Quiz 1

1. This quiz is intended to provide a fair measure of your understanding of the course material to date (Homeworks 1–4 and Lectures 1–7).
2. Do not open this quiz booklet until the quiz begins. Read all the instructions first.
3. When the quiz begins, write your name on every page of this quiz booklet.
4. This quiz booklet contains 14 pages, including this one. An extra sheet of scratch paper is attached. If necessary, you can use it for the continuation of any problem answer.
5. This quiz is open-book, open-notes, open-calculators, open-computers, open-Internet. But no collaboration with anyone else is permitted.
6. Write your solutions in the space provided. If you need more space, write on the back of the sheet containing the problem. Do not put part of the answer to one problem on the back of the sheet for another problem; pages may be separated for grading.
7. Partial credit will be given. You will be graded not only on the correctness of your answer, but also on the clarity with which you express it. Be neat.
8. Good luck!

Problem	Points	Grade	Initials
1	20		
2	20		
3	20		
4	10		
5	10		
6	10		
7	10		
Total	100		

Your Name: _____

Problem Q1-1. Short Answer [20 points]

(a) What is a the US National Birth Certificate Number, and how is it different from the Social Security Number?

Solution: The Birth Certificate Number was a proposed national standard for enumerating each birth in the United States. The standard, proposed in 1948, would have corrected many of the flaws of the social security number — for example, it would include a check digit. But an uproar over enumeration caused the plan to be dropped.

Many student answers implied that there is currently a US National Birth Certificate Number. There is not. Some states have birth certificate numbers, but they are not national in scope. My two sons was recently born in Cambridge and received the Birth Certificate Number #893 and #894. I guess they start renumbering each year.

(b) Do Biometrics eliminate the need for people to memorize difficult-to-remember passwords?

Solution: A perfect biometric would eliminate that need, but there is no perfect biometric. As a result, biometrics are generally paired with PINs or passwords. The use of a unique PIN allows a biometric to be used in a 1-to-1 verification mode, rather than a 1-to-many identification mode. In general, biometrics are more accurate at verification than identification.

Another valid answer is that there will be many applications that cannot be modified to use biometrics, and that any practical biometric system must have some kind of back-door, so at least some person must have some password memorized somewhere.

Several students said that a biometric system could be used with easy-to-remember or simple passwords in case the biometric failed. That would be a bad idea: an attacker could just avoid the biometric and guess the passwords.

A simple "yes" or "no" answer to this question was not acceptable.

(c) Is triple DES as good an encryption algorithm as AES? Why or why not?

Solution: Triple DES is slower than AES but is probably just as secure as AES against brute-force attacks.

3DES has 168 bits of key, giving it more bits than AES-128 but less bits than AES-192 or AES-256. In general, there is no security difference today between a 128-bit key, a 168-bit key, a 192-bit key or a 256-bit key, because they are all resistant against a brute-force key search. That might change in time.

Another advantage of AES over DES is that the encryption block size is larger, making ECB mode more secure.

(d) It is Monday. Your organization is failing, and you have to sanitize 2000 hard drives before your last day of work (this Friday). You are legally required to properly sanitize these drives and you will be held personally responsible if you fail. What do you do?

Solution: Have the drives incinerated or otherwise physically destroyed.
You could hire that Type-1 degausing machine, a photo of which was shown in class. You might try to hire 10 people for a week to sanitize 20 drives/day over the next 5 days, but this approach would be prone to failure.

(e) In completing Assignment #3, the disk forensics assignment, one student obtained from a friend a USB drive that contained a series of files, photos, and documents. Some of the files were deleted and others weren't. One source of confusion, though, was the discovery of several deleted photos that predated the original sale of the drive. These were photos of people who were not known to either the students or his friend. Which probably happened?

Solution: The drive was probably used by a previous customer, returned, and resold to the student's friend. Either that, or the drive was opened in the store, the images were put on the drive, they were deleted, and finally the drive was packed back up.

Problem Q1-2. Multiple Choice [20 points] **Some questions may have multiple correct choices; circle all that apply.**

(a) From which of the following file systems is it possible to recover a file that has been deleted but not overwritten?

(A) FAT-16

(B) FAT-32

(C) NTFS

(D) HFS (Macintosh Hierarchal File System)

(E) UFS (Unix File System)

F None of the above

Solution Note: If the file has not been overwritten, then the file is still there to be recovered.

(b) Which of the following are effective techniques for sanitizing a hard drive?

(A) Destroy the drive in a furnace

B Use the DOS `FORMAT C:\` command with the `/S` option to install a new system.

(C) Delete all of the files on the drive, then create a single large file that uses all available space.

D The second law of thermodynamics states that it is impossible to destroy information, so it is theoretically impossible to completely sanitize a hard drive.

(c) Which of the following are advantages of biometrics?

A Biometrics can't be recorded or copied and played back at a later time.

B Because every person is different, there is no chance that a biometric will mistakingly identify one person as another.

(C) It's generally hard for people to share biometrics.

D None of the above.

(d) Which of the following pieces of information are both sent by your browser to a web server as part of the standard HTTP protocol and could be used to identify you?

(A) A packet containing your IP address.

(B) A cookie from your browser's "cookie jar."

C Your name and social security number

D Your GPS coordinates.

E Your phone number.

F None of the above.

(e) Which of the following techniques or implementation changes could have overcome the usability problems described in the KaZaA paper?

(A) A window that gave the name of every file that was being shared in one long list.

B A website with tutorials and detailed help pages describing how to use KaZaA's various privacy-enhancing features.

(C) An option that allowed the user to specify the specific file types that KaZaA would share, with the default being to share only multi-media files.

D A rule-based language that allowed users to specify regular expressions describing the kinds of files they wanted to share.

(E) A configuration change that caused the system to default to "share no files" rather than "share every file."

F None of the above.

Problem Q1-3. True or False [20 points]

Circle **True** or **False** for each of the following statements. If the statement consists of two parts where one part is true and the other part is false, circle False. No justification is required, but if you think the question is ambiguous, state your clarifying assumptions.

True **False** TEMPEST attacks against video displays are becoming less important as businesses move away from CRTs screens to LCD screens because CRTs have high-voltage transformers while LCD displays do not.

Solution: False; LCDs also have a TEMPEST risk; According to Kuhn et all, some LCDs give off more RF than some CRTs.

True **False** The attack described in *Optical Time-Domain Eavesdropping Risks of CRT Displays* can be trivially prevented by closing a window curtain. Even if the glow from the screen can be seen through the curtain, the attack does not work with light that is reflected or that has passed through a layer of fabric.

Solution: False; the attack works just fine with reflected light.

True **False** The attack described in "Information Leakage from Optical Emanations" is not practical to conduct over long ranges. Indeed, in the article, authors Loughry and Umphress were unable to make the attack work over a distance of more than 50 meters.

Solution: False; they didn't do the attack over more than 50 meters because their lab wasn't big enough.

The following questions draw upon Saltzer and Schroeder's paper The Protection of Information in Computer Systems

True **False** The principle "economy of mechanism" is best summarized by the engineering principle KISS ("keep it simple, stupid.")

Solution: True

True **False** Applying the "fail safe defaults" principle means that a computer with a newly installed operating system that has not been configured should be pretty much unusable until it has been configured.

Solution: True: A fail-safe default would render the system unusable until it was properly configured and secured. That's what is meant by "fail-safe" rather than "fail-usable."

True **False** All web-based applications inherently follow the "Complete mediation" principle because every HTTP GET or POST command is mediated by the web server.

Solution: False; although the web server does mediate every GET and POST, by default these commands are not checked for authorization unless this action has been explicitly coded into the application.

True **False** It is a violation of the "open design" principle for computer cases to be locked or otherwise sealed so that their users cannot open the boxes and see how they work.

Solution: False; "open design" refers to the design of a computer system, rather than to a specific computer case.

True **False** A good example of "open design" is Microsoft's Outlook mail client, since Outlook implements open standards such as HTML, POP and SMTP.

Solution: False; although Outlook uses standard protocols, the underlying design of the system is proprietary. Note: having an "open design" doesn't necessarily mean releasing the source-code: Microsoft could publish its design without revealing its source.

True **False** A biometric authentication system that requires a person to log in with a thumb-print and a secret PIN is an example of the "separation of privilege" principle.

Solution: True; Saltzer's article specifically states that using two passwords, or having two keys to unlock a single door, are examples of this principle.

True **False** The job of a poll worker in an election is to identify voters who show up to vote, check their name off the list, and make sure that they are able to vote without external influence or harassment. An good example of the principle of "least privilege" would be an electronic voting machine that allows a poll worker to change the order of the names on the ballot so that errors can quickly be corrected after they are identified.

Solution: False; according to the principle of least privilege, poll workers do not need such authority to complete their jobs.

Problem Q1-4. Long Answer [10 points]

Moore's law holds that the speed of computers roughly doubles every 18 years. Today an 80-bit cipher is considered to be marginally secure against a brute force attack, while a 128-bit cipher is considered to be quite secure.

Briefly describe what is meant by "a brute force attack."

Solution: A brute force attack is an attack that tries every key until the correct key is found. Any answer that mentioned the word "hash" was incorrect; this was not a hashing problem.

(b) If Moore's law holds, in what year will a 128-bit cipher be considered "marginally secure."?

Solution: This problem contained an error: we meant to write "every 18 months" and not "every 18 years." In either event, if an 80-bit cipher is marginally secure today, than it will take 40 doublings of computational power for a 128-bit cipher will be marginally secure, because there are 2^{40} as many 128-bit keys as 80-bit keys. (Some students wrote that the phrase *marginally secure* was ambiguous, but it was defined in the problem.) That's either going to be in 48×18 months or in 48×18 years, depending on how you read the problem. That is, either in 72 years (the year 2076) or in 864 years (the year 2868).

(c) If Moore's law holds, in what year will a 256-bit cipher be considered "marginally secure" against a brute-force attack?

Solution: A 256-bit cipher would require another 128 doublings. That would be a total of 176 doublings from today. That's either 264 (the year 2268) years from today or 3168 years (the year 5172) from today, depending on if you interpreted this question as 18 months or 18 years.

(d) Is it realistic to think that 256-bit ciphers will be attacked with a brute-force attack?

Solution: Probably not. There's only so much matter in the Universe. Even if computers are getting faster and faster, atoms only vibrate so fast, so you're going to need a lot of mass to build that computer which can try all 2^{256} possible keys.

Problem Q1-5. Long Answer [10 points]

Ben Bitdiddle has created a new hash function called BB-4. The primary advantage of this hash function is its speed. Instead of requiring multiple passes of the data, BB-1 treats the file as an array of 4-byte unsigned integers, which are then multiplied together (mod 2^{32}). There are thus 2^{32} unique hash codes that can be generated.

(a) Is it hard or easy to create hash collisions with BB-1?

Solution: It's pretty easy. Most of the reasons come from the associative and distributive property of multiplication.

(b) If it is easy, then give us a collision:

Solution: The string "ABCDEFGH" and the string "EFGHABCD". Or any two strings that have zero in two different four byte locations.

(c) Describe an important flaw in BB-4 other than the ease of finding hash collisions (if that is, in fact, a problem):

Solution: There are a number of flaws with this algorithm. It is easy to create a digest with a specific code. It is easy to tell what the digest will be from a given file. Changing the file does not necessarily change the digest.

Problem Q1-6. Long Answer [10 points] Ben Bitdiddle has gotten a job working for Dark Side Software, which is creating a piece of spyware that will covertly installed on computers throughout the world. The purpose of this software is to monitor what websites a person views and then report this information back to Dark Side's servers.

Ben has come to you needing help. Although he has created a wonderful user interface for his program, he can't figure out a way to get it installed on people's computers, and he can't figure out a good way for it to secretly and covertly transmit information back to the servers.

Describe a technique for installing this program on a computer system running the Windows operating system.

Solution: There are a variety of ways to do this. Most valid answers included bundling the spyware with another application or sending it via e-mail and convincing the user to install the program. If the computer is unpatched, you could also exploit a known Windows vulnerability.

Describe a covert channel that the software can use for reporting back the list of websites. The channel should not be evident to anyone who is running a client-side firewall or who uses a packet inspection program to examine the data payload area of TCP packets.

Solution: The data can be sent back in an encrypted message over a commonly open port. The biggest mistake here was to send the data back over an open port such as 80 for HTTP, but not to encrypt the data. Remember that even e-mail data must be encrypted in order for the data not to be easily viewable in the TCP packet.

Problem Q1-7. Long Answer [10 points] You are designing a system that will use a biometric (a thumb-print) to automatically validate the information on a driver's license. Here is how the system should work. When a person goes to a bar, an airport, or gets stopped by the cops, that person will present their driver's license and their thumb. The officer will place the driver's license on a hand-held scanner that will read a two-dimensional barcode on the card's back. The officer will place the thumb on the scanner, which will record the thumbprint. If the thumbprint is good, the individual's name is displayed on the officer's hand-held scanner.

Because radio contact is poor in some areas, the system cannot depend upon a real-time connection back to a central database to verify the contents of either the card or the scanner. To solve this problem, public key cryptography will be used. Design a system that uses public key cryptography to verify the identity of a person presenting one of these driver's licenses.

Make sure that the system is secure against the following attacks:

- A crook or terrorist who wishes to make false identification cards and has possession of several thousand valid cards (all obtained from people who have been mugged or murdered.)
- The theft of a police scanner by those same crooks or terrorists.
- An attacker who presents the officer with a card and a severed thumb.

Solution: The identity card contains a two-dimensional barcode that includes a representation of the user's fingerprint and a digital signature of the fingerprint. The signature is signed with a private key that is controlled by the government. The corresponding public key is widely distributed. In particular, it is in the fingerprint reader. When the user presents their card and their finger, the officer first determines that the finger is attached to the user and is alive. The officer then scans the fingerprint and reads the barcode with the 2D barcode reader. The public key in the fingerprint reader verifies the signature on the 2D barcode. The fingerprint from the reader is then compared with the fingerprint in the 2D barcode. No network traffic is created, and no secrets are stored in the fingerprint reader.