employees. Stealing the password of the TOP SECRET employee by looking over their shoulder or waiting for them to leave their computer unlocked and unattended are not examples of an attack where a PUBLIC employee gains the secret documents without the involvement of a TOP SECRET clearance employee.

Error: Using rootkits/backdoors to elevate the privileges of the PUBLIC employee. Answers stating this did not specify the assumptions on the adversary's capabilities and the BLP model implementation or how the attack works specifically.

Error: Tricking a TOP SECRET employee into upgrading the clearance level of the PUBLIC employee. This is contradictory to MAC where there is a central security policy i.e. TOP SECRET employees should not have the authority to upgrade clearance levels of employees.