

Teacher: Prof. Dr. ETH Mathias Payer
 COM-402 Information Security and Privacy – Quiz 04
 16th December 2024
 Duration: 15 minutes

Anon Ymous

SCIPER: **999999**

Do not turn the page before the start of the quiz. This document is double-sided, has 4 pages, the last ones possibly blank. **Do not unstaple.**

- **No other paper materials** are allowed to be used during the quiz.
- Using a **calculator** or any electronic device is not permitted during the quiz.
- For each question, mark the box(es) corresponding to the correct answer(s). Each question has **one or more** correct answers.
- For each question, we give:
 - 3 points by default,
 - 0 points if you give no answer,
 - 1 point per incorrectly checked or missed answer.
 Each question has a minimum of 0 points, we do not award negative points.
- Use a **black or dark blue ballpen** and clearly erase with **correction fluid** if necessary.
- If a question is wrong, we may decide to nullify it.

Respectez les consignes suivantes Observe this guidelines Beachten Sie bitte die unten stehenden Richtlinien		
choisir une réponse select an answer Antwort auswählen	ne PAS choisir une réponse NOT select an answer NICHT Antwort auswählen	Corriger une réponse Correct an answer Antwort korrigieren
     		
ce qu'il ne faut PAS faire what should NOT be done was man NICHT tun sollte		
     		

Question 1

Which of the following statements about side channel attacks is/are correct?

- Side channels are only relevant for square and multiply RSA implementations.
- A timing side channel arises when observed computation time depends on the secret.
- Masking can be used to counteract side channels.
- Spectre is an exploit that affects non secure hardware and is thus not a side channel vulnerability.

Explanation: Side channels can potentially affect any computation. Hiding adds noise, masking splits state into shares and force recomputation. Spectre is a side channel vulnerability, side channel vulnerabilities are also relevant for commodity hardware.

Question 2

Which of the following statements about TPM (Trusted Platform Module) is/are true?

- The TPM cannot do cryptographic operations, but it can be used to store cryptographic keys.
- Secure boot cannot rely on the TPM because the default configuration for any TPM is only generated after the boot process is completed.
- TPM is not part of the main instruction set and therefore cannot be trusted.
- The TPM's platform configuration registers (which are used for attestation) are append-only.

Explanation:

- a TPM chip can do cryptographic operations
- Secure boot can rely on a TPM chip measurements
- TMP is separate from the instruction set but part of the TCB
- the platform configuration registers from TPM which are used for attestation are append-only (see slides)

Question 3

If a mechanism satisfies ϵ -differential privacy, which of the following factors could help an attacker infer private information about individuals in the dataset?

- Access to auxiliary information (e.g., some background knowledge).
- The value of ϵ is very large.
- The value of ϵ is very small.
- Presence of outliers in the dataset.

Explanation: Attackers can combine outputs with auxiliary information to deduce private information. A large ϵ value means less noise is added, making the mechanism less private and more vulnerable to attackers. Outliers can often be uniquely identified and are particularly vulnerable to inference, even when noise is added.

Question 4

Suppose you have the following anonymized database showing the final ISP grades for each student:

Faculty	Cohort	Grade Range
IC	202*	[3-4)
IC	202*	[4-5)
SV	202*	[5-6]
SV	202*	[4-5)
ENAC	202*	[3-4)
ENAC	202*	[5-6]
ENAC	202*	[5-6]

Given the grade (range) is defined as sensitive data, which of the following statement(s) is/are applicable to this database?

- It is K=2-anonymous.
- It is K=3-anonymous.
- It is L=2-diverse.
- It is L=3-diverse.
- The database is vulnerable to a homogeneity attack.
- Irrespective of the initial data, cohort anonymization added no privacy property here.

Explanation: Note that here, our quasi-identifiers are the Faculty and Cohort.

1. K-anonymity is defined by the fact that there exists a set of size at least K for each quasi-identifier group/user. In our case, for the database to be K-anonymous, each {Faculty,Cohort} group should be composed of at least K elements. This is the case for K=2, but not for K=3 (only the {ENAC,202*} group satisfies this)
2. An equivalence class has L-diversity if there are at least L well-represented values for the sensitive attribute. In our case, this means that for each quasi-identifier group, we should have at least L distinct values in the grade range. Luckily, this is the case for each group for L=2. It is impossible for L=3 as some groups don't even have 3 elements.
3. A database is vulnerable to a homogeneity attack when it is not (at least) L=2-diverse, as this means that there exists a group with a unique sensitive attribute - one can then easily infer the attribute directly from the QID.
4. It is easy to find a counterexample showing that cohort anonymization was indeed needed. Suppose that the 2 IC faculty students were from different cohorts (e.g.: one comes from 2022 the other 2023), and this data was left as is in the database (no 202* anonymization). Then, the {IC,202*} group would be split in two groups of one element each, namely {IC,2022} and {IC,2023}, losing K=2-anonymity and L=2-diversity in the process

Question 5

You're asked to review a codebase of a legacy banking application. You have discovered a severe flaw: the RSA implementation includes a square-multiply function which is vulnerable to timing attacks over the network. What precaution(s) might you take to prevent adversaries from extracting the private key through a timing side channel? Your goal is to protect the private key of the banking app. (Your threat model assumes that all other processes and actors are untrusted.)

- I can seal the private key using an SGX enclave.
- I can change the square-multiply with a constant-time alternative.
- I can execute the existing application inside a TPM application.
- I can execute the existing application inside a Hardware Security Module (HSM) which will enforce physical isolation.

Explanation: The only possible solution for this problem is using a constant time implementation. Sealing the private key will not help to prevent timing attacks, since you will need to do an exponentiation operation with a square-multiply, which will have a timing channel. You can not execute custom applications inside of TPM or HSM since they are designed to do just cryptographic operations.

Question 6

You're running compute workloads on sensitive data in the cloud. To protect the data from being read by adversaries, you decide to use Intel SGX as a TEE, and do all the sensitive computations only in the TEE. Which of the following adversaries does this design protect against?

- An outside attacker leveraging a 0-day exploit in the SSH daemon to gain access to your server.
- An outside attacker leveraging a 0-day exploit in the Linux kernel's IPv6 stack to gain code execution in the kernel context on your server.
- An outside attacker leveraging a 0-day exploit in the cloud provider's hypervisor to gain code execution at the hypervisor level.
- A malicious cloud provider sysadmin who's scraping data from customer VMs through his privileged access to the machines.
- A malicious Intel employee implementing a hardware backdoor in the line of CPUs used by your cloud provider.

Explanation: The point of a TEE is to only trust the hardware, not any of the unprivileged (or even privileged) code running on top of it. Therefore, only a hardware flaw can be leveraged to actually exfiltrate data from the TEE.

Question 7

Which of the following is/are true about ML security?

- To steal a 5-dimensional linear model, an adversary needs to query the model at least 6 times.
- To steal a 5-dimensional linear model, an adversary needs to query the model at least 7 times.
- Membership inference can employ machine learning models to determine whether a person is part of a certain social group.
- Adversarial machine learning aims to perturbate inputs such that the queried ML model is fooled but a human in place of the ML model would still classify the input correctly (e.g., an image of a cat being recognized as a guacamole by the ML model).

Explanation: Stealing a linear model is equivalent to solve a set of linear equations, for which you only need $d + 1$ equations to solve a d -dimensional equation system. Membership inference is generally not about social groups but whether a data point is part of the training set of an ML model. However, this can be used in this specific example. Adversarial ML aims to add imperceptible perturbations.

Question 8

Which of the following statements about Privacy is/are true?

- The goal of “privacy as confidentiality” is minimizing data disclosure.
- The goal of “privacy as practice” is minimizing data disclosure.
- Denying privacy to some is denying privacy to all.
- Cryptography (such as TLS) is a PFT to limit/counter, for example, global adversaries.

Explanation:

- Minimize data disclosure is the goal for privacy as Confidential