

Information Assurance: Final Exam Answer Key

117 points total

Multiple Choice (2 points each, 20 points total)

1. The Common Criteria defines what as an implementation-independent set of security requirements for a category of products or systems that meet specific consumer needs?
 - a. Target of Evaluation
 - b. Protection Profile**
 - c. Security Target
 - d. Evaluation Assurance Level

2. Which term matches the following definition? The art and science of hiding secret messages within some other form that is usually visible or out in the open for anyone to see if they knew where to look.
 - a. Encryption
 - b. Digital watermarking
 - c. Signature
 - d. Steganography**

3. Disk scrubbing is necessary because of:
 - a. Magnetic remanance**
 - b. Sticky bits
 - c. File system optimization
 - d. Non-determinism

4. One common attack against web servers using SQL involves passing in an unexpected URL that is passed on to the SQL server. The malformed URL argument can result in an unexpected database operation. This error is most closely described by which of the following RIOS program errors?
 - a. Inadequate identification/authentication/authorization
 - b. Implicit sharing of privileged/confidential data
 - c. Incomplete parameter validation**
 - d. Violable prohibiting/limit

5. Which of the following systems are **not** subject to Communication Assistance for Law Enforcement Act (CALEA) right now and not planned to be subject to CALEA in the near future.

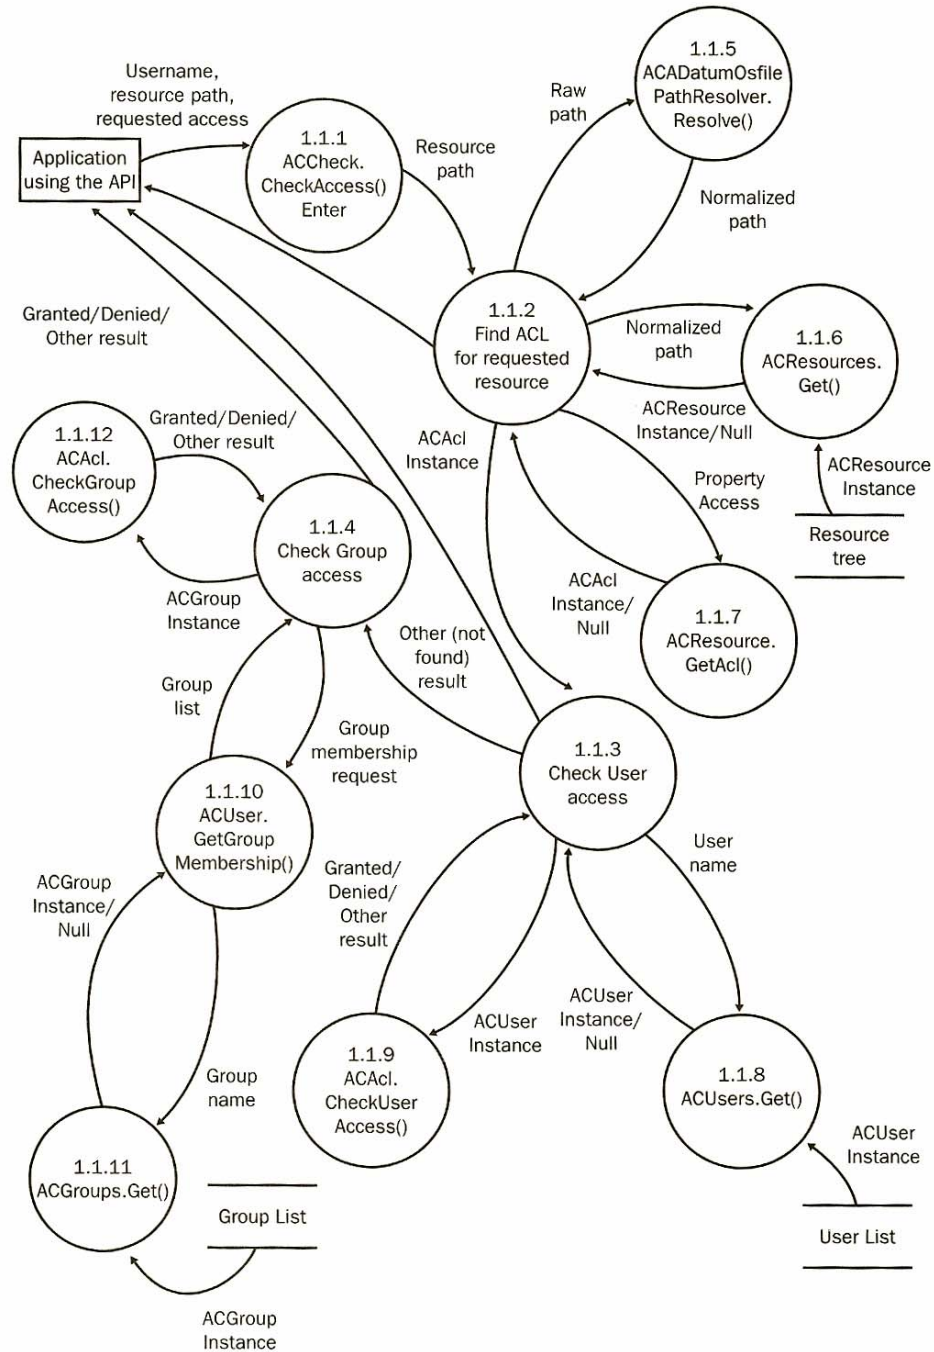
- a. Vonage IP phone system
 - b. Cell phones
 - c. Skype IP phone system**
 - d. Physical System Telephone Network
6. Which of the following definitions best address tamper-proof hardware as it is applied to secure co-processors and smart cards?
- a. **Destroys data when it detects physical tampering.**
 - b. Installed so that it is not physically accessible.
 - c. Logs when it is attacked.
 - d. Makes it difficult to analyze the contents of the processor.
7. Which of the following is **not** a core security service?
- a. Confidentiality
 - b. Availability
 - c. Performance**
 - d. Integrity
8. How does the Seal operation defined by the Trusted Computing Group differ from a regular RSA encrypt operation?
- a. It uses an extra long key.
 - b. It uses a key associated with the Trusted Platform Module (TCPM) to perform the encryption, and it includes information about important aspects of the system configuration in the data to encrypt.**
 - c. It uses the TPM to perform the encryption, and it also signs the data before it encrypts it.
 - d. It encrypts a hash of the data.
9. Which encryption mode does **not** make sense for stream encryption?
- a. Counter Mode
 - b. Electronic Codebook Mode**
 - c. Output Feedback Mode
 - d. Cipher Feedback
10. Which of the following is **not** one of the classes of Authentication information.
- a. What an entity has.
 - b. What an entity does.**
 - c. What an entity is.
 - d. What an entity knows.

Short Answer

11. (10 points total) Match each of the eight design principles used in class with its definition. Two fake principles and definitions have been added. Do not match the fake values.

- a. Principle of Retrofitting Trust
 - b. Principle of Economy of Mechanism
 - c. Principle of Complete Mediation
 - d. Principle of Open Design
 - e. Principle of Evolving Assumptions
 - f. Principle of Separation of Privilege
 - g. Principle of Least Privilege
 - h. Principle of Fail-Safe Defaults
 - i. Principle of Least Common Mechanism
 - j. Principle of Psychological Acceptability
-
- I. f A system should not grant permission based on a single condition.
 - II. Security mechanisms should be added after the initial functionality is designed and implemented.
 - III. c All accesses to objects should be checked to ensure that they are allowed.
 - IV. g A subject should be given only those privileges that it needs in order to complete its task.
 - V. h Unless a subject is given explicit access to an object, it should be denied access to that object.
 - VI. i Mechanisms used to access resources should not be shared.
 - VII. j Security mechanisms should not make the resource more difficult to access than if the security mechanisms were not present.
 - VIII. b Security mechanisms should be as simple as possible.
 - IX. Security mechanisms should automatically adapt to changing security assumptions.
 - X. d Security of a mechanism should not depend on the security of its design or implementation.

12. (16 points, 4 points each) This question involves using Threat Modeling techniques to help design the security of an access control library called “A. Datum Access Control API”. A data flow diagram showing some of the interactions of key library components is shown below



Continued on next page...

Here are the two potential threat profiles for the API.

ID = 10

Name=Replace list that maps users to constant ID's, with a new list file. By replacing the user list file, you could map an unprivileged user name to a privileged user ID.

STRIDE classification=Tampering, Spoofing

Mitigated=?

Entry points=User list file

Assets=Ability to change the evaluation of the access control policy.

ID=11

Name=Use non-standard path to trick API into granting access to user who should not have access. For example, using ".." to confuse the path parsing.

STRIDE classification=Information disclosure.

Mitigated=?

Entry points=1.1.1 ACCheck.CheckAccess()

Assets=Data protected by the access control policy.

- a. For each threat profile, look at the data flow model and determine if the threat is mitigated in the current design. If so, how? If not, how could you mitigate the threat?

Threat 10 does not appear to be mitigated. There are several options for mitigation that people suggested. This threat could be mitigated by using operating system controls to restrict access to the files. Alternatively, system containing the user file could have limited physical access. And the file system could not be networked. Therefore, only cleared administrators could physically access the machine to update the user list. This might be too extreme and make updating the user and group lists too inconvenient. In addition to any OS controls, the API system may want to store a hash of the expected user list file to detect changes. The user information in the file could be encrypted so even if the file was replaced, the attacker would not know what to put in (unless he knew the key).

Threat 11 may be mitigated by the path resolver in subcomponent 1.1.5. Judging from the name, it sounds like this component puts all the paths in a common form before using it as the name of the resource in the access check. This would solve many problems of using odd path names to trick the access control.

- b. For a mitigation you identified for each threat in the previous part, what design principle does this mitigation most closely follow?

By relying on operating system access controls to protect the user list file, we are following the principle of complete mediation.

By normalizing the resource names we are following the principle of economy of mechanism. We break the problem into two parts: normalization and access check. By breaking into two simpler problems, we can use simpler enforcement mechanisms.

- c. Based on these threats, what information should be logged? Which module should be logging the data? Why is this data necessary?

To audit for threat 10, we need to log sufficient information to see if the requested path gets the access requested. I would log this in a couple messages. First from 1.1.2 I would log the requested path, the normalized path and the resource ID. Then from the check user access (1.1.3) and check group access (1.1.4) components I would log the resource ID, the group or user ID, and the access the decision.

Between these two log messages, the auditor will be able to determine if the requested path got the expected access.

To audit for threat 11, I would log a message from check user access (1.1.3) that includes the user name and the user ID. I would also log the second access message that I did for threat 10. By looking at the first message, the auditor can see if the ID mapping is expected. By looking at both messages, the auditor can determine if the user access was expected.

- d. Based on these threats, what test cases should be considered for inclusion in the test suite?

We need to test whether changes to the user list file can be made, if the changes are detected, and how the changes affect the evaluation of access.

We need to test the effectiveness of the path resolving by testing on many variants of non-standard paths.

13. (4 points) You are given responsibility for designing the printing support for the new office. Your boss is particularly concerned with classified data being printed and accidentally seen by people without the appropriate clearance. Identify two controls that can help your design. At least one control should be physical and at least one control should be implemented in software.

physical control: keep high-security printers in locked room

software control: use labeled printers (Bell-LaPadula) so that printing high-security documents to low-security printers is disallowed. Or, require the user to log in at the printer so that he can monitor the printout.

14. (4 points) Identify two additional assurance requirements that would be added from a lower EAL (e.g., EAL2) to a high assurance evaluation level (e.g., EAL7).

The main mistake made here was enumerating functional requirements rather than assurance requirements.

Some additional assurance requirements include

- *More stringent testing*
- *Formal modeling*
- *Formal verification*
- *Covert channel analysis*

15. (4 points) You are in charge of a very sensitive computer lab, and you are concerned with the radio-emanations virus discussed in the Soft Tempest. If your systems get infected with a radio transmitting virus, what are two things you can do to protect against information being transmitted (short of finding and eliminating the virus).

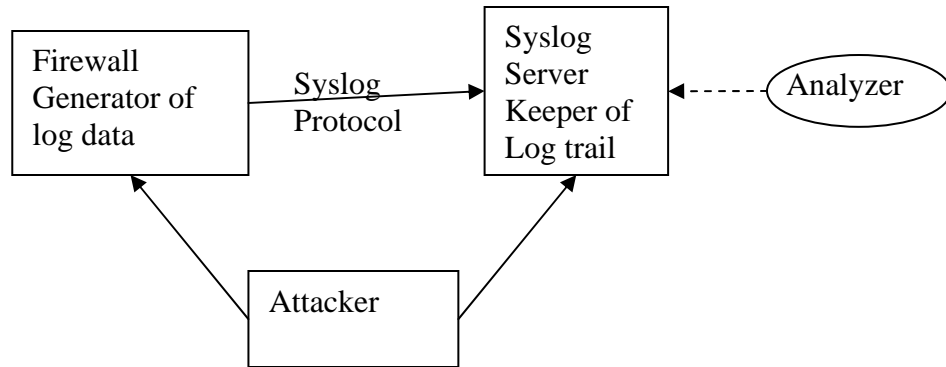
Use RF shielding on monitor or room.

Maintain physical separation of high-security computers.

Broadcast interference noise at frequencies near monitors.

Use Tempest fonts.

16. (8 points) Syslog is a UDP-based protocol that is used by many network devices to emit log information to a physically separate syslog server. The UDP protocol does not ensure reliability. In addition the Syslog protocol does not implement any source authentication. Consider an entity launching an attack against a firewall that uses syslog for its logging, as in the topology shown below



- a. (2 points) Describe how the attacker can eliminate information from the log to hide its attack from the analyzer.

Intercept the UDP packets. Loss will not be detected because UDP is not reliable transport. Downside is that attacker may not be in the interception path.

Attacker sends many of its own packets to the syslog server causing a denial of service attack which will drop legitimate FW log packets.

- b. (2 points) Describe how the attacker could add information to the log to hide or obscure its attack from the analyzer.

The Attacker can send his own messages to the syslog server with the source address set to the firewall. No acks, so the spoof is trivial to do. This would cause the legitimate log of his attack from the firewall to be lost in the noise.

He could intercept fw log packets, change values, and then relay. Again, this assumes that the attacker is in a position in the network to intercept may not be reasonable. You can do things like arp cache poisoning to make this more likely.

- c. (4 points) Describe two mitigations to the syslog protocol that would prevent the attacking program from polluting the audit trail.

One common mitigation suggested was leveraging IPSec or SSL to set up a tunnel between the FW and the syslog server. This gives us authentication and data integrity. Another mitigation would be to use a TCP-based syslog server. Both of these solutions have performance implications. So causing the fw to generate more logs will enable the

launching of a DOS on the firewall. Others suggested directly negotiating a shared key between the firewall and the syslog server and encrypting the log stream.

17. (4 points) Describe two ways to protect a running system from an attacker exploiting a buffer overflow vulnerability.

Use software analysis such as libsafe.

Use hardware support like DEP to prevent execution of code on the stack or heap.

Program with software packages like Java or STL that have boundary-safe constructs

Use a wrapper that does input checking.

Use address space randomization

18. (6 points) The ring memory detection scheme implemented by the Pentium architecture implements access controls based on the privilege levels of the source and destination memory segments. This is very similar to the security levels that direct the access control in the Bell-LaPadula (BLP) model.

a. (4 points) Is the 4 level Intel ring protection scheme just a simplified BLP? If so, describe the rule mappings. If not, identify how the ring protection scheme and BLP differ.

The schemes are similar but differ in the details. The ring scheme rules for write and execute differ. The ring scheme uses more labels than BLP.

b. (2 points) Is the ring memory scheme a mandatory or a discretionary access control? Why?

Mandatory, because the processes cannot change how their memory is labeled to allow different control. The levels are set at boot time.

19. (10 points, 2 each) Compare the following labels as integrity labels in the Strict Biba model. For each pair of subject and object labels and each model determine which access is granted of read or write. For the levels: Truth > Conjecture > Guess > Unknown. In addition the label may include categories A, B, or C. (2 points each, 12 total)

- a. Subject=Conjecture
Object=Guess

write

- b. Subject=Conjecture:A
Object=Conjecture:B

No access

- c. Subject=Truth:AB
Object=Truth:AB

Read and write

- d. Subject=Conjecture:A
Object=Truth:A

Read

- e. Subject=Unknown:C
Object=Truth:C

read

20. (4 points) Describe the key difference between a cryptographic hash and a regular checksum calculation (e.g., CRC or parity bit), and discuss why this difference is important.

The key point is that the cryptographic hash makes it very difficult (computationally infeasible) to find two messages with the same hash. This is not true for standard linear hashes like CRC.

So if an attacker wants to replace data (say a file advertised for download), but he wants his new file to match the original crypto hash, so people don't recognize the change, This is trivial to do for CRC, but not so for a crypto hash like SHA or MD5.

A number of people focused on the need for a key for a crypto hash. While many crypto hashes are keyed (e.g., HMAC-SHA), that is not a fundamental requirement for a crypto hash (e.g. basic SHA is not keyed).

21. (12 points, 3 points each) Given the plaintext “MESSAGE”, compute the cipher text for the following algorithms with the specified keys. In all cases, map the alphabetic characters to their numeric position in the alphabet (e.g., A=1, B=2, ... Z=26).

a. Caesar cipher with key = 8.

UMAAIOM (or 21 13 1 1 9 15 13)

b. Vigenere cipher with key=EXAM

RCTFFEF (or 18 3 20 6 6 5 6)

c. RSA cipher with block length 1 character and public key (3, 899).

399 125 566 566 1 343 125

d. Rail cipher

MSAEESG

22. (6 points) Consider the following situations where you are discussing a potentially embarrassing medical condition with your friend over the telephone. In all cases, someone hears your conversation. For each case, indicate whether this is a loss or a violation of privacy and why.

- a. You are talking on a cell phone in an airport lobby, and a colleague from work overhears you.

Loss of privacy. You should have no expectation of privacy in a public place.

- b. You are sitting at home. You think your friend is also at home, but unknown to you, his phone calls are being forwarded to his cube at work where he is listening to you on his speaker phone. Your friend's cubicle neighbors cannot help but hear the conversation.

Violation of privacy. While the conversation is on a speaker phone you are not aware of it. Some folks argued for loss of privacy because of the speaker phone. If you argued that you would know it is a speaker phone (because you know your friend tends to do this or you can hear the sound quality difference), then loss of privacy is a reasonable answer.

- c. You are sitting at home, and your friend is also taking the call at his house. You are using a cordless phone, and your HAM radio-enthusiast neighbor listens in.

Most people answered violation of privacy because you have an expectation of privacy at home. Some people answered loss, because you should know that the cordless phone could be tapped. This argument depends on the knowledge of the person. Because you know this, then it is a legitimate argument for you.

23. (9 points) You are consulting the on the design of a password system. You have been told that it would be reasonable to assume 10,000 guesses could be made each second. The goal of the system is to age out a password when there is a 50% chance that a brute force attacker could guess the password.

- a. Assume that legal password characters include upper and lower case letters, numbers, space character, underscore character, and dash character. If each password is 5 characters long, what should be the maximum life time of the password?

58,014.5 sec (=966.9 min, =16.1 hours)

- b. If each password is 10 characters long, what should be the maximum life time of the passwords in the system?

*6.7314*10¹³ sec (=2,134,504 years)*

- c. If you add 12 bits of salt to each password, how does that affect password life time if we are concerned only with on-line (or type 2) attack?

salts do not help with type 2 attacks