CS161 Summer 2026 — Introduction to Computer Security — Discussion 8

Q1 *ARP* **(4 points)**

Recall that ARP, the Address Resolution Protocol, translates Layer 3 IP addresses into Layer 2 MAC addresses.

For this part, imagine that Alice has successfully obtained a configuration from the network's router and now she wants to communicate with Bob, whose computer is on the same LAN network.

Alice knows Bob's IP address but wants to learn his MAC address. You want to convince Alice that your MAC address (and not Bob's) corresponds to Bob's IP address, causing messages intended for Bob to be sent to you instead. You, Alice, and Bob are part of the same LAN network and, apart from the router, there are no other machines on the network.

Assume that:

- Your computer has IP address 10.10.10.66 and your MAC address is 66:66:66:66:66:66
- Alice's IP address is 10.10.10.77 and her MAC address is 77:77:77:77:77:77
- Bob's IP address is 10.10.10.88 and his MAC address is 88:88:88:88:88:88
- The network router's IP address is 10.10.10.5 and its MAC address is 99:99:99:99:99:99

Q1.1 (1 point) Alice broadcasts to everyone else on the LAN: "What is the MAC address of 10.10.10.88?" Recall that 10.10.10.88 corresponds to Bob's IP address.

What values for the IP and MAC address could you include in your response to Alice to cause her messages intended for Bob to be sent to you instead?

Solution: You want to convince Alice that your MAC address corresponds to Bob's IP address. Therefore, the only correct answer is 10.10.10.88 as the IP address (Bob's IP) and 66:66:66:66:66:66 as the MAC address (your MAC address).

Q1.2 (1 point) How would your spoofed response to Alice change if Bob was outside the LAN that you and Alice are both part of?

Solution: If Bob is outside of the LAN, Alice would instead generate an ARP request for the gateway router (once Alice's ARP cache entry for the gateway router expires or she must re-query the router's MAC address for some other reason, if she's already obtained the router's MAC). The router would respond with its MAC address and forward messages to some other LAN to get it closer to Bob.

Therefore, you are now instead trying to convince Alice that your MAC address corresponds to the router's IP address. For this reason, the only correct answer is 10.10.10.5 as the IP address (the router's IP) and 66:66:66:66:66:66 as the MAC address (your MAC address).

Q1.3 (1 point) Which defenses exist against such an ARP-spoofing attack?

Solution: A simple defense against ARP spoofing is to use a tool like arpwatch, which tracks the IP address to MAC address pairings across the LAN and makes sure nothing suspicious happens.

Modern wired Ethernet networks defend against ARP spoofing by using switches rather than hubs. Switches have a MAC cache, which keeps track of the IP address to MAC address pairings. If the packet's IP address has a known MAC in the cache, the switch just sends it to the MAC. Otherwise, it broadcasts the packet to everyone. Smarter switches can filter requests so that not every request is broadcast to everyone.

Higher-quality switches include VLANs (Virtual Local Area Networks), which implement isolation by breaking the network into separate virtual networks.

Q1.4 (1 point) You decide to use a network switch to prevent further ARP spoof attacks.

Explain how a network switch prevents such attacks in the following two cases: (1) the switch knows Bob's IP to MAC address mapping and (2) the switch does **not** know Bob's IP to MAC address mapping.

Solution: Switches have a MAC cache, which keeps track of the IP address to MAC address pairings.

(1) If the packet's IP address has a known MAC in the cache, the switch just sends the packet to the known MAC address.

(2) If the switch does not know the mapping, it must broadcast the packet to everyone. Plain switches have no other lines of defense, but advanced switches can, for example, block ARP responses from unauthorized or suspicious devices.

Q2 *Attack on TCP* **(4 points)**

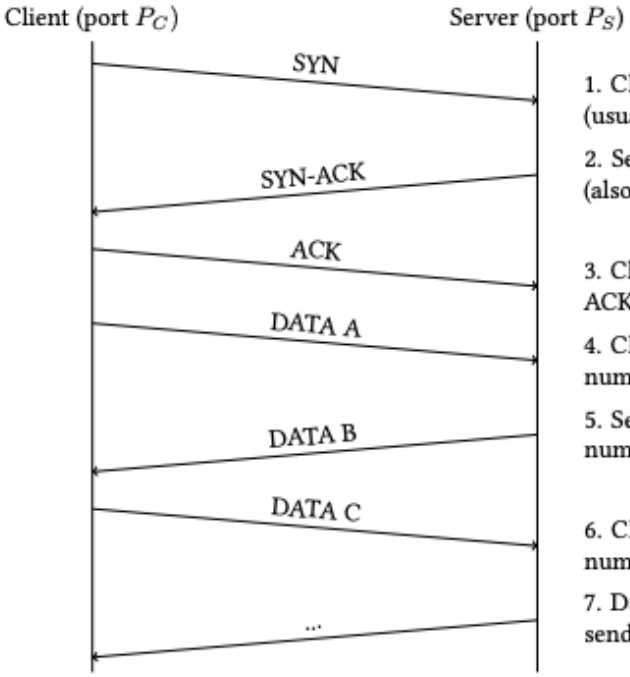

1. Client sends initial SYN with sequence number X (usually random).
2. Server sends SYN-ACK with sequence number Y (also usually random) and ACK $X + 1$.
3. Client sends ACK with sequence number $X + 1$ and ACK $Y + 1$.
4. Client sends DATA A of length L_A with sequence number $X + 1$ and ACK $Y + 1$.
5. Server sends DATA B of length L_B with sequence number $Y + 1$ and ACK $X + 1 + L_A$.
6. Client sends DATA C of length L_C with sequence number $X + 1 + L_A$ and ACK $Y + 1 + L_B$.
7. Data exchange continues until both sides are done sending data.

Q2.1 (1 point) Assume that the next transmission in this connection will be DATA D from the server to the client. What will this packet look like?

Sequence Number:	_____	ACK:	_____
Source Port:	P_S	Destination Port:	P_C
Length:	L_D	Flags:	ACK

Solution: Source Port: $Y + 1 + L_B$

ACK: $X + 1 + L_A + L_C$

Q2.2 (1 point) You should be familiar with the concept and capabilities of a *man-in-the-middle* as an attacker who **can observe** and **can modify** traffic. There are two other types of relevant attackers in this scenario:

1. *On-path* attacker: **can observe** traffic but **cannot modify** it.
2. *Off-path* attacker: **cannot observe** traffic and **cannot modify** it.

Carol is an *on-path* attacker. Can Carol do anything malicious to the connection? If so, what can she do?

Solution: Yes, Carol can leverage the information she learns from your traffic to hijack the session.

In part (a), we identified the values of all of the fields of concern expected in the next data transmission in the connection. Say Carol wants to spoof a packet from the server to the client; Carol can create a packet with the source IP as the server's IP, the destination IP as the client's IP, and the payload as whatever she wants. To spoof traffic in the other direction, she swaps the sequence number/ACK, source port/destination port, and source IP/destination IP. The recipient of this data cannot distinguish it from legitimate traffic, so she has effectively hijacked the session from her victim, allowing her to inject arbitrary data.

Q2.3 (1 point) David is an *off-path* attacker. Can David do anything malicious to the connection? If so, what can he do?

Solution: No, there isn't much he can do.

In part (b), we demonstrated that we are effectively defenseless against an attacker that knows the *sequence numbers* and the *port numbers* of the connection. An off-path attacker, however, does not have the power to observe the traffic and find these parameters.

Even without prior knowledge of these parameters, though, an *off-path* attacker may attempt to guess them. In a typical TCP client-server connection, the client's port is an *ephemeral* port, with a maximum potential range of $\left[0, 2^{\{16\}} - 1\right]$ (this varies, so we make an overestimation). The server's port is usually a *well-known* port for a specific service, such as port 80 for HTTP, which makes it much easier to guess. The sequence number and acknowledgement numbers are in the range $\left[0, 2^{\{32\}} - 1\right]$. However, an attacker can just ignore the acknowledgement number or use a random number. The TCP connection will not be reset with an invalid acknowledgement number as long as it's within the window (out of scope). Thus, an attacker has a rough $\frac{1}{2^{48}}$ chance of successfully brute forcing the correct parameters.

If, for some reason, the initial sequence numbers are not properly randomized, David may be able to make educated guesses on the sequence numbers and significantly decrease the range of possibilities. However, assuming that they are properly randomized, this attack is theoretically possible but largely improbable.

We call this attack *blind hijacking*, as David has no concrete information when attempting to hijack the session.

Q2.4 (1 point) The client starts getting responses from the server that don't make any sense. Inferring that David is attempting to hijack the connection, the client then immediately sends the server a **RST** packet, which terminates the ongoing connection. David wants to impersonate the client by establishing a new connection. How would he go about doing this?

Solution: If David attempts to start a new connection, he can *choose* the source port (the *ephemeral* port) and the source sequence number to be whatever values he wants. The values of these parameters for any subsequent transmissions in the connection will then be predictable. The server's port remains a well-known port; the only remaining unknown is the server's sequence number. Because David is still an off-path attacker, he still has to guess this field, with an overall probability of $\frac{1}{2^{32}}$ of success, which we note is higher than the emph{blind hijacking} approach.

Note that there's now a time constraint on David's attack: if the client receives a response from the server based on his spoofed *SYN*, it will send a **RST** and terminate the connection, putting David back at step 1.

We call this attack *blind spoofing*, as David has no concrete information when attempting to spoof a new session.

Q3 *TLS Protocol Details* **(4 points)**

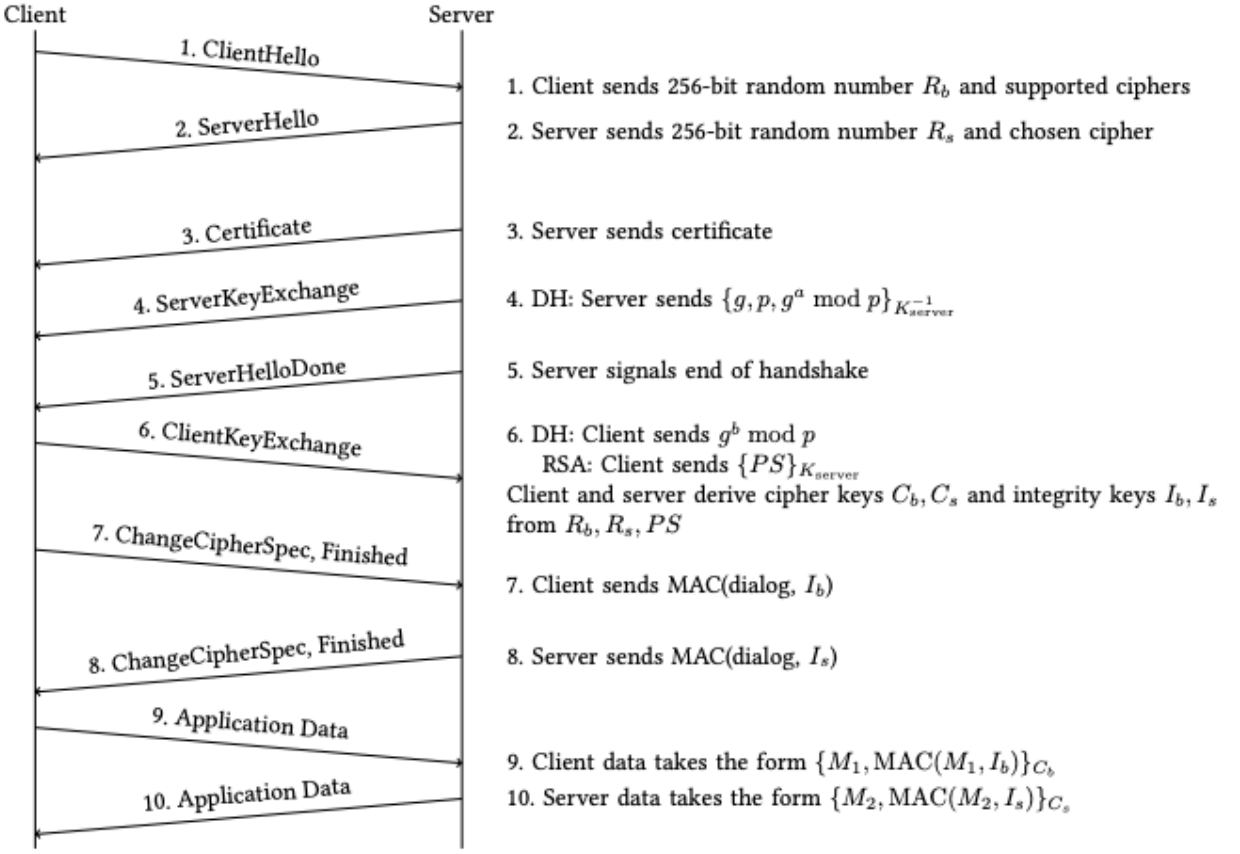

Figure 1: TLS 1.2 Key Exchange

Q3.1 (1 point) What is the purpose of the *client random* and *server random* fields?

Solution: They act as nonces to prevent replay attacks.

Q3.2 (1 point)

ClientHello and ServerHello are not encrypted or authenticated. Explain why a man-in-the-middle cannot exploit this. (Consider both the Diffie-Hellman and RSA case.)

Solution: The use of either public key encryption (RSA handshake) or a Diffie-Hellman exchange prevents an eavesdropper from learning the Premaster Secret.

A MITM attacker who alters any of the values will be exposed, as follows.

For the RSA handshake case, the MITM will be unable to read the Premaster Secret sent by the client because it is encrypted using the server's public key. When the client and server exchange MACs over the the handshake dialog, the MITM will be unable to compute the correct MACs for their altered dialog because they will not know the corresponding integrity keys derived from the Master secret.

For the Diffie-Hellman case, the MITM will be unable to alter the value of $g^a \bmod p$, because the client requires that the value have a correct signature using the server's public key. Because the MITM cannot alter the value, they cannot substitute $g^{a'} \bmod p$ for which they know a'. Without knowledge of the exponent, the MITM cannot compute $g^{ab} \bmod p$ to obtain the Premaster Secret.

Q3.3 (1 point) Note that in the TLS protocol presented above, there are two cipher keys C_b and C_s. One key is used only by the client, and the other is used only by the server. Likewise, there are two integrity keys I_b and I_s. Alice proposes that both the server and the client should simply share one cipher key C and one integrity key I. Why might this be a bad idea?

Solution: Vulnerable to **reflection** attacks: a man-in-the-middle can send a client's application data back to them. It will still verify the appropriate checks, but the user will think that this is the response from the server. Likewise, an attacker can reflect a server's response back to the client. Note that due to the existence of sequence numbers in the TLS specification, the actual attack would be a little more complicated.

Q3.4 (1 point)

The protocol given above is a simplified form of what actually happens. After step 8 (ChangeCipherSpec), the protocol as described above is still vulnerable. What is the vulnerability and how could you fix this?

Solution: An attacker can perform a **replay** attack, where they send the same application data twice. The solution is to add sequence numbers (which is what the actual TLS specification does, with some extra details involved).