Question 1 ***Intrusion Detection*** **(15 min)**

FooCorp is deciding which intrusion detection method to employ in a few target scenarios. In each part, consider which of the intrusion detection methods learned in class would be most appropriate, and justify why. Try to be as specific as possible.

(a) FooCorp wants to detect attacks for a specific vulnerability that may exist in some of their web servers.

> **Solution:** FooCorp is probably best off using vulnerability signatures to try to detect attacks for the specific vulnerability. If there are key elements to all attacks on the vulnerability, these can be listened for and detected.

(b) FooCorp is using HTTPS, but all of their services use the same modular web framework. They are interested in detecting any time their servers receive arguments that are suspicious, in real-time.

> **Solution:** Some form of HIDS would probably be most appropriate here. A HIDS can check arguments and detect suspicious ones in real-time, and since the web framework used is modular and extendible, it wouldn't be very difficult to deploy HIDS software for every service.

(c) FooCorp is a diverse company, with a wide variety of different web servers built on top of different web frameworks, offering different services. They wish to detect suspicious arguments for all of their services. Every service uses HTTP and not HTTPS, and FooCorp has a low budget for security, but they want real-time detection.

> **Solution:** A NIDS would probably be best here. A NIDS is endpoint-agnostic, so the variety of servers that FooCorp is using doesn't affect a NIDS. It also only involves writing one piece of software and deploying at one point, so it would be a cheap solution.

(d) FooCorp again has many different web servers built on different web frameworks, but each uses the same logging format. They are using HTTPS, and do not need real-time detection.

Solution: Logfile analysis would probably work best. Since FooCorp does not need to detect attacks in real-time, a HIDS would probably be overkill here, while a NIDS wouldn't be as powerful.

Question 2 ***Detection Tradeoffs*** **(15 min)**

Suppose that S is a network-based intrusion detector that works by passively analyzing individual UDP and TCP packets. Suppose that A is a host-based intrusion detector that is a component of the browser that processes and analyzes individual URLs before they are loaded by the browser. Suppose S has false positive rate S_P and false negative rate S_N, and A has false positive rate A_P and false negative rate A_N.

Your company decides to build a hybrid scheme for detecting malicious URLs. The hybrid scheme works by combining scheme S and scheme A, running both in parallel on the same traffic. The combination could be done in one of two ways. Scheme H_E would generate an alert if for a given network connection either scheme S or scheme A generates an alert. Scheme H_B would generate an alert only if both scheme S and scheme A generate an alert for the same connection. (Assume that there is only one URL in each network connection.)

(a) Assuming that decisions made by S and A are well-modeled as independent processes, and ignoring any concerns regarding evasion, what can you say about the false positives and false negatives of H_B and H_E? In terms of S_P, S_N, A_P, A_N, what are the false positive and false negative rates for H_B and H_E.

Solution: The key insight here is that alarms by H_B will be a subset of the alarms generated by H_E. Since H_B will generate fewer alarms for non-malicious activities, it will have less false positives. On the other hand, because it generates fewer alarms, it might miss more malicious activity, implying more false negatives. The false positive rate for H_E would be $S_P + A_P - S_P A_P$, and for H_B would be $S_P A_P$. Similarly, the false negative rate for H_E would be $S_N A_N$ and for H_B would be $S_N + A_N - S_N A_N$

(b) If deploying the hybrid scheme in a new environment, is one of H_E and H_B clearly better? If not, what environment parameters would help determine whether H_E or H_B is better, and for each parameter p, increasing p favors which hybrid scheme?

Solution: In the absence of more data, particularly the cost of false positive and false negatives, as well as the rate of malicious and non-malicious activity, it is impossible to make any decision.

Increasing the cost of false-positives and the rate of non-malicious activity favors scheme H_B, while increasing the cost of false-negatives, and the rate of malicious activity favors scheme H_E.

Question 3 ***Proof of Work*** **(15 min)**

(a) Eve is buying a penguin from Alice. Eve generates and then sends Alice a valid transaction message, which transfers 100BTC to Alice's wallet. The signature is correct, and Eve has enough funds to make this transaction. Immediately upon receiving and verifying the transaction, Alice gives Eve the penguin. What attack could Eve do to avoid paying Alice the 100BTC?

Solution: Eve can do a double spend attack, where she generates another transaction sending the 100BTC to herself, and races Alice to get hers confirmed first.

(b) What can Alice do to make sure she's actually received the money?

Solution: Alice should wait for some blocks to be mined with her transaction in it. The more blocks she waits for, the more unlikely it is for her transaction to be orphaned (ignored).

(c) Alice tells Eve she will wait until the next block is mined to determine if the transaction went through. Given that the last block was mined 9 minutes ago, and a block is mined every 10 minutes on average, how long does Eve expect to wait?

Solution: 10 minutes. The last 9 minutes of mining do not affect the probability of a block being found in the future (think about the probability of flipping a heads, given you've flipped several tails in a row).

(d) Alice is unsure if waiting for the next block is secure enough. Let's say Eve controls a mining pool with a large fraction of the total network hash rate, and is trying to perform an attack similar to the one from part 1, even though Alice is now waiting for the next block to be mined with her transaction in it before releasing the penguin. What does Eve need to do to pull off the attack?

Solution: In order to perform the double spend, Eve needs to mine a chain where she sends the money to herself, and have it overtake the rest of the network. One option is to mine 2 consecutive blocks faster than the rest of the network does.

Whenever the rest of the network mines one block containing the transaction to Alice, this gets broadcast over the network, Alice is satisfied and gives Eve the penguin. Before the rest of the network is able to mine a second block, Eve's mining pool might have already mined two blocks, the first of which includes a transaction giving the money to Eve instead. When Eve broadcasts those two

blocks, they become the longest chain and the whole network will start working on this new chain and abandon the previous chain.