Final Exam

CS642: Computer Security
May 8, 2016

NAME:

UW ID:

It is a dark time for the Silicon Valley startup Hoolibits. Although it's previous security vulnerabilities have been patched, hackers and competition have driven daily uniques from their previous highs and venture capital is running low in the Valley.

Evading the dreaded "down-round", a group of clever engineers led by you, the proven security guru, are pitching investors on ways to pivot Hoolibits into a security tech company. You are looking to establish a middle-out level breakthrough in new security technology ...

This exam is open book, open note. You may use any printed materials you like on this exam. You may not use a computer, phone, any electronic devices, or the internet to help you complete this exam.

Please ensure answers are neat and legible. Illegible answers may be given no points.

If you choose, for any question (that is not multiple choice) you may write "I don't know" and circle this phrase as the answer. Questions answered this way will be given 30% of the points for that question (including fractional points). Nonsense or incorrect answers may be given no points when appropriate.

NOTE: This option is NOT available to multiple choice questions.

DO NOT START UNTIL THE CORRECT TIME

Problem 1 — [12 points]

NOTE: The “I don’t know” option is not available for this question. Unanswered items will be awarded no points.

The following list of protocols or technologies are used in the Hoolibits technical infrastructure. Identify which of the following are using cryptography (and which types) by default by circling the correct answers. Please note that some technologies use **both** types of cryptography in which case one must circle **both** to get full points.
The term “pubkey crypto” means public key cryptography or asymmetric cryptography.

UDP	Symmetric Crypto	Pubkey Crypto	No Crypto
HMAC	Symmetric Crypto	Pubkey Crypto	No Crypto
DNS	Symmetric Crypto	Pubkey Crypto	No Crypto
PBKDF	Symmetric Crypto	Pubkey Crypto	No Crypto
TLS	Symmetric Crypto	Pubkey Crypto	No Crypto
Port scanning	Symmetric Crypto	Pubkey Crypto	No Crypto
ARP	Symmetric Crypto	Pubkey Crypto	No Crypto
TCP SYN flood	Symmetric Crypto	Pubkey Crypto	No Crypto
Bitcoin	Symmetric Crypto	Pubkey Crypto	No Crypto
HTTP	Symmetric Crypto	Pubkey Crypto	No Crypto
ASLR	Symmetric Crypto	Pubkey Crypto	No Crypto
VM Co-residence detection using cache side channels	Symmetric Crypto	Pubkey Crypto	No Crypto

Problem 2 — Many Bothans died to bring us this information [12 points]

Hoolibits is considering a pivot from its core business to building an anonymizing network for internet traffic. The service must be both fast and secure. It will compete with the popular Tor network.

A. (3 points) What is the fundamental difference between a one-hop anonymizing proxy server and the Tor anonymity network?

In a one-hop proxy, the proxy server knows both the client's IP address and the destination. In Tor, circuits are comprised of 3 relays: no single relay knows both simultaneously.

B. (3 points) Assuming no cooperation between any relay nodes, what is the minimum number of Tor relays that must be used in a circuit to provide anonymity even if relay nodes are spying on traffic passing through that node? Explain why this is the minimum number.

Two nodes. Assuming no cooperation: one node knows the client's IP address but not the destination, the other knows the destination but not the client's IP address.

C. (3 points) One proposal for the Hoolibits anonymizing service is that a client installs a program on their computer and, when anonymizing mode is activated, all TCP traffic is routed through the anonymizing network; other IP traffic is routed normally. Does this design provide anonymity for web browsing? Under what threats does it provide anonymity and under what threats does it fail?

If only TCP traffic is anonymized then DNS queries made by the web browser to resolve URLs to IP addresses (sent over UDP) will not be anonymized. If an eavesdropper is watching all traffic from a user's computer, then the eavesdropper will see these DNS queries. However, if someone inspects a particular web server's log files, the user's IP will not appear, instead, an IP address from an anonymizing node will appear.

Problem 3 — Plop and Go [12 points]

(Cotton candy sweet to go, let me see that plopsicle.)

The Hoolibits CTO wants to build a new file sharing service called PlopIt! Users upload files to the PlopIt website (over TLS) and then the service assigns a URL that looks like:
`https://plopit.net/username/HASH`
where HASH is the SHA256 digest computed over the contents of the uploaded file. The user is given the URL (called a *plopsicle*) when the upload completes, and can use the URL or send it to any of her friends. Anyone with the URL can download the file.
A+B = 6 points; C+D = 6 points

A. Is this design secure? Why or why not?

Many possible answers. Here's an attack against the privacy of PlopIt users:
Say the MPAA suspects that Alice is illegally sharing copyrighted material. They can probe the URL by computing the hash of some file then test the URL. If it succeeds, then the MPAA knows that Alice has uploaded the file.

B. If you think it is insecure, describe how you will change the design to make it more secure. Be specific.

The design can be secured against this attack by making the URLs hard to guess. Examples: generate random URLs, or use a non-public salt when computing the hash.

C. Some of the investors think the large URLs are too ugly. They also want to provide, by default, short URLs that look like: `plop.it/xyzab` where "xyzab" is the hash shortened to just 5 lower-case alphabetic characters. Is this design secure? Why or why not?

This is not secure. The search space is too small (26^5 ~ 12M). Iterating through a large number of all possible URLs is feasible. This would reveal lots (possibly all) private content uploaded to PlopIt.

D. If you think it is insecure, describe how you will change the design to make it more secure. Be specific; if necessary, include some basic calculations to justify your solution.

"Short" URLs must not be too short even if they are random. One fix: make the URLs 10 characters, uppercase, lowercase, and numeric. This gives a search space of 62^{10} ~ 2^{80}. Even if there are 1M uploads to PlopIt (2^{20}), the search space is still very large and sparse (chance of discovering a file with one query: $1/2^{60}$).

Problem 4 — Never Gonna Give Your Password Up

[12 points]

One of your competitors, CodeSwap, recently had their website breached; the password database was stolen and posted publicly on pastebin.com. The entries in this database are of the form: `[user_email,SHA256(password)]`. There are approximately 20 million entries in this password database.

A. (4 points) If an attacker wants to crack the password from a single account in this database, does the lack of salting make the attack any easier? Explain why or why not.

Salting doesn't make cracking a single account any harder. Running a dictionary attack on a single account is no different (in terms of performance) with or without a salt.

B. (4 points) If an attacker wants to recover at least 50% of the passwords from this database, does the lack of salting make the attack any easier? Explain why or why not.

Salting makes this much harder. Without salting, an attacker can run a dictionary attack one time, compute the hash of each password guess just once, and check if any hash in the database matches. With salting, the attacker must compute hashes separately for each account because each account has a distinct salt. This makes the attack roughly 20M times harder with salts.

C. (4 points) It turns out that that 30% of the users at CodeSwap also have accounts on the Hoolibits website. Unlike CodeSwap, the Hoolibits site uses salts in its passwords. Should Hoolibits be concerned about this breach? Explain why or why not.

Hoolibits should be concerned about the security of its users' accounts. Cracking unsalted passwords is easy and users often re-use the same password across multiple accounts. Attackers may use cracked passwords to access user accounts on the Hoolibits website. Testing a cracked password is easy: just try to login using the cracked password.

Problem 5 — PC Load Letter [12 points]

The (very expensive) printers in the Hoolibits office are from PrintAsylum. These printers use a protocol for automatic discovery and configuration on the local network.

In this protocol the user's computer automatically broadcasts a *Printer Discovery* packet over UDP. The packet has a destination address as the broadcast address (it will be routed to every endpoint on the network). The source and destination ports are both 1848. Any time a printer sees a discovery packet, it responds with a *Printer Announcement* packet broadcast to the entire network. The packet payload contains the name of the printer, it's IP address, and configuration options (model number, double-sided printing, etc.).

When a machine receives a Printer Announcement packet, it verifies that the source IP address matches the IP address in the payload. In case of a mismatch, the packet is ignored. Otherwise, each machine automatically adds this printer to it's list of known printers. If a previous entry for this printer name exists, it is overwritten. All machines automatically process Printer Announcement packets whether or not that machine sent the initial Printer Discovery packet.

The Hoolibits CEO, Veronica, is negotiating another round of VC funding. The negotiations are complete and she is printing a final copy of the contract to sign and mail. Unknown to Hoolibits, a competitor, ODIN, has compromised another **computer** on the network and wants to read this contract. If possible, ODIN would also like to modify the document to generate distrust between Hoolibits and its investors.

Veronica will connect her laptop to the local network; it will run the printer discovery protocol; then she will print the document. The ODIN spy, Malory, has the ability to receive all broadcast packets and inject packets with spoofed IP addresses. Malory has not been able to hack the printers, any routers or switches, or any other machine in the network; her only hope is to abuse this printer discovery protocol.

A. (6 points) Can Malory arrange to learn the contents of this contract without physically accessing any of the printers? Describe the attack or explain why the attack isn't possible.

This is the ARP man-in-the-middle attack adapted for UDP. When V(eronica) plugs in her laptop, M(alory) will see the printer discovery and and announcement messages. Then V sends a spoofed announcement packet announcing her IP address as the printer. When her announcement packet is received by V, her laptop will overwrite the printer details with M as the IP for the printer. When V sends the document, it will go to M's computer where she can read it, and if she chooses, M can then send the document to the IP address of the printer if she wants it to be printed for V to find.

B. (6 points) Can Malory modify the document to be printed by replacing the original document with a modified one? It's not acceptable to allow the original to be printed and also print a modified version: this creates a dangerous *observable* and might compromise Malory's access to the Hoolibits local network for future espionage.
Describe the attack or explain why the attack isn't possible.

Same as above except M modifies the document then sends the modified version to the printer.

Problem 6 — The Ballad of Alice & Bobby [12 points]

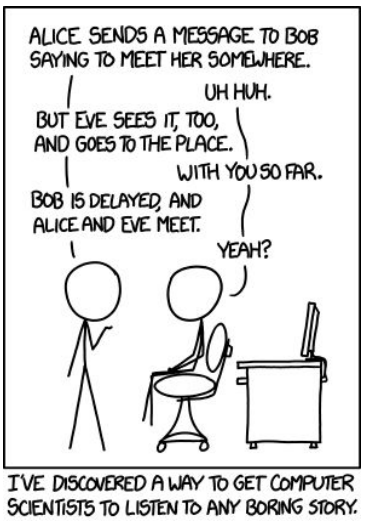

Alice wants to send a large document as an encrypted attachment to an email to Bob over the internet. Alice also wants Bob to know that this attachment was sent by her (and not a forged attachment sent by someone else). Assume Alice and Bob have each other's public keys.

In the questions below, use the cryptographic primitives we've discussed in class. Define any cryptographic functions that you use (example: one could say, H is cryptographic hash function, or H is MD5).

A. (6 points) Give the steps for Alice to prepare the attachment that will be sent.

PKa, SKa, PKb, SKb - Alice and Bob's public/private key pairs
Ek()/Dk() -- Authenticated encryption/decryption scheme using key k (say, AES-GCM)
Enc()/Dec() -- Public key encryption/Decryption (say, RSA encryption)
Sign/Verify -- Digital signature/verification algorithm (say, RSA signature)

// Choose random, one-time symmetric key, then encrypt the attachment M
k = rand_key()
C_1 = Ek(M)

// Encrypt k with Bob's **public key** and sign the entire ciphertext with Alice's **private key**
C_2 = Enc(PKb, k)
sig = Sign(SKa; C_1,C_2)

Send: C_1,C_2,sig

B. (6 points) Give the steps for Bob to decrypt Alice's attachment and verify that the message is valid and authentic (it has not been tampered and it was definitely sent by Alice).

// Verify the digital signature using Alice's **public key**
if Verify(PKa; sig; C_1,C_2) = INVALID: return ERROR

// Decrypt the symmetric key with Bob's **private key**, then decrypt the attachment
k' = Dec(SKb, C_2)
M' = Dk'(C_1)
if M' = ERROR: return ERROR
else: return M'