CS 361S - Network Security and Privacy
Spring 2014

FINAL

May 12, 2014

DO NOT OPEN UNTIL INSTRUCTED

YOUR NAME: _____

Collaboration policy

No collaboration is permitted on this exam. Any cheating (*e.g.*, submitting another person's work as your own, or permitting your work to be copied) will automatically result in a failing grade. The UTCS code of conduct can be found at `http://www.cs.utexas.edu/undergraduate-program/code-conduct`

Final (125 points)

Problem 1 (28 points)

Circle only one of the choices **(4 points each)**.

1. **TRUE** **FALSE** BGP route updates are authenticated by 16-bit random transaction IDs.
2. **TRUE** **FALSE** To use a session-filtering firewalls, applications such as Web browsers and FTP clients must be modified accordingly.
3. **TRUE** **FALSE** It is possible to create an anomaly detector which produces no false positives.
4. **TRUE** **FALSE** Recall that an RSA modulus n is a product of two large primes. If someone discovers an efficient algorithm for computing the greatest common divisor of two numbers, then breaking RSA will become feasible.
5. **TRUE** **FALSE** The main reason why WEP is insecure that the attacker who knows the plaintext can easily recover the keystream by XORing this plaintext with the ciphertext.
6. **TRUE** **FALSE** With properly validated, unforged certificates, HTTPS security guarantees hold even if DNS has been poisoned.
7. **TRUE** **FALSE** You are attempting to connect to a server which presents a self-signed certificate. If you accept it, your security is the same as if you accepted a certificate issued by an untrusted certificate authority.

Problem 2

Problem 2a (4 points)

What is the difference between a Web attacker and a network attacker?

Problem 2b (4 points)

What capabilities do Web attackers have that allow them to perform cross-site request forgery attacks against websites that rely only on cookies for authentication?

Problem 2c (4 points)

What capabilities do Web attackers have that network attackers do *not* have?

Problem 3

Recall `PwdHash`, a browser extension that automatically converts user's password into website-specific, random-looking values. When the user visits a website like `chase.com` and types in a plaintext password like `monkey`, the browser instead sends `Hash(monkey,chase.com)`, where `Hash` is a cryptographic hash function. This hashed value is what the website sees as the user's password.

Problem 3a (4 points)

Suppose that the user uses the same weak password (eg, `monkey`) at multiple bank websites.

Does `PwdHash` protect the user's password from being cracked by a brute-force attack? If so, how?

Problem 3b (4 points)

Suppose that the user visits a fake website set up by a Web attacker—for example, a website that looks `citi.com` but is hosted at `citi.com.cn`—and mistakenly types his true `citi.com` password into his modified browser.

Does `PwdHash` protect the user's password from being stolen and used by the Web attacker? If so, how?

Problem 3c (4 points)

Suppose that the user accesses a banking website over HTTP via a malicious Wi-Fi access point.

Does `PwdHash` protect the user's password from being stolen and used by the network attacker? If so, how?

Problem 4 (5 points)

Describe at least **two** changes that could be made to the C compiler to prevent buffer overflow attacks. Explain why these defenses would be effective.

Problem 5 (4 points)

SPF and DKIM are two defenses against spam.

With SPF, the receiving email server performs a DNS lookup on the "From:" domain name in the received message. As part of the DNS response, it receives the list of IP addresses authorized to send email from that domain. If the IP address from which the message arrived is not on the list, the message is rejected.

With DKIM, the sending email server digitally signs the message with its private key. The receiving email server performs a DNS lookup on the "From:" domain name in the received message. As part of the DNS response, it receives the public key of that domain.

It uses this public key to verify the signature on the received message. If verification fails, the message is rejected.

Describe an attack on SPF that does not work against DKIM.

Problem 6

Consider a stateless packet filtering firewall installed at the gateway of a corporate network. Assume that all traffic to and from the network flows through the firewall. The format of a firewall rule is as follows:

```
Interface Action SourceIP SourcePort DestIP DestPort
```

Problem 6a (4 points)

Can a packet filter block all external attempts to connect using HTTP to a Web server located at a particular address within the corporate network, but permit HTTPS access to the same server? If yes, what would the firewall rule(s) look like? If no, why not?

Problem 6b (4 points)

Can a packet filter block all incoming email messages containing the word `V1AGRA`? If yes, what would the firewall rule look like? If no, why not?

Problem 6c (6 points)

List **three** different network attacks that even a stateful firewall cannot protected against.

Problem 7 (10 points)

Suppose that every packet observed by a network-based intrusion detection system (NIDS) belongs to one of the following mutually exclusive categories: legitimate (88% of all traffic), known worm (4%), distributed denial of service (4%) or port scan (4%).

The NIDS correctly classifies all known-worm packets. A legitimate packet is classified as legitimate with probability 91%, and misclassified as belonging to any of the three attack categories with equal probability. A DDoS packet is classified as DDoS with probability 50%, as a known worm with probability 40%, and as a legitimate packet with probability 10%. A port-scan packet is classified correctly with probability 85%, and misclassified as a legitimate packet with probability 15%.

If the NIDS announces that a particular packet belongs to a known worm, what is the probability that this packet is **not** a legitimate packet? Show your calculations.

Problem 8 (4 points)

How does a network telescope work? Give an example of a real-world worm or virus that would have been difficult to analyze using a network telescope and explain why.

Problem 9 (4 points)

The Bank of Molvanîa website uses authentication cookies in which the username and timestamp are encrypted with RC4(K), where K is the fixed key known only to the website itself.

If you are a customer of the bank, explain how you can log in under the username of any other customer.

Problem 10

For each of the following threats, explain in detail what mechanism is used in SSL/TLS to provide protection, and how it is used. Do not make any assumptions about the specific encryption or signature scheme used by SSL/TLS, as it is supposed to be compatible with multiple schemes.

Problem 10a (4 points)

A network eavesdropper records all of the server's messages in an SSL handshake. Later, he impersonates the server by replaying these messages to a client.

Problem 10b (4 points)

A network attacker poisons DNS and tricks the client into thinking that the domain to which the client is connected is hosted at an attacker-controlled IP address.

Problem 10c (4 points)

A network attacker modifies the client's "Hello" message in transit and tricks the server into thinking that the client supports only relatively weak cryptographic algorithms.

Problem 11

Explain the differences between an HTTPS session protected using a regular certificate and an HTTPS session protected using an extended-validation certificate.

Problem 11a (4 points)

What does the browser do differently?

Problem 11b (4 points)

What is the difference from the viewpoint of a Web attacker?

Problem 11c (4 points)

What is the difference from the viewpoint of a network attacker?

Problem 12

Sotirov et al. used MD5 hash collisions to forge a rogue SSL certificate for the name "MD5 Collisions Inc." This certificate can be used to issue additional fraudulent certificates for arbitrary Web domains, such as `gmail.com`

Problem 12a (4 points)

Why would anyone believe that "MD5 Collisions Inc." is authorized to issue certificates for other domains?

Problem 12b (4 points)

The forged "MD5 Collisions Inc." certificate is not one of the root certificates stored by the Web browser. Would the browser accept a `gmail.com` certificate issued using the forged certificate? Why or why not?