Weaver
Fall 2017

CS 161
Computer Security

Homework 4

Due: Monday, December 4th, at 11:59pm

Instructions: This homework is due **Monday, December 4th, at 11:59pm** on Gradescope. No late homeworks will be accepted unless you have prior accommodations from us. This assignment must be done on your own.

Make sure you have a Gradescope account and are joined in this course. **The homework must be submitted electronically via Gradescope** (not by any other method). Your answer for each question, when submitted on Gradescope, should be a single file with each question's answer on a separate page. Please don't forget to assign pages to questions on your submission.

This will be your final homework assignment. We hope that you've enjoyed this course and its very applicable content. Best of luck on all your finals!

Problem 1 ***Finding Outis*** **(20 points)**

(a) The following question will require you to download wireshark. Search online on how to perform the tasks below.

It's time to apply everything you learned in this class to discover who Outis is.

We've intercepted traffic from a forum that Nick and Outis frequently converse on. However, the site uses TLS!

What is the cipher suite that was agreed on? (Hint: What step of the handshake does the client and server agree on crypto functions. You should copy this EXACTLY as shown on wireshark.)

(b) Using your mastery on security, you managed to find a command injection vulnerability on the forum to obtain their private key.

Decrypt the communication to find out who Outis is! Post the EXACT message you find[1].

(c) Would you have been able to decrypt the message had the client and server agreed on ephemeral Diffie-Hellman rather than RSA? Why or why not?

[1] Note: Nick is not actually Outis. Also as a reminder, do not wiretap unless it's your own network. I created the fake message in the pcap file.

Problem 2 ***DNSSEC & DOS*** **(35 points)**

DNSSEC (DNS Security Extensions) is designed to prevent network attacks such as DNS record spoofing and cache poisoning. When queried about a record that it possesses, such as when the DNSSEC server for `berkeley.edu` is queried about the IP address of `oskibear2k17.berkeley.edu`, the DNSSEC server returns with its answer an associated *signature*.

(a) Suppose you make a DNSSEC-enabled request for the A record of oskibear2k17.berkeley.edu. To your dismay and disappointment, you receive the response: NXDOMAIN. The response also contains NSEC and RRSIG records.

What further DNS queries, if any, must you make to validate this response? Assume all caches are empty. Provide both the domain(s) to be queried and the necessary record type(s) for each domain.

(b) For the following, suppose that a user R (a resolver, in DNS parlance) sends a query Q to a DNSSEC server S, but all of the network traffic between R and S is visible to a network attacker N. The attacker N may send packets to R that appear to originate from S.

Suppose that when queried for names that do not exist, DNSSEC servers such as S simply return "No Such Domain," the same as today's non-DNSSEC servers do. This reply is not signed.

Describe an attack that N can launch given this situation.

(c) Suppose now that when queried for a name Q that does not exist, S returns a signed statement "Q does not exist."

Describe a DoS attack that N can launch given this situation. Assume that signatures are costly.

(d) One approach to address the above considerations is to use NSEC Records. As mentioned in class, when using NSEC S can return a signed statement to the effect of "when sorted alphabetically, between the names N_1 and N_2 there are no other names." Then if the name represented by the query Q lexicographically falls between N_1 and N_2, this statement serves to confirm to R that there's no information associated with the name in Q.

NSEC has a shortcoming, which is that an attacker can use it to *enumerate* all of the names in the given domain that do indeed exist. To counter this threat, in the April 17 section we introduced the NSEC3 Record, which is designed to prevent DNS responses from revealing other names in the domain. NSEC3 uses the lexicographic order of *hashed* names, instead of their unhashed order. In response to a query without a matching record, NSEC3 returns a signed statement of the hashed names that come just before and just after the hash of the query.

Suppose that the server S has records for `a.cs161.com`, `b.cs161.com`, and `c.cs161.com`, but *not* for `abc.cs161.com`. In addition, assume that `a.cs161.com` hashes to

`dee60f2e...`, `b.cs161.com` to `80a4cb36...`, `c.cs161.com` to `c218f96a...`, and `abc.cs161.com` to `99f3e2ba...`.

If R queries S for `abc.cs161.com`, what will S return in response? Describe how R uses this to validate that `abc.cs161.com` does not exist.

(e) The hashes in NSEC3 are computed as a function of the original name plus a *salt* and an *iteration parameter*, as follows:

```
Define H(x) to be the hash of x using the Hash Algorithm selected by
the NSEC3 RR, k to be the number of Iterations, and || to indicate
concatenation. Then define:

   IH(salt, x, 0) = H(x || salt), and

   IH(salt, x, k) = H(IH(salt, x, k-1) || salt), if k > 0

Then the calculated hash of a name is

   IH(salt, name, iterations)
```

The name of the hash function, the salt and the number of iterations are all included in an NSEC3 reply (that is, they are *visible* and assumed to be easily known). All replies from a given server use the same salt value and the same number of iterations.

Suppose an attacker has a list of "names of interest," i.e., names for which they want to know whether the given name is in a particular domain. If the attacker can get all of the NSEC3 responses for that domain, can they determine whether these names exist? If so, sketch how. If not, describe why not.

(f) Why would a domain change their *salt* value in NSEC3 replies?

(g) What is the purpose of the *iteration parameter* in NSEC3 replies?

(h) The specification of NSEC3 also sets an upper bound on the *iteration parameter*. What could happen if that upper bound did not exist?

Problem 3 ***Under Attack from the APT Team*** **(20 points)**

You are in charge of `fu.co`, a new startup which is unfortunately become the target of the latest APT (Annoyingly Persistent Teenager) attacker team, the Suspended RedBand which is believed to operate out of the enclave known as Shallow Alto.

This attack group is using multiple techniques to try to break into your systems. Fortunately, you also have multiple layers of defense, including passive network monitoring (NIDS) on the network border, host-based monitoring (HIDS) on your web server and your mail server, and log based monitoring where the centralized file server records all file accesss by all systems.

(a) One primary tactic that the Suspended RedBand uses is to send malicious word documents through email. Your email server is properly configured to receive email using StartTLS. Fortunately these emails are easy to recognize because of various features. Which defensive system can detect these attacks? Can the defensive system also block these attacks?

(b) The Suspended RedBand successfully compromised a user's workstation and began exfiltrating (stealing and sending outside the network) information from the file server through an HTTPS connection. How would you determine what data the attacker accessed? Can you estimate how much data was actually exfiltrated rather than just examined on the compromised system?

(c) Your security posture doesn't just stop external threats but also internal damages. Human Resources reports that a worker has complained about another worker viewing inappropriate web pages at work. Which system can substantiate if this claim is correct?

(d) One of your vendors offers an IPS (Intrusion Prevention System) solution which you consider installing. The vendor claims that it has a zero false negative rate when detecting attacks by the RedBand. and a false positive rate of less that .1%, however false positives are very disruptive and will cause a lot of collateral damage. The vendor is honest but will provide no other information before you decide to deploy it. Is this a good system? Why or why not?

Problem 4 ***Detecting Worms*** **(20 points)**

Assume that you are working for a security company that has to monitor a network link for worm traffic. The link connects a large site with the rest of the Internet, and always has lots of traffic on it. Your company sells a monitoring box that can scan individual packets for fixed strings at very high speeds.

(a) Suppose that after careful analysis you discover that a particular TCP-based worm that you need to protect against generates traffic that always contains a fixed 4-byte sequence. You program your company's specialized hardware to generate an alarm whenever it detects a packet containing this 4-byte pattern.

Explain how this detector can exhibit false negatives.

(b) Propose an alternative architecture for your company's monitoring box that fixes the problem from part (a). Your alternative should not increase false positives by more than a modest amount. Does your revised approach completely eliminate false negatives? Explain why or why not.

(c) Suppose benign network traffic is uniformly distributed in terms of its content. That is, every payload byte is chosen uniformly at random. Calculate the expected time until a 150 MB/sec (megabytes per second) link will cause a false alarm.

(d) Assume now that the signature length is 8 bytes. Calculate the expected time until a 20 GB/sec (gigabytes per second) link will cause a false alarm.

(e) What could a worm author do to try to ensure that their worms do not have many fixed-byte sequences?

Problem 5 ***Intro to Tor*** **(20 points)**

(a) True or False: When a client sends a message over the Tor network, Tors onion routing works because each intermediary encrypts the message they receive with the public key of the following intermediary, so no one else can decrypt the messages

(b) True or False: Even if a MITM managed to be placed in between the last onion router and its destination, that MITM will still not be able to decrypt traffic originally sent in plaintext.

(c) True or False: Bob wants to message Alice using Tor. However, Mallory controls 160 of the 161 intermediaries (i.e. these intermediaries are dishonest). As a result, Mallory can determine that Bob and Alice are messaging.

(d) The remaining questions will require you to download TorBrowser

After downloading Tor, login to facebook on Tor (create a temporary account if you don't have one). Once logged in, take a screenshot and include it in your submissions

(e) One common observation is that using Tor is extremely slow. Explain why.

(f) Now go to youtube on Tor and include a screenshot in this submission. What do you notice when you visit this site other than its speed?

Problem 6 ***Feedback*** **(0 points)**

Optionally, feel free to include feedback. What's the single thing we could do to make the class better? Or, what did you find most difficult or confusing from lectures or the rest of class, and what would you like to see explained better? If you have feedback, submit your comments as your answer to Q6.