CSC474 Fall 2023 - Homework 2*

Due Friday November 10, 2023, 11:59pm ET

100 points

Prof. William Enck

1 Crypto Protocols {30 points}

1.1 *PayMe!* {10 points}

Hoping to become the next dotcom millionaire, Bob decides to create an online money payment service similar to PayPal. His service, *PayMe!*, allows users to transfer money to other users of the system.

To ensure that no fraudulent activity takes places, the *PayMe!* service stores the public key of each user. (You should assume that the sharing of the public key is secure; that is, the server has each user's correct public key.)

If Alice ("A") wishes to give X dollars to Bob ("B"), she sends the following message to the *PayMe!* service ("S"):

$$A \rightarrow S : A, B, X, n, Sig(A^-, [X|n])$$

where n is a nonce, A^- is Alice's private key, and $Sig(K^-, M)$ denotes a digital signature over M computed using the private key K^-.

(a) {5 points} What is a nonce, and why does Bob include one in his protocol? Does it prevent any type of attack?

(b) {5 points} Explain how an active adversary can exploit a weakness in the *PayMe!* protocol to steal money from an honest user, Alice.

1.2 PompousPass {20 points}

Bob believes he has a come up with a simple way of performing authentication. His system, PompousPass$^{\text{TM}}$, uses RSA signatures. Let Id_x and Pw_x respectively be the username and password for user x, and (x^+, x^-) be the public/private keypair associated with user x. Assume that the

*Last revised on October 5, 2023.

server S knows the user's username (Id_x), password (Pw_x), and his public key (x^+). To authenticate to the server, x sends:

$$x \rightarrow S : Id_x, r, Sig(x^-, [Id_x|Pw_x|r])$$

where r is a nonce.

The server should only authenticate the user iff (1) the transmitted password matches the password stored in the server's database and (2) the nonce is fresh.

(a) {5 points} Describe two attacks on this protocol. You should assume the attacker cannot get access to the password database.

(b) {10 points} Fix the protocol to defend against these attacks.

(c) {5 points} In practice, public-key cryptography is often used to distribute session keys, which are then used with symmetric algorithms. Why is this approach preferred over using solely public-key operations? (1-2 sentences)

2 Firewalls {30 points}

You are given the following "informal firewall policy" details to be implemented using a firewall setup for the following network.

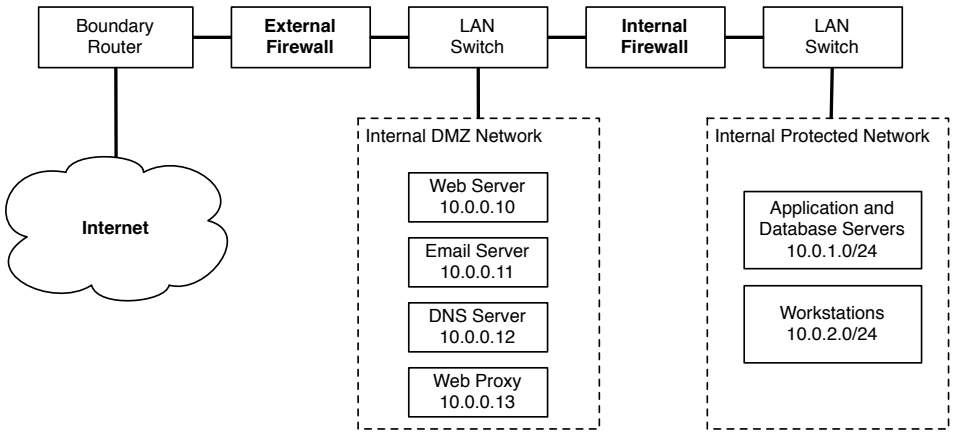

1. E-mail may be sent using SMTP in both directions through the firewall, but it must be relayed via the DMZ mail gateway that provides header sanitization and content filtering. External e-mail must be destined for the DMZ mail server.

2. Users inside may retrieve their e-mail from the DMZ mail gateway, using either IMAP or IMAPS, and authenticate themselves.

3. Users outside may retrieve their e-mail from the DMZ mail gateway, but only if they use the secure IMAP protocol, and authenticate themselves

4. Web requests (both insecure and secure) are allowed from any internal user out through the firewall but must be relayed via the DMZ Web proxy, which provides content filtering (noting this is not possible for secure requests), and users must authenticate with the proxy for logging. The DMZ Web proxy must be allowed to make Web requests to anywhere on the Internet.

5. Web requests (both insecure and secure) are allowed from anywhere on the Internet to the DMZ Web server

6. DNS lookup requests by internal users allowed via the DMZ DNS server, which queries to the Internet.

7. External DNS requests are provided by the DMZ DNS server.

8. Management and update of information on the DMZ servers is allowed using secure shell connections from relevant authorized internal users (may have different sets of users on each system as appropriate), but only from the internal network and the DMZ.

Design suitable packet filter rulesets to be implemented on the "External Firewall" and the "Internal Firewall" to satisfy the aforementioned policy requirements.

For this question, you should assume a **stateless** firewall. Recall that this means that you need to define rules for both directions, as discussed in class. Note that this question requires you to define the firewall policy for two different firewalls ("External Firewall" and "Internal Firewall"). Use the following table as an example of how define your rules. This example assumes a single firewall that is protecting hosts in the 10.0.0.0/24 subnet. It defines a stateless ruleset that allows the hosts in the subnet to access HTTP on the Internet.

Action	Src IP	Src Port	Dest IP	Dest Port	Protocol	Flags	Comment
allow	10.0.0.0/24	*	*	80	TCP		allow hosts to access the web (outbound)
allow	*	80	10.0.0.0/24	*	TCP	ACK	allow hosts to access the web (reply traffic)
deny	*	*	*	*	*		Default deny

3 Intrusion Detection {20 points}

(a) {10 points} This problem considers the base rate fallacy discussed in class. Let $Pr(M)$ be the probability that a given packet is malware (i.e., ground truth). Let Pr(A) be the probability that there is an alarm raised by the IDS. Your IDS is 99.9% accurate at detecting intrusions [i.e., $Pr(A|M) = 0.999$] and it is 99.9% accurate at detecting when an event is not an intrusion [i.e., $Pr(!A|!M) = 0.999$]. An intrusion occurs once every one million events (i.e., the base rate of incidence is 1 / 1,000,000). What is the "true alarm" rate [i.e., determine $Pr(M|A)$]? Show your work.

(b) {10 points} This problem considers the creating of ROC curves discussed in class. Assume the same trivial detection algorithm as the slides. $D(k, T) \rightarrow [0, 1]$, takes a package of length k and a threshold T. If the packet length $k \leq T$, then an alarm is raised. Produce a table similar to that in the lecture slides showing the TP% and FP%. From this table, draw an ROC curve. Use the following traffic classifications:

- Attack packet lengths: 1, 2, 2, 3, 3, 6, 6, 10
- Non-attack packet lengths: 3, 3, 5, 6, 7, 7, 8, 8, 8, 9

4 Routing {10 points}

(a) {5 points} BGP hijack attacks can be classified into two types: *prefix* and *subprefix*. Which is more dangerous and why?

(b) {5 points} AS relationships can be classified as *customer-provider* and *peer-peer*. A customer AS pays the provider AS to both send and receive traffic. In contrast, peer ASes commonly have a settlement-free peering arrangement, meaning they transit each others traffic for free. ASes typically avoid forwarding traffic from one neighbor to another if it cannot generate revenue for doing so.

BGP prefix filtering rules are a allowlisting technique used to filter out bogus BGP announcements. Rules are commonly based on an economically motivated rule of thumb: AS a will typically announce a route to a neighboring AS n only if: (1) n is a customer of a; (2) the route for a prefix originated by a; or (3) the route is through a customer of a.

For what type of sources (e.g., customer, provider, peer) does prefix filtering work well? For what type of sources does it not work well? For both cases, explain why.

5 DNS {10 points}

(a) {5 points} The *identifier* field in DNS requests and responses is 16 bits long. Consider an adversary, Edward, who is located in a far off region of the Internet from the victim, Victor. Edward is positioned such that he cannot intercept Victor's DNS requests (or the returned responses).

Further suppose that Edward can inject 1024 forged DNS responses per second (i.e., cause Victor to receive 1024 forged responses a second). If it takes one second for Victor to receive the correct DNS response from his resolver, what is the probability that Edward will be able to poison a particular request? Assume the requested hostname is not locally cached.

(b) {5 points} Now, assume Edward and Victor are on the same subnet and share a switch. Edward is still confined to forging 1024 DNS responses per second. Explain how Edward can effectively increase the probability of a cache forgery attack to 100%.