

CS 378 - Network Security and Privacy
Spring 2009

FINAL

May 13, 2009

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 Computer Sciences department code of conduct can be found at <http://www.cs.utexas.edu/academics/conduct/>

Final (125 points)

Problem 1 (28 points)

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

1. **TRUE** **FALSE** An encryption scheme can be secure against the chosen-plaintext attack even if there exists a key K such that all messages encrypted with K result in the same ciphertext.
2. **TRUE** **FALSE** Same-origin policy in modern Web browsers says that a script can access Web resources only from the same server, protocol, and port as the script itself.
3. **TRUE** **FALSE** It is not known if computing $\phi(n)$ is harder than computing the prime factors of an RSA modulus n .
4. **TRUE** **FALSE** Recall that an RSA modulus n is a composite number, which is a product of two large primes. If someone discovers an efficient algorithm for computing the greatest common divisor of two composite numbers, then breaking RSA will become feasible.
5. **TRUE** **FALSE** It is impossible to create an anomaly detector which produces no false positives.
6. **TRUE** **FALSE** To take advantage of the authentication and confidentiality services provided by IPsec, applications such as Web browsers and FTP clients must be modified accordingly.
7. **TRUE** **FALSE** One of the benefits of Tor as opposed to basic onion routing is that sending messages over an established Tor circuit does not require any public-key operations.

Problem 2

You are eavesdropping on encrypted messages between Alice and Bob. You notice that many ciphertexts have the same prefix. When two ciphertexts have the same prefix, it consists of several hundred bytes, and the number of common bytes is always a multiple of 16. The parts of the ciphertexts that follow the common prefix never seem to have any common byte sequences of any significant length.

Problem 2a (6 points)

What cipher is likely being used (be specific and explain your reasoning)? In what mode of operation?

Problem 2b (5 points)

What are Alice and Bob doing wrong?

Problem 3

Kerberos 4 involves three message exchanges: (1) between the client C and the key distribution center D , (2) between C and the ticket-granting service T , and (3) between C and some server V .

The communication between C and D goes like this:

1. C sends a ticket request containing its own name and the name of the ticket-granting service T .
2. D checks the both C and T are known to the system. If they are, D creates a ticket containing C 's and T 's names, C 's network address, the current time, the lifetime of the ticket and a fresh session key K_{CT} . The ticket is encrypted with a secret key K_{DT} , which is known only to D and T .
3. The reply from D to C consists of the above ticket, T 's name, the current time, the lifetime of the ticket and K_{CT} , all encrypted with C 's secret key K_C . To prevent con-

fusion between messages, the plaintext of this encrypted reply also contains a constant string `krbtgt`, identifying this reply as a ticket-granting ticket.

Problem 3a (6 points)

Explain concisely the purpose of each of the three message exchanges: between C and D , C and T , and C and V .

Problem 3b (5 points)

It seems that the protocol would have been simpler if the ticket-granting service were eliminated, and the client would request tickets directly from D for each server connection. What are the disadvantage(s) of this design?

Problem 3c (3 points)

Suppose the attacker guessed the client's password. Explain how he can use D 's message to C to verify his guess.

Problem 4 (6 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

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 5a (5 points)

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

Problem 5b (5 points)

Can the packet filter block all email messages containing the word **V1AGRA** to a particular client within the corporate network? If yes, what would the firewall rule look like? If no, why not?

Problem 5c (6 points)

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

Problem 6 (5 points)

DKIM and SPF are two defenses against spam. Describe an attack on SPF that does not work against DKIM.

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 (6 points)

Tripwire is a software tool intended to assure integrity of system files by detecting unexpected modifications (such modifications are often a sign of rootkit activity). One version of Tripwire reads the names of the directories to be protected from a configuration file. For each file in the specified directories, Tripwire uses the host machine's cryptographic library to compute SHA-1 of the file. It then stores a database on the machine's hard drive which, for each file name, contains the associated SHA-1 value.

A Molvanian IT administrator configures Tripwire to protect all files in the `/usr/bin` directory. A month later, he executes Tripwire again to see if any of the files have been modified. The newly generated database of SHA-1 values perfectly matches the database that was generated when Tripwire was first executed. The administrator concludes that the files have not been modified.

In fact, his computer has been infected by a rootkit, and all system programs substituted by their rootkitted versions.

Describe at least **three** attacks that are consistent with this scenario.

Problem 9

Consider the following variant of RSA encryption. Recall that an RSA public key is a pair (n, e) . To encrypt some message m , first generate a fresh random value r of the same length as m . Use r as if it were a one-time pad to encrypt m (*i.e.*, let $s = m \oplus r$), and then encrypt r using plain RSA (*i.e.*, let $t = r^e \pmod n$). The ciphertext is the (s, t) pair.

Problem 9a (4 points)

How does decryption work in this scheme?

Problem 9b (5 points)

Is this encryption scheme secure against the chosen-plaintext attack? Explain.

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 (5 points)

Protection against replay attacks:

Problem 10b (5 points)

Protection against man-in-the-middle attacks:

Problem 10c (5 points)

Protection against known-plaintext attacks based on pre-computation:

Problem 11 (5 points)

Do you think Tor should be run over IPsec, or IPsec over Tor? Explain.