Network Security Transport Layer

Target Course

Networks

Learning Goals

A student shall be able to:

1. Describe foundational security concepts in securing networks and systems.
2. Describe security design principles and identify security issues associated with common threats and attacks.

IAS Outcomes

IAS Knowledge Topic	Outcome
Network Security	3. Describe virtues and limitations of security technologies at each layer of the network stack. [Familiarity] 4. 4. Identify the appropriate defense mechanism(s) and its limitations given a network threat. [Familiarity]

Dependencies

- Cover after the ***Network Security Concepts*** module.

Summary

Describe how the transport layer may be used to support the security goals of CIA and the fundamental concepts of assurance, authentication, anonymity, and non-repudiation.

Estimated Time

This module took approximately one lecture hour to cover.

Materials

How does this layer affect the security goal of confidentiality?

- TCP and UDP do not automatically encrypt their payload.
- TLS (Transport Layer Security), and the older protocol SSL (Secure Sockets Layer), both use cryptography to have both hosts agree to a shared secret which is then used to generate a unique symmetric key used by both hosts.

How does this layer affect the security goal of integrity?

- TCP and UDP uses 16-bit checksum designed to catch transmission errors.
    - This checksum is not cryptographically secure, so this checksum does not provide integrity from the perspective of computer security.
- TLS will ensure integrity of the data that is encrypted before being sent and can be decrypted only by the receiving host.

How does this layer affect the security goal of availability?

- TCP has flow control using a sliding window protocol.
- TCP has congestion control to reduce possibility of overwhelming a network.
- UDP does not have flow control or congestion control.

How does this layer affect the fundamental security concept of assurance?

- TCP and UDP protocols allow packets to be sent between any two devices.
- TCP and UDP protocols do not include any permissions or security policies (e.g., similar to firewall capabilities).
- TCP session hijacking allows an attacker to pretend they are someone else.

How does this layer affect the fundamental security concept of authenticity?

- TCP and UDP protocols do not include any type of digital signature. These protocols have no notion of user identity. While an IP/port is associated with a device, any type of user could be using this device.
- TCP session hijacking allows an attacker to pretend they are someone else.

How does this layer affect the fundamental security concept of anonymity?

- TCP and UDP protocols do not include any type of digital signature. These protocols have no notion of user identity. Thus, transport layer supports anonymity - which is a two-edged sword since an attacker may pretend they are someone else without attribution.

How does this layer affect the fundamental security concept of non-repudiation?

- Since TCP and UDP have no notion of user identity, non-repudiation is not supported.

What type of risks are known about the Transport layer?

The information below is from Chapter 14 in [1] and Chapter 1 in [2].

The Transport layer general risks include the following:

- Transport layer hijacking
  - Attacker focuses on sequence numbers and port numbers
  - Performs some type of network layer compromise e.g., using promiscuous mode, simple address impersonation, or MitM (Man-in-the-middle)
  - Must identify the transport sequencing
  - Must impersonate network layer traffic
- Servers should have minimum number of ports “open”
  - TCP & UDP: port #’s below 1024 are reserved
- Static vs dynamic port assignment
  - Client connection initially made to known server port number
  - Client port number may be dynamic (selected from range of numbers); client must include dynamic port number in its request
  - Dynamic ports used when app spawn’s processes for managing network traffic. This creates a security risk since large port range must be accessible to network
  - Firewalls
    - Configure port numbers to allow or prevent network access
    - Can permit all higher (> 1023) port numbers
    - When using only static ports, can prevent access to unused ports
- Port scans
  - Attacker looking for an “open” port by doing either targeted port scans or a port sweep
    - Targeted port scans - Scan same port number across range of IP addresses
    - Port sweep - Scan all port numbers for same IP address
  - To mitigate a port scan:
    - Use nonstandard port numbers
    - Use a “no reply” defense e.g., BSD systems do not reply to packet requests when port is inactive
    - Use an “always reply” defense i.e., Have system reply to every packet request, whether port is active or not

The UDP risks include the following:

- An un-validated inbound source. Any host can connect to UDP server e.g., any type of UDP packet can potentially flood a server.
- UDP hijacking. Since UDP packets do not have sequence numbers, can guess port number (only 65,535; takes a few seconds)

- UDP keep-alive attack. UDP server ports closed after period of inactivity e.g., attacker can hold open a UDP port; tries to keep open many ports, possibly preventing other ports from opening
- UDP smurf attack. Flood remote network with packets e.g., attacker forges victim's network address as sender

The TCP risks include the following:

- TCP reconnaissance. This may include doing any of the following.
  - Operating system profiling, to determine the OS and its patch level.
  - Port scans, to attempt to connect to a port. A host can reply in one of four ways:
    - SYN-ACK - positive identification that service running on port.
    - RST - typically confirms no service on port.
    - ICMP unreachable - indicates failure to reach host/server.
    - Reply with nothing - cannot determine status of port.
  - The mitigation for this is to log network activity e.g., connection requests.
- TCP hijacking, which is any attack that interferes with a TCP connection.
  - Full session hijacking. The attacker tells client to disconnect but then acts as client to the server. This is fairly rare, typically requires attacker to have direct link layer access
  - ICMP (Internet Control Message Protocol) is used to report unsuccessful connections. This can be used maliciously to redirect TCP connections to different ports.
- TCP DoS (Denial of Service). This include any of the following.
  - SYN attacks. Send large number of SYN packets in order to consume all available memory on server.
  - RST and FIN attacks. Abnormally terminate a connection.
  - ICMP attacks. Use to terminate a connection.
  - LAND attacks. Send SYN to server where the packet source IP address and port matches server's address and port i.es, Server is in a feedback loop.
- To mitigate these attacks, one or more of the following options may be deployed.
  - Alter the system profile e.g., change SYN timeout, retry counts, retry durations, initial window size, available TCP options, initial sequence values
  - Block attack vectors by using a firewall.
  - Identify network devices since some devices may be more vulnerable to certain types of attacks.
  - Stateful packet inspection. Track state of TCP connections; reject packets that do not match known state e.g., silently drop an RST sent to a closed port.
  - Use an intrusion detection system to monitor network for nonstandard or unexpected packets.
  - Use an intrusion prevention system to actively disable attack vectors.
  - Ensure the app-layer should authenticate traffic and detect potential attacks.

Assessment Methods

None used.

References

[1] Krawetz, N. (2007). Introduction to Network Security. Cengage Charles River Media. Accessed via Books 24x7 Digital Library.

[2] Xiao, Y. & Pan, Y, eds, (2007). Security in Distributed and Networking Systems: Computer and Network Security, Vol. 1. World Scientific Publishing Company. Accessed via Books 24x7 Digital Library.