

Q1 *Intrusion Detection*

(8 points)

FooCorp is deciding which intrusion detection *method* to employ in a few target scenarios. In the following parts, consider which of the intrusion detection methods learned in class would be most appropriate (NIDS, HIDS, or logging), and justify why.

Q1.1 (1 point) FooCorp is hosting a web application over HTTPS and needs to detect any use of black-listed characters in real time.

Solution: FooCorp should use a HIDS. While the style of detection is ideal for a NIDS, it is useless in this specific scenario because TLS hinders a NIDS from reading the application data.

Q1.2 (1 point) FooCorp is hosting a web application over HTTP and wants to pass all user traffic through an anomaly detection algorithm (which uses some computationally expensive machine learning). The web application needs to have low latency when many users are online during the day.

Solution: FooCorp should use logging with overnight analysis. Using a NIDS or HIDS would cause the web application to have increased latency since running the analysis is expensive.

Q1.3 (1 point) FooCorp uses the Simple Mail Transfer Protocol (SMTP) for email and wants to be able to quickly detect phishing attacks against any of their internal computers. SMTP runs on port 25 and is unencrypted.

Solution: FooCorp should use a NIDS here. They only need one device to be able to monitor the whole corporation and it will be real-time detection.

Q1.4 (1 point) FooCorp doesn't trust its employees and sets-up a NIDS to monitor their traffic. However, many employees use TLS, hindering what can be monitored.

FooCorp decides to turn their NIDS into a *Man-in-the-Middle*, giving it a certificate that all the employee's computers trust. Whenever an employee visits a website they complete a TLS handshake with the NIDS, the NIDS connects to the requested website using TLS, and any traffic between the employee and website is forwarded across the two TLS links by the NIDS.

Which security principle does this violate? Describe everything an attacker can do if they compromise the NIDS.

Solution: Separation of privilege. The NIDS has complete control over the contents, integrity, and authenticity of any connection.

The only certificate being verified by the client is the NIDS's so an attacker can reroute hosts to a malicious website and they wouldn't be able to tell. In other words, all the guarantees of TLS are lost.

The NIDS essentially has the same power as a MiTM with HTTP. Except it's even worse since there's a false sense of security due to the presence of TLS.



(Question 1 continued...)

FooCorp now needs to decide which intrusion detection *technique* to employ in a few target scenarios. In the following parts, consider which technique would be most appropriate (signature-based, anomaly-based, specification-based, or behavioral), and justify why.

Q1.5 (1 point) FooCorp wants to detect script kiddies (hackers who primarily use publicly available tools or exploits).

Solution: Signature-based detection would be best here since you can compile a list of all the things the hacker might do.

Q1.6 (1 point) FooCorp wants to detect a seasoned hacker who crafts custom exploits for each attack.

Solution: Behavioral/anomaly-based detection will be best here since the hacker will most likely use clever, possibly novel, techniques to pull off the attack. They can most likely avoid writing code that violates specifications or matches known signatures.

Q1.7 (1 point) FooCorp wants to detect publicly-available malware that a hacker manually tweaks to avoid signature checks.

Solution: Behavioral detection will be best here since you have a good idea of how the attack works.

Q1.8 (1 point) FooCorp wants to detect any attempts by their employees to access the protected `/etc/passwd` file.

Solution: Specification detection is best here. Simply set up rules that flag any syscalls which access `/etc/passwd`.

Q2 *Low-level Denial of Service*

(2 points)

In this question, you will help Mallory develop new ways to conduct denial-of-service (DoS) attacks.

CHARGEN and ECHO are services provided by some UNIX servers. For every UDP packet arriving at port 19, CHARGEN sends back a packet with 0 to 512 random characters. For every UDP packet arriving at port 7, ECHO sends back a packet with the same content.

Mallory wants to perform a DoS attack on two servers. One with IP address A supports CHARGEN, and another with IP address B supports ECHO. Mallory can spoof IP addresses.

Q2.1 (1 point) Is it possible to create a single UDP packet with no content which will cause both servers to consume a large amount of bandwidth?

- If yes, mark 'Possible' and fill in the fields below to create this packet.
- If no, mark 'Impossible' and explain within the provided lines.

Possible

Impossible

If possible, fill in the fields:

Source IP: _____ Destination IP: _____

Source Port: _____ Destination Port: _____

If impossible, why?

Solution: Source IP: B , port: 7. Destination IP: A , port: 19. Source and destination can be flipped. Notice this will create a chain of CHARGEN and ECHO that will generate a lot of network traffic.

(Question 2 continued...)

Q2.2 (1 point) Assume now that CHARGEN and ECHO are now modified to only respond to TCP packets (post-handshake) and not UDP. Is it possible to create a single TCP SYN packet with no content which will cause both servers to consume a large amount of bandwidth? Assume Mallory is off-path from the two servers.

- If yes, mark 'Possible' and fill in the fields below to create this packet.
- If no, mark 'Impossible' and explain within the provided lines.

Possible

Impossible

If possible, fill in the fields:

Source IP: _____ Destination IP: _____

Source Port: _____ Destination Port: _____

Sequence #: _____ Ack #: N/A

If impossible, why?

Solution: Impossible. As seen in previous question, source/destination IP has to be B/A for the chain to work. If you send a SYN packet to A pretending to be B, A will send SYN-ACK to B, which won't respond since it never sent a SYN. The connection won't be established.