Network Security

Can we trust the connection?

ITN 4252 Advanced Topics in Networking

Partially based on PowerPoint slides for Computer Networking: A Top-Down Approach, 8th edition, Jim Kurose, Keith Ross, Pearson, 2020



What makes a network connection secure?

- A secure connection protects four things
  - Confidentiality: Others cannot read the data
  - Integrity: Changes can be detected
  - Authentication: Communicating parties can establish who they are
  - Availability: Services remain accessible when required
- **Security is not one control, these properties work together.**

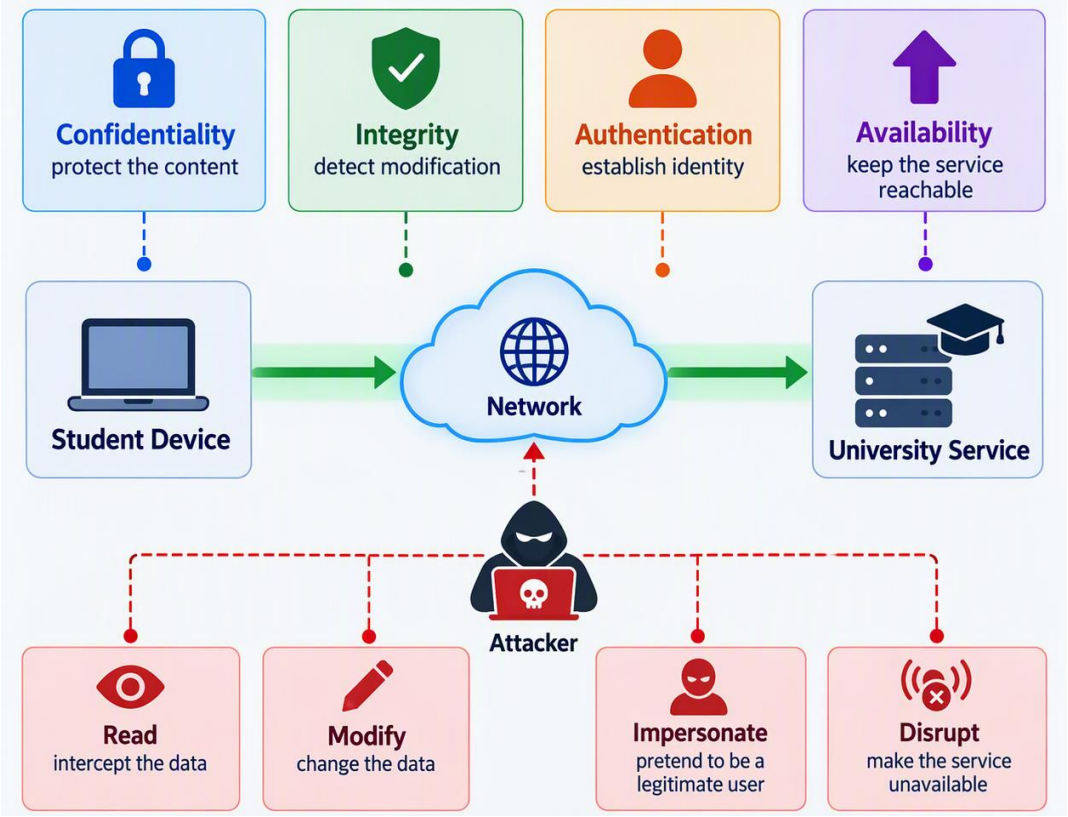

The cryptographic building blocks

- Modern network security combines several tools
  - Symmetric encryption: Protects data efficiently
  - Public-key cryptography: Supports authentication and key establishment
  - Hashing & MACs: Help protect integrity
  - Digital signatures: Prove origin and detect modification
  - Certificates: Bind identities to public keys
- **Secure protocols combine these building blocks, they do not rely on one alone.**

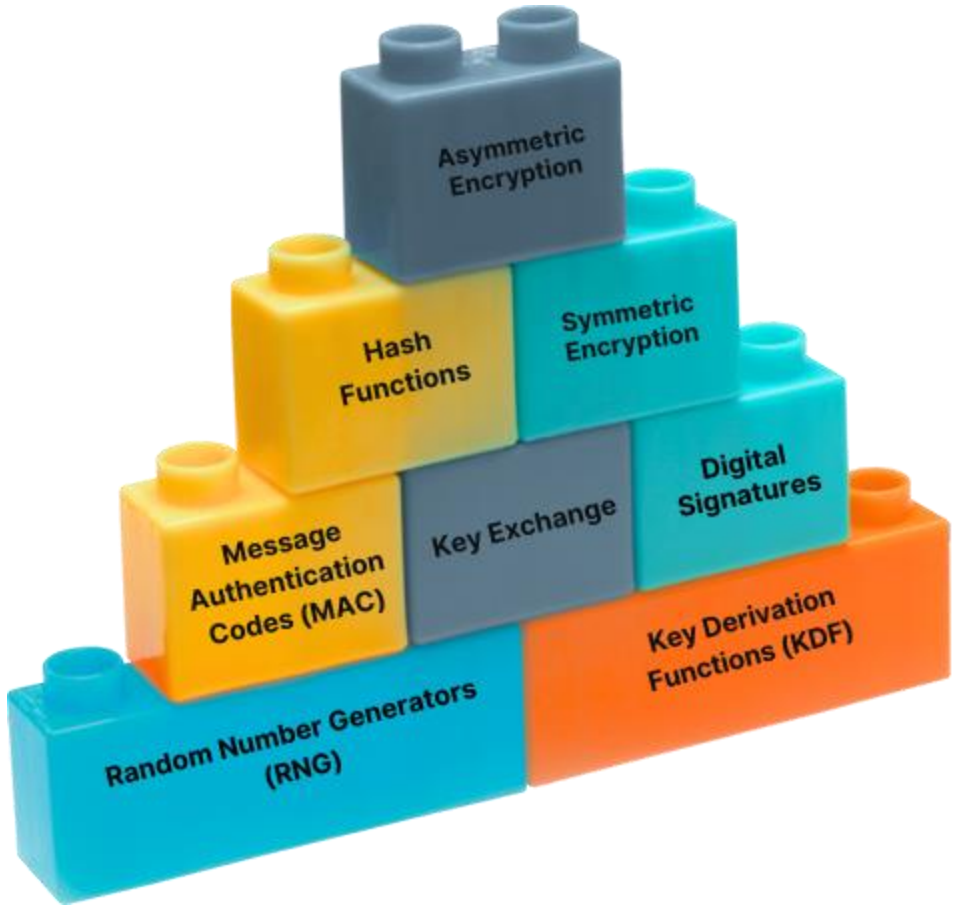

Certificates establish trust in public keys

- How does your browser trust the university website?
  - The server presents a digital certificate
  - The certificate binds a domain name to a public key
  - A trusted Certificate Authority (CA) signs the certificate
  - The browser validates the certificate before trusting the connection
- **A public key is useful only when you can trust who it belongs to.**

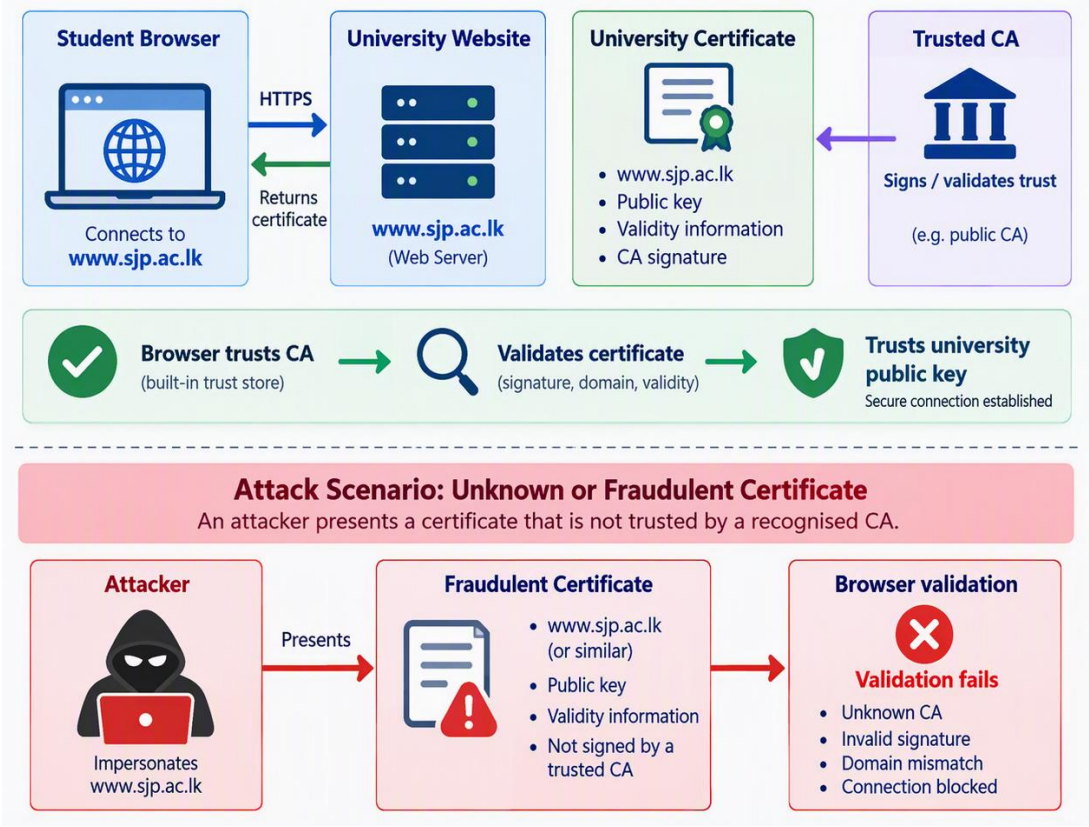

TLS protects data in transit

- TLS creates a secure channel over an untrusted network
- TLS provides:
  - Authentication: Verify the server
  - Key establishment: Create shared session keys
  - Confidentiality: Encrypt application data
  - Integrity: Detect unauthorised modification
- **HTTPS = HTTP protected by TLS**

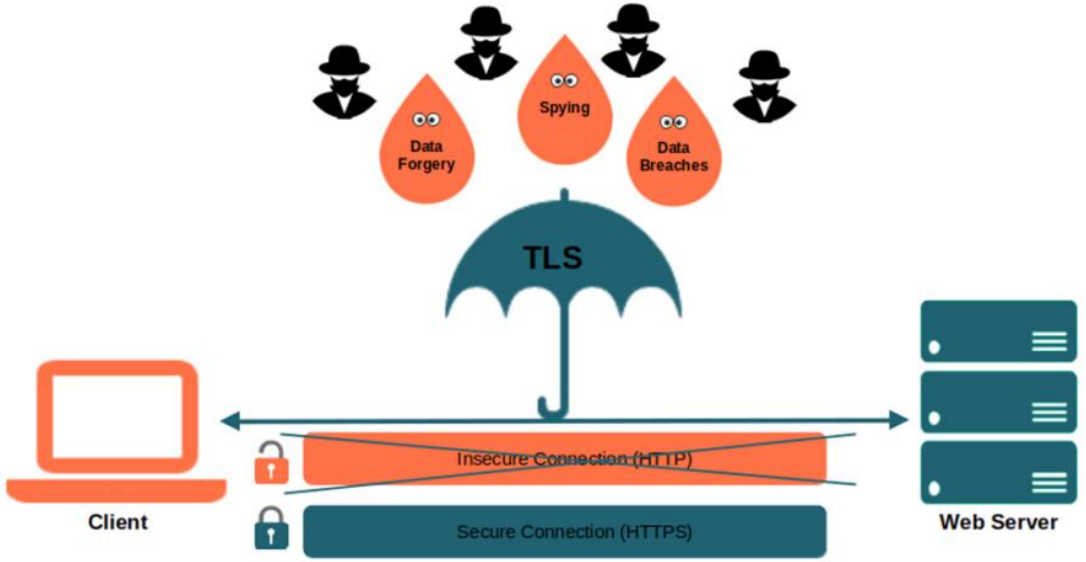

How TLS 1.3 establishes a secure connection

- TLS establishes trust before protecting data
  1. Client Hello: Proposes TLS options and key-share information
  2. Server Hello: Selects parameters and contributes its key share
  3. Certificate: Server proves its identity
  4. Key derivation: Both sides derive shared session keys
  5. Encrypted traffic: Application data is protected
- **Authenticate → Establish keys → Protect the session**

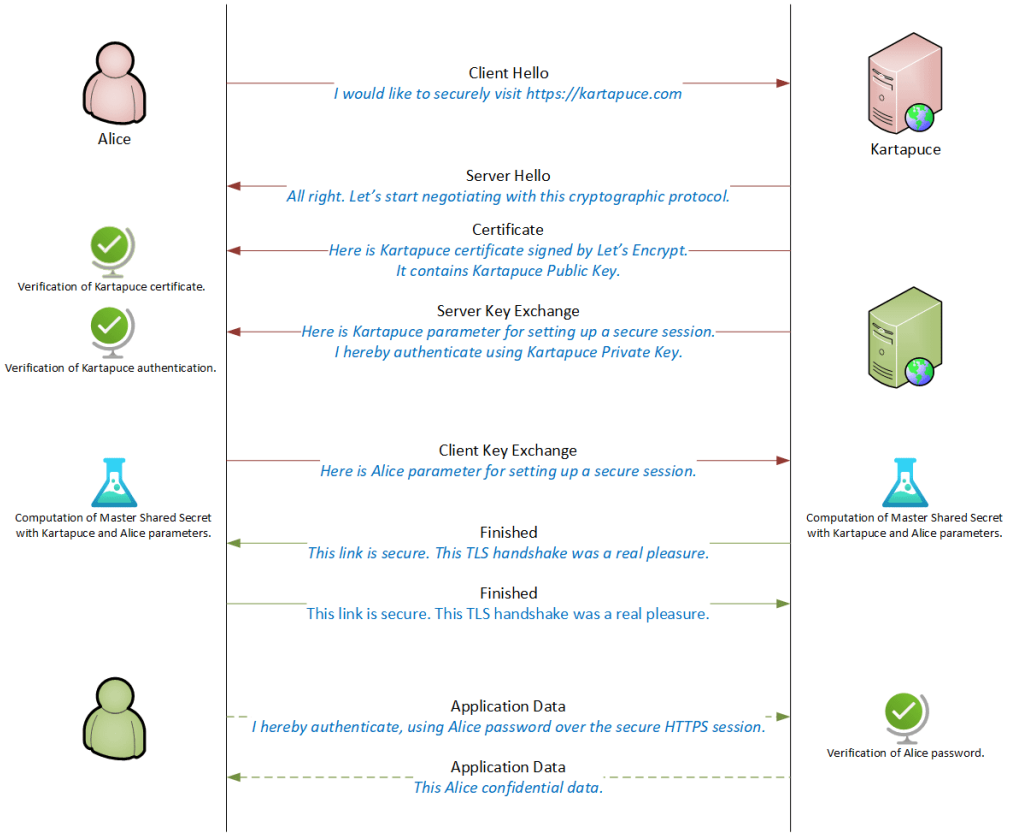

What TLS does, and does not, protect

- TLS protects the connection, not everything around it
- TLS protects
  - Data between TLS endpoints
  - Data from eavesdropping and modification
  - The authenticated server identity
- TLS does not protect
  - A compromised endpoint or application
  - Data before encryption or after decryption
  - Users from every malicious or deceptive website
- **A secure connection does not guarantee a secure service.**

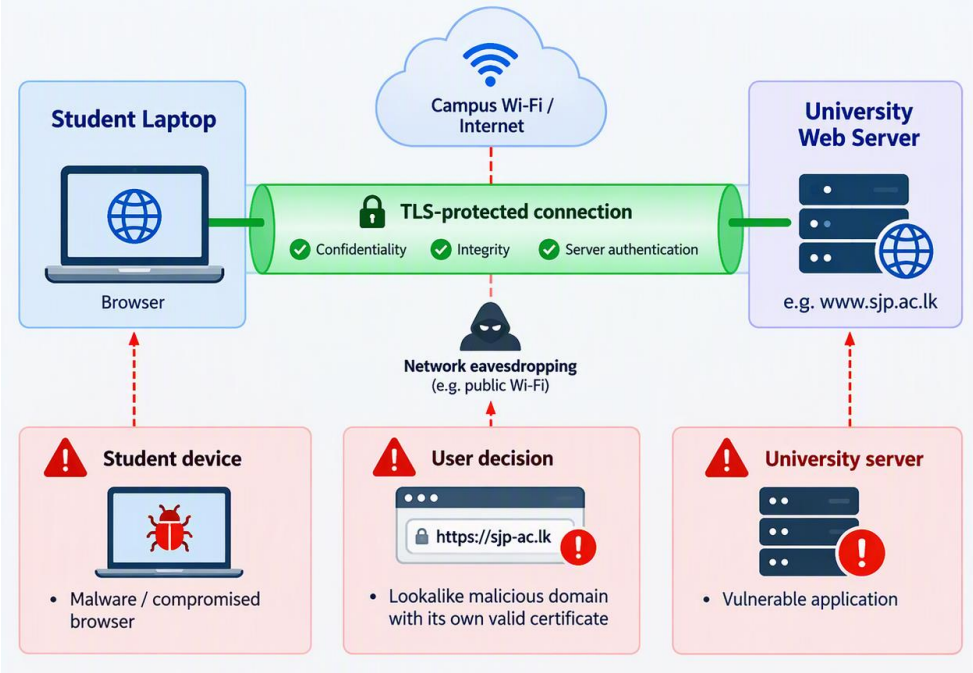

TLS is not the only place to protect traffic

- Security can operate at different network layers

Protection	Where it operates	Typical use
TLS	Application transport	HTTPS, APIs, email
IPsec	Network layer	VPNs, site-to-site links
Wi-Fi security	Wireless link	Campus wireless access

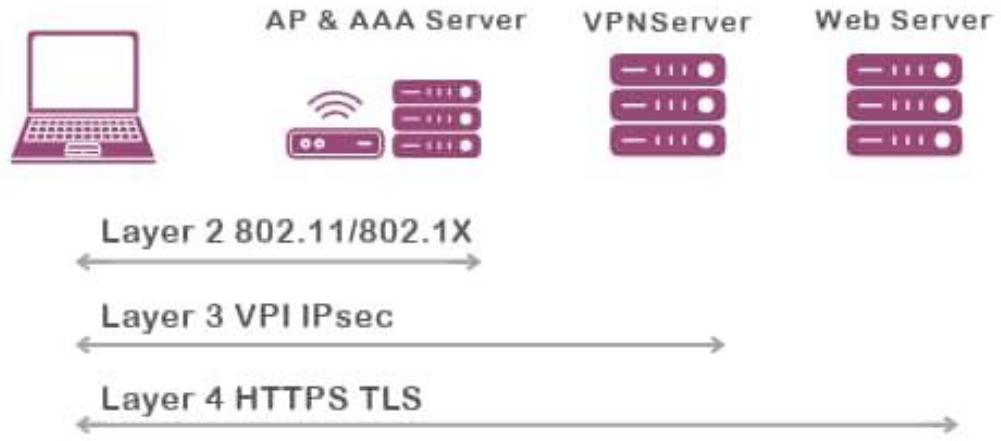

- **Choose the protection based on what part of the communication path you need to secure.**

Email has a different trust problem

- Email crosses multiple trust boundaries
  - Sender → Sending mail server → Internet → Receiving mail server → Recipient
- Security questions:
  - Who really sent the message?
  - Was the sending domain authorised?
  - Was the message altered?
  - Was transport between mail servers protected?
- **Receiving an email does not prove the sender is genuine.**

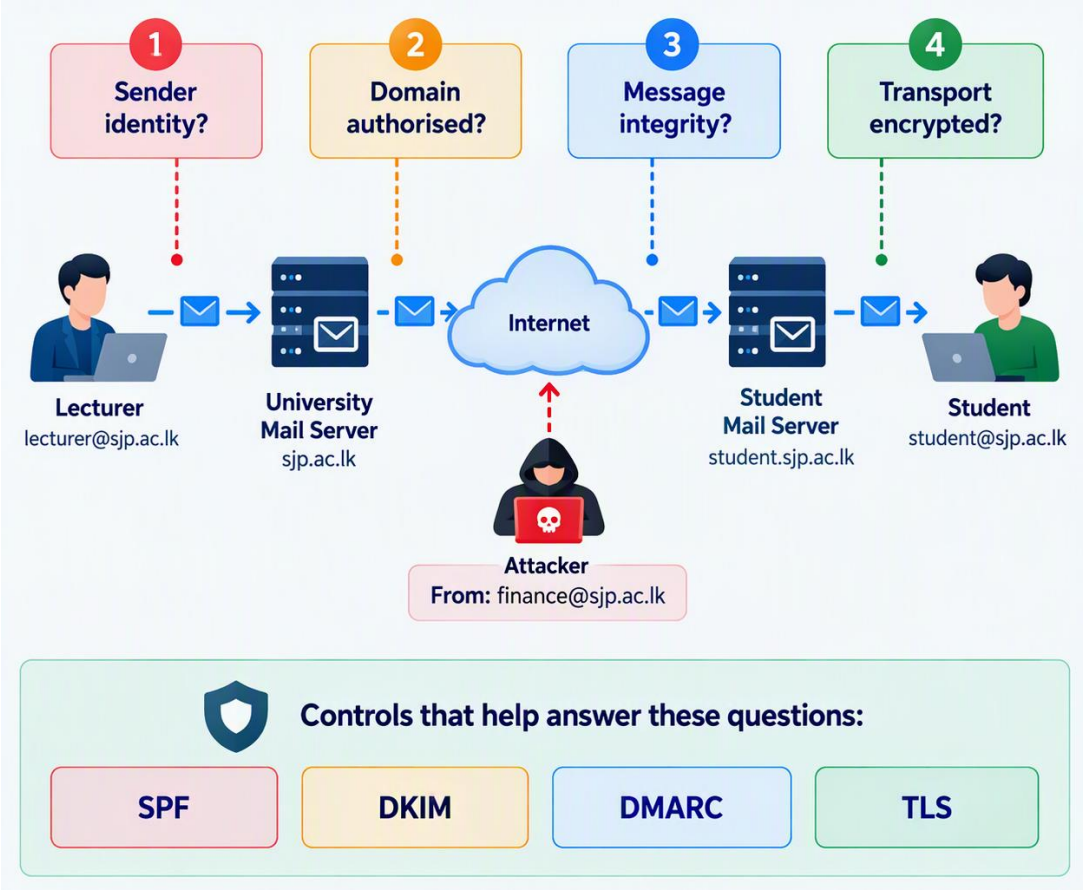

TLS protects email in transit, not the sender

- Email encryption and email authenticity are different problems
- TLS protects transport
  - Encrypts SMTP traffic between mail systems
  - Reduces eavesdropping and modification in transit
- But TLS does not prove
  - The visible From: address is genuine
  - The sender was authorised by that domain
- **A protected delivery path does not make the message trustworthy.**



SPF: Who is allowed to send for this domain?

- SPF identifies authorised sending systems
- The domain owner publishes an SPF record in DNS
- It tells receiving mail systems:
  - Which servers may send for the domain
  - Whether the sending server is authorised
  - How unauthorised sources should be treated
- **SPF asks: “Is this server allowed to send for this domain?”**

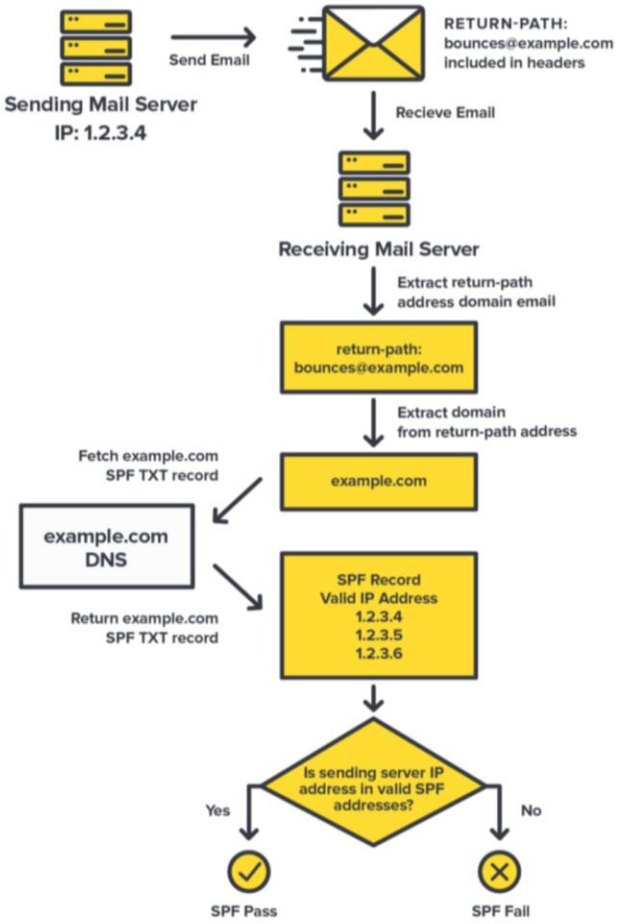

Diagram from postmarkapp.com

DKIM: Was the message signed by the domain?

- DKIM adds a cryptographic signature to email
- The sending mail system:
  - Signs selected message content with a private key
  - Publishes the matching public key in DNS
  - Allows the receiver to verify the signature
  - Helps detect modification of signed content
- **DKIM asks: “Did a domain take responsibility for signing this message?”**

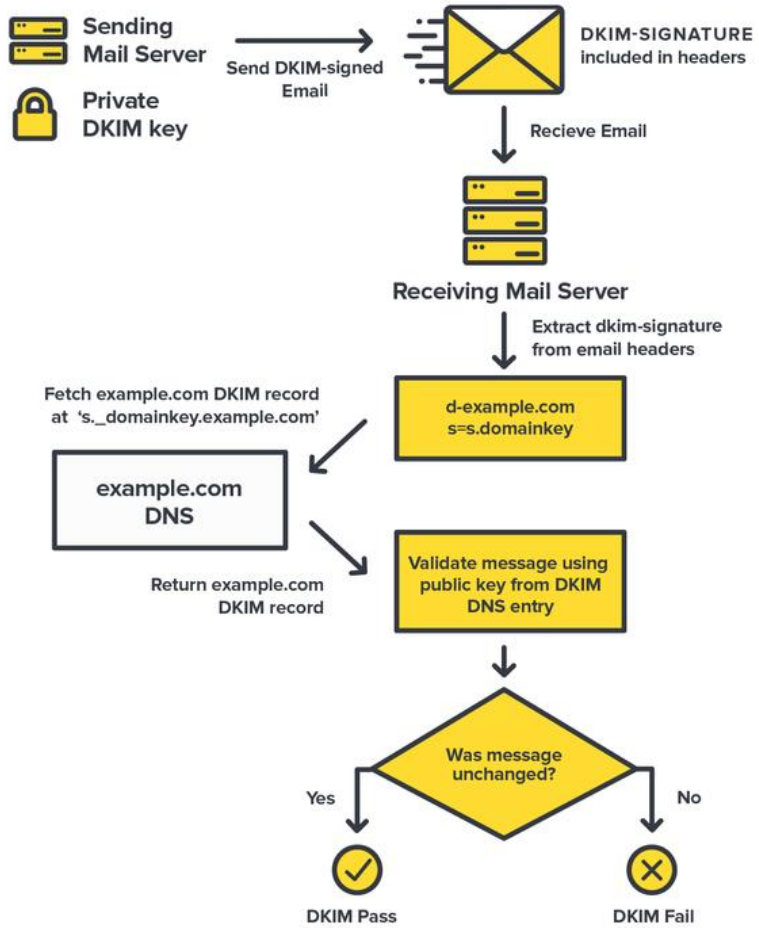

Diagram from postmarkapp.com

DMARC: does the authentication match the visible sender?

- DMARC connects authentication to the sender users see
- DMARC checks whether the From: domain aligns with:
  - SPF: authorised sending infrastructure
  - DKIM: authenticated signing domain
- The domain owner can publish a policy:
  - Monitor
  - Quarantine
  - Reject
- **DMARC asks: "Does the authenticated domain match the claimed sender?"**

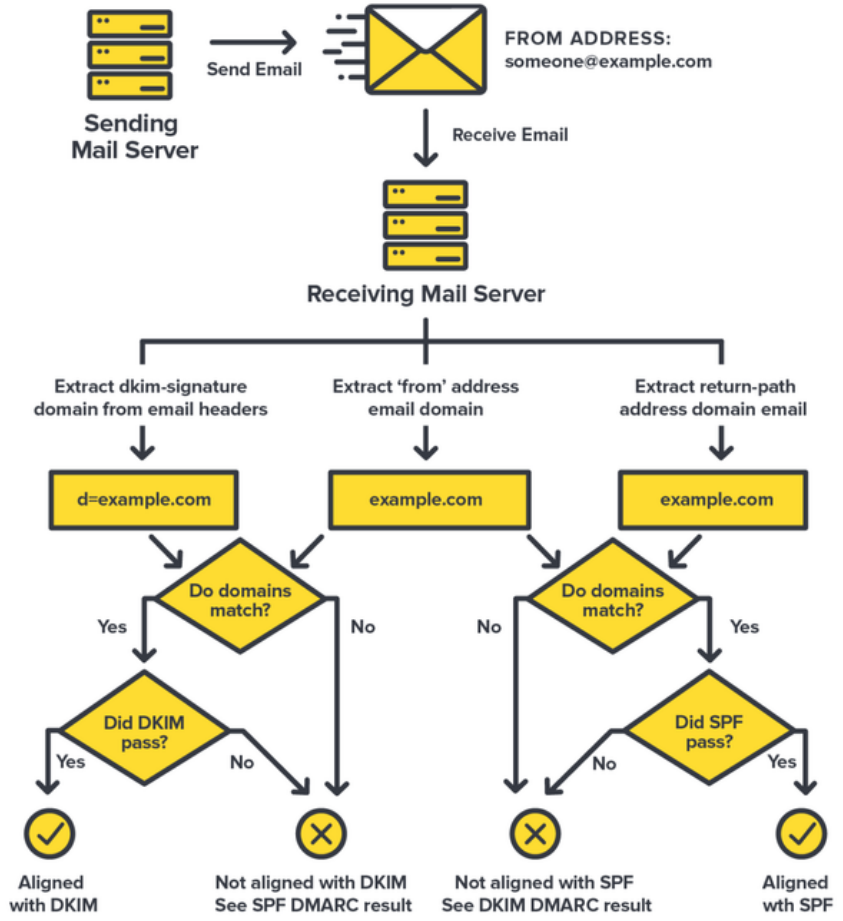

Diagram from postmarkapp.com

What email authentication does not solve

- Authenticated email can still be malicious
- SPF, DKIM and DMARC help reduce domain spoofing, but they cannot stop:
  - **Compromised accounts**: legitimate account, malicious sender
  - **Lookalike domains**: university-support.lk ≠ university.lk
  - **Malicious links or attachments**: authenticated but harmful content
  - **Social engineering**: trusted identity, deceptive request
- **Email authentication verifies domains, not intent.**

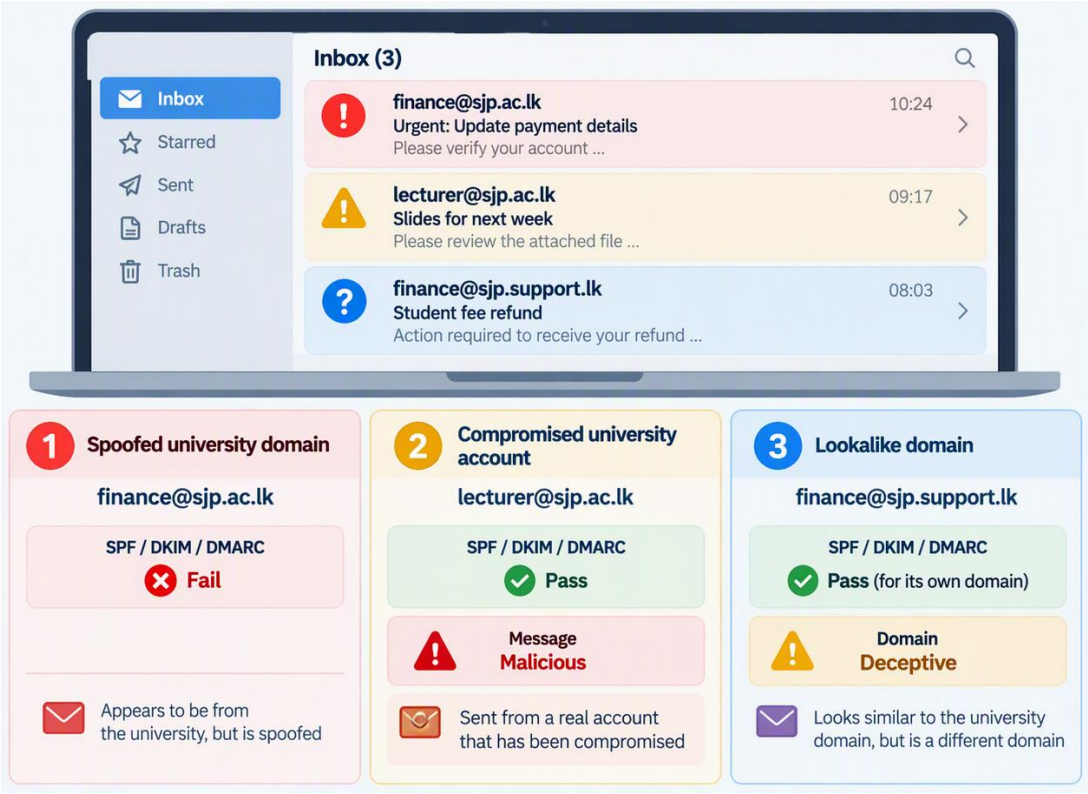

IPsec protects traffic at the network layer

- IPsec protects IP traffic between trusted endpoints
  - Operates at the network layer
  - ESP can provide confidentiality, integrity and authentication
  - Transport mode protects the IP payload
  - Tunnel mode protects an encapsulated IP packet
  - Commonly used for VPN and site-to-site connectivity
- **TLS protects an application connection; IPsec can protect IP traffic between networks or hosts.**

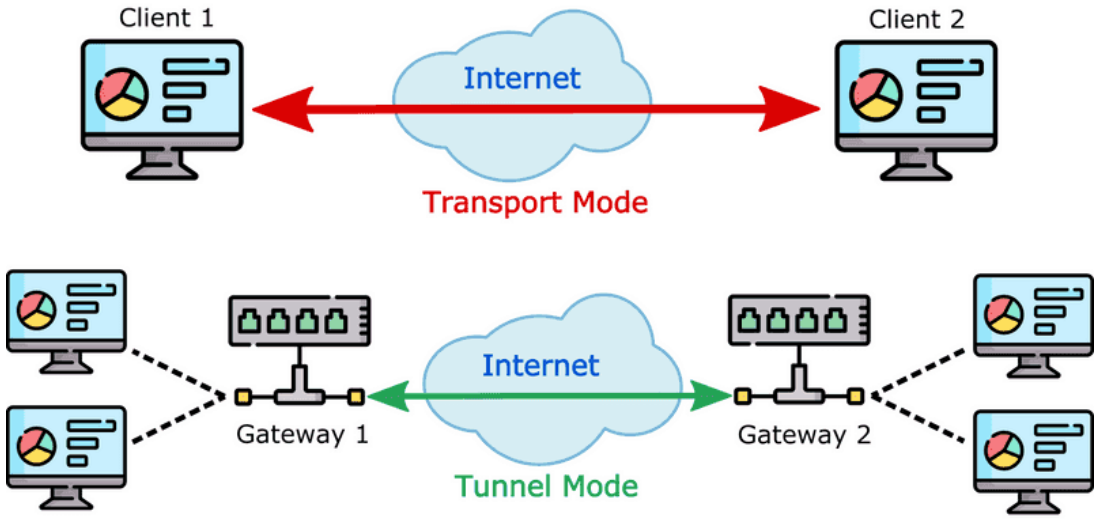

IKEv2 establishes the IPsec tunnel

- IPsec needs trusted endpoints and shared keys
- IKEv2 helps the endpoints:
  - Authenticate each other
  - Negotiate security parameters
  - Establish shared key material
  - Create and maintain IPsec Security Associations
- **IKEv2 establishes the trust; ESP protects the traffic.**

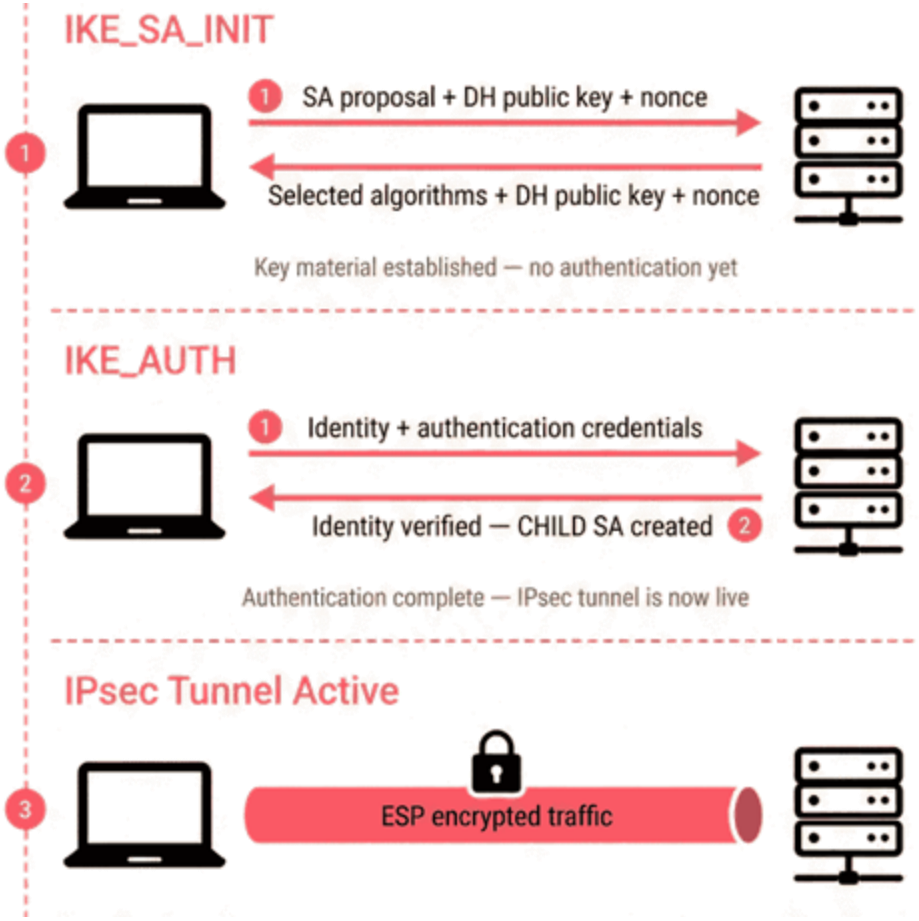

Wi-Fi extends the network beyond the cable

- Wireless changes the security boundary
- Unlike a wired connection, Wi-Fi signals extend beyond the physical building.
- This creates risks such as:
  - **Eavesdropping**: Wireless traffic can be captured nearby
  - **Unauthorised access**: Attackers may attempt to join the network
  - **Rogue access points**: Unapproved APs create new entry points
  - **Evil twins**: Fake APs imitate trusted university Wi-Fi
- **Being near the network should not mean being trusted by the network.**



WPA3 strengthens personal Wi-Fi security

- WPA3 improves how shared-password Wi-Fi establishes trust
- WPA2-Personal
  - Uses a shared Wi-Fi password
  - Captured handshakes can enable offline password guessing
- WPA3-Personal
  - Uses SAE authentication
  - Resists offline dictionary attacks
  - Requires Protected Management Frames
- **A stronger protocol cannot compensate for a weak or poorly managed network.**

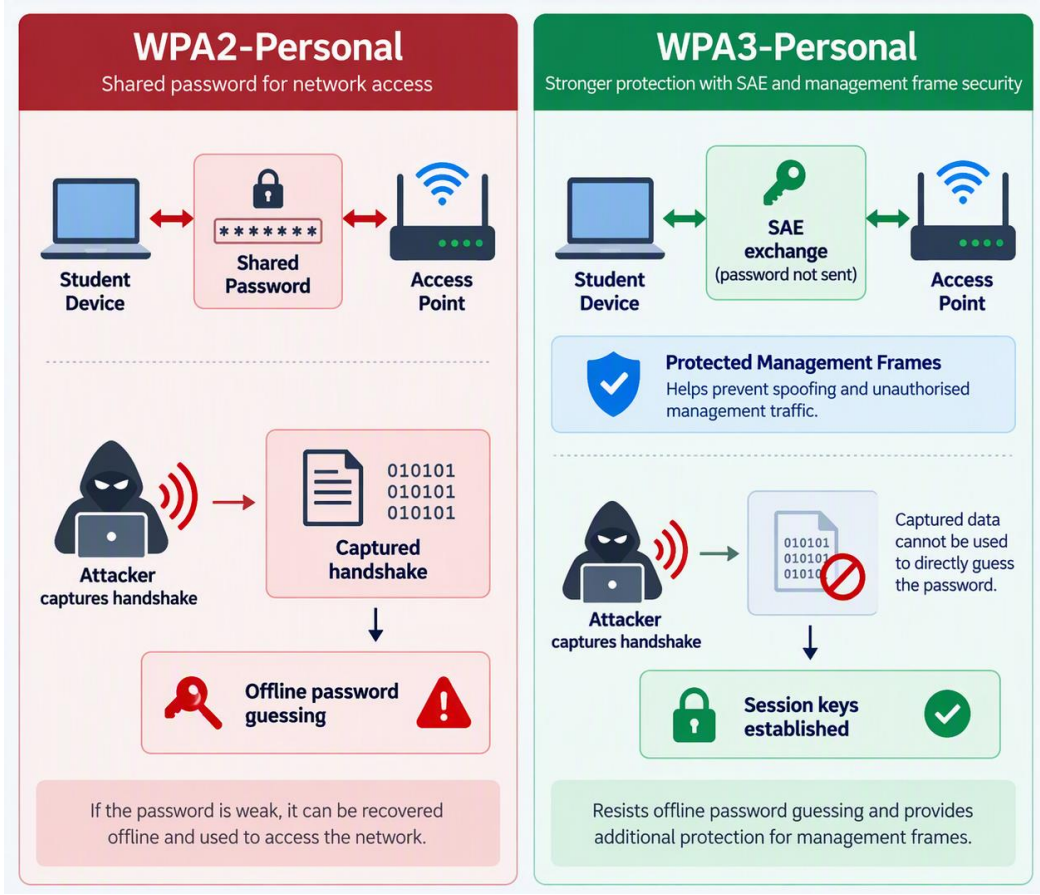

Enterprise Wi-Fi authenticates each device

- Campus Wi-Fi needs individual authentication
- WPA3-Enterprise commonly uses 802.1X:
  - Supplicant: Student or staff device
  - Authenticator: Wireless access point
  - Authentication server: Validates the connection
  - EAP: Carries the authentication exchange
- **The Wi-Fi password is replaced by individual trust.**

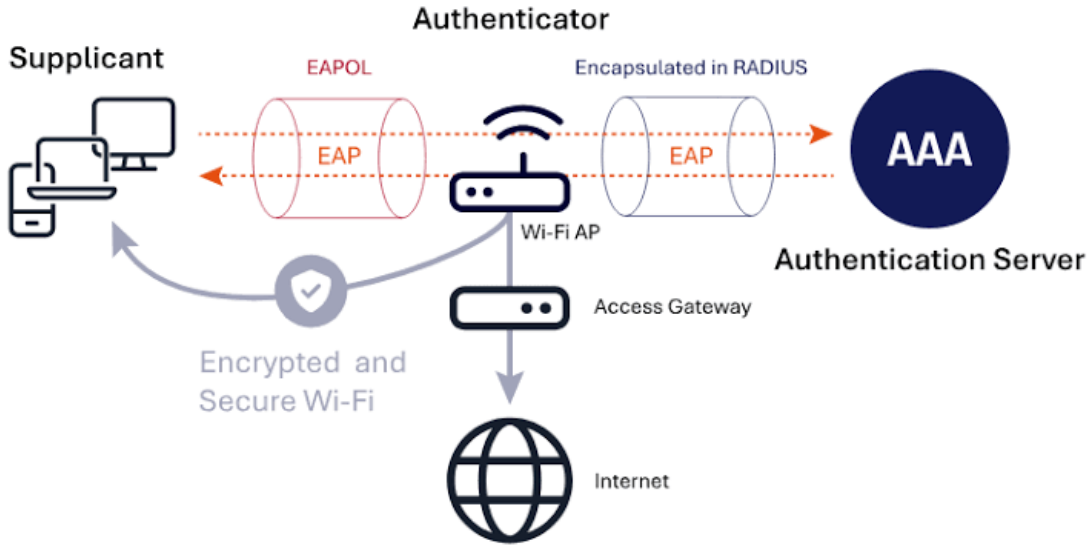

Trust the network, not just the SSID

- The SSID does not prove the network is genuine
- An attacker can advertise the same Wi-Fi name.
- Enterprise devices should:
  - Validate the authentication server certificate
  - Trust only the expected CA and server identity
  - Avoid accepting unexpected certificate warnings
  - Use centrally managed Wi-Fi profiles where possible
- **Mutual authentication helps defeat the evil-twin problem.**

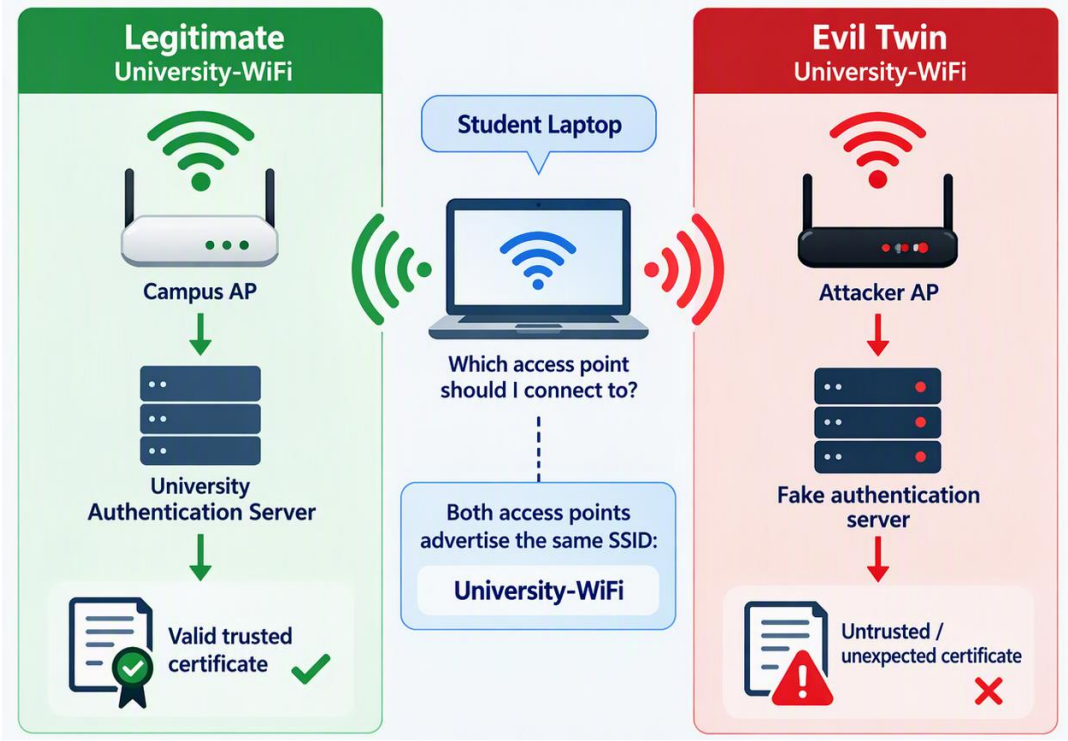

Wi-Fi security is evolving with the standard

- Newer Wi-Fi generations strengthen the security baseline
- Wi-Fi 6E
  - Extends Wi-Fi into the 6 GHz band
  - Requires WPA3 or Enhanced Open
  - Removes legacy WPA2-only operation in 6 GHz
- Wi-Fi 7
  - Continues the WPA3 security baseline
  - Adds new performance capabilities
  - Security still depends on correct authentication and configuration
- **Faster Wi-Fi does not automatically mean safer Wi-Fi.**

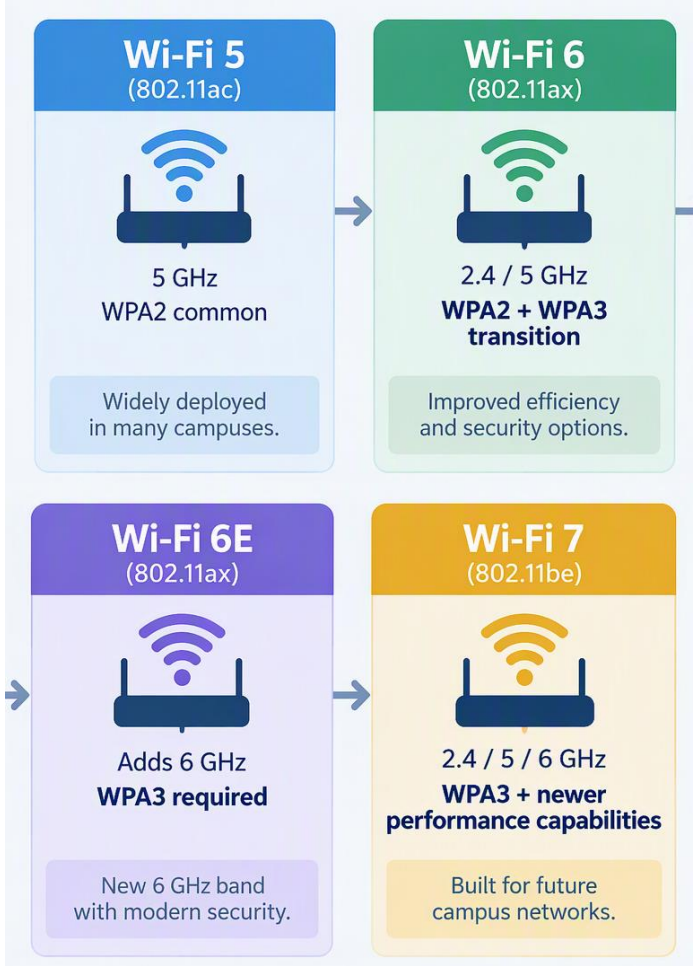

Mobile networks authenticate the subscriber

- Mobile access uses subscriber-based trust
  1. USIM securely stores the subscriber's long-term credentials
  2. 4G EPS-AKA / 5G-AKA authenticate the subscription and network
  3. Successful authentication derives session-specific keys
  4. These keys enable protection of signalling and radio traffic
  5. TLS still protects university email and web sessions
- **Cellular security protects network access; TLS protects the university service.**

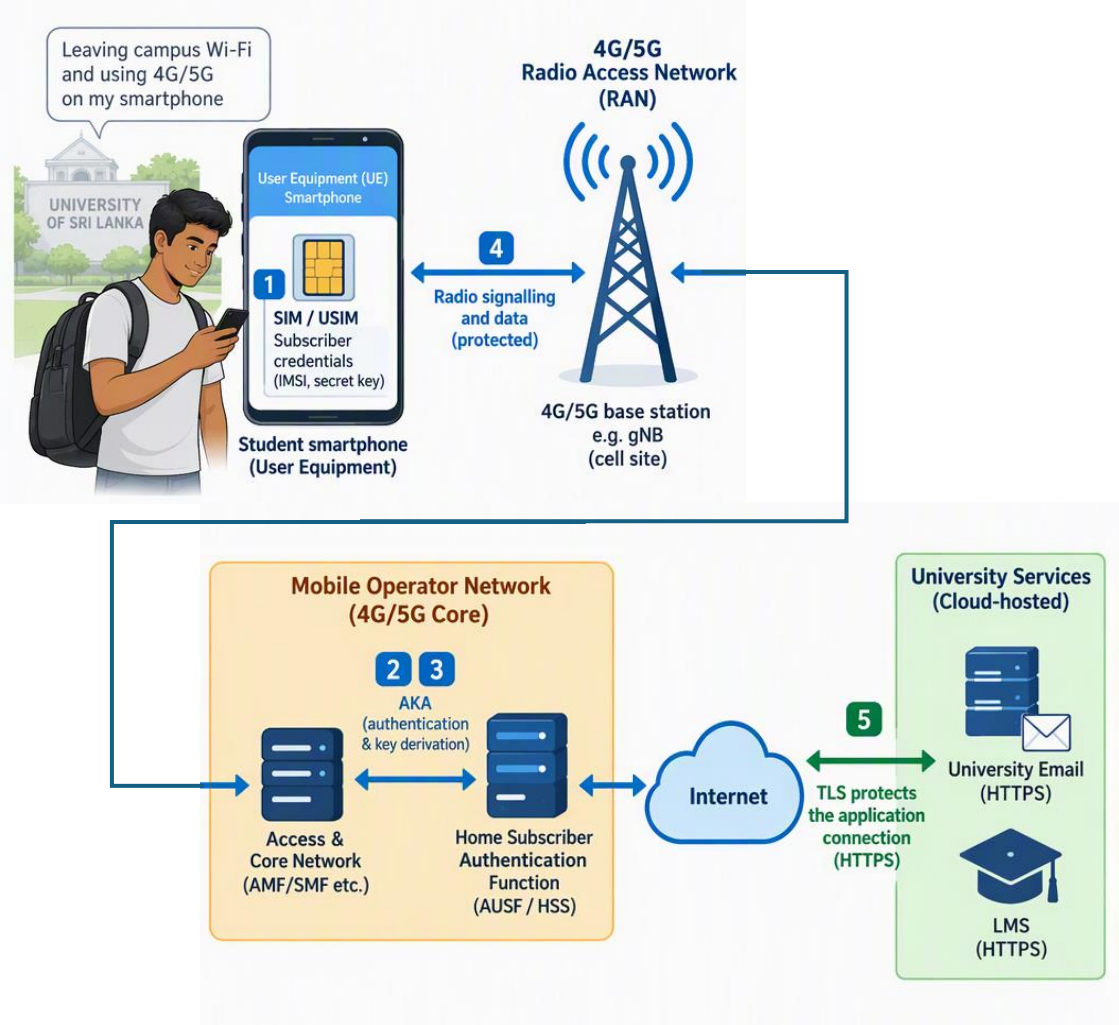

5G conceals the subscriber identity

- 5G improves subscriber-identity privacy
  - SUPI: Permanent identifier associated with the subscription
  - SUCI: Concealed identifier used for initial identification
  - The UE/USIM creates the SUCI using home-network information
  - The home network resolves SUCI back to SUPI for authentication
- **The permanent identity need not be exposed in clear text over the radio interface.**

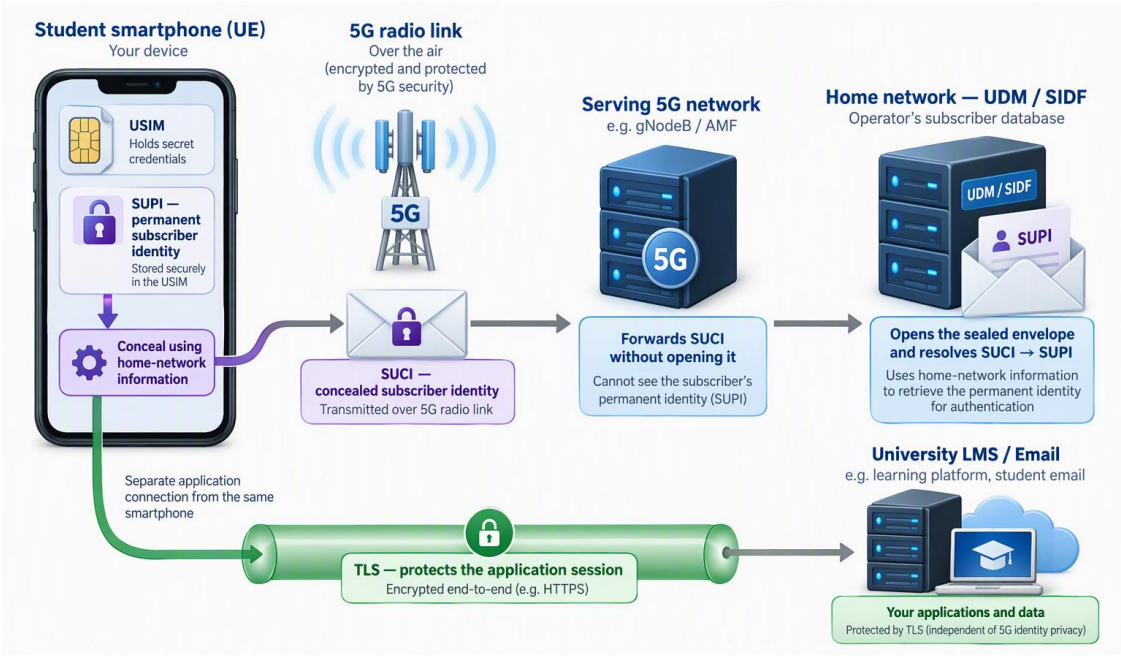

5G improves security, not end-to-end trust

- Mobile security still has trust boundaries
  - Roaming: Security depends on relationships between operators
  - Fake base stations: Devices may encounter untrusted radio infrastructure
  - Legacy fallback: Older network generations may offer fewer protections
  - Device compromise: Cellular encryption cannot secure an infected endpoint
- **5G strengthens the access network. TLS still protects the university service.**

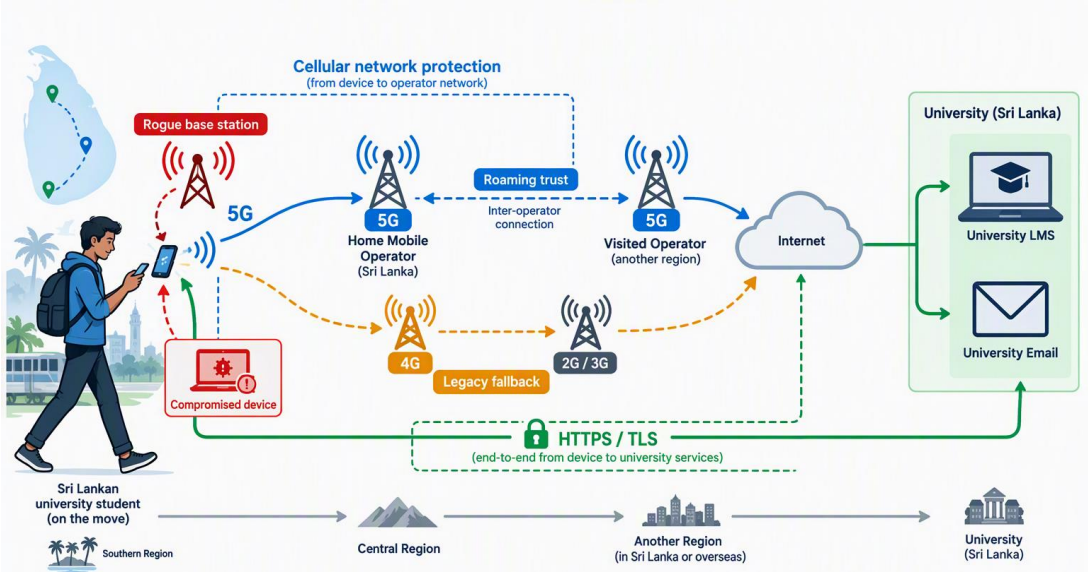

One student, many security boundaries

- Follow one student through a typical university day:

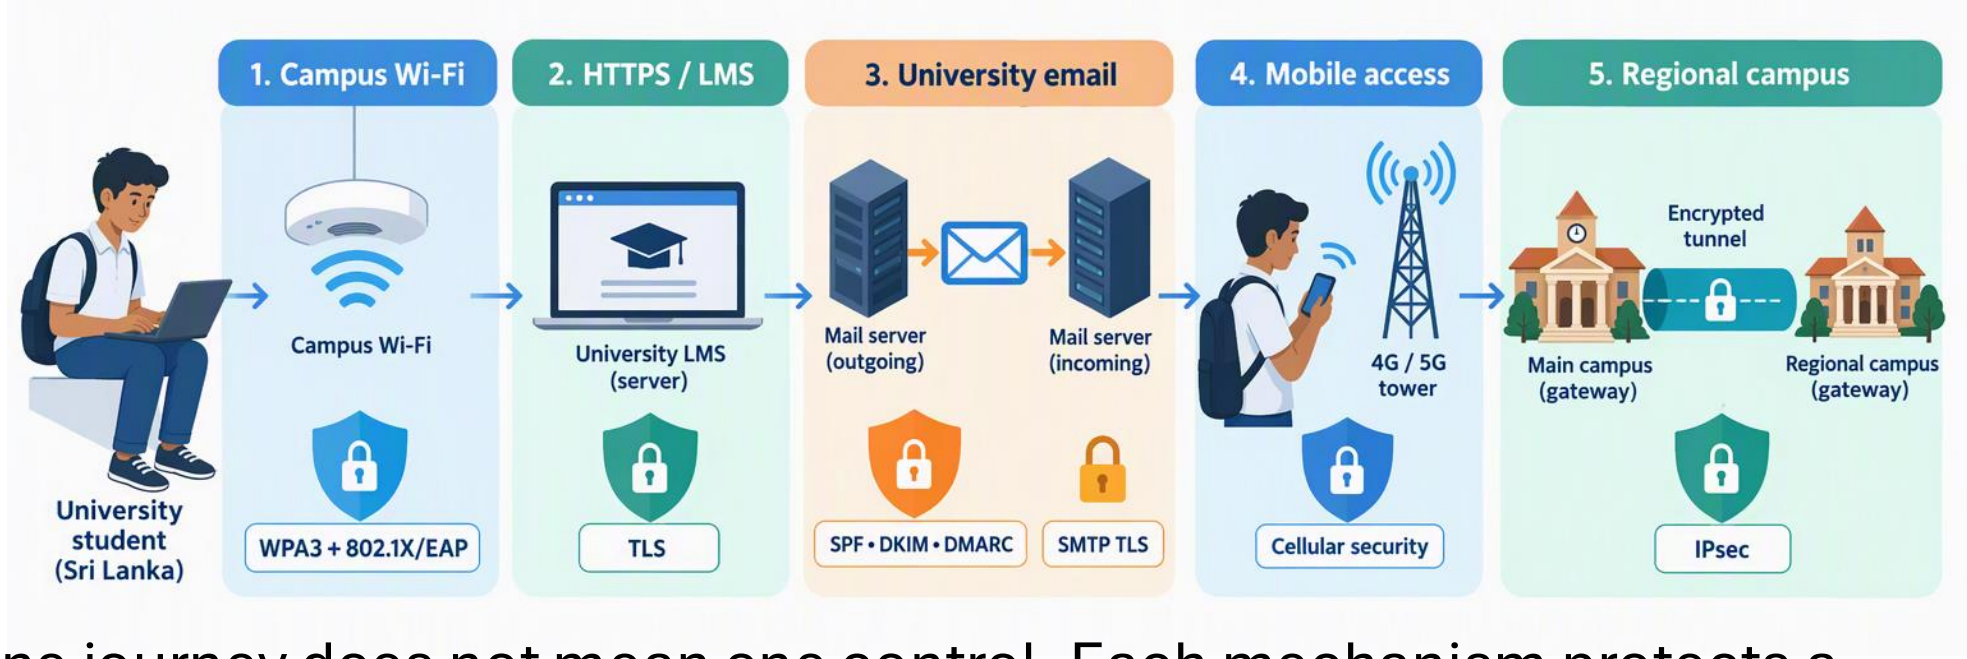

- One journey does not mean one control. Each mechanism protects a different layer, path or trust boundary.

Defence Across Layers

- No single control protects everything. Security comes from overlapping trust boundaries.

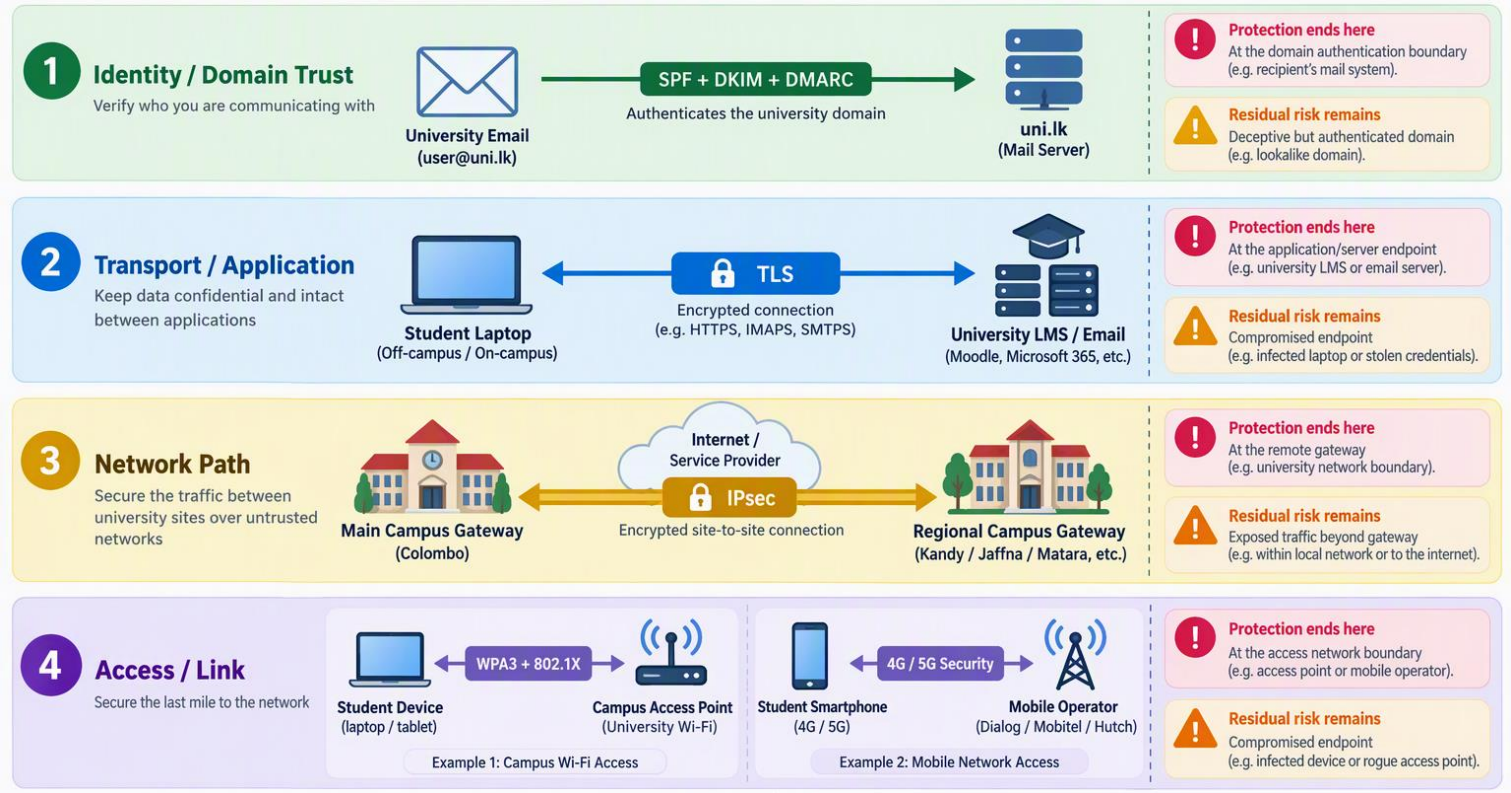

Five network security takeaways

- Five engineering takeaways
  1. Authenticate before you trust
  2. Encrypt the right communication boundary
  3. Know where protection starts and ends
  4. Use complementary controls across layers
  5. Always ask what risk remains
- **Secure networking is not about one protocol, it is about building trustworthy communication paths.**

