Cloud Networking

Where is the network boundary now?



ITN 4252 Advanced Topics in Networking

Partially based on Cloud Computing and Cloud Networking by Adel Nadjaran Toosi, Introduction to Cloud Computing by Indranil Gupta, and resources from AWS.

The university LMS is moving to the cloud

- It must move before enrolment opens
- The new environment must support:
  - Students and staff across Sri Lanka
  - Secure access from the main and regional campuses
  - Public access to learning services
  - Private access to admin and student data
  - Continued service during an infrastructure failure
- Moving the application is easy. Designing the network around it is harder.



The LMS is online, but is it safe and resilient?

- All components run in one Availability Zone (AZ)
- The database is reachable through an Internet-routed subnet
- A broad rule allows access from any IP address
- Regional campuses have no private cloud connection
- Network logging and monitoring are not enabled
- **Working is not the same as well designed.**

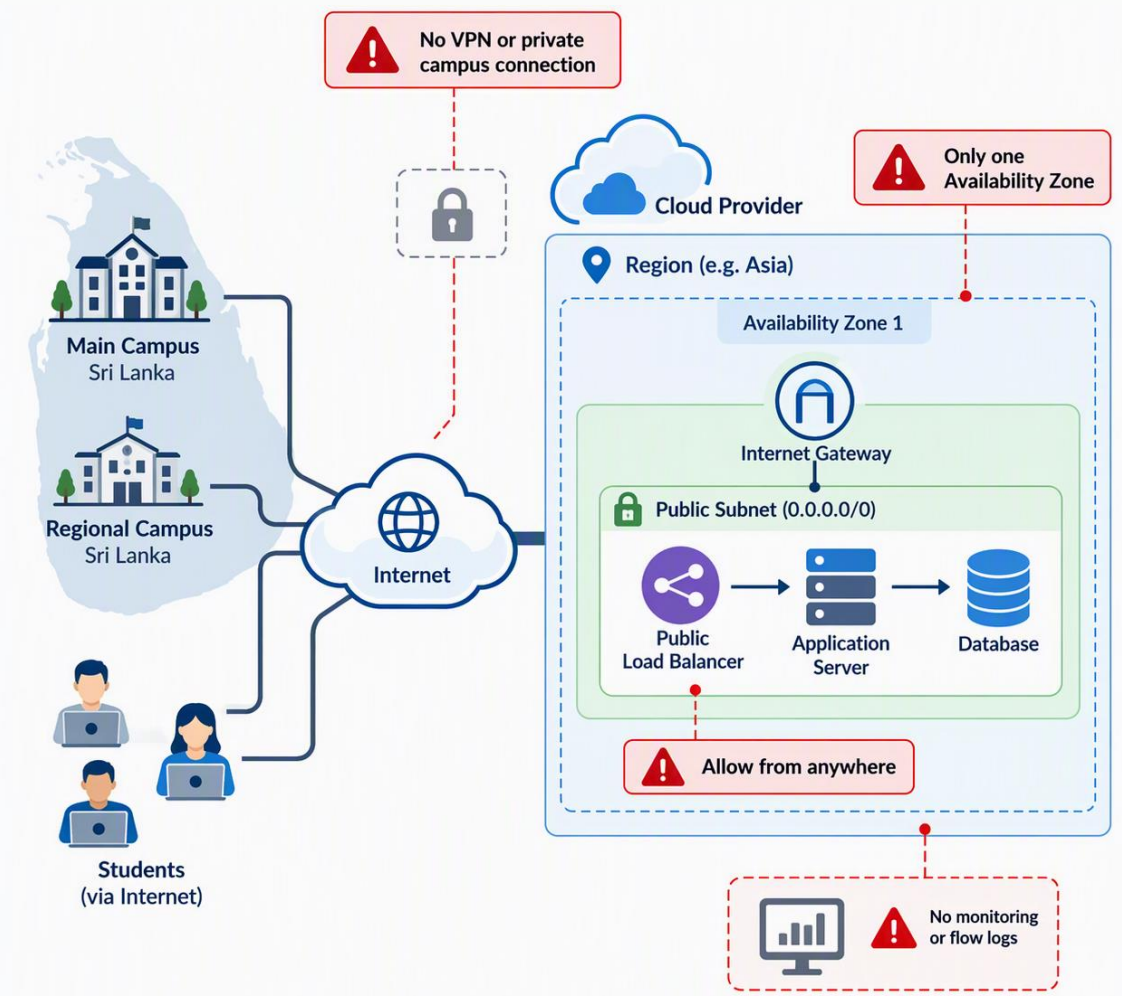

What makes an environment a cloud?

- Cloud is more than remote hosting.
- A cloud environment provides:
  - On-demand access to shared resources
  - Rapid provisioning and release
  - Elastic capacity
  - Access through networks and APIs
  - Usage that can be measured
- **Virtualisation provides resources.**
  Cloud makes them available as an on-demand service.

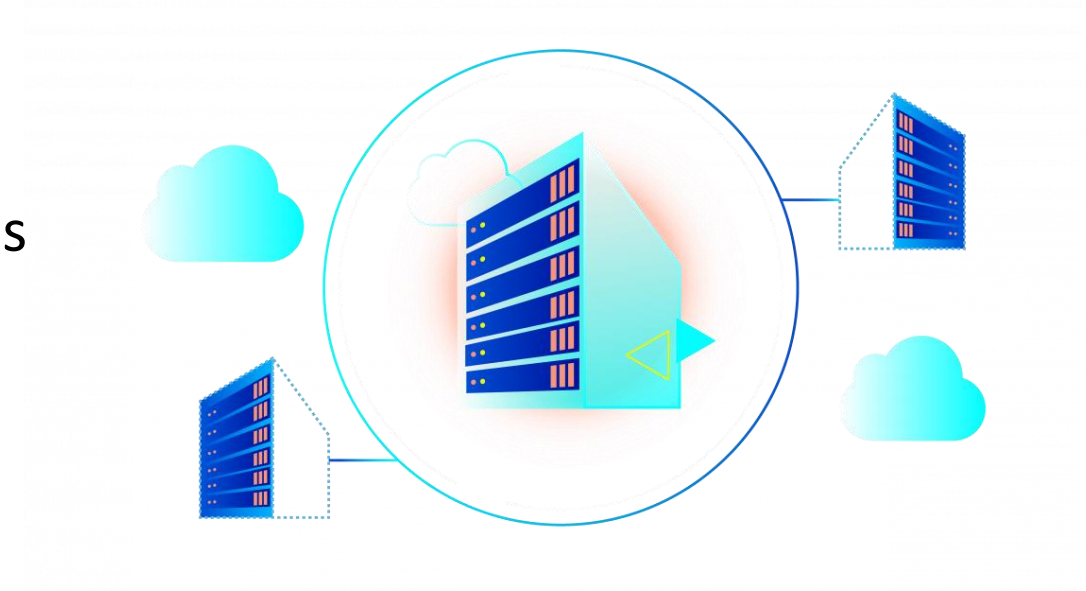

Five characteristics make cloud different

- On-demand self-service
  - Provision resources when needed
- Broad network access
  - Access services through standard networks
- Resource pooling
  - Share provider resources across customers
- Rapid elasticity
  - Scale capacity with changing demand
- Measured service
  - Monitor and charge based on usage
- **Cloud is defined by how resources are delivered, not where the servers are located.**

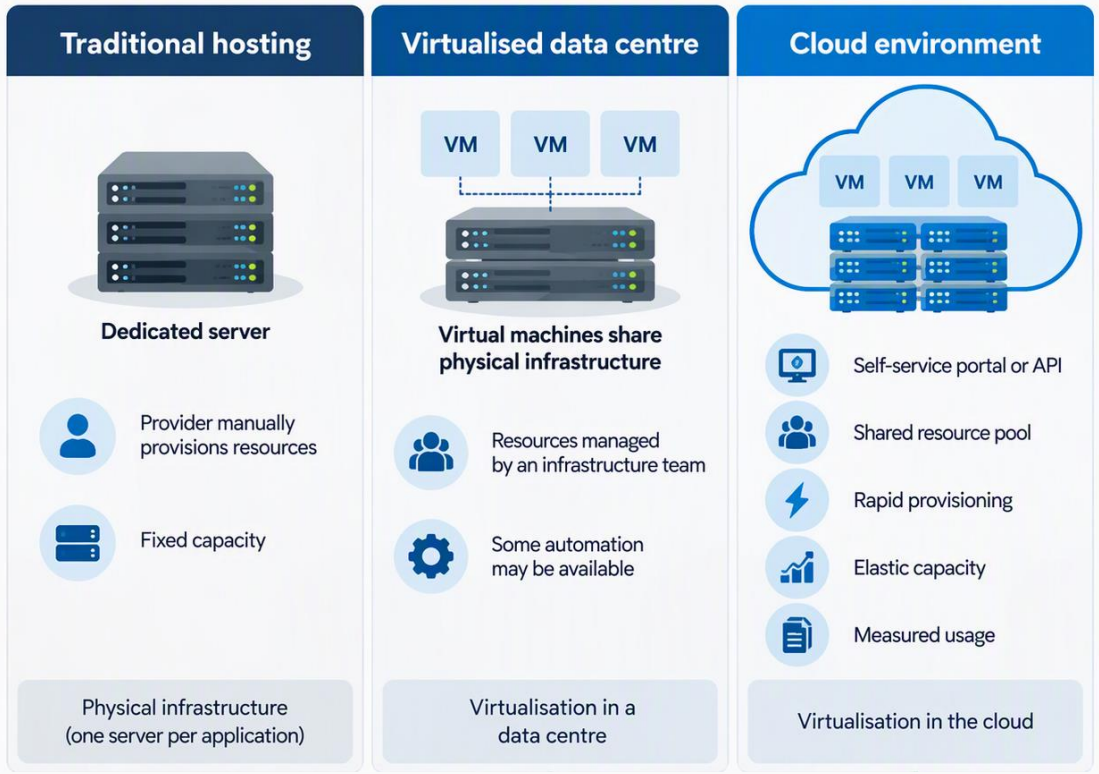

Service models change customer responsibility

- More managed service, less infrastructure responsibility.
- IaaS
  - Customer manages the operating system, applications and virtual network
- PaaS
  - Customer manages the application, data and service configuration
- SaaS
  - Customer uses the application and manages users, data and settings
- **The provider manages more of the stack, but the university always retains responsibilities.**

Technology layer	IaaS	PaaS	SaaS
Application and data	University	University	Shared
Runtime and operating system	University	Provider	Provider
Virtual networking	University	Shared	Provider
Physical infrastructure	Provider	Provider	Provider

Shared responsibility, you cannot outsource risk

- Cloud security is a shared responsibility.
- Cloud provider
  - Protects the underlying cloud infrastructure
- University
  - Protects its data, identities and configurations
- Shared
  - Both parties contribute to security and service resilience
- **Moving to the cloud transfers some responsibilities, not the university's risk.**

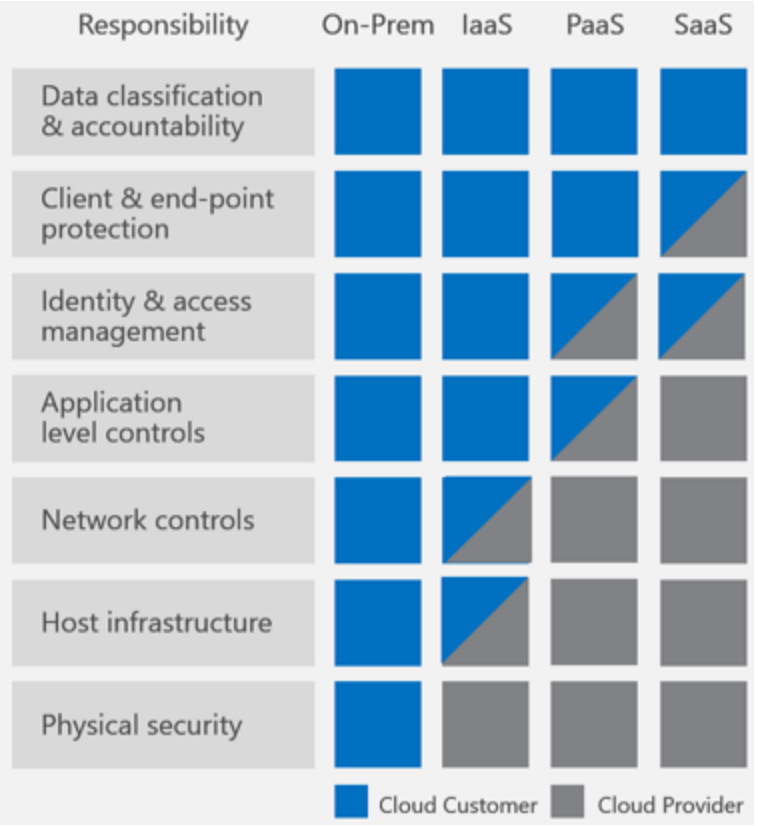

Public, private, hybrid and multi-cloud

- Where should the LMS services run?
- Public cloud
  - Provider-operated, shared cloud infrastructure
- Private cloud
  - Dedicated cloud environment for one organisation
- Hybrid cloud
  - Connected public-cloud and private environments
- Multi-cloud
  - Services from more than one cloud provider
- **Choose the model according to control, risk, integration and operational capability.**

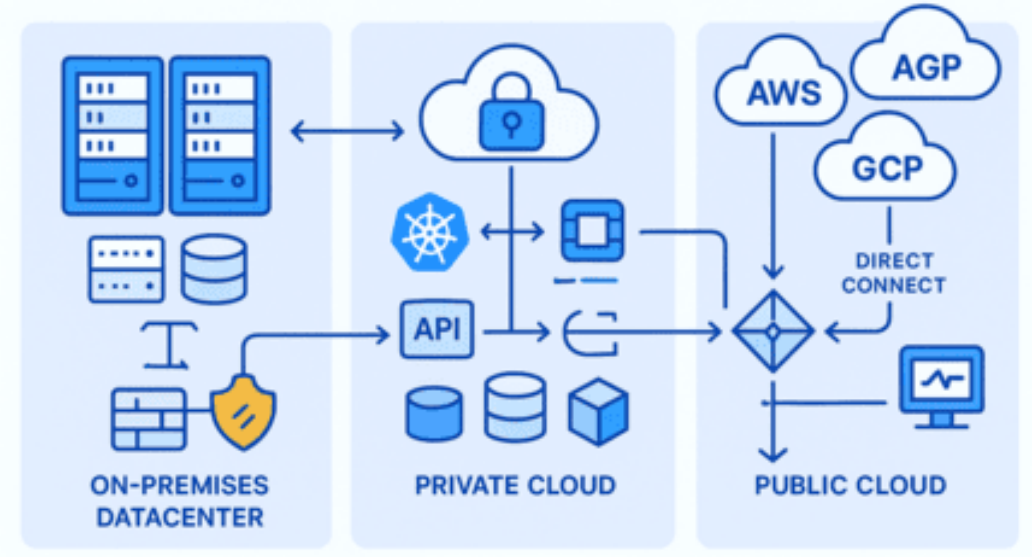

Cloud geography, region, zone and edge

- Cloud services operate at different geographic layers.
- Region
  - A geographic area containing cloud infrastructure
- Availability Zone
  - An isolated failure domain within a region
- Edge location
  - A network location closer to users for content delivery
- **Region affects location. Zones affect resilience. Edge affects performance.**

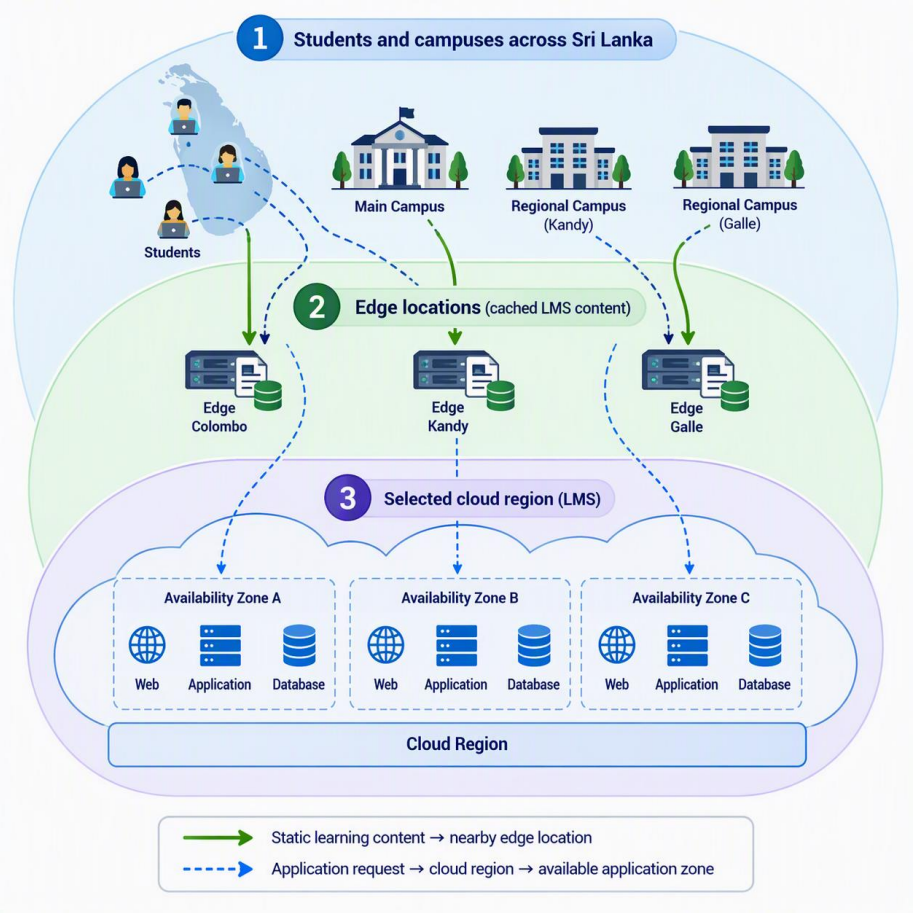

Which cloud region should host the LMS?

- Region selection is an architecture decision.
- Requirements
  - Data location and compliance obligations
- Performance
  - Latency from users and campuses across Sri Lanka
- Capabilities
  - Required services, zones and capacity
- Cost
  - Resources, connectivity and data transfer
- **The closest region is a candidate, not an automatic choice.**



One zone is still one failure domain

- If the LMS runs in only one zone:
  - Application servers may become unavailable
  - Database access may be interrupted
  - New capacity cannot be added elsewhere
  - The entire LMS path may fail
- **Cloud-hosted does not automatically mean highly available.**

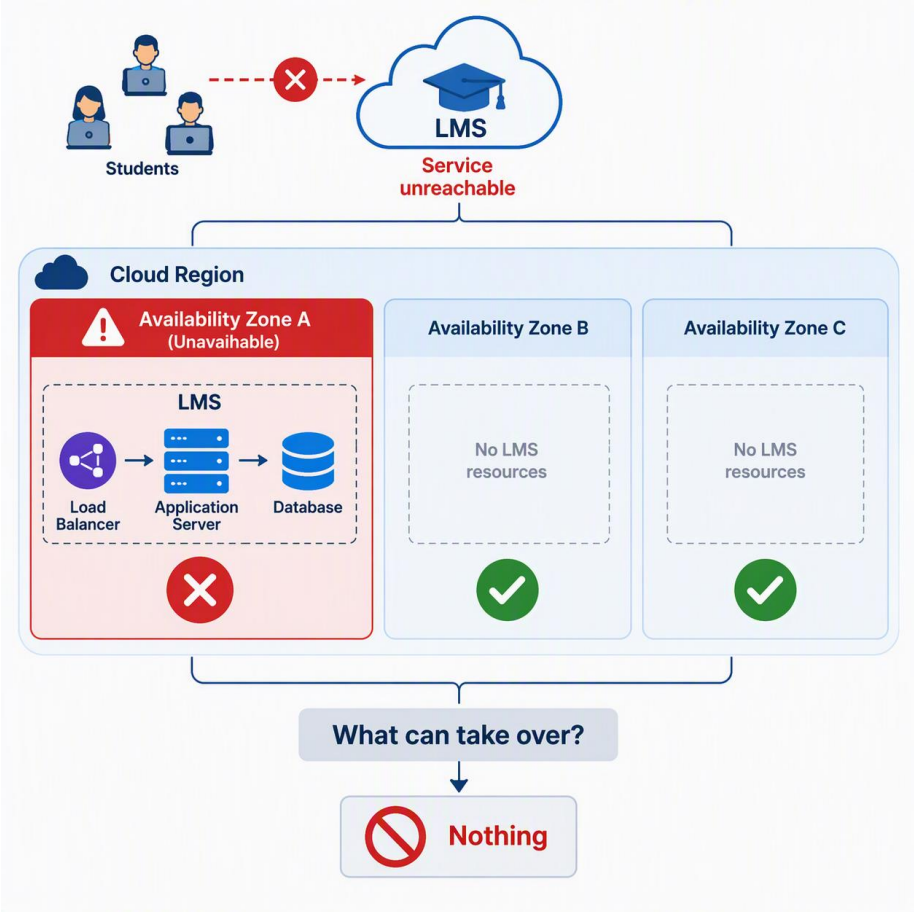

High availability must be designed

- Resilience must cover every critical layer.
  - Distribute application instances across zones
  - Route requests only to healthy instances
  - Replicate data across failure domains
  - Maintain capacity when one zone is unavailable
  - Test failover and recovery regularly
- **Redundancy without failover is only duplication.**

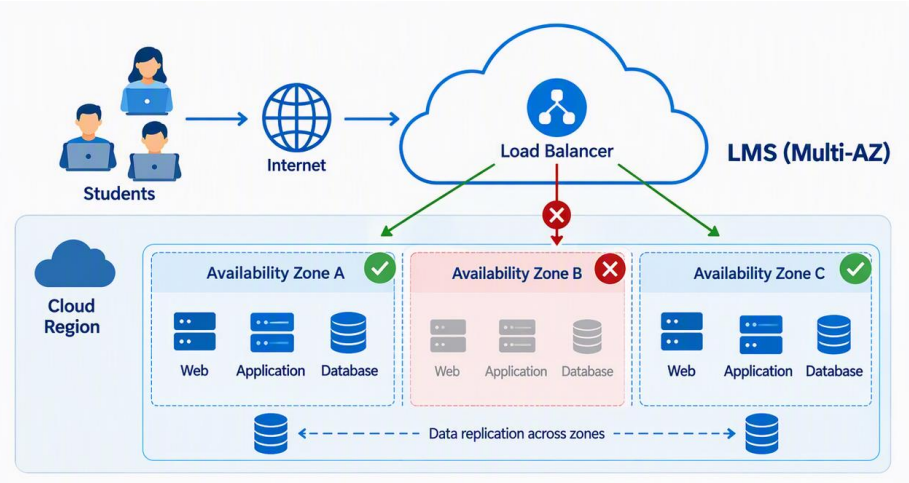

The VPC, a logical network boundary

- The VPC creates the LMS network boundary.
- Within the virtual network, the university defines:
  - IP address ranges
  - Subnets and workload placement
  - Routes and network gateways
  - Security and access controls
  - Connections to campuses and other networks
- The provider supplies the infrastructure. The university designs the virtual network.

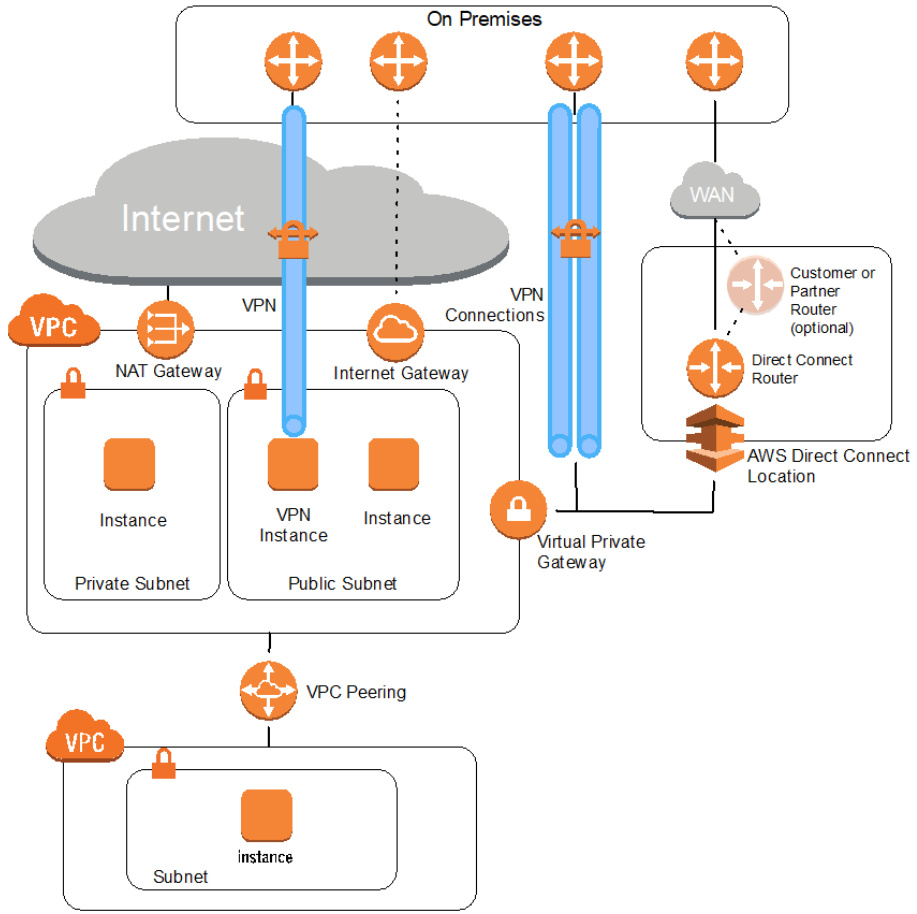

Addressing and CIDR planning

- Plan the address space before connecting the cloud.
  - Use a private, non-overlapping CIDR range
  - Allow capacity for future workloads and campuses
  - Divide the range into logical subnet tiers
  - Document allocations across cloud and campus networks
  - Consider both IPv4 and IPv6 requirements
- Addressing decisions made today determine which networks can connect tomorrow.

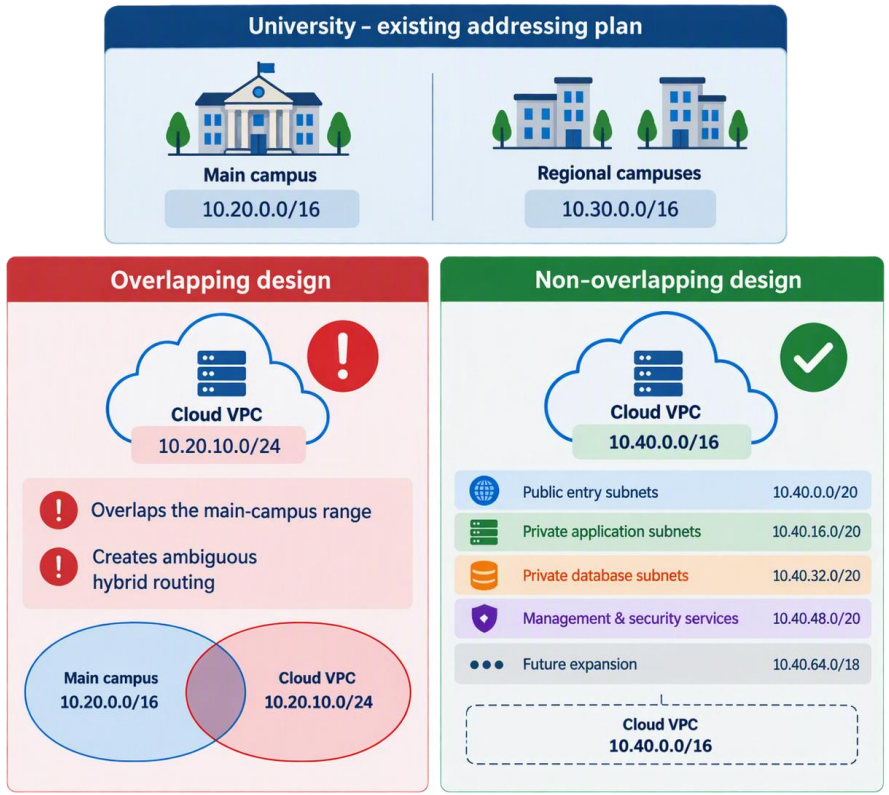

Subnets form placement, routing and security

- Subnets divide the cloud network into logical tiers.
- Public entry tier
  - Internet-facing load balancers and gateways
- Private application tier
  - LMS application servers and services
- Private data tier
  - Databases and sensitive information
- Across failure domains
  - Create equivalent subnet tiers in multiple zones
- **Place each component according to who needs to reach it.**

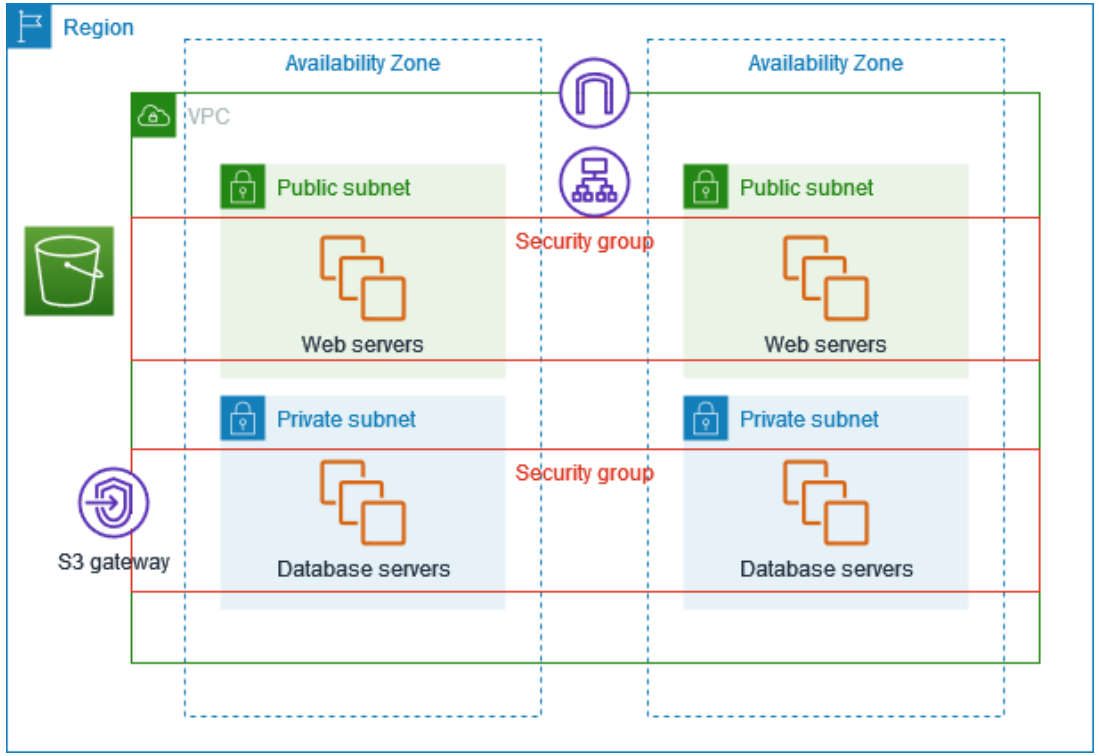

Public, private and isolated routing outcomes

- A subnet’s route determines its role.
- Public subnet
  - Direct route to an Internet Gateway
- Private subnet
  - No direct Internet Gateway route
- Isolated subnet
  - No route beyond the virtual network
- **Names do not create isolation. Routes and controls do.**

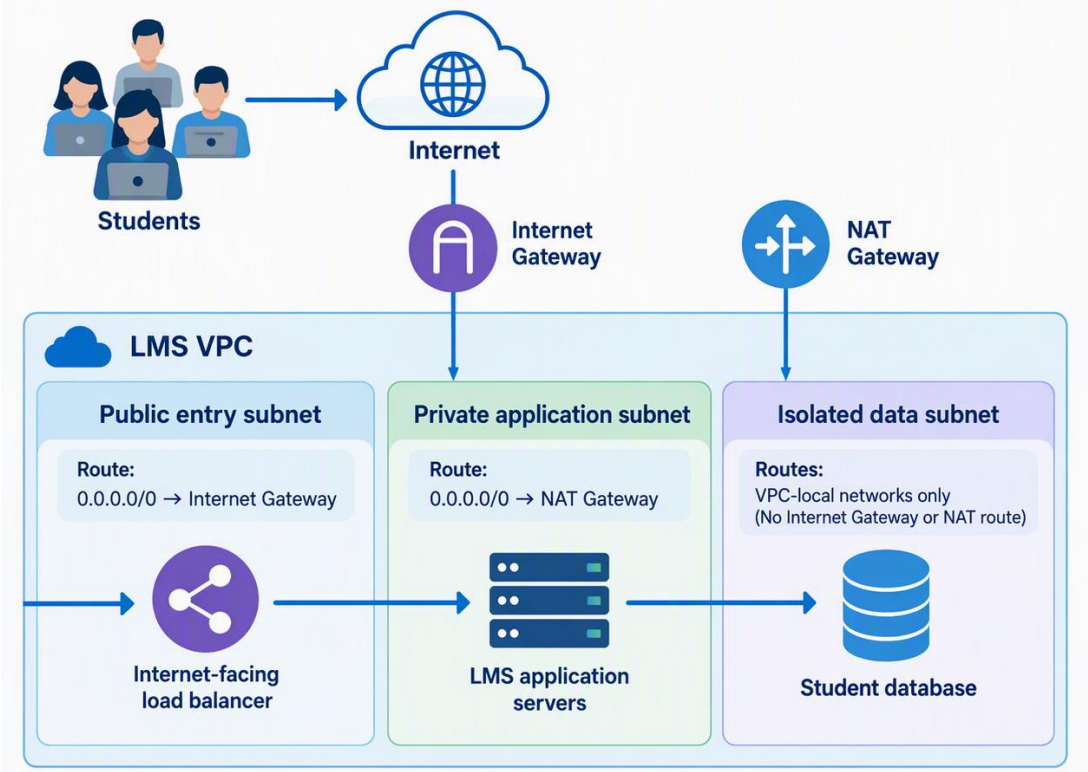

Route tables decide the next hop

- Every route answers two questions:
  1. Destination: Which address range does the traffic match?
  2. Next hop: Where should matching traffic be sent?
- Route selection
  - The most specific matching route is preferred
- **A correct security rule cannot fix an incorrect route.**

Public entry subnet

Destination	Next hop
VPC address range	Local
0.0.0.0/0	Internet Gateway

Private application subnet

Destination	Next hop
VPC address range	Local
Campus networks	VPN Gateway
0.0.0.0/0	NAT Gateway

Isolated data subnet

Destination	Next hop
VPC address range	Local
Approved campus administration network	Private connection

Internet gateway, NAT and private endpoints

- Different destinations need different connection paths.
- Internet Gateway
  - Direct Internet connectivity for public resources
- NAT Gateway
  - Outbound access for private resources
- Private endpoint
  - Private access to supported cloud services
- **Use the least exposed path that meets the application requirement.**

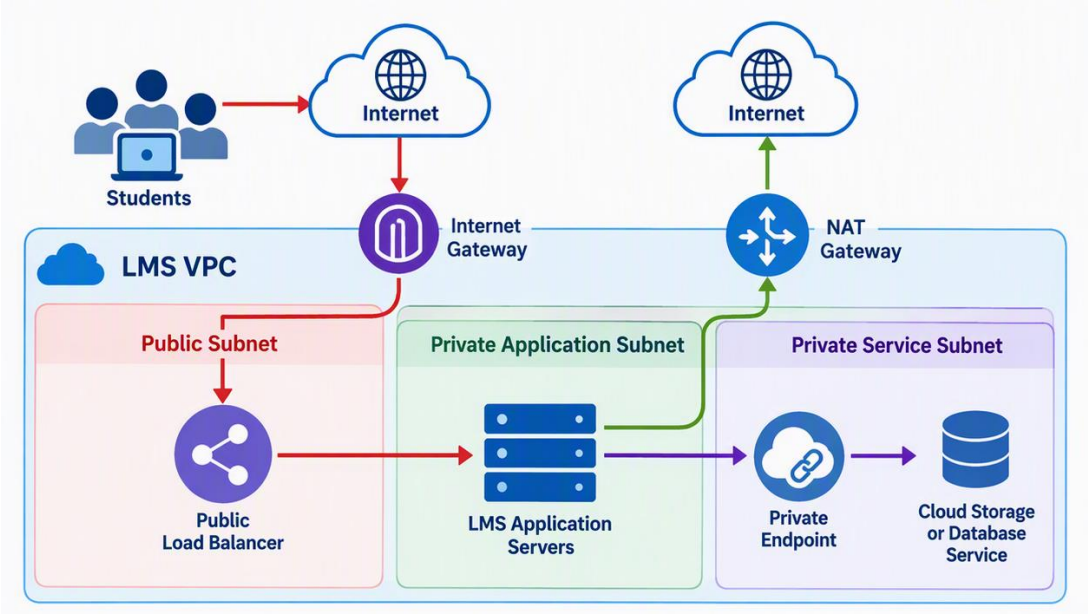

Public entry path

- Students → Internet Gateway → Public load balancer

Private outbound path

- Application servers → NAT Gateway → Internet

Private service path

- Application servers → Private endpoint → Cloud database

Connect the university to the cloud

- Campus systems need a private cloud path
- Site-to-site VPN
  - Encrypted connection across the public Internet
- Routing
  - Exchanges or configures campus and cloud network prefixes
- Redundancy
  - Uses multiple tunnels, gateways or connectivity paths
- Security
  - Restricts which campus systems can reach cloud resources
- **Private connectivity extends the network boundary, not the trust boundary.**

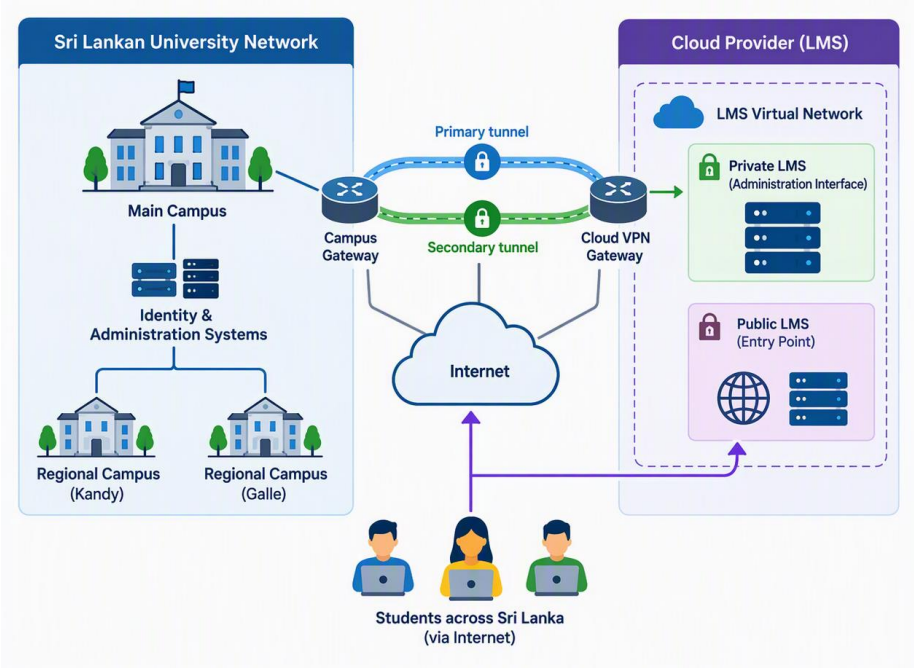

VPN or dedicated private connectivity?

- Choose connectivity according to the workload.
- Site-to-site VPN
  - Faster to establish
  - Encrypted over the public Internet
  - Suitable for pilots, backup and moderate traffic
- Dedicated private connection
  - More predictable performance
  - Higher-capacity options
  - Requires provider coordination and additional cost
- **Use VPN for agility. Use dedicated connectivity for predictable production demand.**

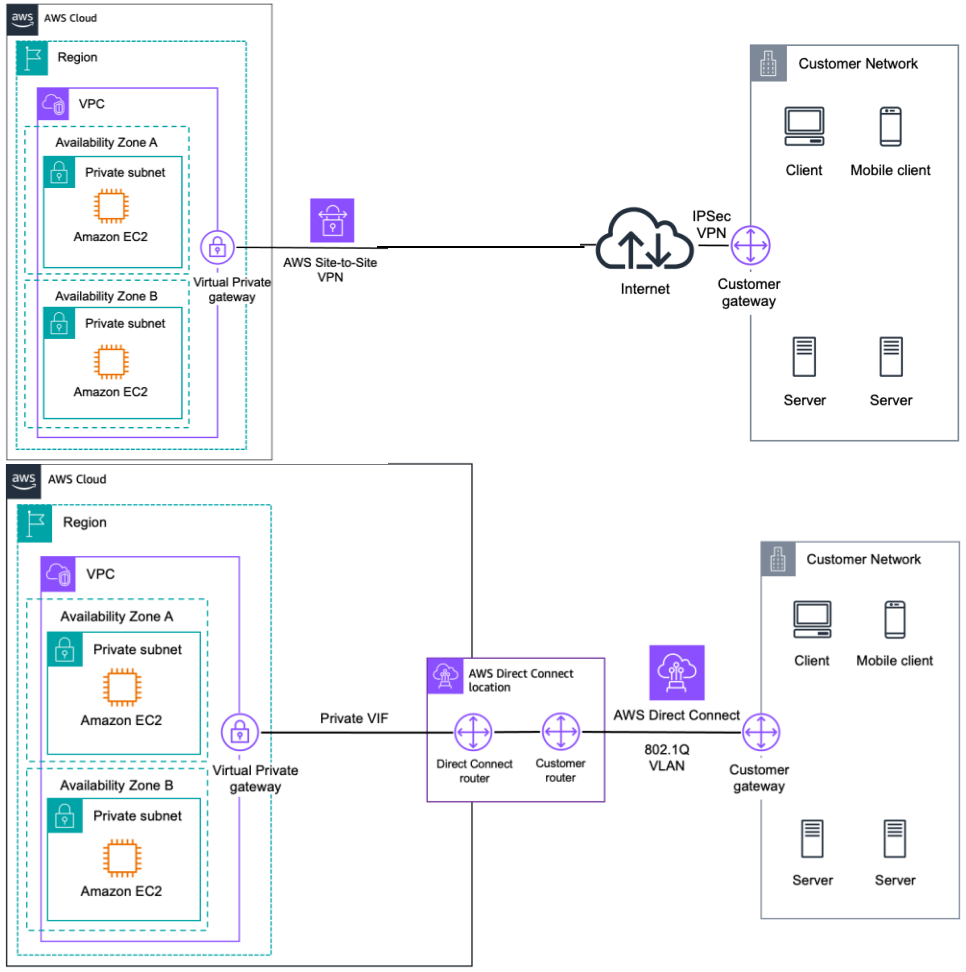

Follow a student request to the LMS

- Follow one request through the cloud network.
    1. DNS resolves the LMS name
    2. Public entry point receives the connection
    3. Load balancer selects a healthy application instance
    4. Network controls evaluate the traffic
    5. Application accesses the private data tier
- **Every successful request depends on several independent decisions.**

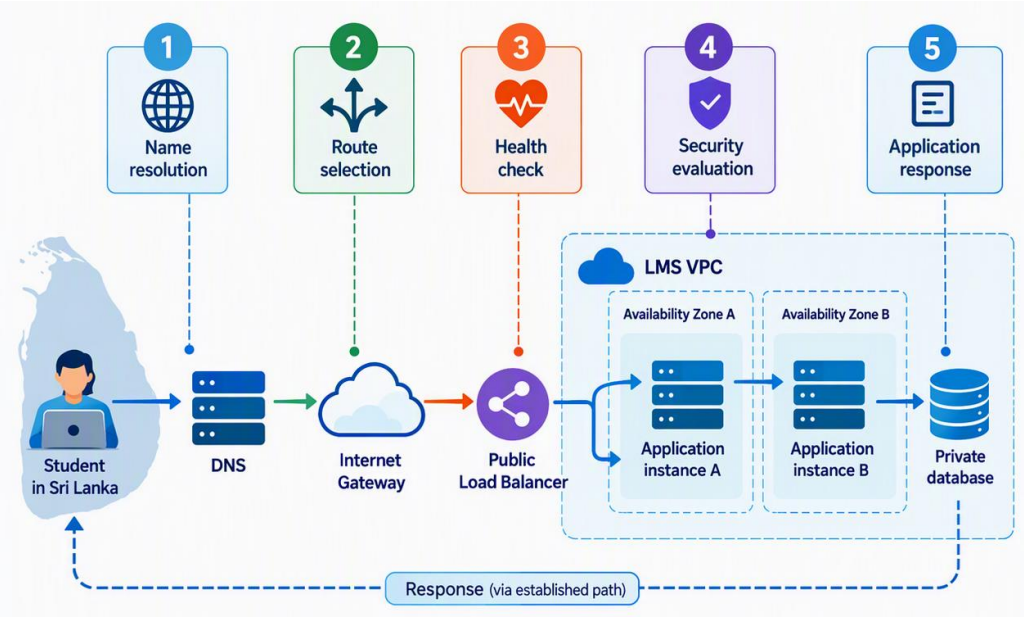

Does the firewall remember the connection?

- Stateless filtering
  - Evaluates inbound and outbound traffic separately
  - Requires rules for both request and response
  - Does not track connection state
- Stateful filtering
  - Tracks established connections
  - Automatically recognises permitted return traffic
  - Simplifies resource-level access control
- **Stateless checks each direction. Stateful remembers the conversation.**

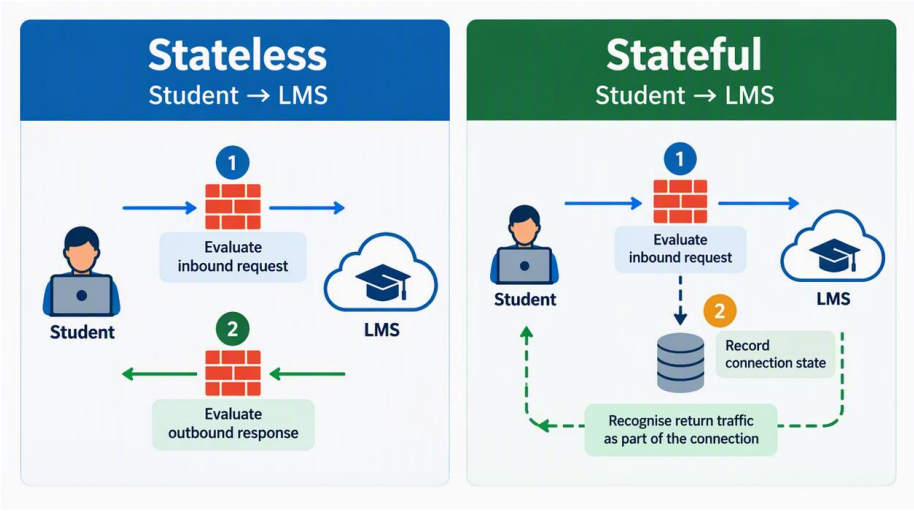

Layered cloud network security

- Apply security controls at multiple layers.
- Network boundary
  - Controls traffic entering or leaving the cloud network
- Subnet or segment
  - Restricts traffic between application tiers
- Workload
  - Allows only required connections to each resource
- Application
  - Inspects requests using application context
- **Each layer reduces exposure if another fails.**

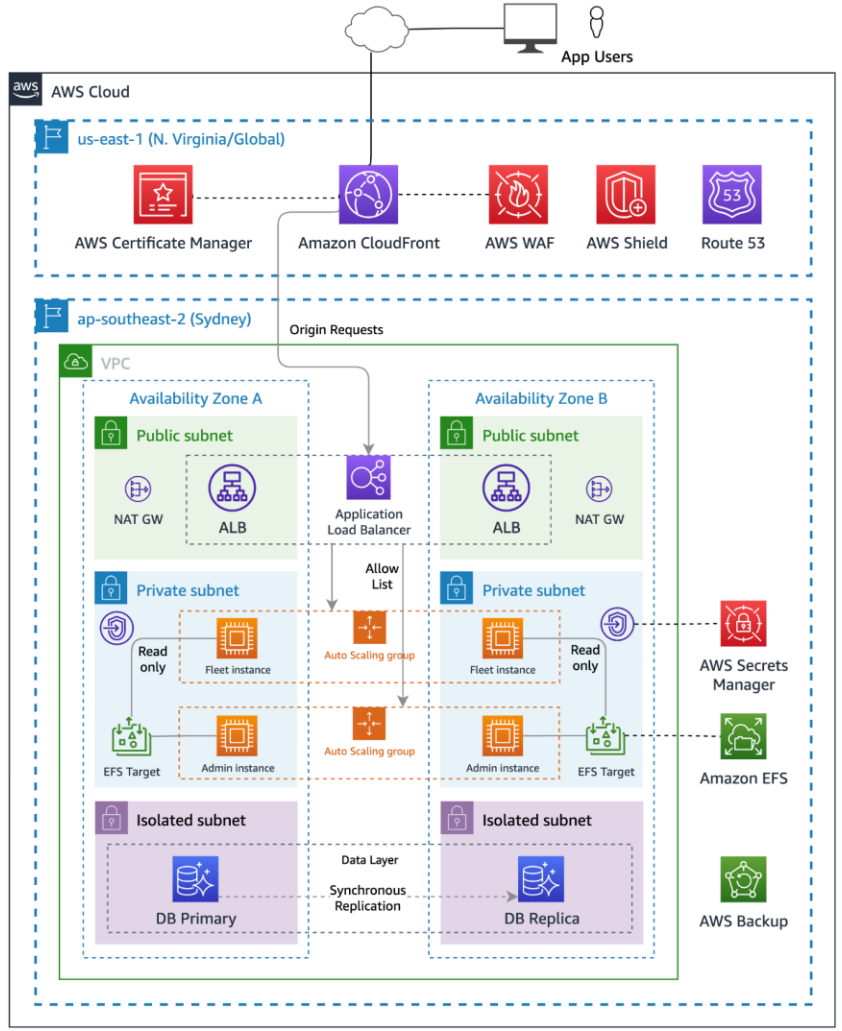

Cloud security models differ

- The security outcome is transferable.
- Across cloud platforms, understand:
  - **Scope**: Network, segment, workload or application
  - **Enforcement**: Where traffic is evaluated
  - **State**: Whether connection state is tracked
  - **Rules**: How allow, deny and priority are processed
  - **Identity**: Whether policy can follow workload roles
- **Compare behaviour, not product names.**

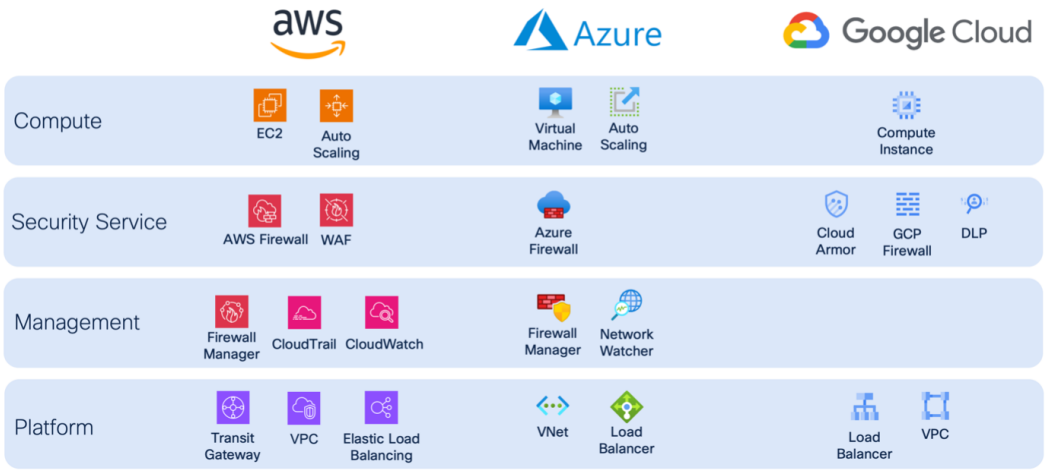

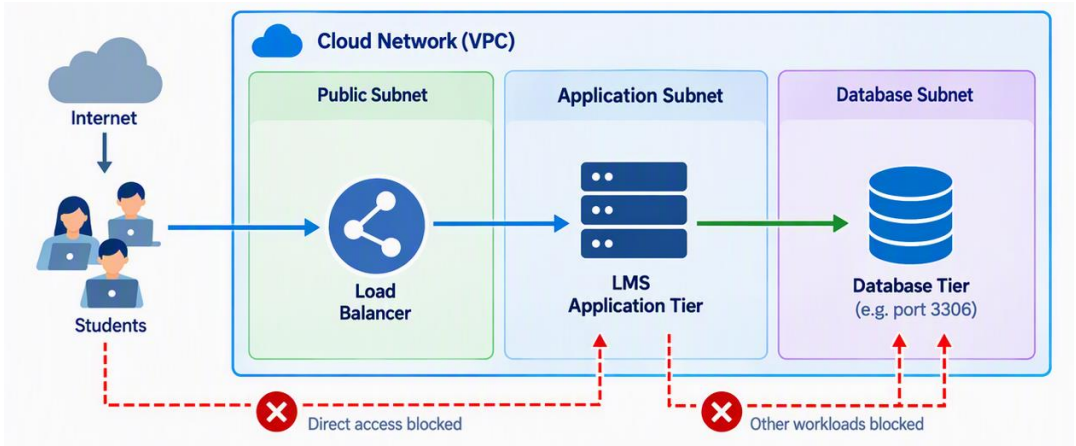

Least privilege and shared security responsibility

- Allow only what the LMS requires:
  - Expose only the public entry point
  - Permit only approved traffic between tiers
  - Restrict administration to authorised users and paths
  - Review rules, identities and exceptions regularly
  - Log changes and denied connections
- **Connected does not mean trusted.**

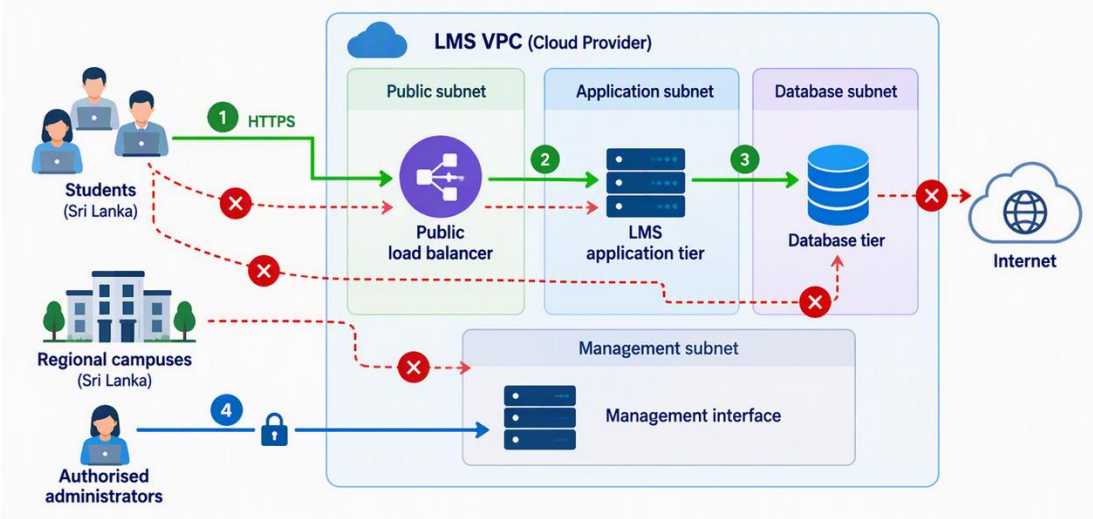

1. **Students → Public load balancer**
   Allow HTTPS
2. **Load balancer → Application tier**
   Allow the LMS application port
3. **Application tier → Database tier**
   Allow only the required database port
4. **Authorised administrators → Private management path**
   Allow approved administration access

Cloud network observability

- You cannot troubleshoot what you cannot see.
- Monitor:
  - **Flows**: Who connected to what?
  - **Routes**: Which next hop was selected?
  - **Security**: What was allowed or denied?
  - **Health**: Are gateways, tunnels and application targets available?
  - **Performance**: Where are latency, loss or capacity problems?
- **Observability turns a network diagram into operational evidence.**

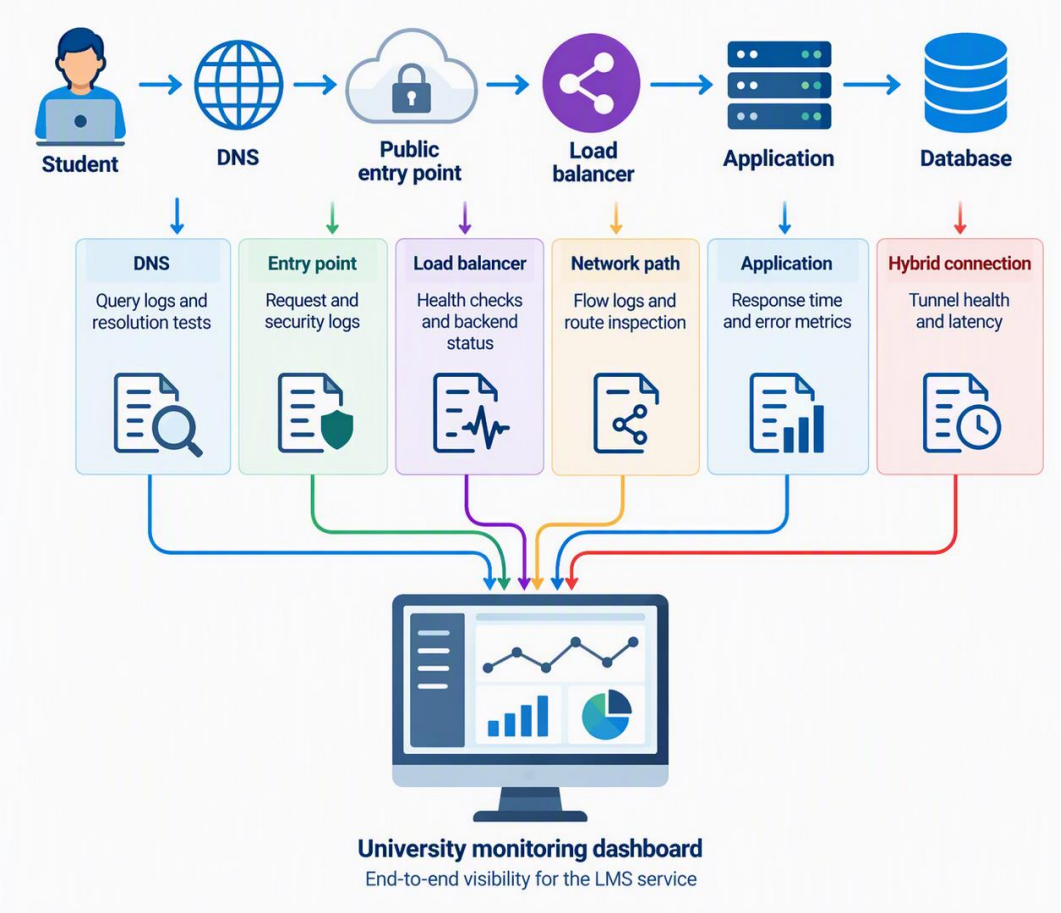

DNS, load balancing and application health

- Direct each request to a healthy service.
- DNS
  - Resolves the LMS name to an available entry point
- Load balancer
  - Distributes requests across application instances
- Health checks
  - Remove unhealthy instances from service
- Capacity
  - Add or remove instances as demand changes
- **Reachability is not enough. The destination must also be healthy.**

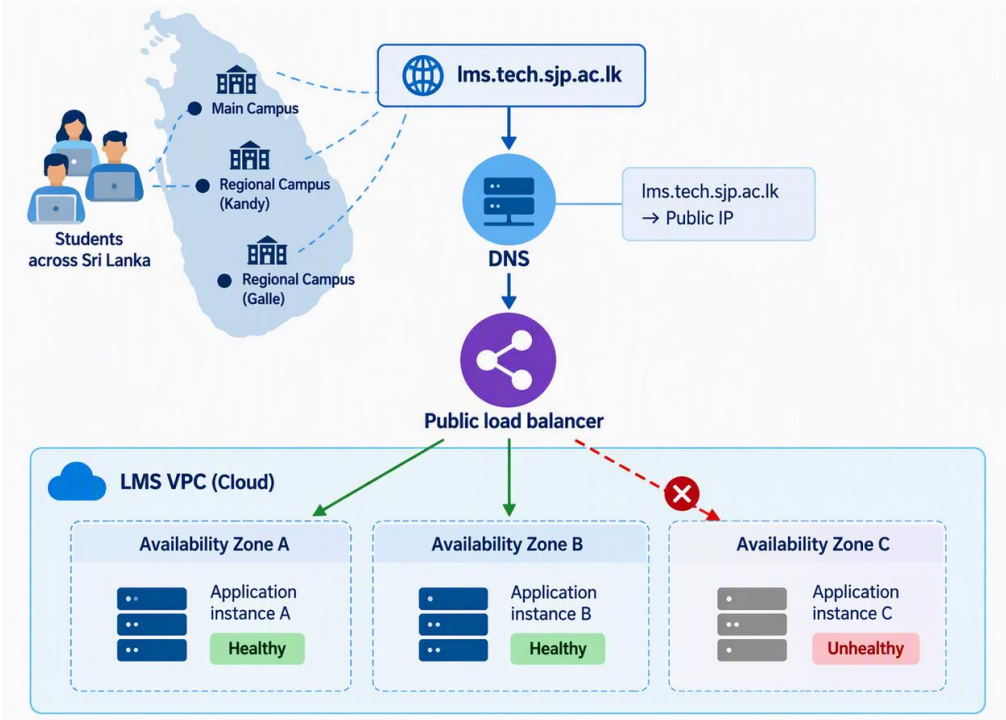

CDN, bring content closer to students

- Deliver static content from the network edge.
- Cache hit
  - The edge returns a stored copy
- Cache miss
  - The edge retrieves content from the origin
- Benefit
  - Lower latency and reduced origin demand
- Control
  - Caching rules determine what is stored and for how long
- **Move content closer to users, not the complete application.**

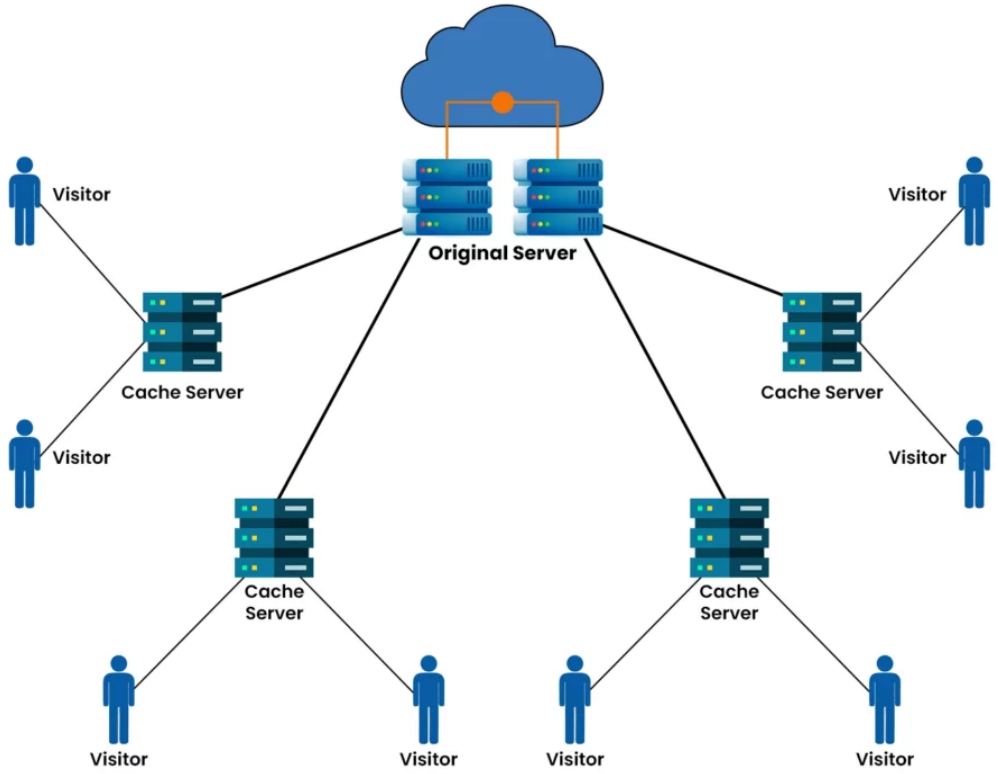

Infrastructure as Code

- Treat the cloud network as code.
  - Define the desired architecture in configuration files
  - Store every change in version control
  - Review and test changes before deployment
  - Deploy environments consistently
  - Detect and correct configuration drift
- **If the architecture matters, it should be repeatable and reviewable.**

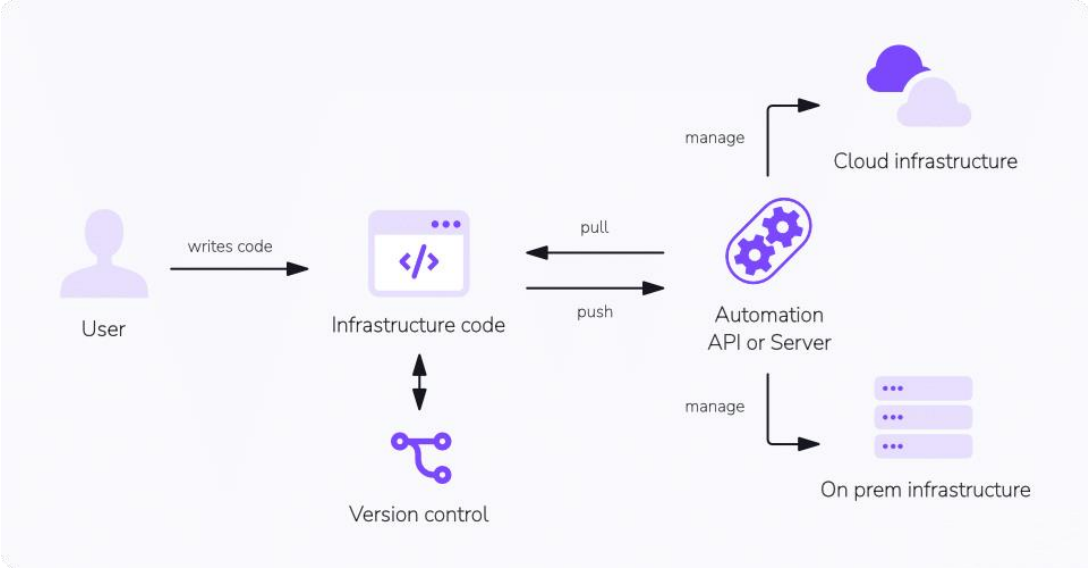

Final Design Challenge, Defend the Architecture

- Can this LMS design survive real-world conditions?
- Validate the complete architecture:
  - **Reachability**: Students can access the LMS
  - **Privacy**: Application and data tiers remain private
  - **Resilience**: Service continues during a zone failure
  - **Security**: Only required traffic is permitted
  - **Operations**: Failures can be detected and investigated
- **Every connection, control and dependency must have a reason.**

Questions?

- Cloud provides capabilities, not automatic resilience
- Routes determine reachability
- Layered controls protect each application tier
- Observability provides operational evidence
- The university remains responsible for its design and configuration
- **Cloud changes where the network runs, not the need for network engineering.**

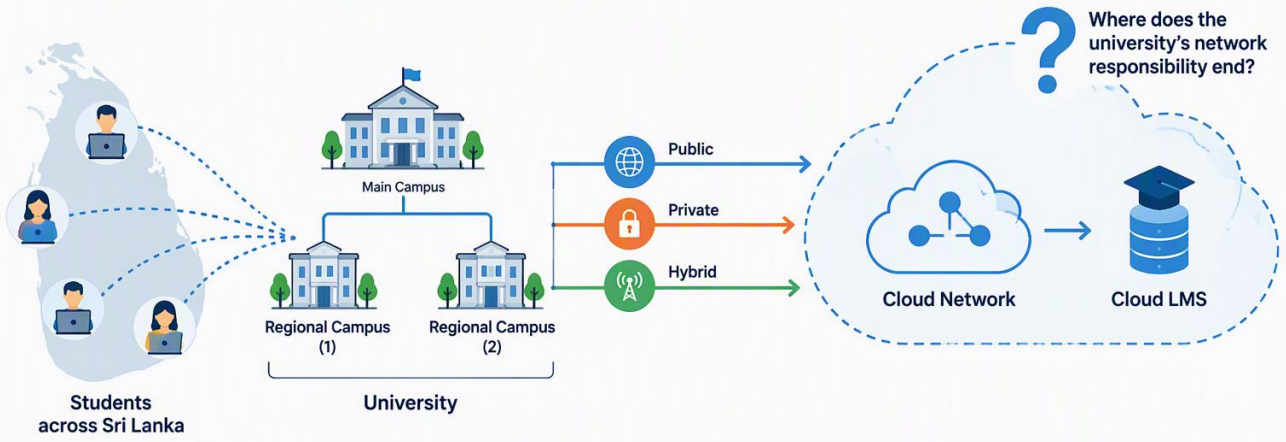

Additional resources

- AWS Educate
  - Free cloud learning platform with labs, courses, and certifications.
  - https://aws.amazon.com/education/awseducate/
- Microsoft Learn
  - Interactive modules and certifications for Azure networking services.
  - https://learn.microsoft.com/en-us/training/azure/networking
- Google Cloud Training
  - Hands-on labs and learning paths for cloud networking and security.
  - https://cloud.google.com/learn/training/networking-security
- Oracle Cloud Infrastructure Training
  - Courses and certifications for Oracle Cloud networking and infrastructure.
  - https://www.oracle.com/au/education/training/oracle-cloud-infrastructure/