Operational Network Security:

An overview of Network Security Mechanisms – Firewalls, DPI, IDS, IPS and NGFW

ITN 4252 Advanced Topics in Networking

Partially based on resources from Palo Alto Networks, SolarWinds, Fortinet, and Cloudflare; and PowerPoint slides from Computer Networking: A Top-Down Approach, 8th edition, Jim Kurose, Keith Ross; Implementing Intrusion Prevention, CCNA Security v2.0

A suspicious connection reaches the university

- The public service must be reachable, but not trusted:
  - Internet access is required for students and staff
  - The public service must accept legitimate connections
  - Suspicious traffic may look similar to normal traffic
  - A compromised public service must not expose internal systems
- **The challenge is not simply blocking traffic, it is controlling where traffic can go next.**

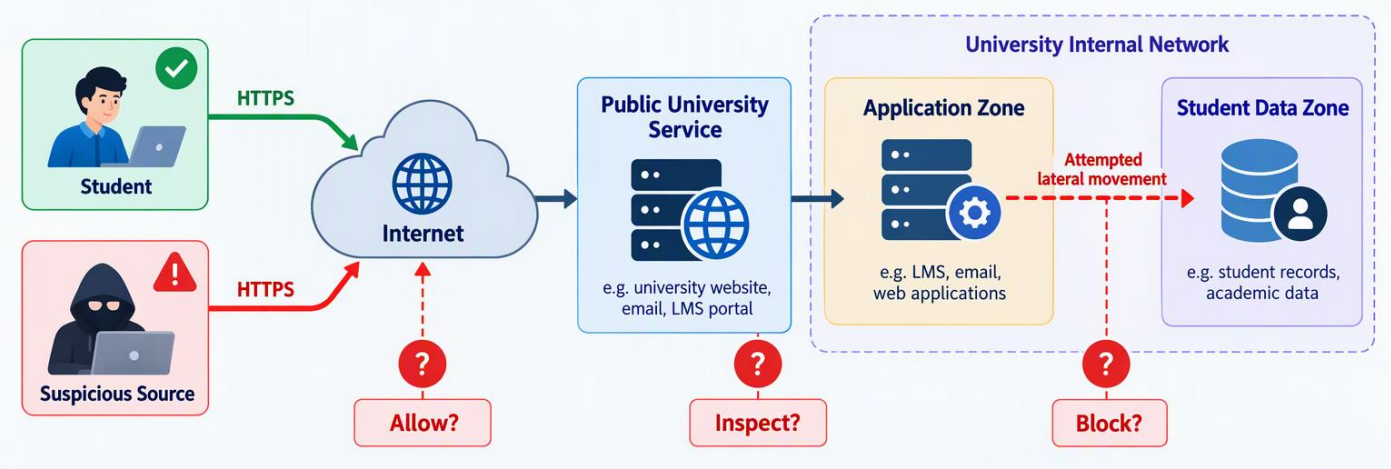

The attack path: enter, explore, move

- One allowed connection can become a path inward.
- A network intrusion may progress through different zones:
  - Enter: Reach an exposed university service
  - Explore: Discover what the service can communicate with
  - Move: Attempt to reach other systems
  - Access: Target sensitive services or data
- **The goal is to break the path before one compromise becomes many.**

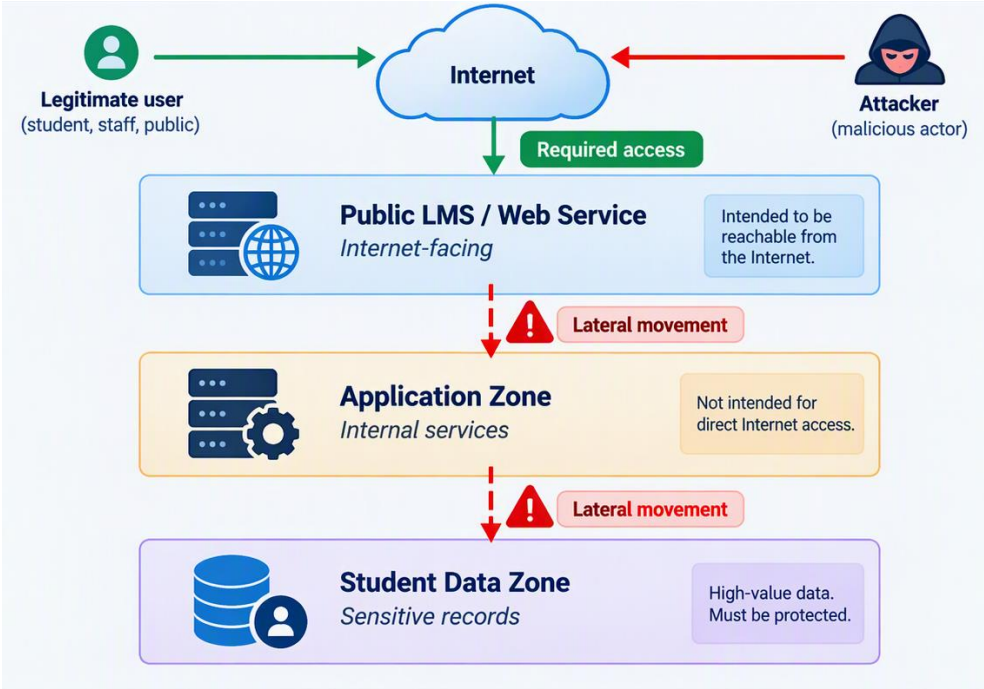

Four defensive decisions

- Every security control makes a different decision:
  - **Allow**: Permit required communication
  - **Restrict**: Limit where traffic can travel
  - **Detect**: Identify suspicious behaviour
  - **Block**: Stop traffic when confidence is sufficient
- **Good security is not “block everything”, it is making the right decision at the right point.**

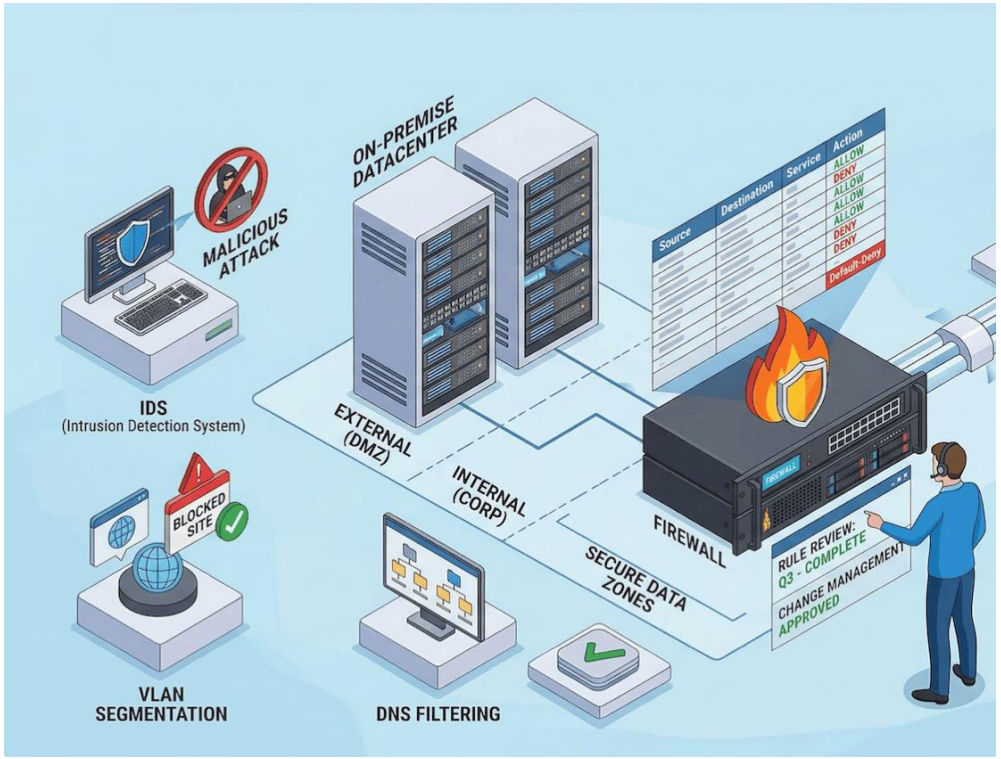

A firewall enforces communication policy

- A firewall turns policy into traffic decisions:
  - Controls traffic between hosts, networks or security zones
  - Evaluates communication against defined allow and deny rules
  - Permits only the paths required by university services
  - Blocks direct access to protected internal zones
  - Records decisions for review and investigation
- **Allowed by policy does not mean proven safe.**

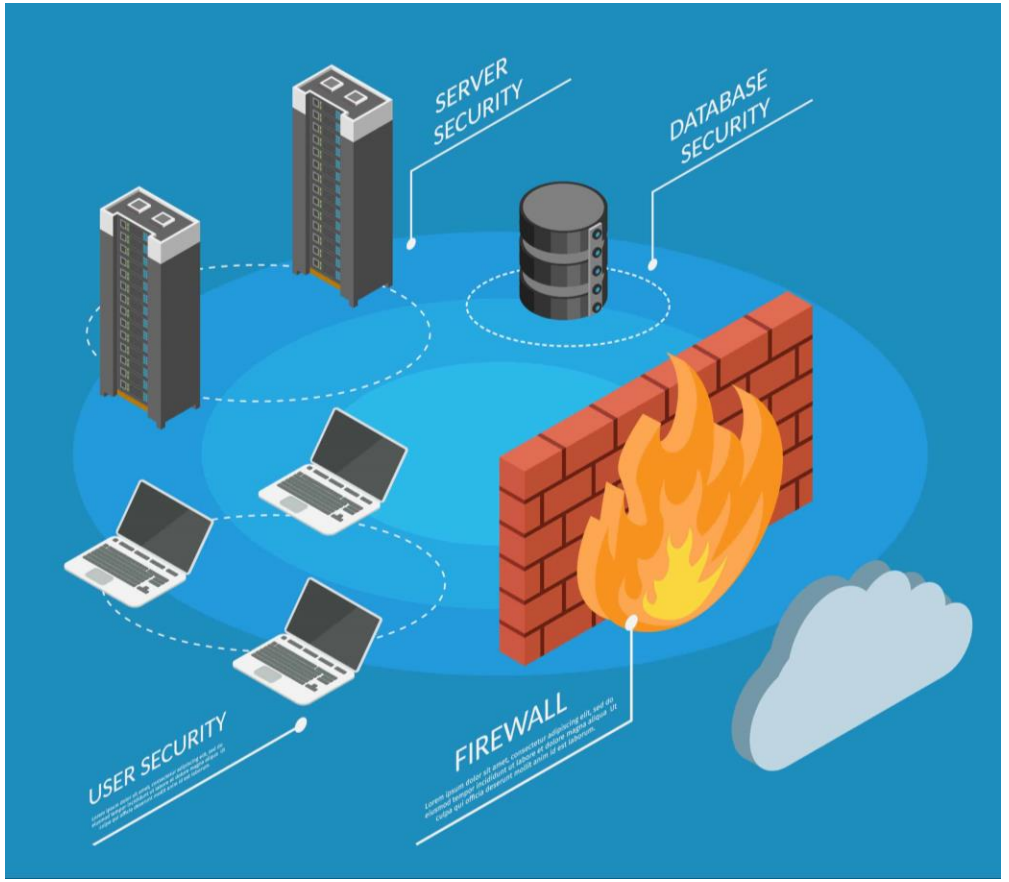

Default deny, permit only required flows

- Start with deny, then add justified exceptions.
- For each permitted flow, define:
  - **Source**: Who or what initiates the connection?
  - **Destination**: Which service is required?
  - **Service**: Which protocol and port are necessary?
  - **Direction**: Who may initiate the communication?
- **If the communication is not required, it should not be allowed.**

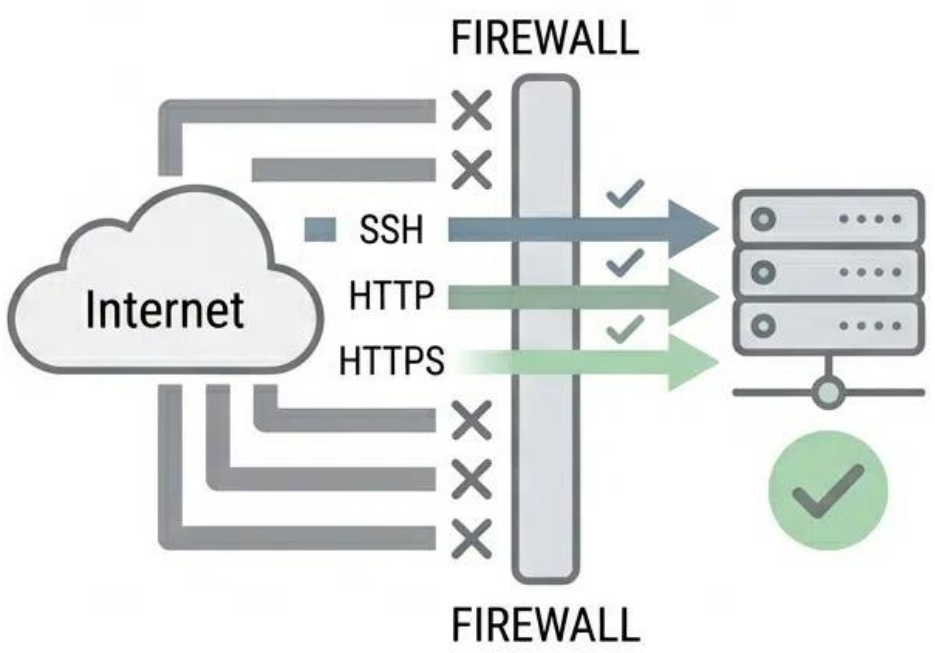

Stateless filtering, decide packet by packet

- Stateless filtering evaluates each packet independently.
- Rules examine:
  - Source and destination IP address
  - Protocol
  - Source and destination port
  - Direction or interface
- **It applies the rule, but does not remember the connection.**

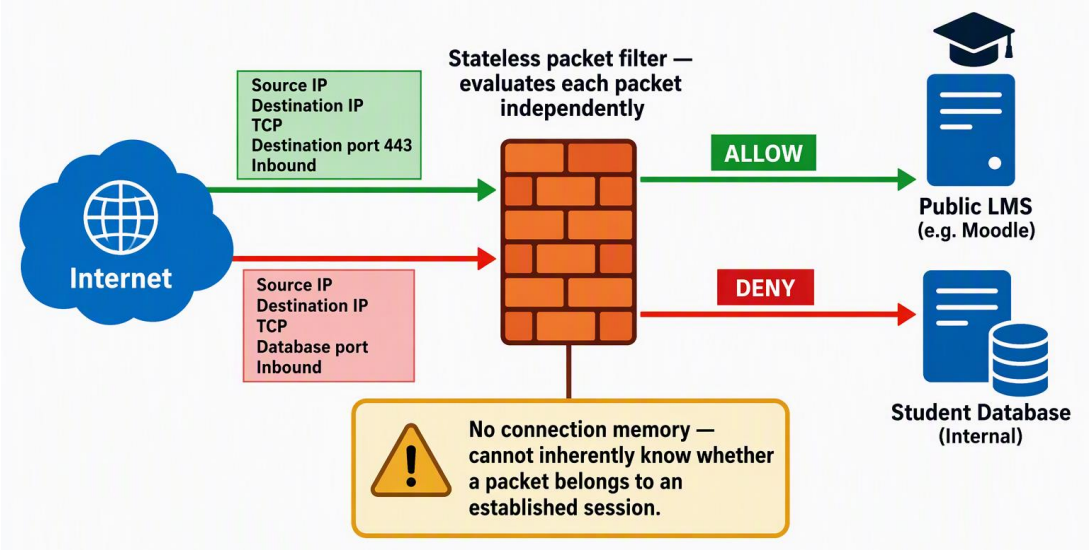

Stateful inspection, track the connection

- Stateful inspection remembers active connections.
- The firewall tracks:
  - Who initiated the connection
  - Source, destination and ports
  - Connection state
  - Whether new packets belong to an expected flow
- **A packet is evaluated in the context of the connection, not in isolation.**

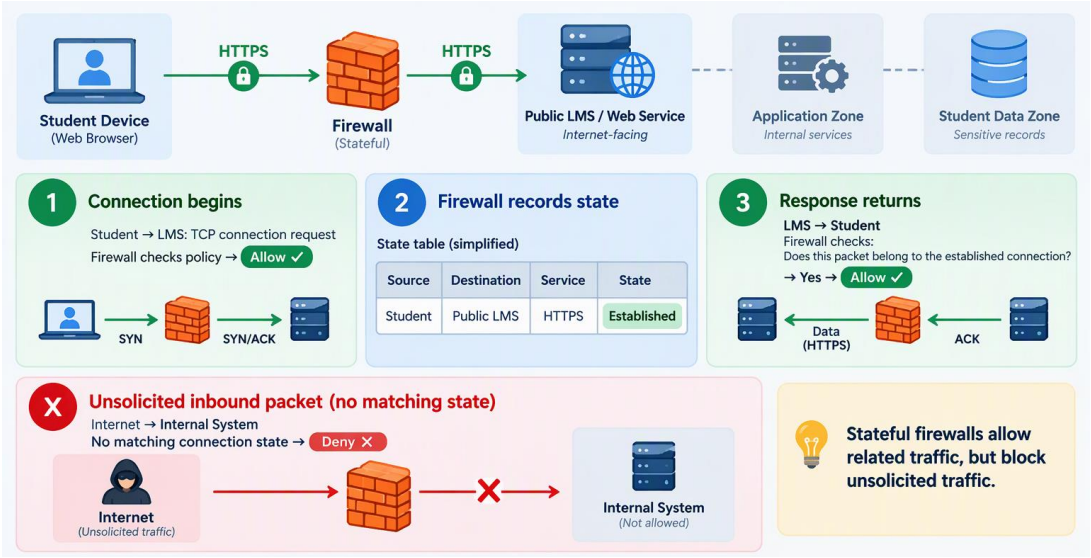

Rule order changes the outcome

- Specific rules must come before broad rules.
- Intended order
  1. Allow Public LMS → Required Application Service
  2. Deny other traffic → Application Zone
  3. Deny anything not explicitly permitted
- Reversed order
  1. Deny all traffic → Application Zone
  2. Required LMS connection never reaches its allow rule
- **The same rules in a different order can produce a different result.**

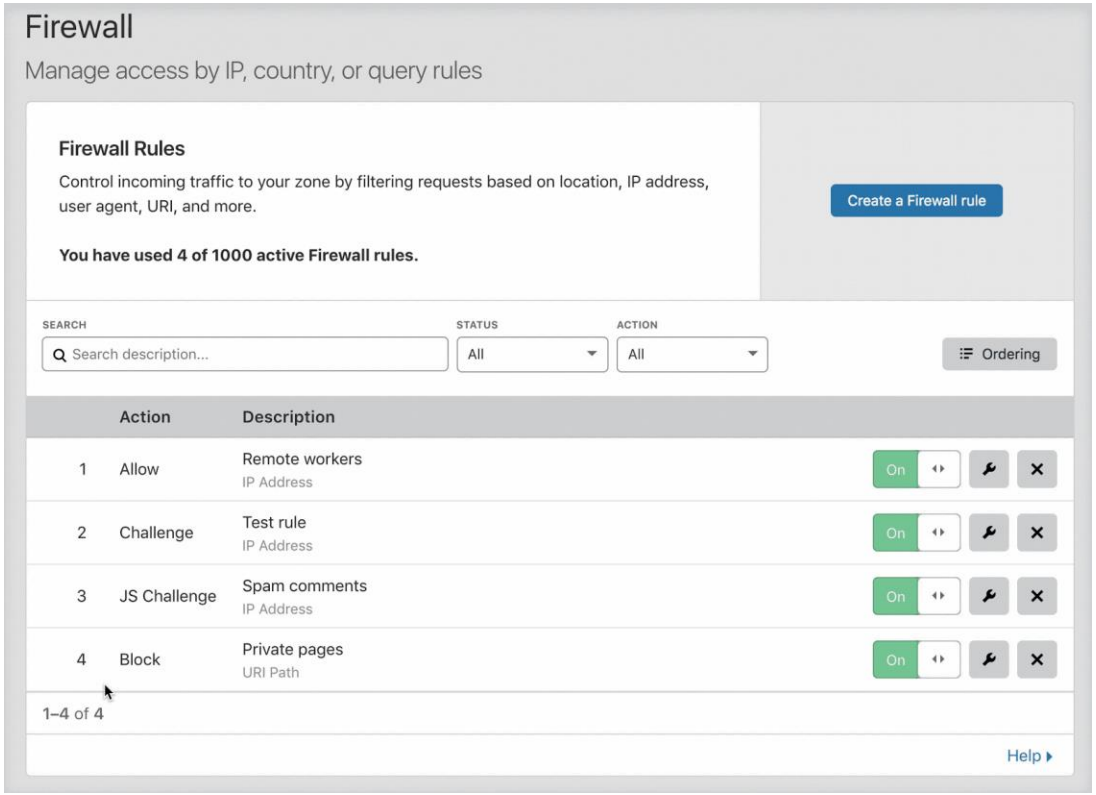

Firewall decision activity

- Would you allow or deny this traffic?

Communication	Decision
Internet → Public LMS HTTPS	Allow ✓
Internet → Student Database	Deny ✕
Public LMS → Required Application Service	Allow ✓
Guest Wi-Fi → Administration Systems	Deny ✕
Application Service → Student Database	Allow only if required ✓

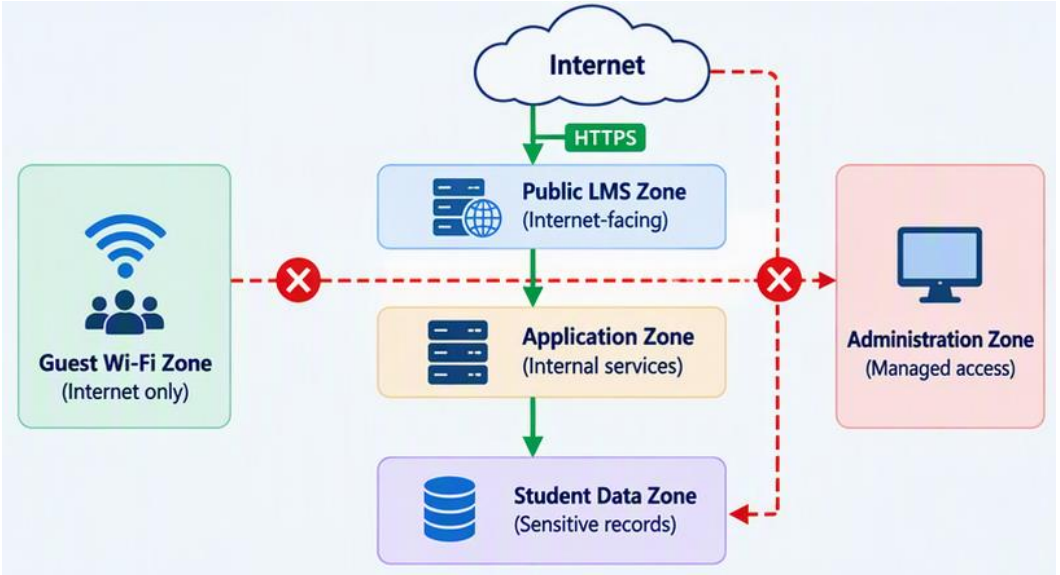

- **Every allowed path should have a clear purpose.**

Inside does not automatically mean trusted

- Internal does not mean trusted.
- Threats can originate from:
  - Compromised devices already inside the network
  - Guest or unmanaged devices
  - Misconfigured internal systems
  - Compromised public-facing services
- **Trust should come from policy and context, not network location alone.**

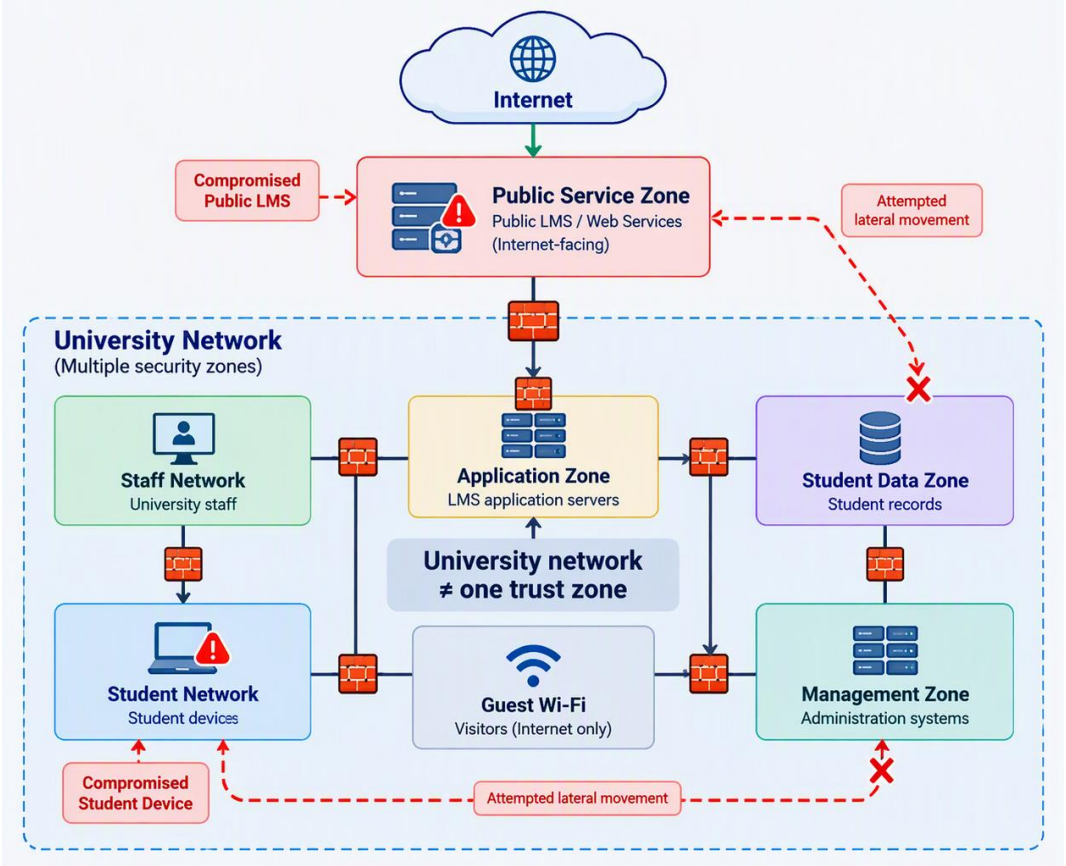

Segmentation: Divide into security zones

- Different systems need different security zones.
- Separate the university by purpose and risk:
  - **Public Services**: Internet-facing LMS and web services
  - **Application**: Internal application services and APIs
  - **Student Data**: Sensitive academic and personal records
  - **User Networks**: Staff, students and guests
  - **Management**: Administration and security systems
- **Segmentation creates boundaries where communication can be controlled.**

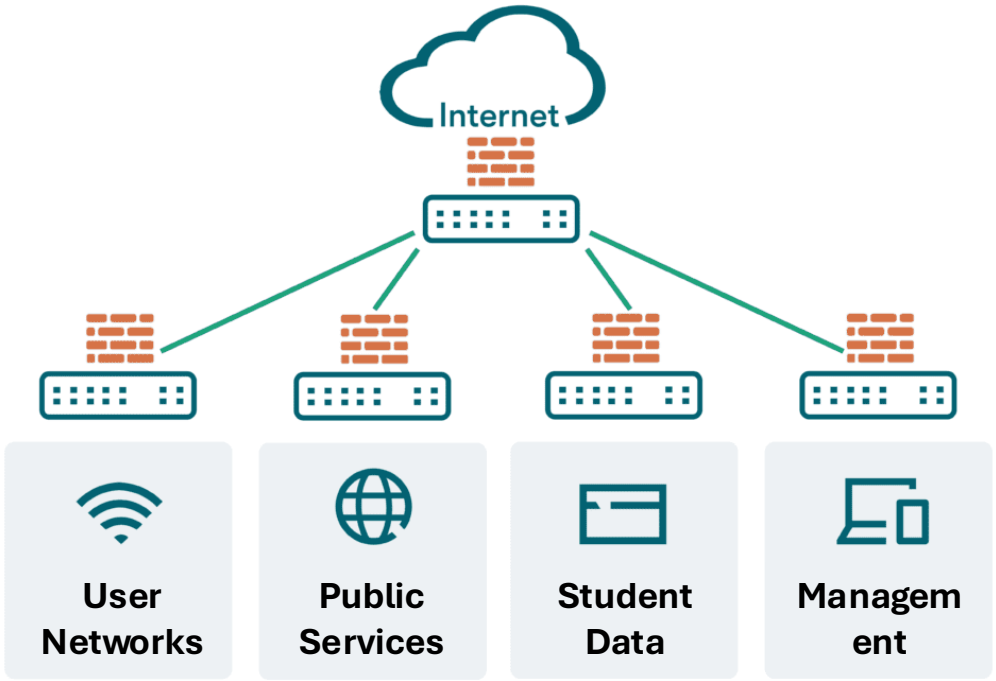

DMZ: Expose the service, not the network

- Public services need their own security zone.
- A DMZ separates Internet-facing services from sensitive internal systems:
  - Internet users reach the public LMS/web service
  - The DMZ permits only required inbound traffic
  - DMZ services receive only specific internal connections
  - Internal data and management systems remain separately protected
- **Expose the service, not the university network behind it.**

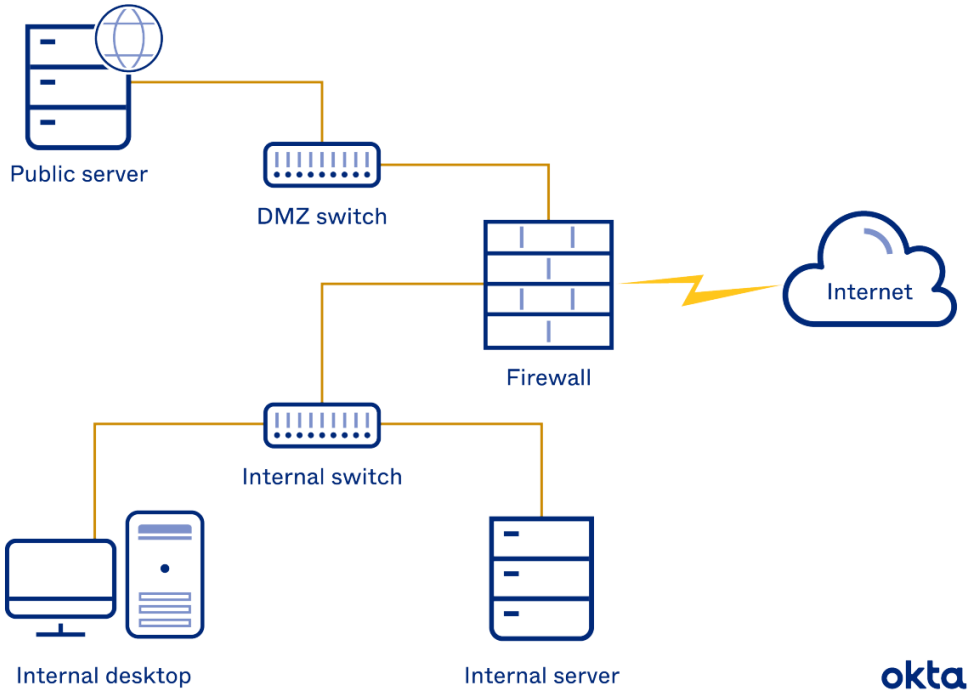

Permit only required inter-zone flows

- Zones are useful only when traffic between them is controlled.
  - Public → DMZ: Only required public services
  - DMZ → Application: Only specific application flows
  - Application → Data: Only required data services
  - User → Management: Deny unless explicitly authorised
- **Every zone boundary should reduce unnecessary communication paths.**

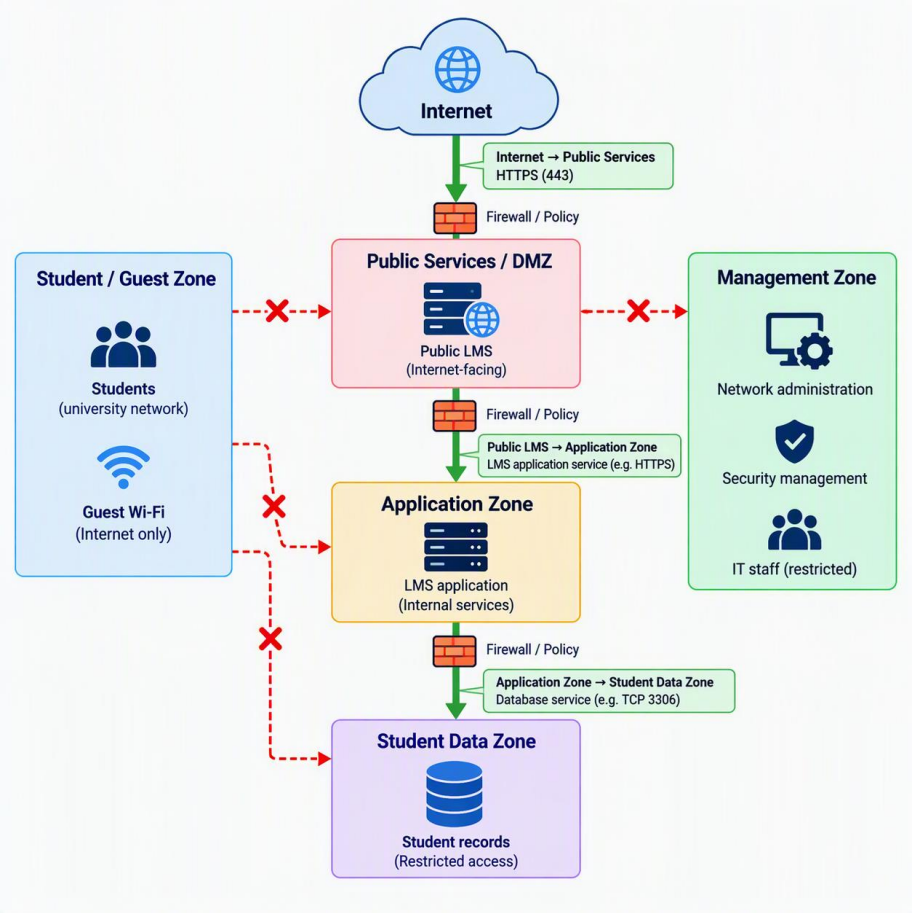

Segmentation limits lateral movement

- Segmentation limits how far a compromise can spread.
- If the public LMS is compromised:
  - DMZ boundary limits access to internal systems
  - Application zone permits only required services
  - Data zone blocks unnecessary direct access
  - Management zone remains separately protected
- **Assume one system may fail, design so the whole network does not fail with it.**

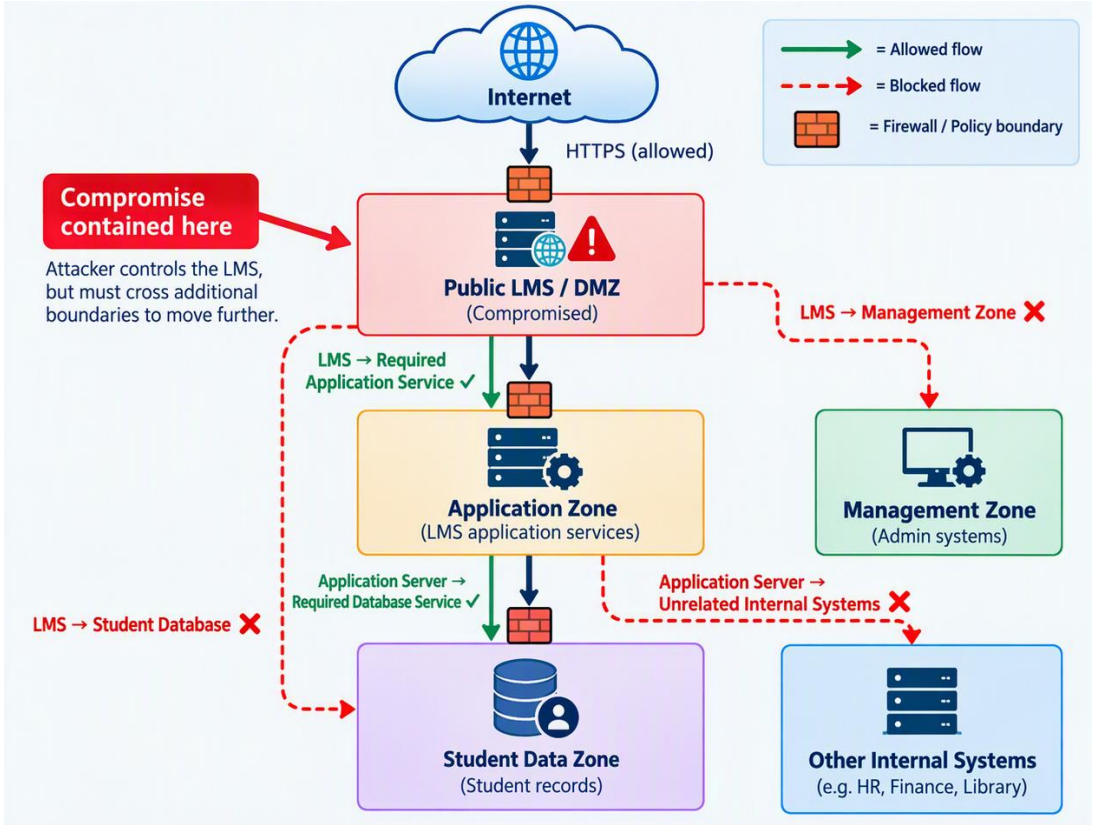

From header inspection to application context

- Not all traffic inspection looks equally deep.
  - **Packet filtering**: IP addresses, protocols and ports
  - **Stateful inspection**: Adds connection context
  - **Application awareness**: Identifies the protocol or application
  - **Content inspection**: Examines visible payload data
- **More inspection can provide more context, but requires more visibility and processing.**

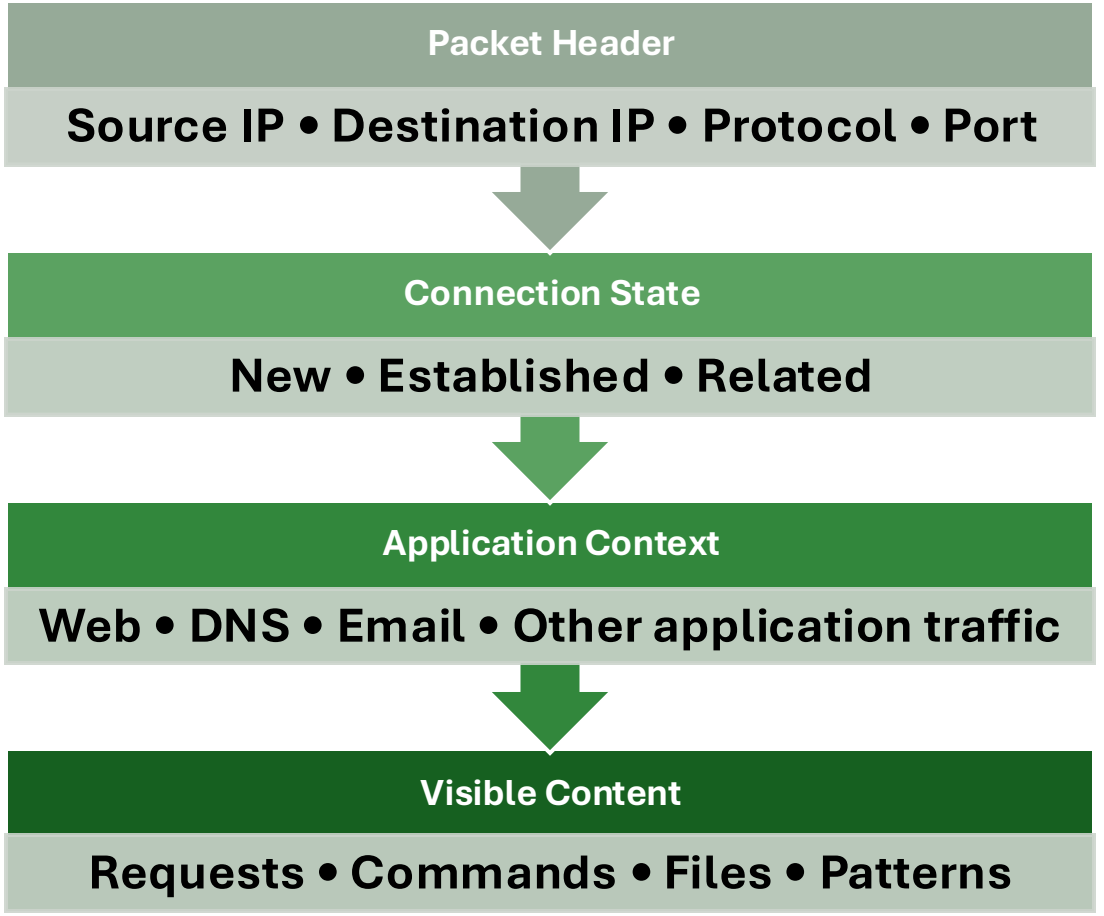

DPI sees only what is visible

- Encryption changes what the network can inspect.
- Without encryption
  - Headers, application data and payload may be visible
- With encryption
  - Some metadata remains visible
  - Application content is protected between encryption endpoints
  - DPI cannot simply read the encrypted payload
- **Deep inspection requires visibility, encryption deliberately removes much of it.**

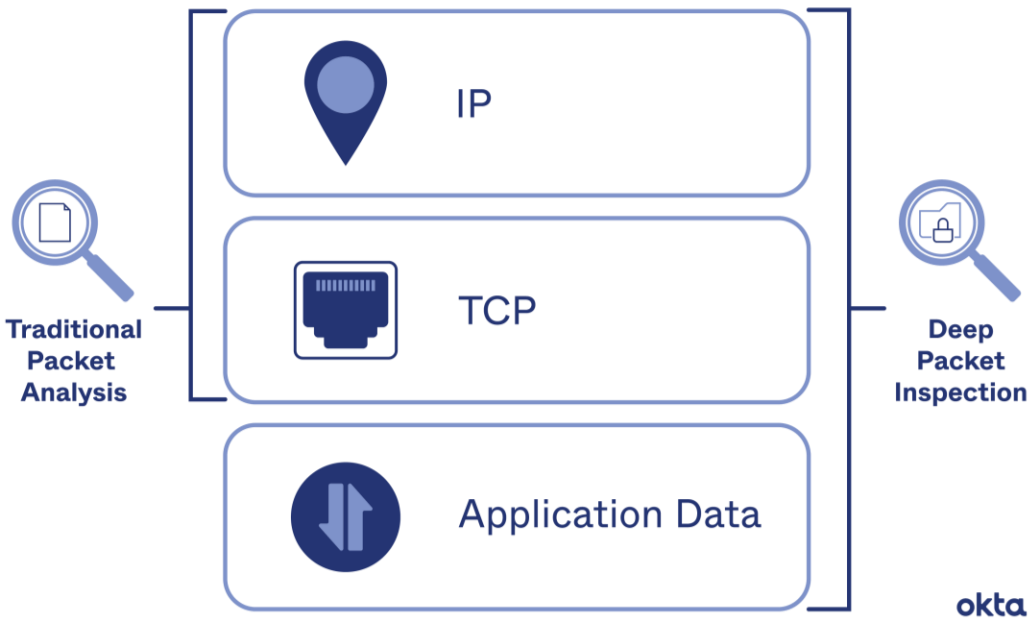

Encrypted-traffic inspection is a risk decision

- Decrypting traffic improves visibility, but creates new risk.
- TLS inspection can:
  - Decrypt selected traffic at an authorised inspection point
  - Inspect the plaintext for threats or policy violations
  - Re-encrypt traffic towards the destination
- But it introduces:
  - Privacy and trust concerns
  - Performance and operational overhead
  - A new high-value security point
- **Inspect selectively, not simply because you can.**



Known patterns and unexpected behaviour

- Detection uses two complementary approaches.
- Signature-based
  - Matches traffic against known patterns or indicators
  - Effective when the threat is already recognised
  - May miss new or significantly modified activity
- Behaviour / anomaly-based
  - Compares activity with an expected baseline
  - May reveal previously unseen suspicious behaviour
  - Requires tuning and may generate more false positives
- **Signatures recognise what we know. Anomalies highlight what does not look normal.**

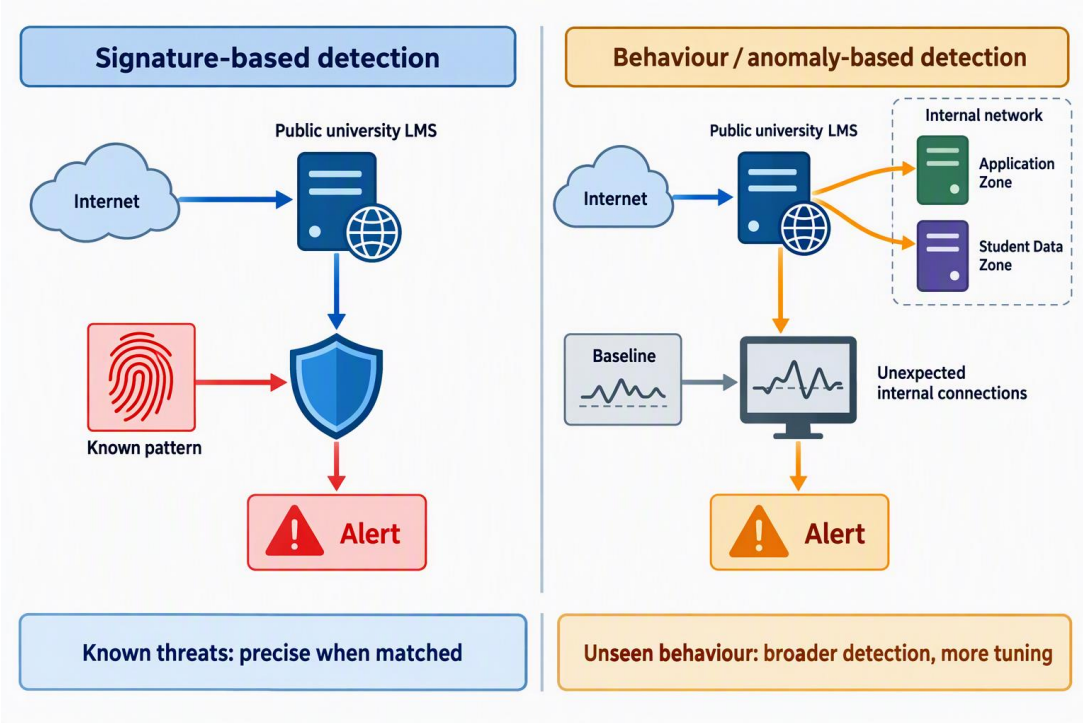

IDS observes and alerts

- An IDS watches traffic without sitting in the path.
  - Receives a copy of network traffic
  - Analyses traffic for signatures and unusual behaviour
  - Generates alerts when suspicious activity is detected
  - Does not normally block the triggering packet
- **IDS provides visibility without directly interrupting the traffic flow.**



IPS sits inline and can block

- An IPS can detect and stop traffic in real time.
  - Sits inline with the traffic path
  - Analyses traffic using detection rules and behaviour
  - Can drop or reject suspicious traffic
  - Incorrect decisions can disrupt legitimate services
- **IPS adds prevention, but also becomes part of the production path.**



Where should IDS and IPS be placed?

- Placement determines what the control can see, and affect.
  - Internet edge: Observe external traffic entering the university
  - DMZ boundary: Monitor public-facing services
  - Internal zone boundaries: Detect attempted lateral movement
  - Inline IPS: Use where confident blocking justifies the availability risk
- **Place controls where important traffic crosses a meaningful trust boundary.**

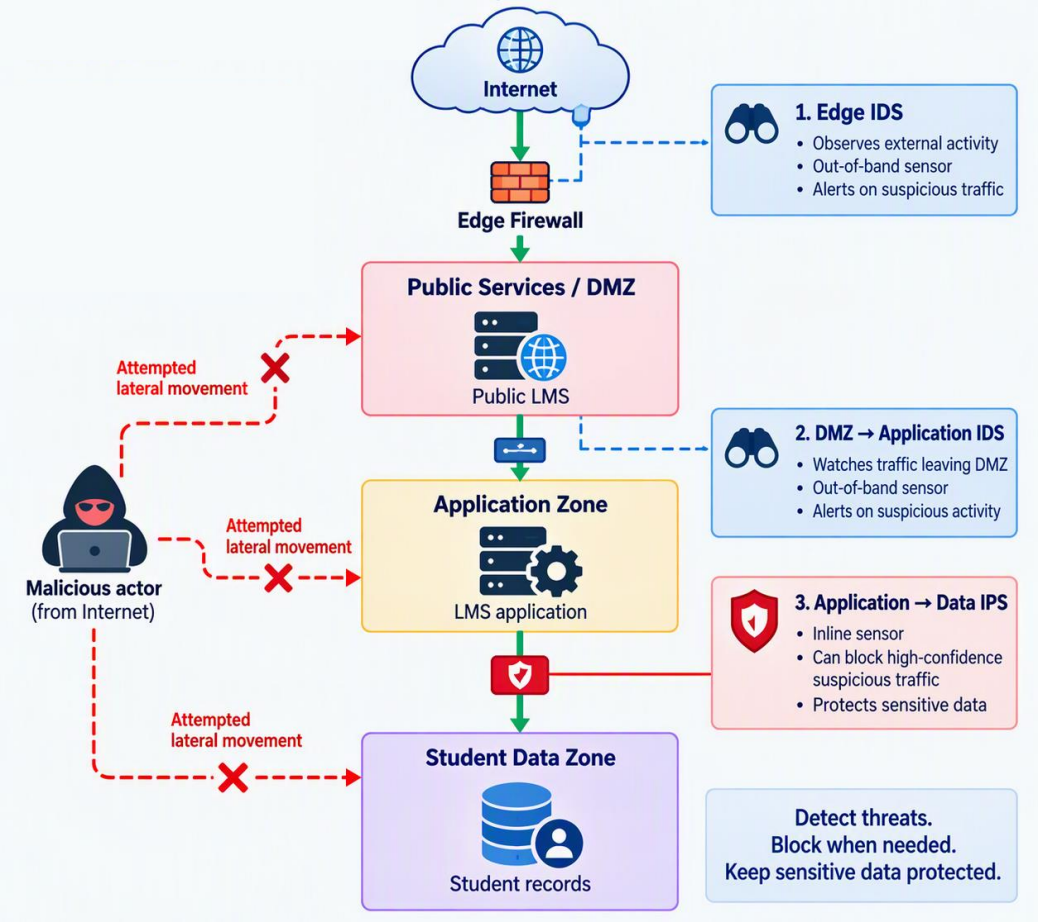

Detection is never perfect

- Every detection system makes mistakes.
  - True positive: Threat correctly detected
  - False positive: Legitimate activity incorrectly flagged
  - True negative: Legitimate activity correctly ignored
  - False negative: Threat is missed
- **Tuning balances detection coverage against operational disruption.**

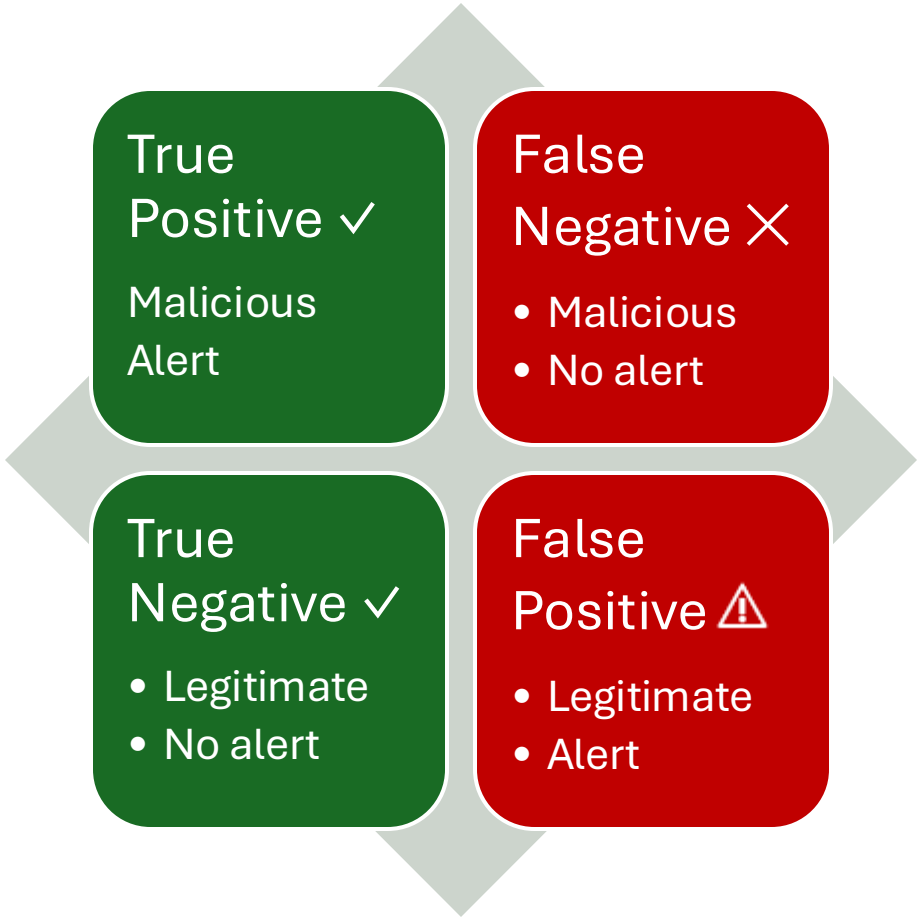

NGFW: An integrated capability bundle

- NGFW combines several security functions.
- Common capabilities include:
  - Stateful firewalling: Enforce connection policy
  - Application awareness: Identify and control application traffic
  - Content inspection: Inspect visible or authorised decrypted traffic
  - Integrated IPS: Detect and block suspicious activity
  - Context and intelligence: Add identity, threat and risk information
- **Integration improves context, not certainty.**

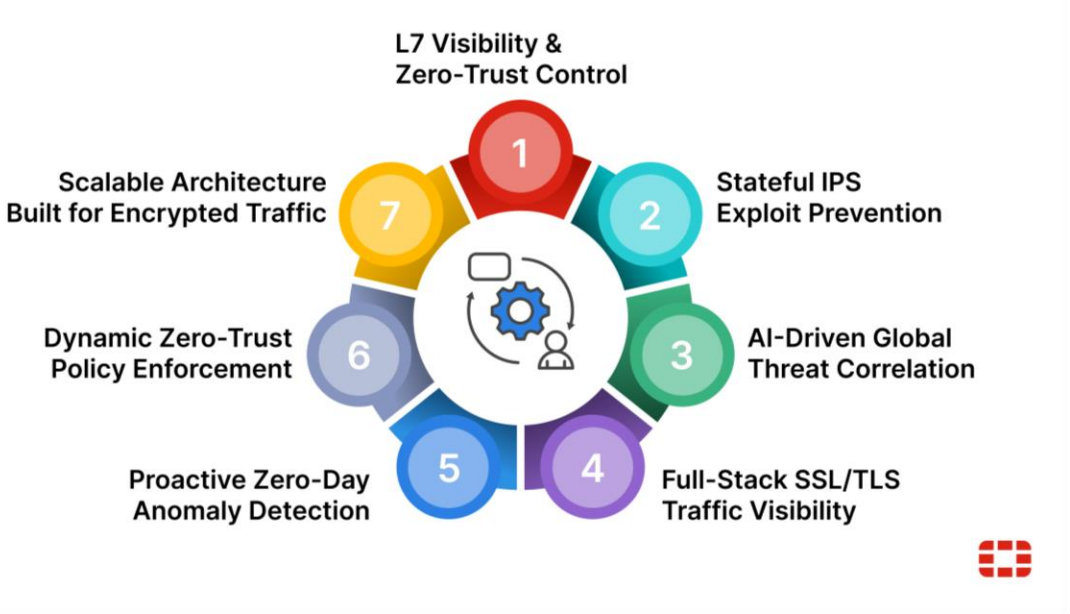

What an NGFW still cannot guarantee

- Integration does not eliminate security gaps.
- An NGFW cannot guarantee:
  - Complete visibility: Encryption can hide application content
  - Perfect detection: False positives and false negatives remain
  - Safe endpoints: Compromised systems can use permitted connections
  - Correct policy: Poor rules still create unnecessary access
  - Unlimited capacity: Inspection can affect performance and availability
- **A powerful control is still only one layer of defence.**

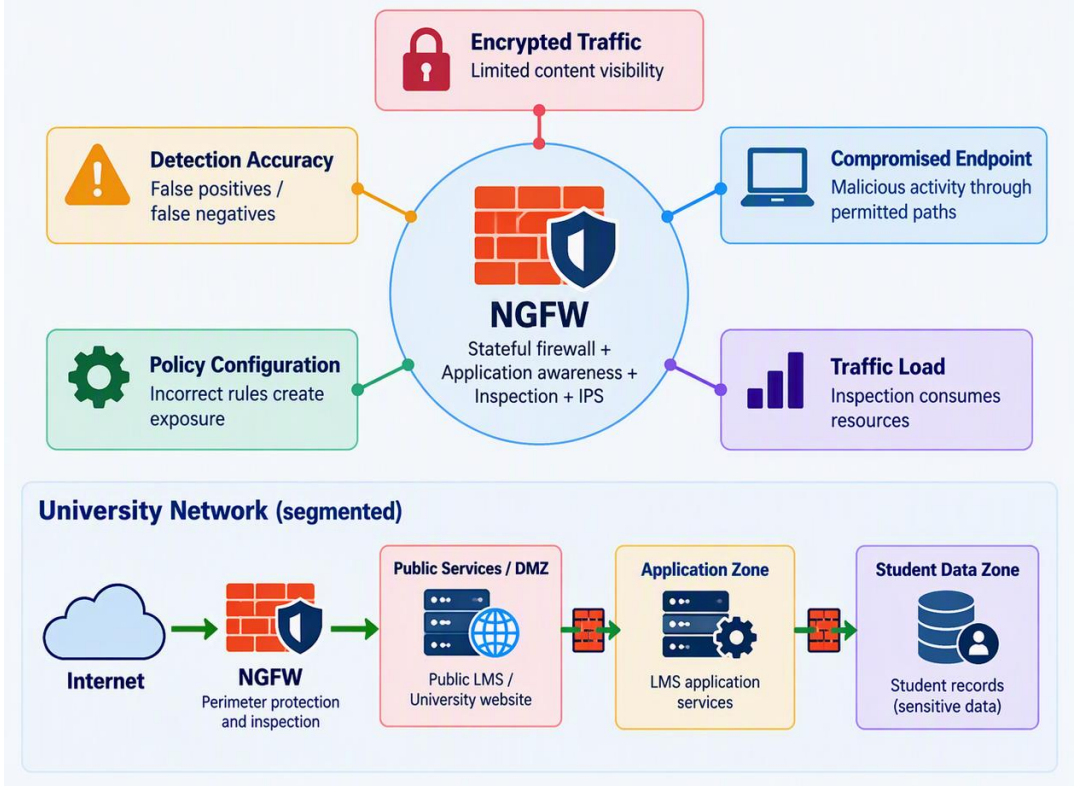

SASE: Security beyond the campus perimeter

- The university perimeter is becoming distributed.
- SASE combines networking and cloud-delivered security to support:
  - Main and regional campuses
  - Remote students and staff
  - Cloud applications and services
  - Consistent security policy across locations
- **Bring security enforcement closer to the user and application, not everything back to one campus.**

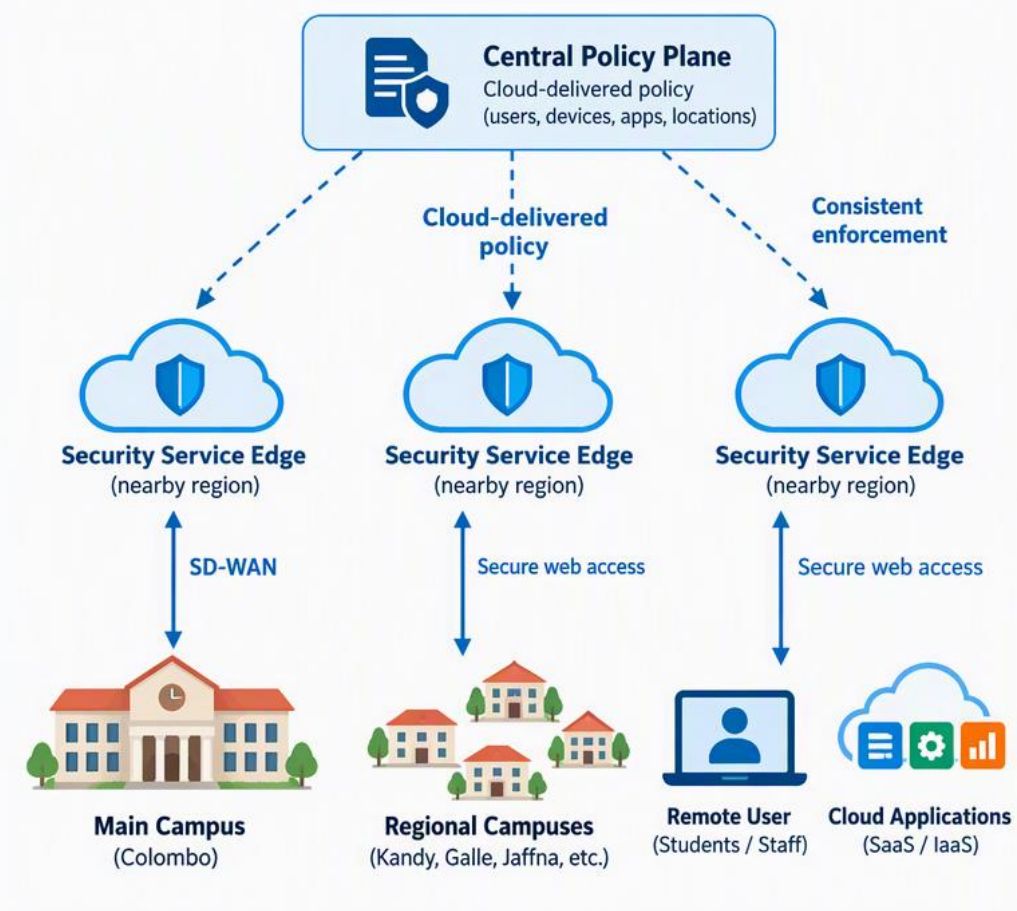

XDR: Connect security evidence across domains

- One incident may appear across several security systems.
- XDR can bring together signals from:
  - Network
  - Endpoints
  - Email
  - Cloud services
- **The goal is broader context, not another network firewall.**

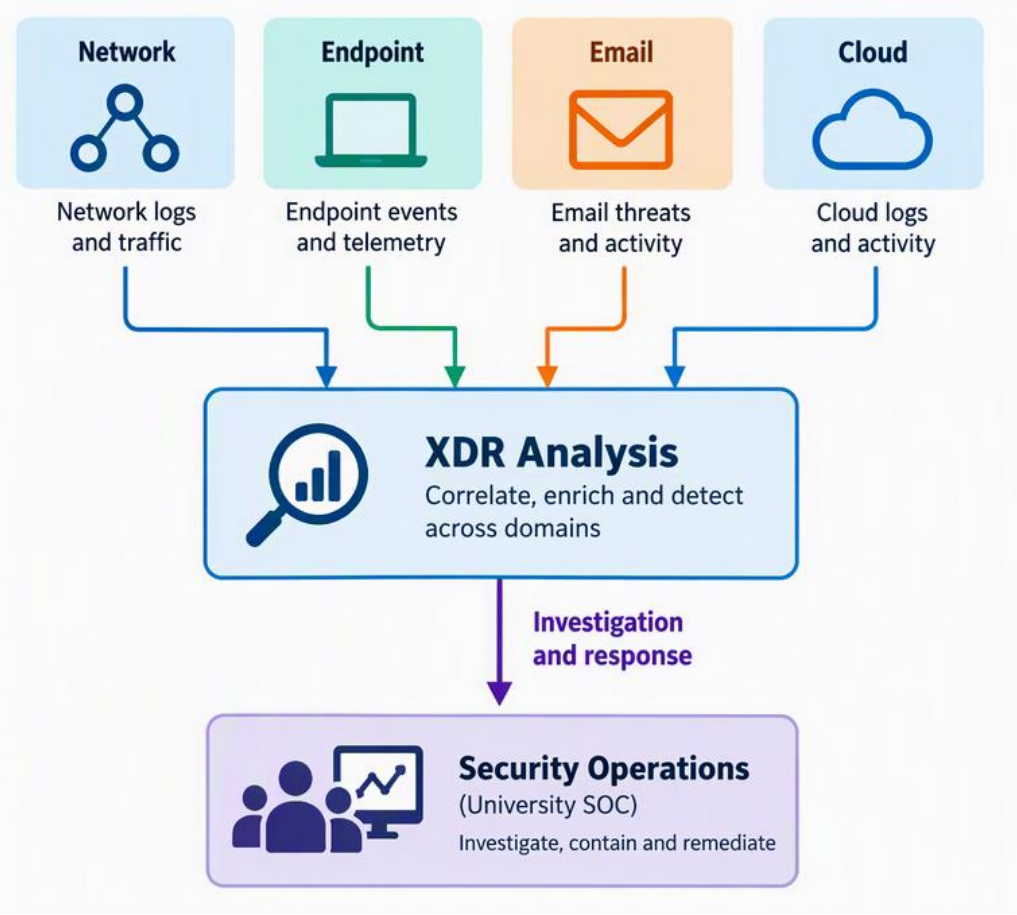

SASE and XDR solve different problems

- Do not confuse enforcement with detection context.

	SASE	XDR
Problem	Distributed access	Fragmented security evidence
Focus	Networking + security enforcement	Detection + response context
Connects	Users, sites and applications	Security signals across domains
Think	Where do we enforce?	What evidence connects?

- **SASE protects distributed access. XDR helps understand activity across security domains.**

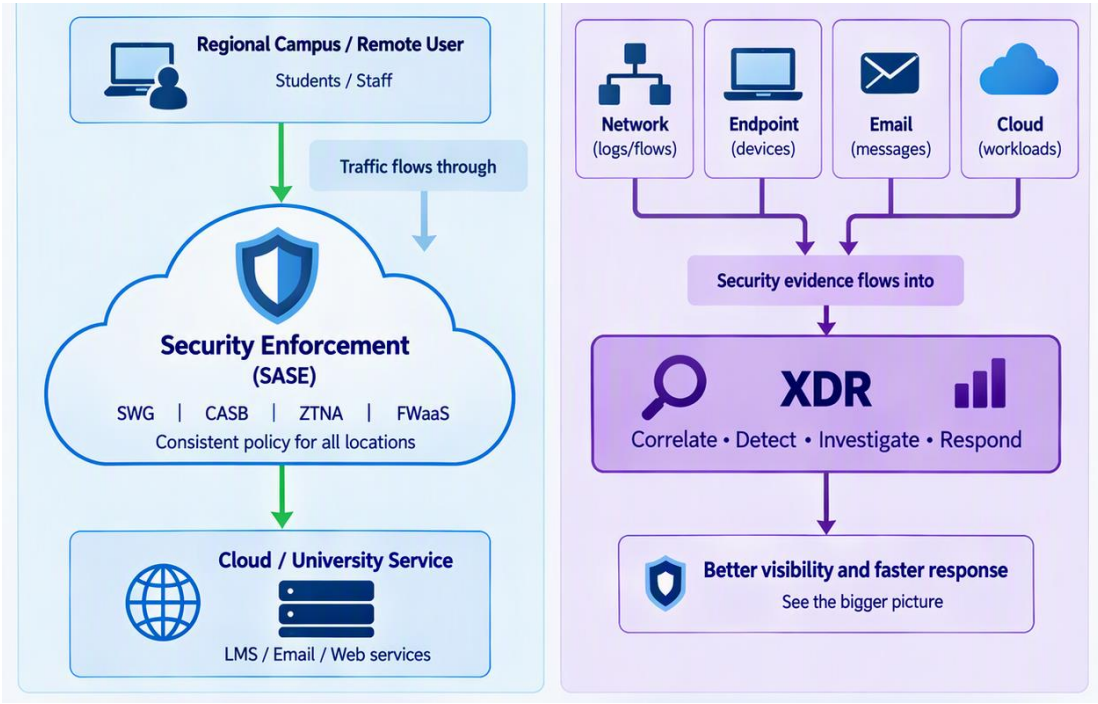

Final design challenge: Stop the lateral move

- Where would you break the attack path?
  - Internet → Public LMS → Application Zone → Student Data
- For each boundary, decide:
  - Allow: Which communication is genuinely required?
  - Restrict: Which paths should not exist?
  - Detect: Where should suspicious behaviour be observed?
  - Block: Where is inline prevention justified?
- **Design for compromise, then limit how far it can spread.**

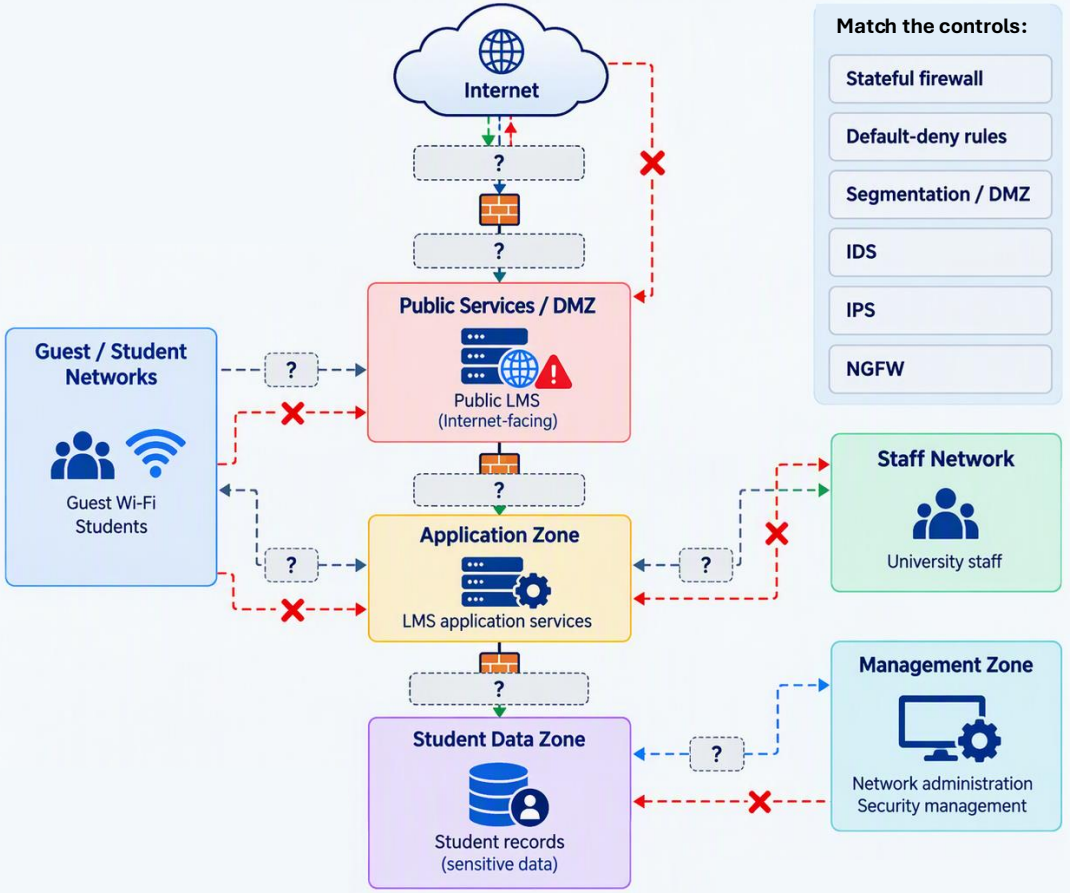

Five operational network-security lessons

- Five engineering lessons:
  1. Permit only required communication
  2. Segment to limit lateral movement
  3. Inspect only what is visible and justified
  4. Detect broadly — block carefully
  5. Assume controls can fail
- **Effective security comes from layered policy, visibility and containment, not one device.**

