CS161 Summer 2025 — Introduction to Computer Security — Exam Prep 13

Q1 *Networking: A TORrible Mistake* **(7 points)**

Q1.1 (1 point) An $n > 1$-node Tor circuit provides anonymity (i.e. no node learns who both the user and server are) when at least _____ node(s) are honest. Assume that malicious nodes can collude, but they do not correlate traffic. Fill in the blank.

○ 0 ● 1 ○ n-1 ○ n

Solution: As seen in lecture, a Tor circuit is secure if at least one node is honest. Anonymity is only broken if every node in the circuit colludes, so that together they can reconstruct the entire circuit that messages are being routed through. Without traffic correlation, the honest node will cryptographically protect the Tor request.

For the next 3 subparts, a user is using Tor to send a message to a server. Assume that there is no collusion between any Tor nodes, and that the user choses exactly 3 nodes for their Tor circuit.

Q1.2 (1 point) Which values can a malicious **entry** node learn? Select all that apply.

■ The IP address of the user ☐ The list of all nodes in the circuit

☐ The IP address of the server ☐ None of the above

Solution: The user sends messages to the entry node, telling the entry node to forward those messages to the next node.

The IP address of the server is wrapped in many layers of encryption inside the message sent to the entry node, so the entry node cannot see that value.

The entry node knows about the second node in the circuit, but not the entire list of nodes.

Q1.3 (1 point) Which values can a malicious **exit** node learn? Select all that apply.

☐ The IP address of the user ☐ The list of all nodes in the circuit

■ The IP address of the server ☐ None of the above

Solution: The exit node is the last node in the circuit, who needs to know the server's identity so that they can forward the message to the server.

By the time the message reaches the exit node, all information about the original user's identity has been stripped away (the entry node removed all traces of the original user's identity when forwarding the packet to the second node).

The exit node knows about the second-to-last node in the circuit, but not the entire list of nodes.

Q1.4 (1 point) Which values can an on-path attacker on the user's local network learn? Select all that apply.

- [x] The IP address of the user
- [ ] The list of all nodes in the circuit
- [ ] The IP address of the server
- [ ] None of the above

Solution: The on-path attacker in the local network can see the user sending messages into the Tor network (to the entry node).

However, the IP address of the server is encrypted inside the message sent to the entry node, so the on-path attacker cannot see that value.

The on-path attacker only knows about the entry node, not the entire list of nodes in the circuit.

When a new user first downloads Tor, they need to download a list of nodes from a trusted directory server.

A malicious, on-path attacker on the user's local network wishes to eavesdrop on the new user's Tor connection. Assume that the attacker controls 3 nodes out of 100 total Tor nodes, and can win any data race.

For the next three subparts, select the approximate probability that the attacker can learn the identity of the server.

Q1.5 (1 point) User connects to the directory via TLS, attacker is on-path.

- ○ Exactly 0%
- ○ Greater than 50%, less than 100%
- ● Greater than 0%, less than 50%
- ○ Exactly 100%

Solution: Because the directory connection is made over TLS, and TLS has end-to-end security, the on-path attacker cannot tamper with the list of nodes.

Therefore, the on-path attacker can only hope that the user randomly selects the three nodes controlled by the attacker.

The probability of selecting the 3 attacker-controlled nodes out of 100 nodes is intuitively less than 50%, but it's not 0%.

Formally, you can calculate this probability to be $\frac{3!}{100 \cdot 99 \cdot 98}$, where the numerator is the number of ordered ways to choose the 3 attacker nodes (counting all possible orders, since order doesn't matter), and the denominator is the number of ordered ways to choose any 3 nodes.

Q1.6 (1 point) User connects to the directory via TCP, attacker is on-path.

- ○ Exactly 0%
- ○ Greater than 50%, less than 100%
- ○ Greater than 0%, less than 50%
- ● Exactly 100%

Solution: Unlike the last subpart, the user is now using just TCP to connect to the directory, so the attacker can tamper with the response from the directory.

Specifically, the attacker can trick the user into thinking that the list of nodes only has 3 nodes: the attacker-controlled nodes.

Now, the user is forced to always choose the attacker-controlled nodes, and the attacker will always be able to break anonymity by controlling every node in the resulting circuit.

Note that we don't have to worry about data races, since the question says the attacker can win any data race.

Q1.7 (1 point) User connects to the directory via TCP, attacker is off-path.

- ○ Exactly 0%
- ○ Greater than 50%, less than 100%
- ● Greater than 0%, less than 50%
- ○ Exactly 100%

Solution: As in the previous subpart, the attacker can trick the user into using the attacker's nodes.

However, because the attacker is now off-path, they need to guess the sequence number in order to inject a malicious message into the TCP connection. The probability of the attacker guessing a valid 32-bit sequence number is under 50% (but not 0%).

Q2 *Making New Friends* **(9 points)**

Consider two local broadcast networks, as shown in the diagram below.

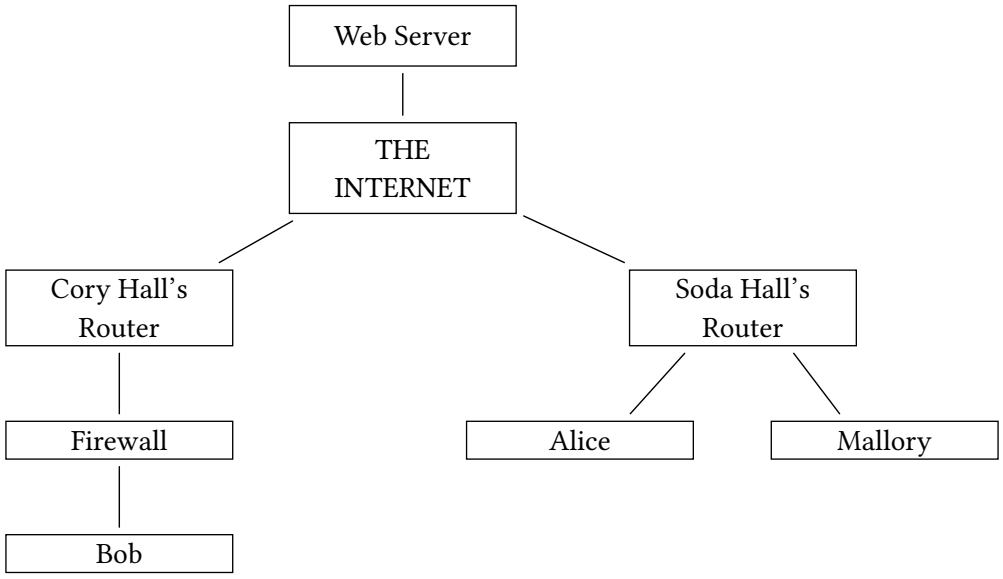

Q2.1 (2 points) Alice broadcasts an ARP request for Mallory's MAC address.

Which of these entities, if malicious, can poison Alice's ARP cache? Select all that apply.

- [x] Mallory
- [ ] Bob
- [ ] None of the above
- [x] Soda Hall's router
- [ ] Cory Hall's router

Solution: ARP is a local network protocol. The ARP broadcast is only sent to users on the local network, so only users on the local network can spoof an ARP response.

Q2.2 (4 points) Mallory and Bob form a TLS connection. Then, Bob adds a rule to the firewall disallowing all inbound packets from Mallory.

EvanBot argues that TLS messages are encrypted, so the firewall cannot stop Mallory from sending more TLS messages to Bob. Is EvanBot correct? Justify your answer in 10 words or fewer.

- ○ Yes
- ● No

Solution: No, because the IP header is not encrypted. TLS does not provide anonymity/availability.

Q2.3 (3 points) Bob adds a rule to the firewall disallowing all inbound packets from anybody in Soda Hall's local network.

Which of the following attacks can Mallory still perform on Bob? Assume that Mallory cannot spoof packets. Select all that apply.

- [x] DoS
- [ ] TLS Hijacking
- [x] XSS
- [ ] None of the above

Solution: Mallory could DoS Bob by overwhelming the firewall.

Mallory could perform a stored XSS attack on Bob by storing malicious JavaScript on an external web server. Bob then loads the JavaScript from the web server, not Mallory.

Mallory cannot hijack TLS regardless of the firewall.