System Security 2023
Exam (with solutions guidelines) - June 30

Solutions guidelines in red. Note that for many questions, other answers would also be valid, these are just short examples of solutions.

Question 1 (10 Points)

b. [4 points] What are the advantages and disadvantages of static analysis vs. dynamic analysis in the context of firmware? Provide a comparison, mentioning at least two advantages and disadvantages of each approach.

Advantages static analysis: no code emulation, fast, can in theory scan all execution paths.
Disadvantages static analysis: custom formats, manual loading, unknown architectures.
Advantages dynamic analysis: handle code obfuscation and unpacking, observable results.
Disadvantages dynamic analysis: hard to emulate, limited code coverage, rely on external peripherals.

c. [3 points] What are the differences between Type I and Type III firmware? When is it more favorable to use a Type III firmware image?

Type I = Linux-based; Type III = no OS abstraction (favorable for low power devices).

d. [3 points] What security features/technologies we usually do not find in the firmware for IoT devices and why? Name at least 3.

CFI, Seccomp, ASLR. Why? Performance overhead and scalability reasons.

Question 2 (10 Points)

A new company released a service named ChatIoT. The service allows users to manage all connected IoT devices in a smart home via a chat-based mobile app. ChatIoT receives text-based commands (e.g, “turn off the A/C in the living room when the dog is outside”) or queries (e.g., “what is my dog doing”), processes them, translates them into low-level commands for IoT devices, collects the results, and provides text-based answers to the users (e.g., “your dog is currently playing in the garden, the living room A/C is off”). Consider the following scenario related to the smart home of a ChatIoT customer.

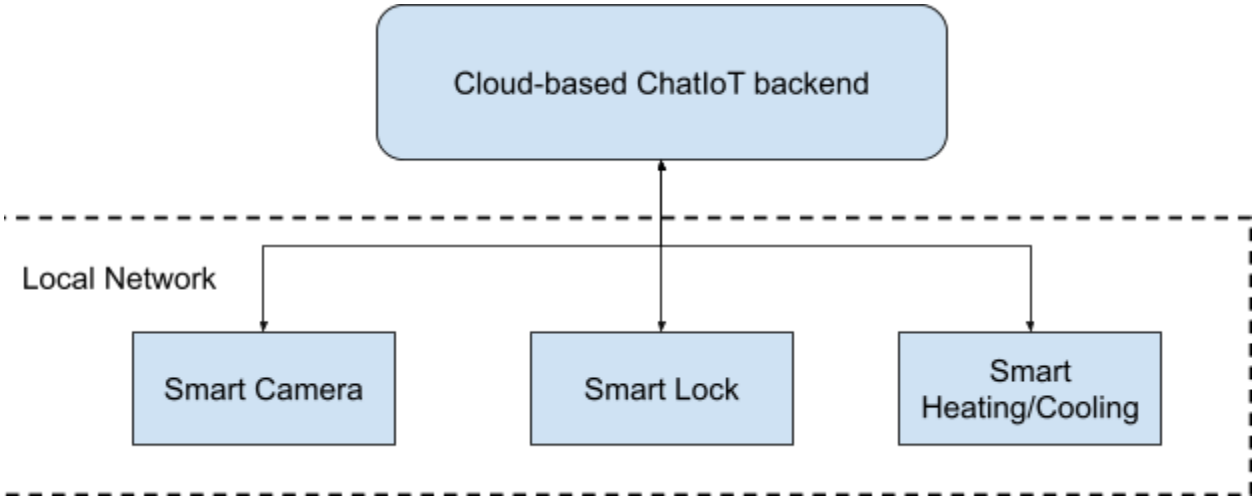

a. [2 points]. Describe two potential threats to this system discussing their risks, how attackers could execute the attacks, their motivation, and what attackers could gain.

#1 Malware infecting the mobile device and compromising the companion app. Motivation: Disable smart lock and enter the property
#2 Passive MITM attack on the network side to fingerprint users' activity and reveal info about whether they are home.

b. [2 points]. The update mechanism of the smart camera works as follows: The camera downloads the updated image via HTTP and a certificate via HTTPS. The camera then verifies the integrity of the new image using the downloaded certificate. Instead, the update mechanism of the smart cooling system works as follows: The device downloads both the updated image and the certificate via HTTPS. Then, it verifies the integrity of the new image using the downloaded certificate. Imagine that an attacker can intercept the network traffic of the whole smart home during firmware updates. What can the attacker do to compromise the smart camera and/or cooling system? Which of the two mechanisms is more secure?

They are equally secure, the attacker cannot tamper with the updates. For the smart camera, they can however obtain the new image and look for potential vulns.

c. [6 points]. Imagine that you extracted the firmware running on the smart lock and loaded it in Ghidra. You see the following output.

```
s_00101038 = "Sending output"
s_00102000 = "Executing command"
s_0010202a = "unauthorized command"

char * FUN_0010117a(undefined8 param_1) {
  FUN_00101149(s_00102000);
  return library_func_2(param_1);
```

```
}

void FUN_001011a4(char *param1) {
  FUN_00101149(s_00101038);
  library_func_1(param1);
}

int FUN_001011c3(char *param_1, char *param_2) {
  size_t sVar1;
  undefined8 uVar2;
  int local_1c;

  sVar1 = strlen(param_1);
  if (sVar1 == 8) {
    uVar2 = FUN_0010117a(param_2);
    for (local_1c = 0; local_1c < 8; local_1c = local_1c + 1) {
      if (param_1[local_1c] != "seeecret"[local_1c]) {
        FUN_001011a4(s_0010202a);
        return 0;
      }
    }
    FUN_001011a4(uVar2);
    uVar2 = 1;
  }
  else {
    FUN_001011a4(s_0010202a);
    uVar2 = 0;
  }
  return uVar2;
}
```

The firmware running on the smart lock presents at least two security flaws. Describe them and explain what you think `FUN_001011c3` does.

FUN_001011c3 authenticates and executes commands, returning the command output.
Command is executed before authentication, although if authentication fails, the output is not sent back
Short and harcoded password, which suggests it is reused for all devices

Question 3 (9 Points)

a. [3 pts] What is the difference between covert channels and side channels?

Side channels: unintentional leakage, one party leaks as a side effect of what it does
Covert channels: intentional leakage where two parties (sender and receiver) agree on a hidden protocol.

b. [3 pts] Consider a newly designed digital modulation scheme, named 10+10 QAM, which works as follows. 10+10 QAM allows for encoding 10 symbols in the complex plane. One of the symbols is however special, and, when sent, it tells the receiver that the subsequent transmission will enable 20 symbols (only for one transmission). In comparison with 16 QAM, which scheme facilitates the implementation of a radio-based covert channel and why?

Assuming the special symbol is not used significantly more frequently than others, 10+10 QAM facilitates the implementation of a radio-based cover channel because it has more room in between symbols (less symbols in the complex plane)

c [3 pts]. How can you detect such potential radio-based covert channels?

Anomaly detection. Provide some examples of features to look at, e.g., average noise level, average distribution of the symbols.

Question 4 (10 Points)

a. [2 pts] Describe what happens when an application installs a seccomp filter, i.e., the filter installation process.

Prctl syscall to move compiled filters in kernel space.

b. [3 points] To overcome the limitations of seccomp, a group of elite cyber-experts proposes sseccomp (super-seccomp). To significantly reduce the performance overhead and improve compatibility, sseccomp directly executes within the standard libraries, which check whether invoked library functions would produce allowed syscalls. A high-level example is shown in the figure below. Compare super-secomp with seccomp from a security perspective.

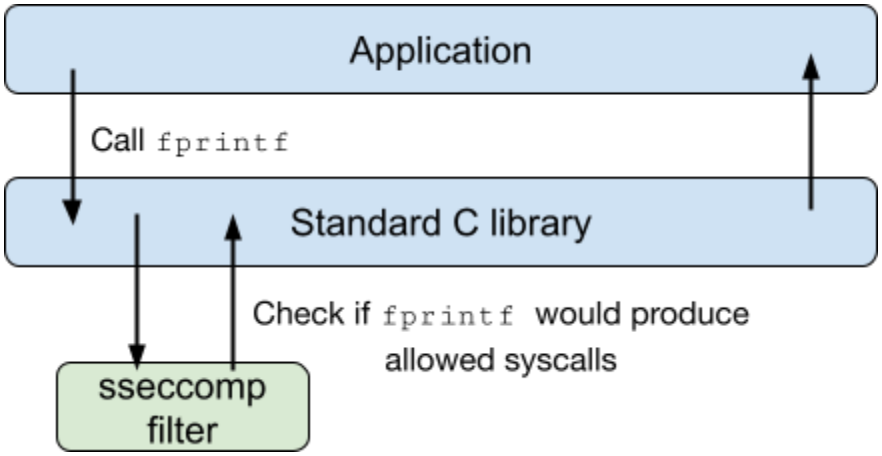

Weaker, checks enforced in userspace

c. [2 points]. The cyber-experts also claim that sseccomp's design can *securely* dereference pointers passed as arguments to syscalls. Is this true? Why?

No, this scheme still opens room for a race condition

d. [3 points]. Ssecomp 2.0 has another mode: it automatically deploys libraries in a separate, isolated Docker container. When an application invokes a library function, the function call is forwarded to the isolated container (see figure below). How does this improve the previous scenario? Discuss the security of this mechanism.

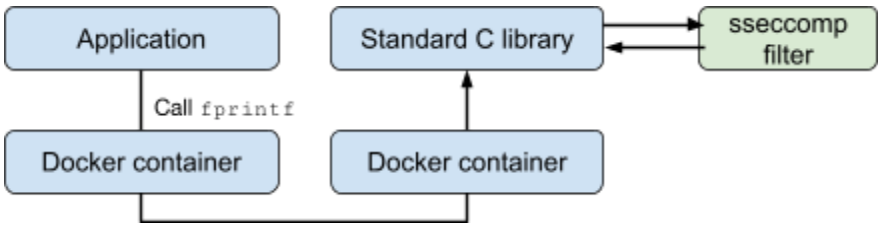

Still weaker, apps can directly invoke syscalls without going through libraries.

Question 5 (11 Points)

a. [3 pts] What are the differences between Docker containers and Xen Virtual Machines, also from a security perspective?

Docker, userspace, shared kernel. A kernel vulnerability Compromises all containers
Xen, hypervisor, separate kernels, stronger/more isolated

b. [2 pts] What are Linux control groups (cgroups) and how are they useful?

Set limits on resource usage (e.g., memory, CPU time). Mitigate certain types of attacks, e.g., DoS, or limit attacker capabilities.

c. [3 pts] What is a Trusted Execution Environment and why is it useful?

Separate hardware-level environment with separate secure-world OS.
Useful when the normal OS is fully compromised (kernel/root level), e.g., avoid leaking sensitive info.

d. [3 pts] What is meant by code integrity? Can you mention and describe at least one technique to guarantee code integrity?

Making sure the only right code is executed. Secure/Trusted Boot

Question 6 - PUF (10 Points)

Consider a setup where you are asked to use a Physically Unclonable Function, and you are required to explain to the project manager what its advantages are.

a) [2 points]. Explain an application of a weak and strong PUF.

Weak PUF can generate a small number of responses so it can be used for authentication. Strong PUF gives many responses so it could be used for random number generation.

b) [2 points] How do we use the linear additive model in PUFs, and why is that model good?

We use it by adding separate stages of PUF. The model is good as it approximates reality well.

c) [3 points] Consider a setup where we use a custom form of an interpose PUF. More precisely, we have three chains instead of two chains, as we saw during the lecture/lab. The output bit of the first chain is fed into the middle position of the middle chain, and the output bit of the middle chain is fed into the middle position of the bottom chain. What would be the theoretical max accuracy one can achieve and why (neglect effects of environment and modeling)? Hint: consider what happens if a bit is correct or wrong.

Theoretical max accuracy (so, disregarding influence of any specific classifier used to attack or number of CRPs), the result is 75%. The reasoning is simple: let us assume the upper bit is wrong. That one will influence 50% of lower ones. And from the 50% rest, on average, 50% will be wrong. So, 50%+25%=75%.

d) [3 points] What are the differences between the challenge-response pairs attack and the reliability attack?

Challenge-response pairs attack does not account for noise and only tries to find the delays in PUFs that would coincide with the real-world example. Reliability attack considers multiple (same) challenges and whether they end up with different responses.

Question 7 - SCA (12 Points)

In this question, we consider side-channel attacks.

a) [2 points] Explain two scenarios when CPA will not work.

CPA will not work if we do not have enough measurements to conduct statistical analysis. Also, it will not work if the target is protected with certain masking schemes.

b) [3 points] Consider the side-channel attack on AES S-box in the Hamming weight leakage model. What would happen if the S-box would be linear?

If S-box is linear, two inputs that have the same Hamming weight would also result in same Hamming weight at output. Which means we could not differentiate them and attack is more difficult.

c) [2 points] What is the difference between CPA and online CPA? Would those techniques give different results, and why?

CPA runs with all measurements from start while for online CPA, we can add measurements. They will give the same results.

d) [2 points] How many different output values do we have if we use the Hamming distance leakage model and we attack the PRESENT S-box that has four input and four output bits?

5.

e) [3 points] Consider the 1-bit DPA attack on the AES-256. What is the attack complexity, and why?

2^{13} as you still need to do key enumeration for all possible keys and 32 S-boxes.

Question 8 – SCA/FI Countermeasures (8 Points)

This question considers the countermeasures for SCA and FI.

a) [2 points] Explain a countermeasure that works against SCA and glitches.

Threshold implementation.

b) [3 points] Let us consider an SCA countermeasure that first involves masking the sensitive variable with Boolean masking and then duplicating the masking. What would be the main advantages and disadvantages of such a combined countermeasure compared with independent countermeasures (only duplication or only masking)?

Advantage is combined countermeasure, which can provide stronger resilience against SCA while disadvantage is increased area to implement.

c) [3 points] What are the necessary conditions for the threshold implementation? What are the main advantages and disadvantages of threshold implementation compared to Boolean masking?

Uniformity, correctness, non-completeness. Advantage is provable security and disadvantage is high implementation cost.

Question 9 - FI (10 Points)

This question considers fault injection attacks.

a) [2 points] What would be the main advantage and disadvantage of laser fault injection over voltage glitching?

Laser fault injection is more precise but has more parameters to consider.

b) [3 points] A countermeasure developed against a given attack, if not carefully examined, may benefit another physical attack. Let us consider a type of attack called computational safe-error attack (C safe-error attack), which can be done against the classical, side-channel protected exponentiation algorithms given below. The C safe-error attack is developed by inducing any temporary random computational fault(s) inside the ALU. Consider where and how you would induce an error into the code below to reveal secret information (note that both algorithms can be attacked in the same way).

Observe that since the algorithm runs in constant time, an attacker can more easily locate the exact moment of the second multiplication “Rb ← Rb R2” for each iteration. When the current exponent bit (e.g., kj) is equal to 0, then this multiplication is a dummy operation and so has no influence on the final result. Therefore, if an attacker induces any kind of computational fault into the ALU during the operation of Rb ← Rb R2 at jth iteration, then according to whether the final result is incorrect or not, the attacker may deduce if kj = 1 or kj = 0. Note that this attack only reveal one bit of exponent k.

c) [2 points] Explain how a fault injection attack can be made and prevented on a system where sensitive data has a limited set of values (like phase and state variables). Hint: consider the Hamming distance between the values.

The attack could be made with laser fault injection that has bit granularity. If the data has limited number of values, then random change of a bit could still result in allowed (but not correct value). We can prevent it by ensuring maximal distance among allowed values.

d) [3 points] Consider the fault analysis where the goal is to find such a fault that does not change the intermediate result and, as such, leads to a correct ciphertext. How can this attack be used to break the security of a cipher?

If the attack does not make any change in the output, then we know that the attacked computation does not contribute to the solution, which immediately reduces the security of the implementation.

Question 10 - HWT, CFI, Microarchitectural attacks (10 Points)

You are in the role of security evaluator and need to explain how to protect against various attacks.

a) [2 points] On what does CFI's performance and code size overhead depend?

On the number of checks that need to be implemented.

b) [2 points] What are the two parts each hardware trojan needs to have and provide an example of those parts in a security scenario of your choice.

Trigger and payload. Trigger could be a combination of events (e.g., a specific path of gates). Payload could be a part that stops working or that always responds in the same way.

c) [3 points] Describe the pre-requirement, initial phase, waiting phase, and measurement phase for the Flush+Reload attack.

Pre-requirement: shared memory between attacker and victim. Availability of clflush.

Initial phase: attacker flushes shared critical data from cache.

Waiting phase: victim executes his operations.

Measurement phase: attacker reloads critical data. Lower access=data has been accessed by victim.

d) [3 points] Consider the pseudocode for the left-to-right binary algorithm given above. How would you mount a Prime+Probe attack on it?

For instance, the attacker would prime the line with R0. Then, he probes the line and checks the time required to load the line for each cache set. Depending on time, the attacker decides if the cache set was accessed or not.