Network Control Plane

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

The LMS is running, but students can’t reach it

- The application is healthy
- DNS resolves correctly
- Main campus access is working
- Remote and other campus students cannot connect
- **If the destination has not moved, what changed?**



Where would you start investigating?

- Application
  - Is the LMS service responding?
- Name resolution
  - Does DNS return the correct address?
- Local network
  - Can affected users reach their default gateway?
- Network path
  - Where does traffic stop or change direction?
- Good troubleshooting follows the path, not the assumption.

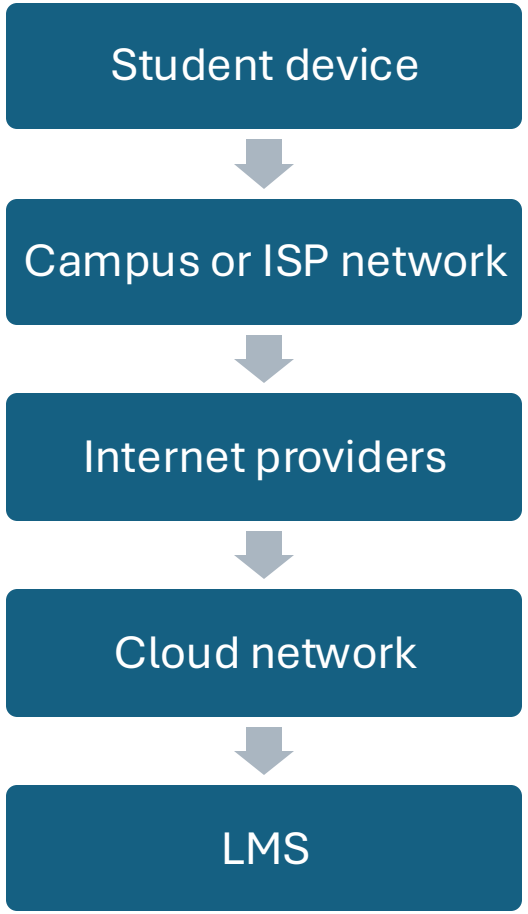

One packet, two different responsibilities

Forwarding

- Moves the packet to the correct output interface
  - Happens at each router
  - Uses an existing forwarding table
  - Operates packet by packet

Routing

- Determines which paths should be used
  - Learns network reachability
  - Calculates or selects routes
  - Updates the forwarding table



Forwarding moves packets. Routing decides the path.

Three planes, three responsibilities

Data plane

- **Forwards traffic**
- Examines packet headers
- Applies forwarding rules
- Sends or drops packets

Control plane

- **Determines paths**
- Exchanges reachability information
- Calculates or selects routes
- Populates forwarding tables

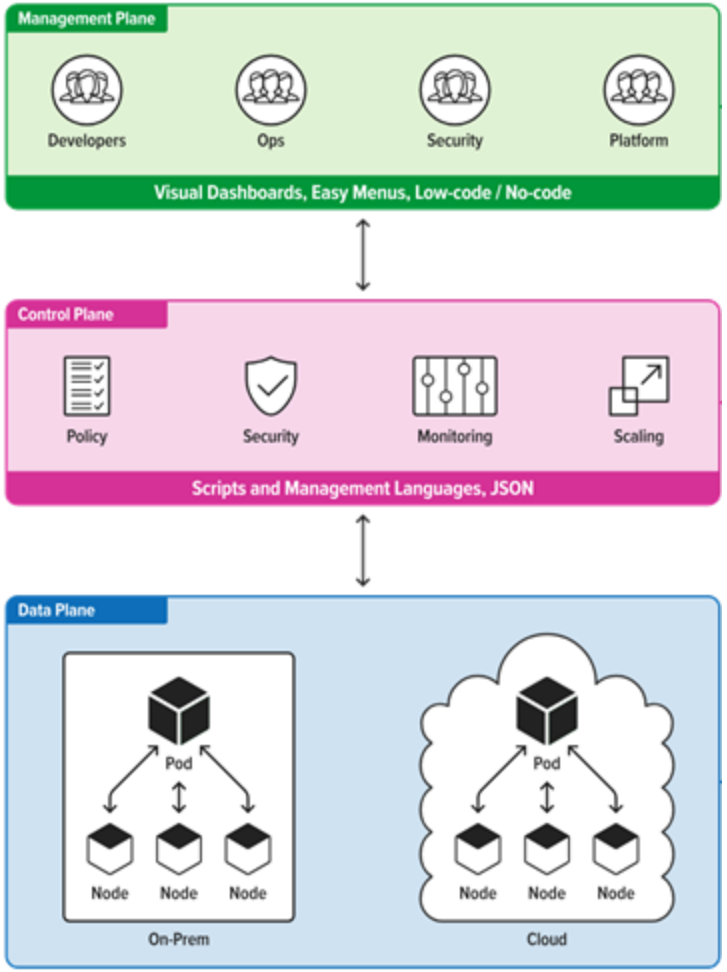

Management plane

- **C**onfigures and observes
- Applies configuration and policy
- Collects logs and telemetry
- Supports administration and troubleshooting

From routing decision to packet forwarding

1. **Learn** which destinations are reachable
2. **Select** the preferred route
3. **Install** the route in the forwarding table
4. **Forward** packets using that entry
5. **Recalculate** when the network changes

The control plane creates the decision. The data plane applies it.

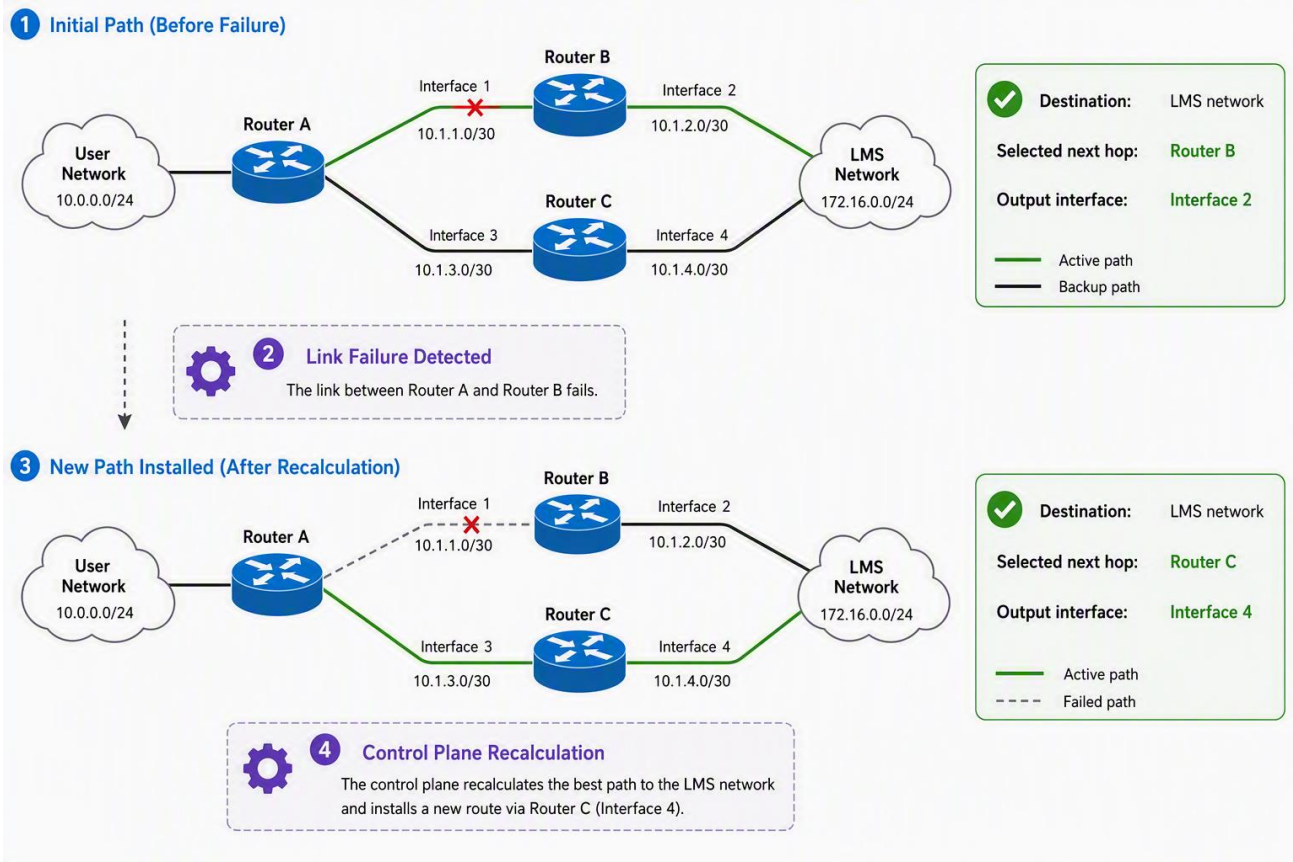

Turn the network into a graph

- **Vertices**
  - Routers or networks
- **Edges**
  - Links between them
- **Costs**
  - Values representing the preference for each link

Routing algorithms calculate paths through this model.

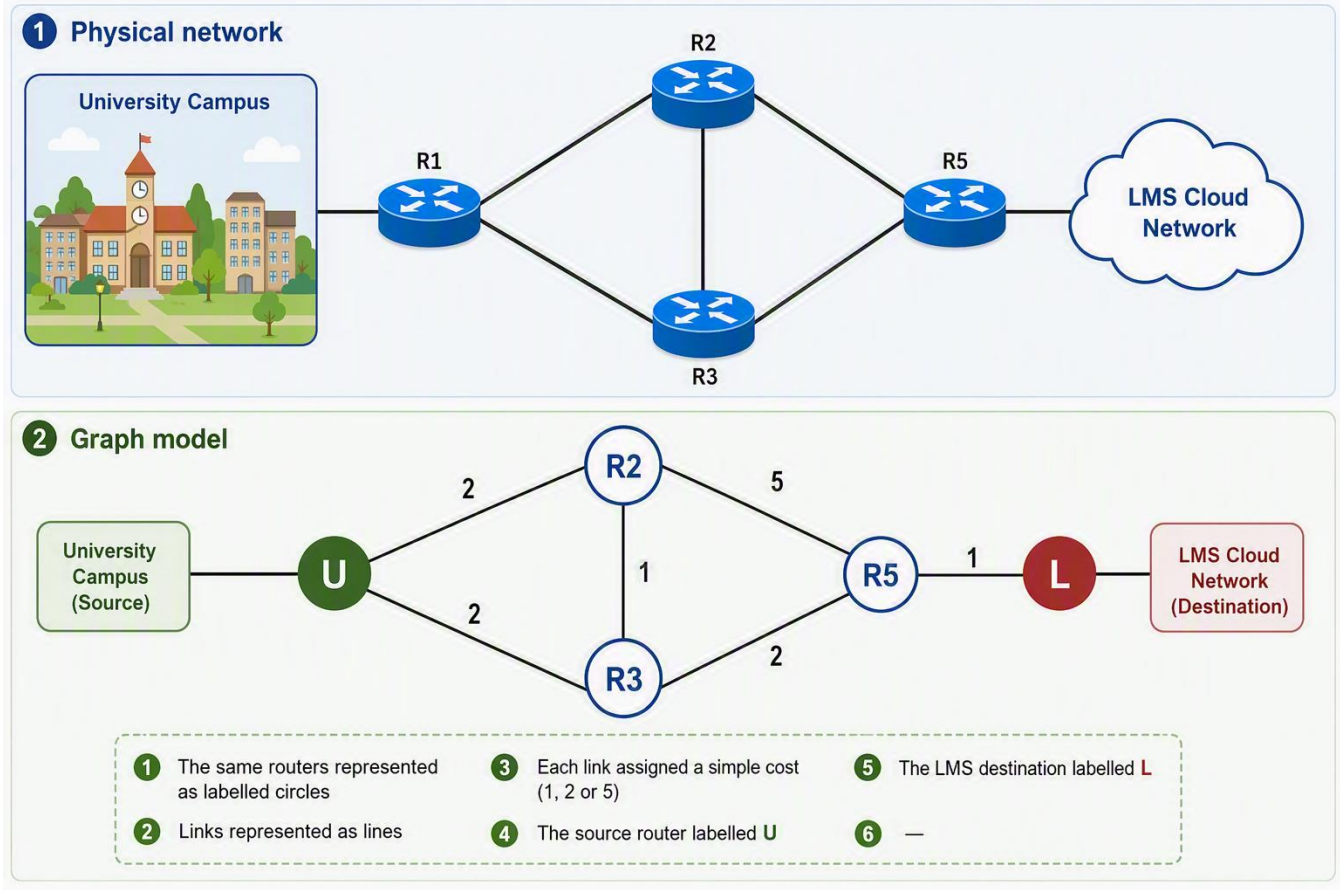

What makes one path “better”?

- **Hop count**
  - How many routers must be crossed?
- **Link cost**
  - What preference has been assigned to each link?
- **Performance**
  - What bandwidth, delay or congestion exists?
- **Policy**
  - Are there organisational reasons to prefer one path?

The best path depends on what the network is designed to optimise.

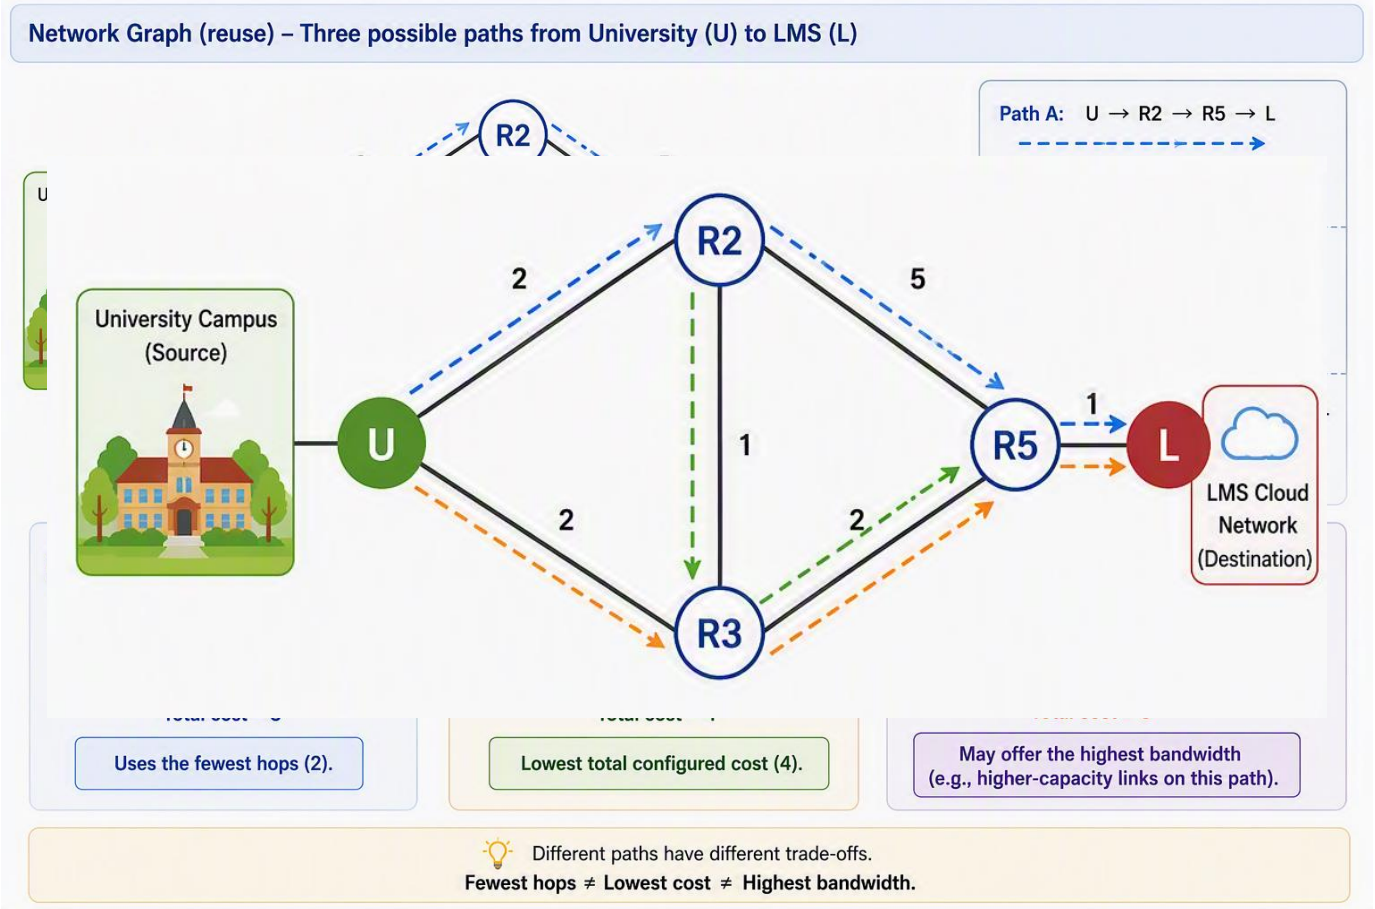

The preferred path has failed

1. **Detect** the topology change
2. **Share** updated reachability information
3. **Recalculate** the best available path
4. **Update** the forwarding table
5. **Resume** packet delivery

The time required to agree on the new path is convergence.

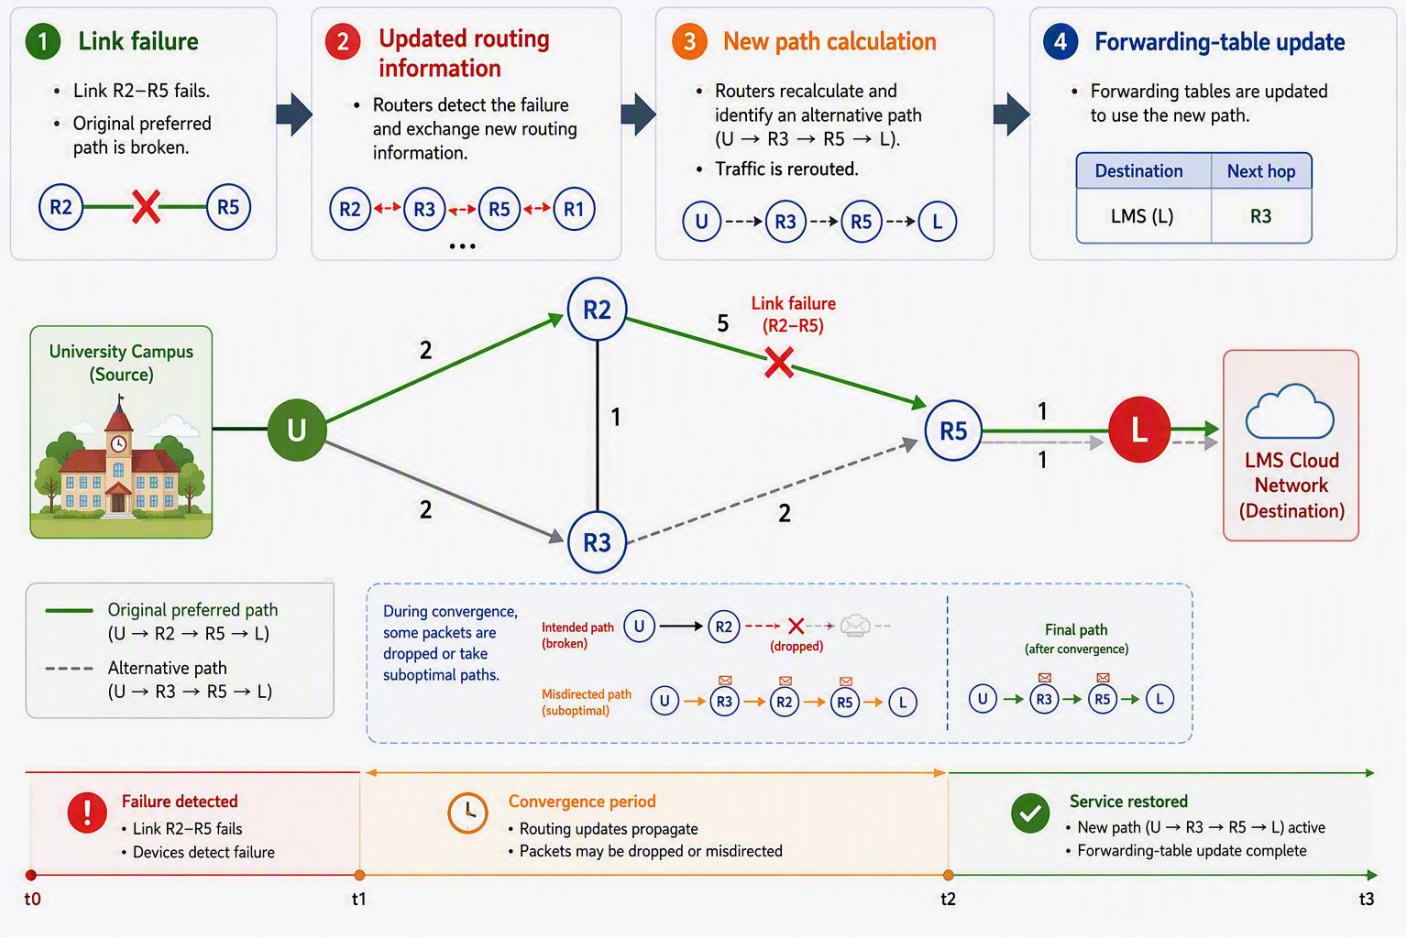

Link-state routing: Build a map, then calculate

1. **Discover** directly connected neighbours and link costs
2. **Advertise** local link-state information
3. **Build** a shared topology database
4. **Calculate** the least-cost paths
5. **Install** the preferred routes

Each router calculates independently from a common network map.

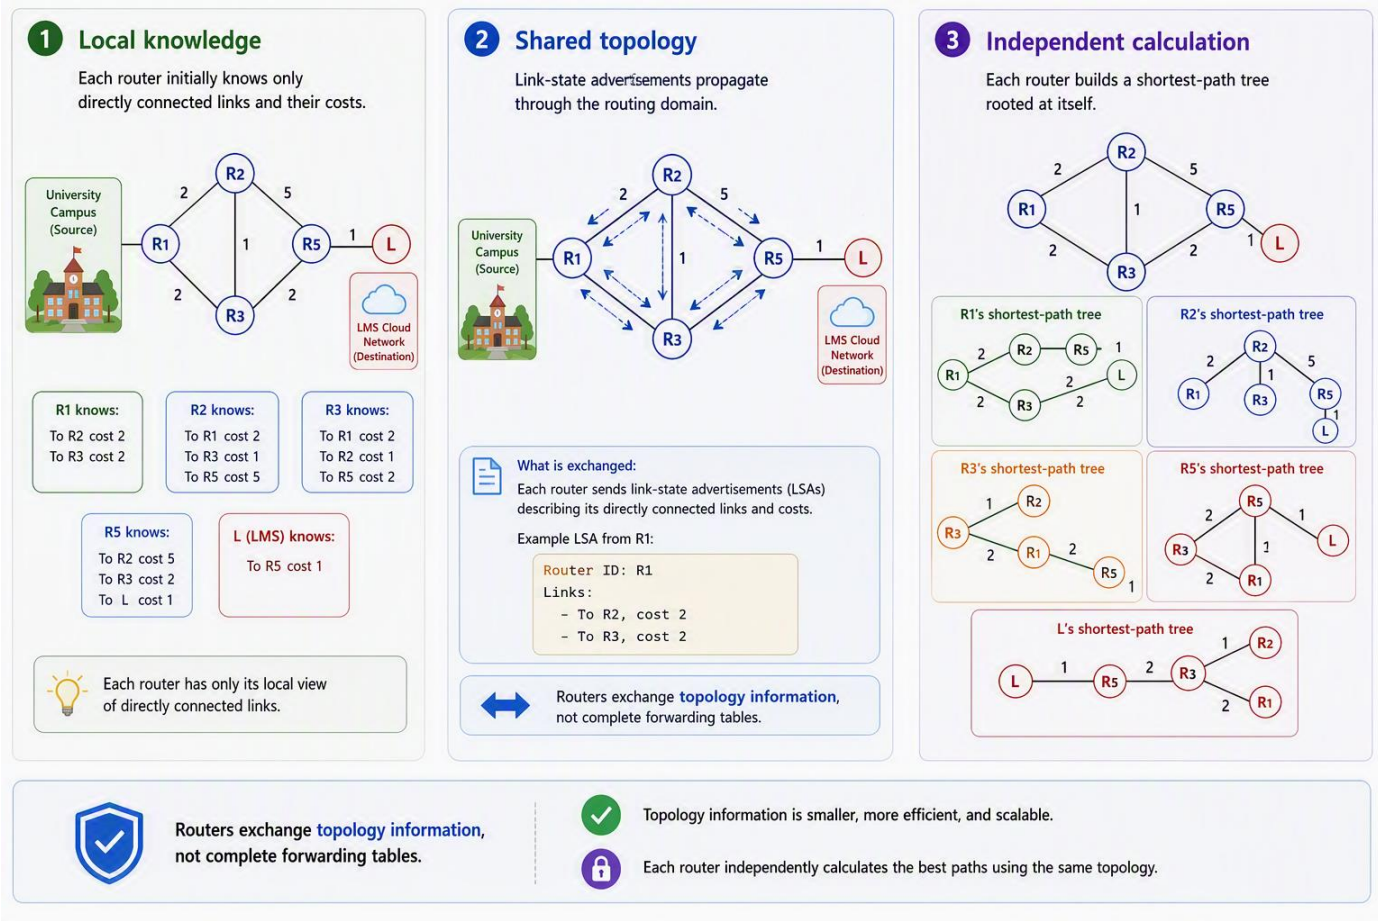

Dijkstra's algorithm: Find the least-cost path

1. Start with the source router
2. Record the cost to each neighbour
3. Select the lowest-cost unvisited router
4. Update the costs to other destinations
5. Repeat until the shortest-path tree is complete

Choose the next lowest-cost step, then improve what is known.

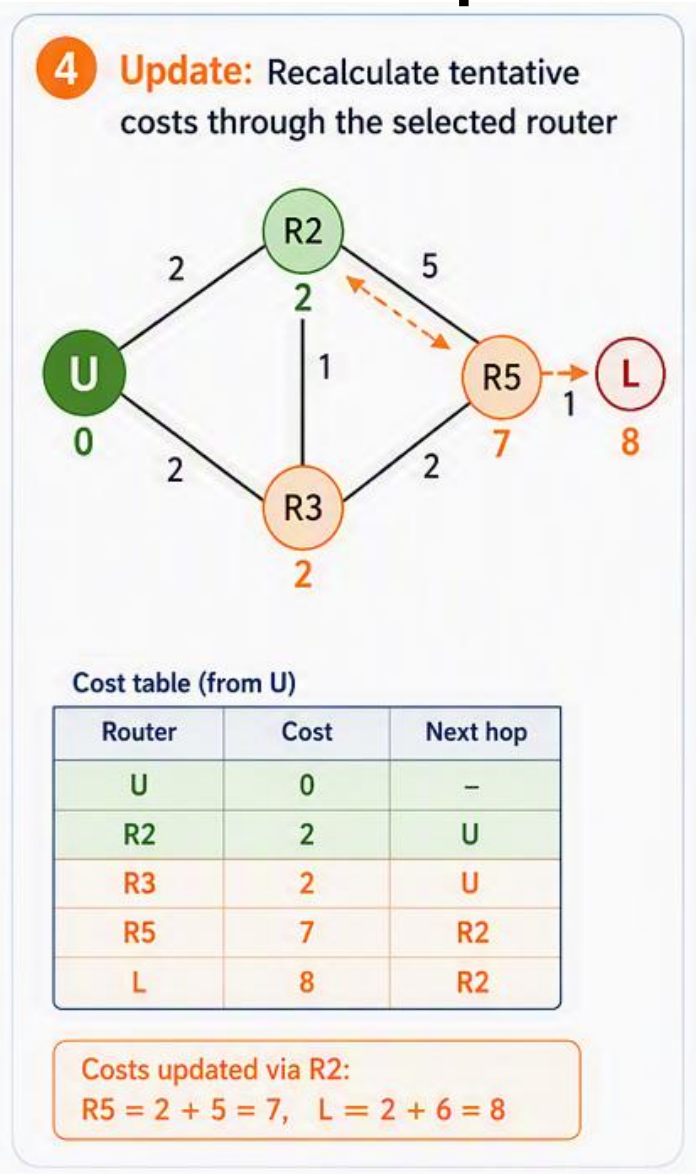

Distance vector: Learn from your neighbours

1. Know the cost to directly connected neighbours
2. Receive each neighbour's known distances
3. Add the local link cost
4. Select the lowest-cost route
5. Repeat until routes stabilise

No router needs a complete map of the network.

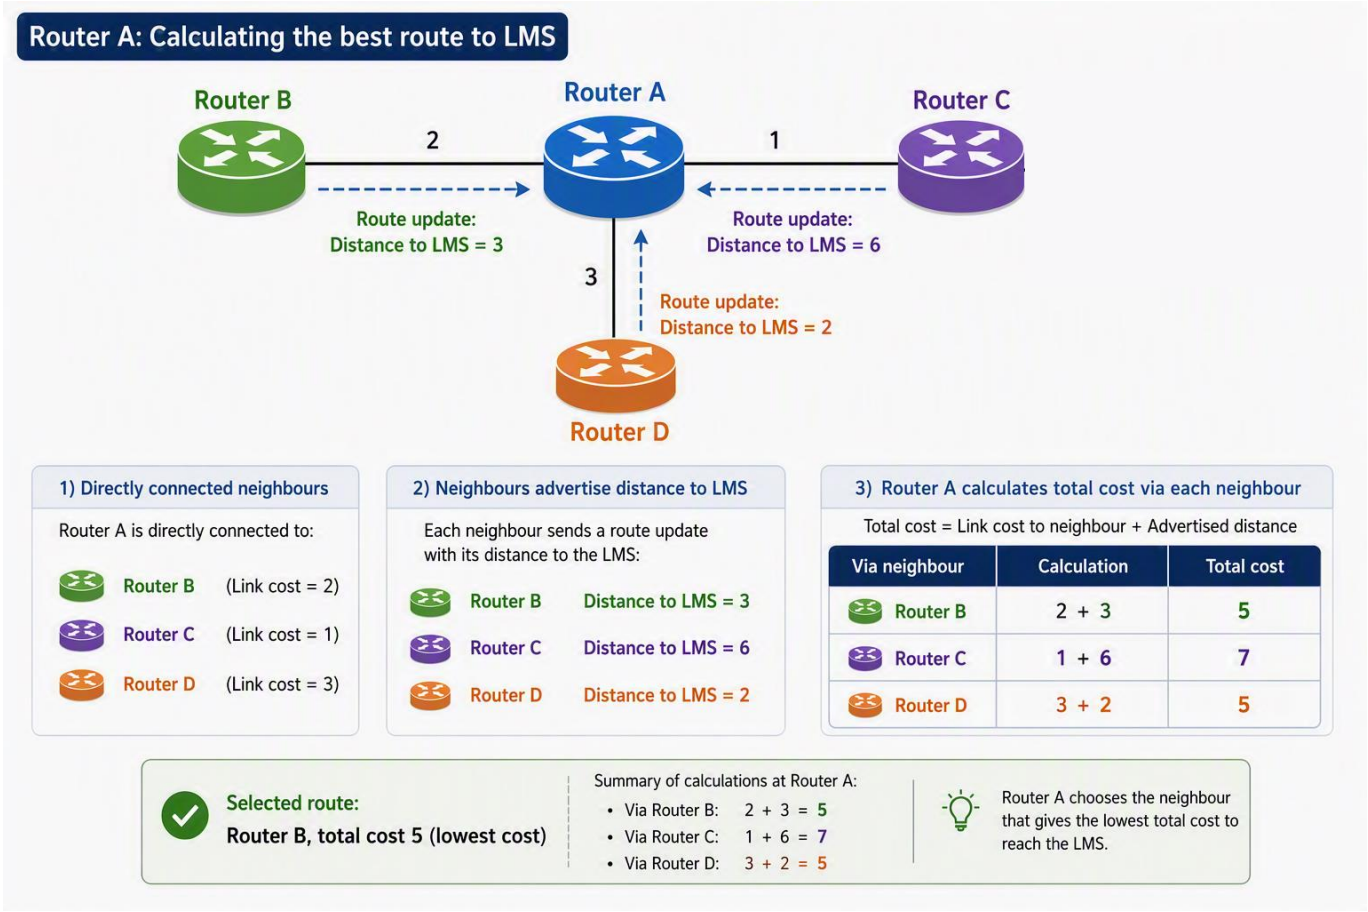

When neighbours repeat outdated information

1. A destination becomes unreachable
2. Some routers still advertise the old route
3. Neighbours believe each other has a valid path
4. Packets circulate between routers
5. The metric increases towards “infinity”

Outdated reachability can create a route that leads nowhere.

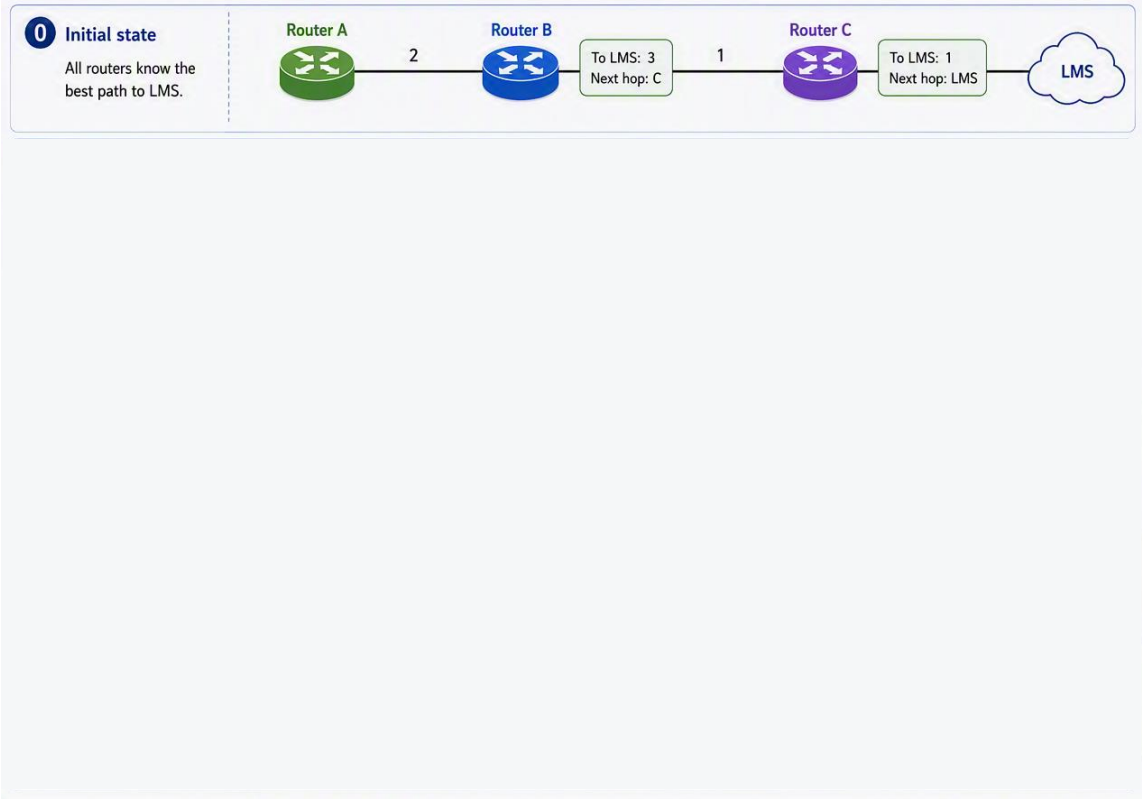

Two ways to learn the network

Link-state

- Builds a complete topology map
- Calculates paths independently
- Responds quickly to topology changes

Distance-vector

- Learns routes from neighbours
- Calculates paths iteratively
- Uses less topology information

Different knowledge creates different trade-offs.

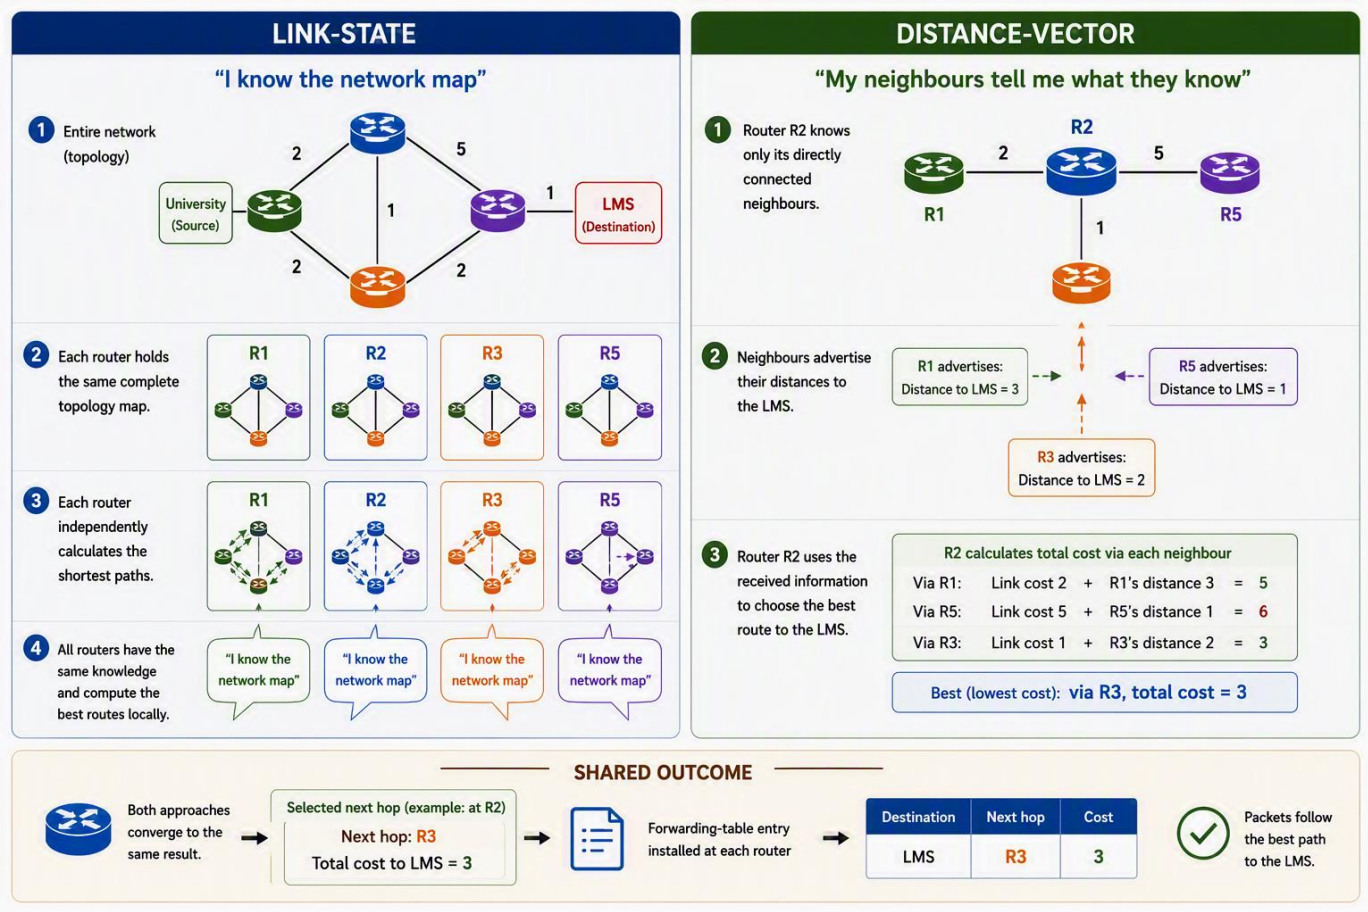

Why not use one global routing system?

- **Scale**
  - Millions of network prefixes must remain reachable
- **Change**
  - Links, routes and organisations constantly change
- **Autonomy**
  - Each organisation controls its own network
- **Policy**
  - Business and security decisions influence path selection

The Internet is a network of independently managed networks.

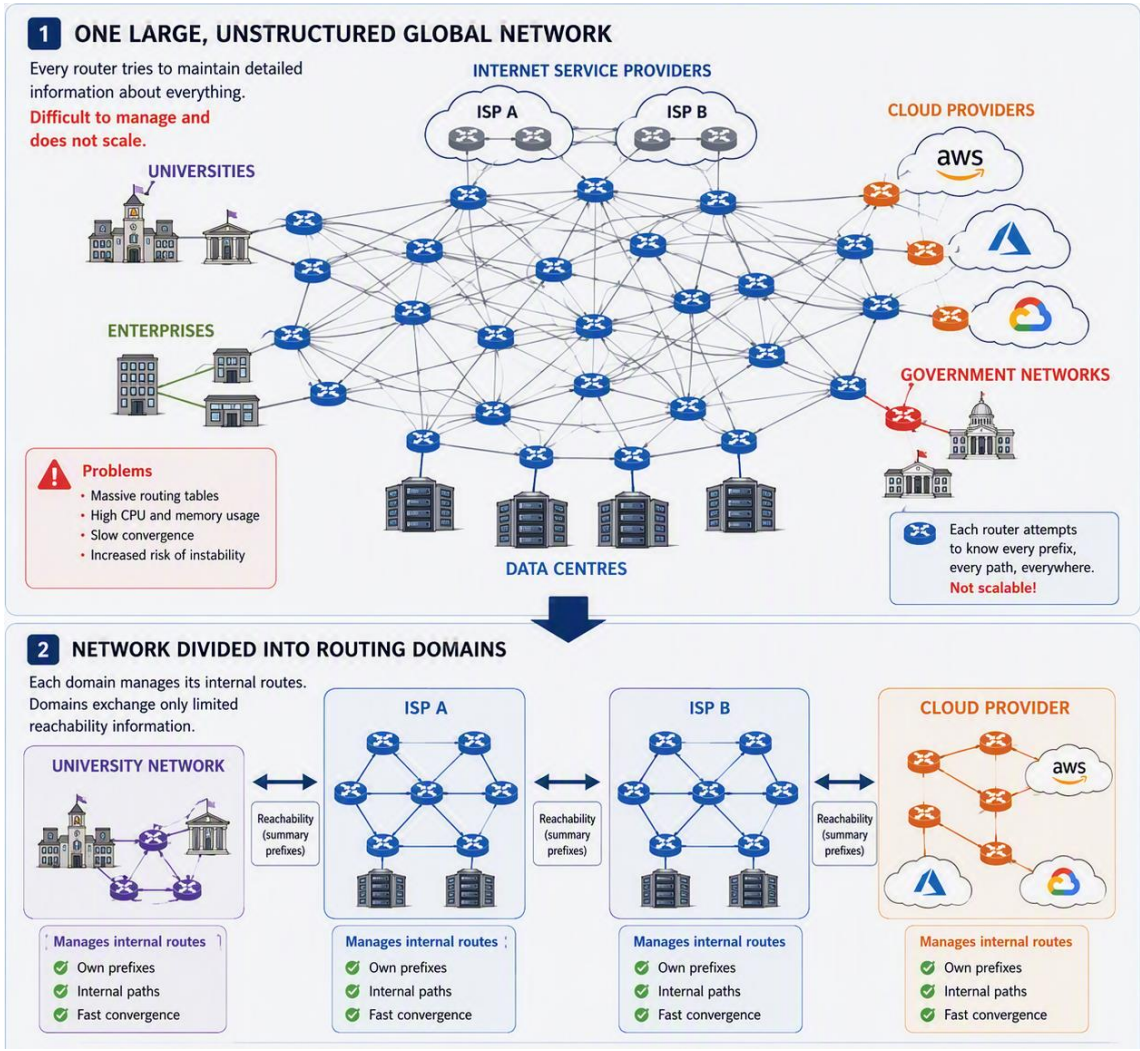

The Internet is a network of networks

- **Autonomous System**
  - A collection of IP prefixes managed under a defined routing policy
- **Autonomous System Number**
  - A unique identifier used to exchange routing information
- **Administrative control**
  - Each AS manages its own internal network and external relationships

Networks cooperate without giving up operational control.

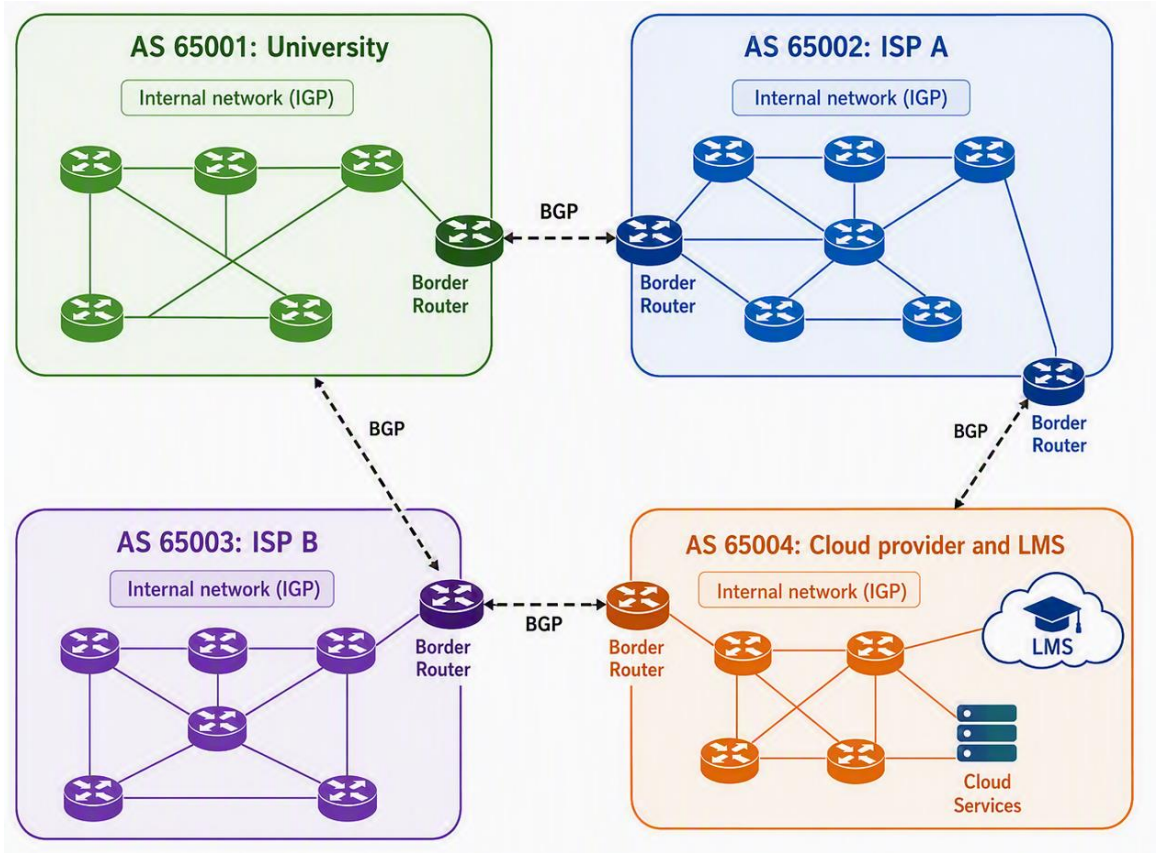

Two levels of routing

Intra-AS routing

- **Within one Autonomous System**
- Optimises internal paths
- Managed by one organisation
- Uses protocols such as OSPF or IS-IS

Inter-AS routing

- **Between Autonomous Systems**
- Exchanges network reachability
- Applies organisational policy
- Uses BGP

Inside the network: optimise the path.

Between networks: control the relationship.

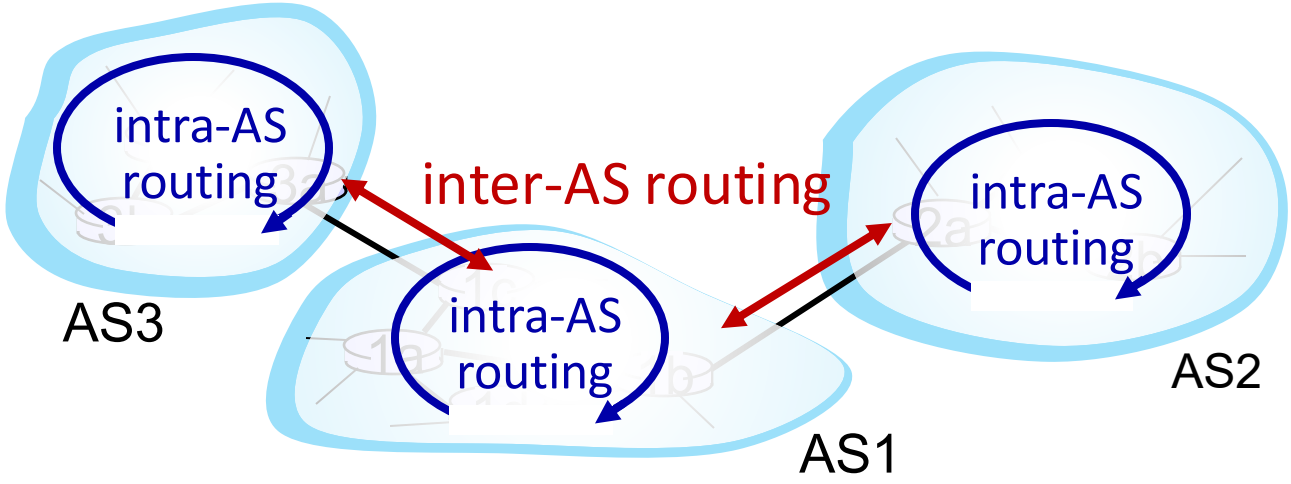

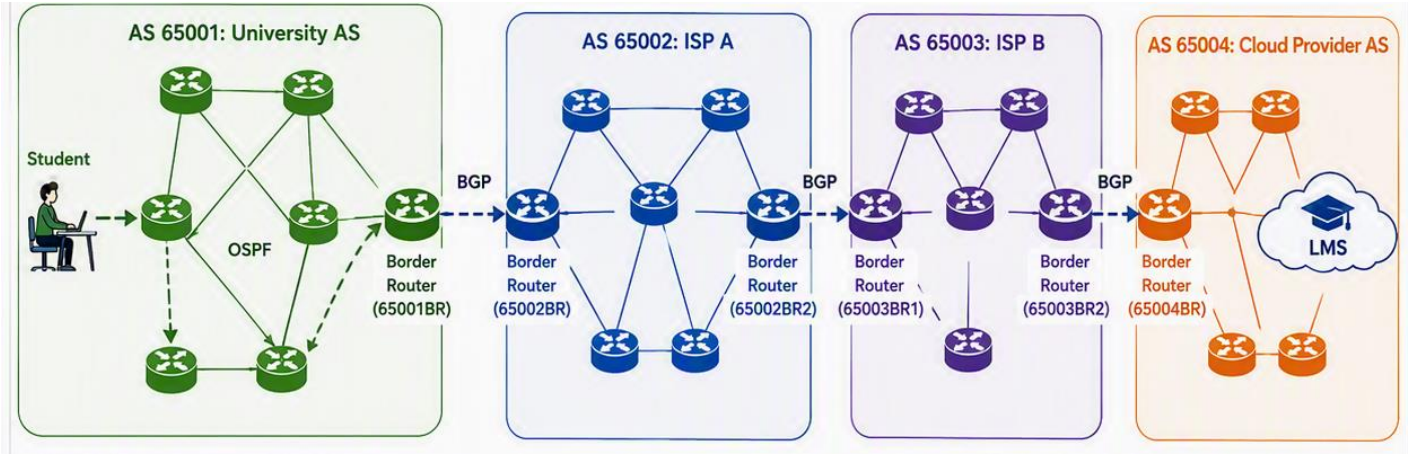

BGP tells networks what is reachable

A BGP route advertises:

- **Prefix**: Which destination network can be reached?
- **Next hop**: Where should traffic be sent?
- **AS path**: Which Autonomous Systems does the route cross?
- **Attributes**: How should the route be evaluated?

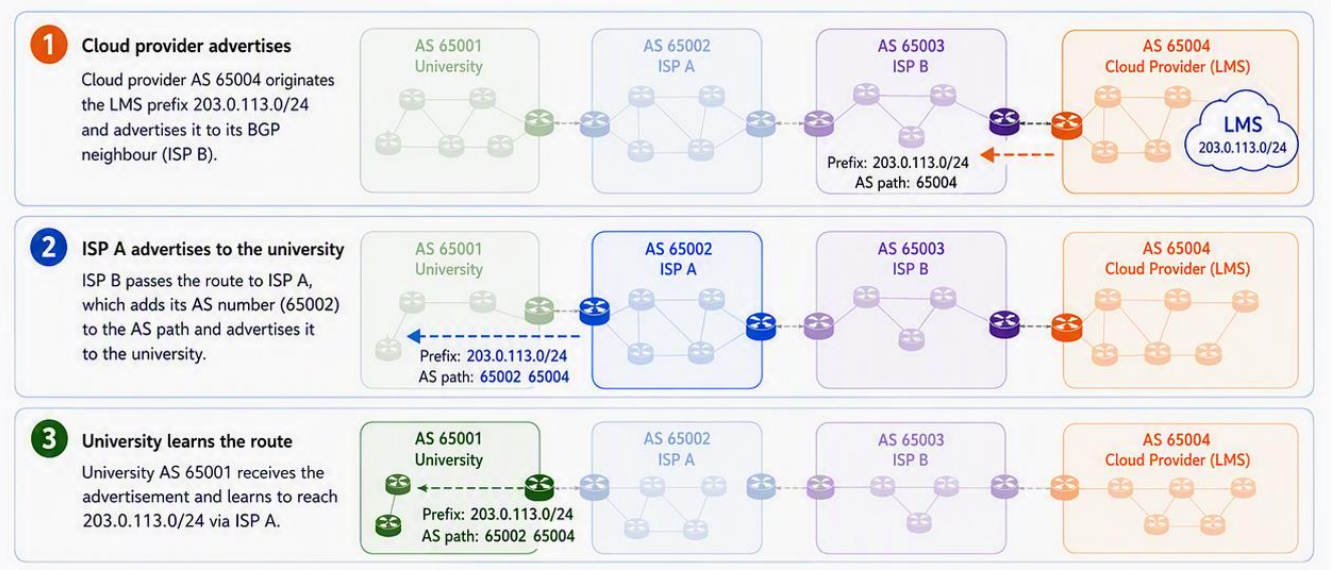

BGP shares reachability, not the complete internal topology.

Best BGP path is not always the shortest

- **Policy**
  - Which provider or relationship should be preferred?
- **Path attributes**
  - What information accompanies each route?
- **AS path**
  - How many Autonomous Systems are traversed?
- **Operational conditions**
  - Is the route valid, reachable and permitted?

BGP asks: “Which acceptable route best matches our policy?”

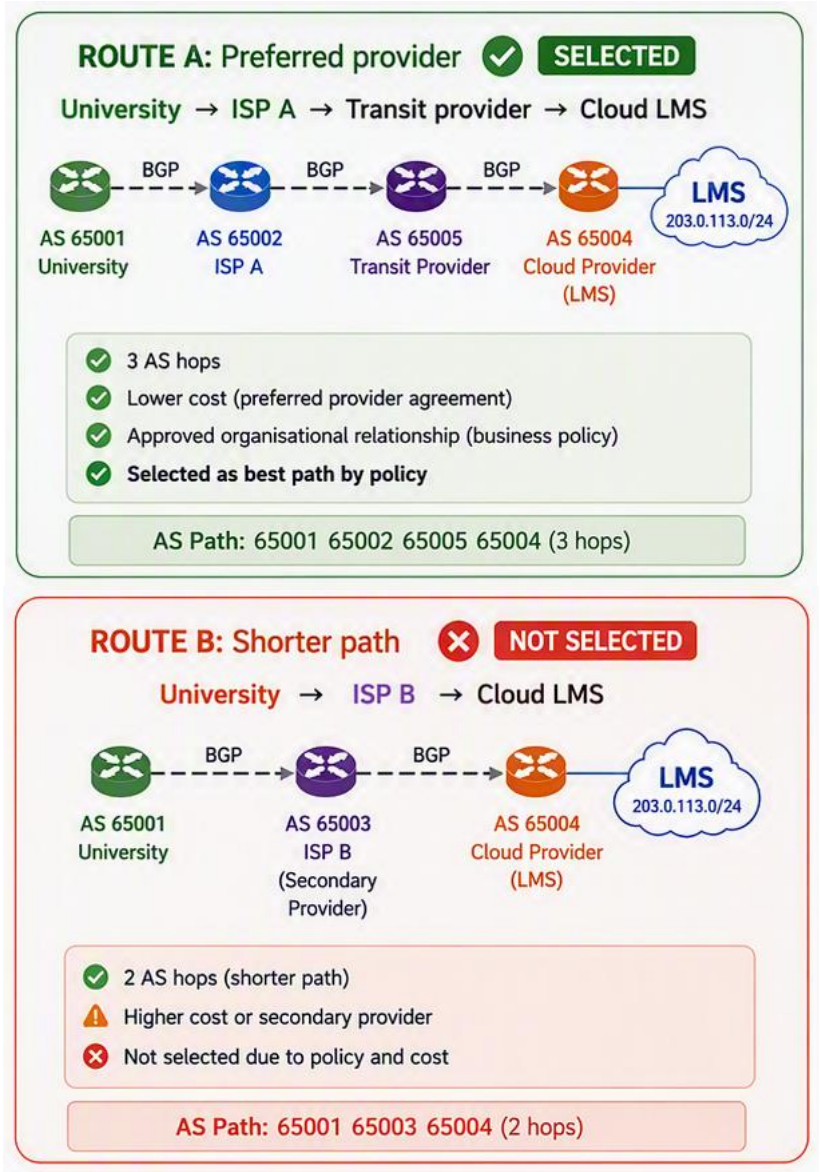

When an incorrect route spreads

Pakistan Telecom–YouTube incident (2008)

1. Pakistan Telecom announced a more-specific YouTube prefix
2. An upstream provider propagated it globally
3. Networks preferred the more-specific route
4. YouTube traffic was redirected towards Pakistan
5. Users worldwide lost access

A local routing decision created a global impact.

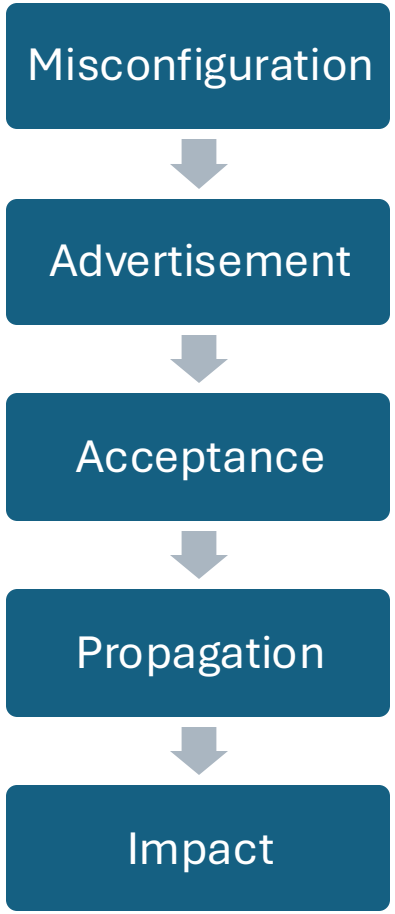

Do not trust every route automatically

- **Prefix Filter**
  - Accept only expected prefixes and paths
- **RPKI validation**
  - Check whether the origin AS is authorised
- **Maximum-prefix limit**
  - Restrict the number of routes accepted from a neighbour
- **BGP monitoring**
  - Detect unexpected announcements and path changes

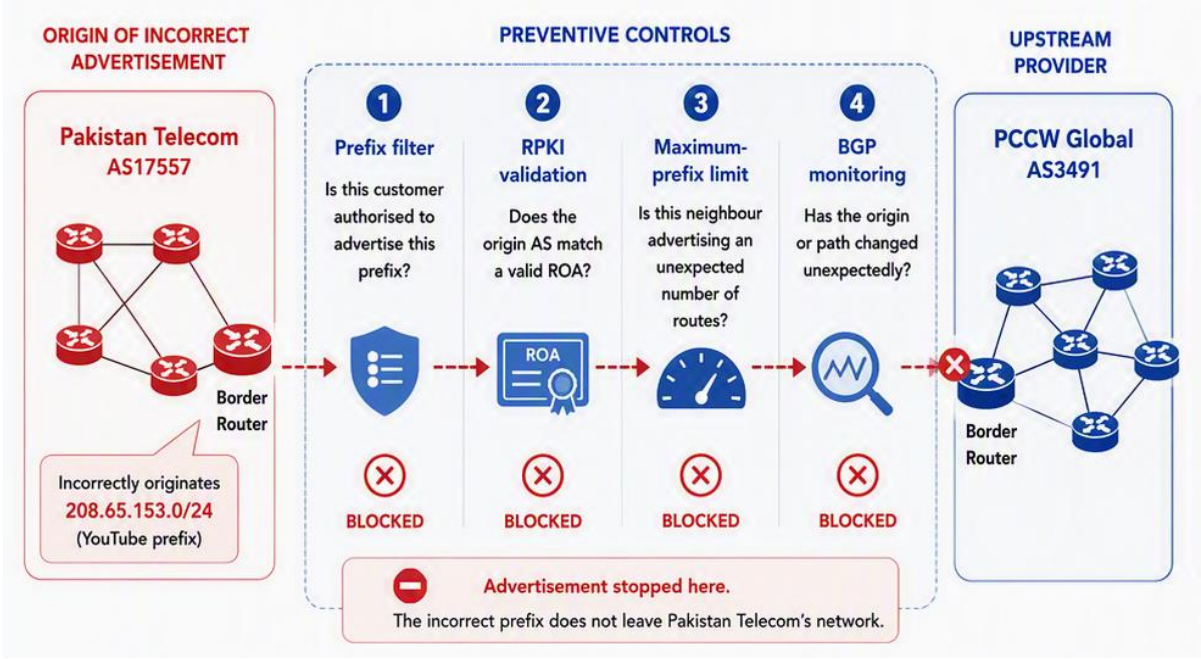

Secure routing requires prevention, detection and coordination.

SDN: Programmable network control

- **Applications and policy**
  - Define the desired network behaviour
- **SDN controller**
  - Maintains network state and translates policy into rules
- **Network devices**
  - Apply the rules and forward traffic

Logically centralised control, physically distributed forwarding

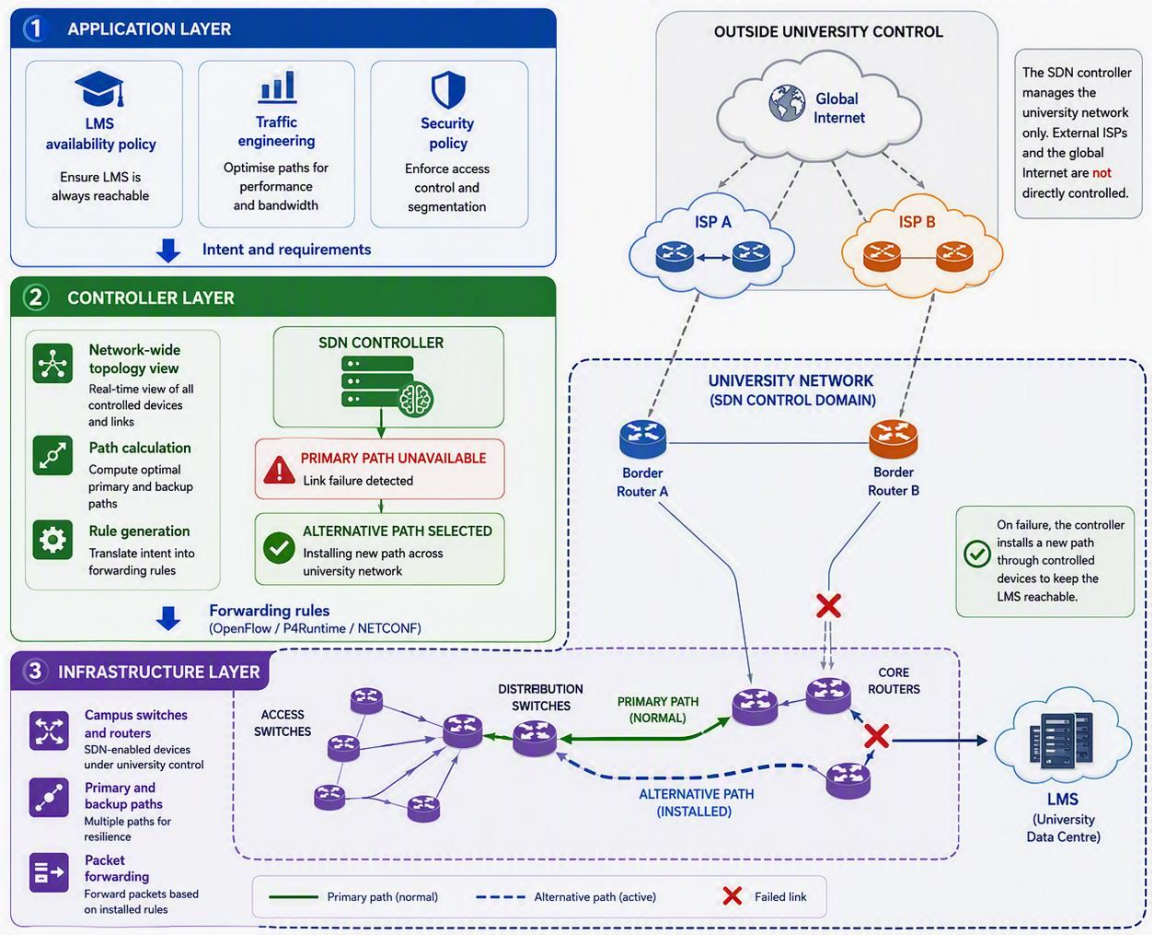

The control plane must be trustworthy

- **Correct**
  - Select accurate and policy-compliant routes
- **Resilient**
  - Adapt safely when links, devices or controllers fail
- **Secure**
  - Reject unauthorised or unexpected routing information
- **Observable**
  - Make routing decisions and changes visible to operators

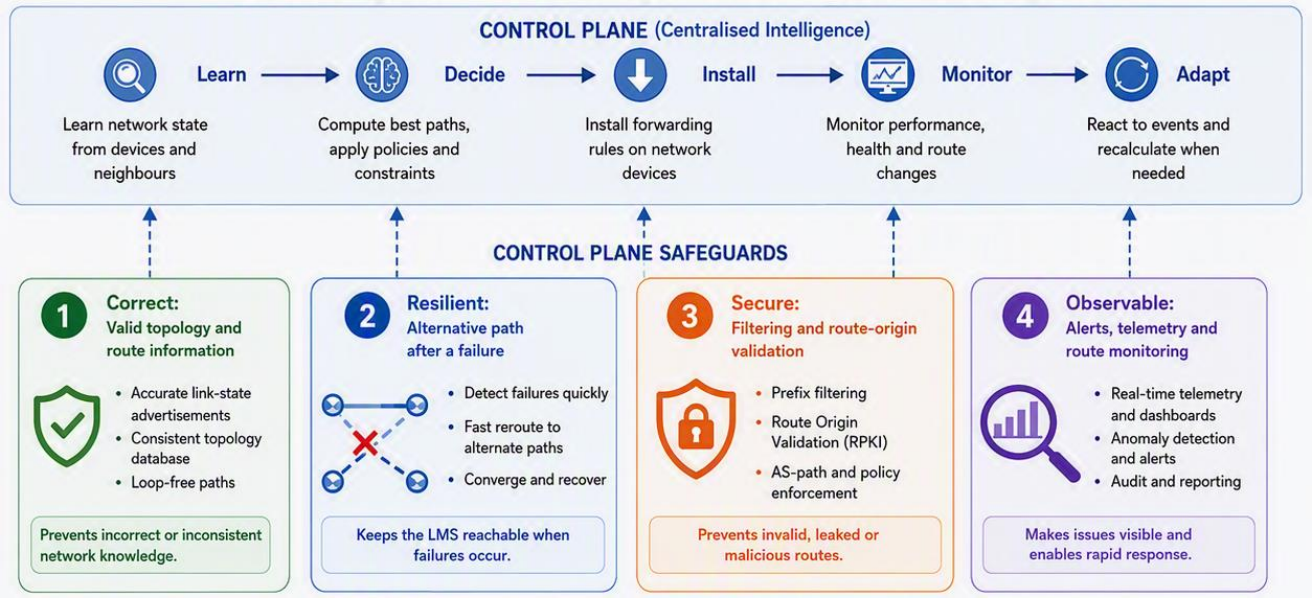

Packets can only follow the decisions the control plane provides.

From routing decisions to programmable paths

- **Segment routing**
  - Steer traffic through an ordered list of network instructions
- **Intent-based networking**
  - Translate desired outcomes into network policy
- **Programmable data planes**
  - Define how devices process selected packets
- **AI-assisted operations**
  - Analyse telemetry and recommend network changes

Programmability increases capability, but also increases the need for validation and governance.

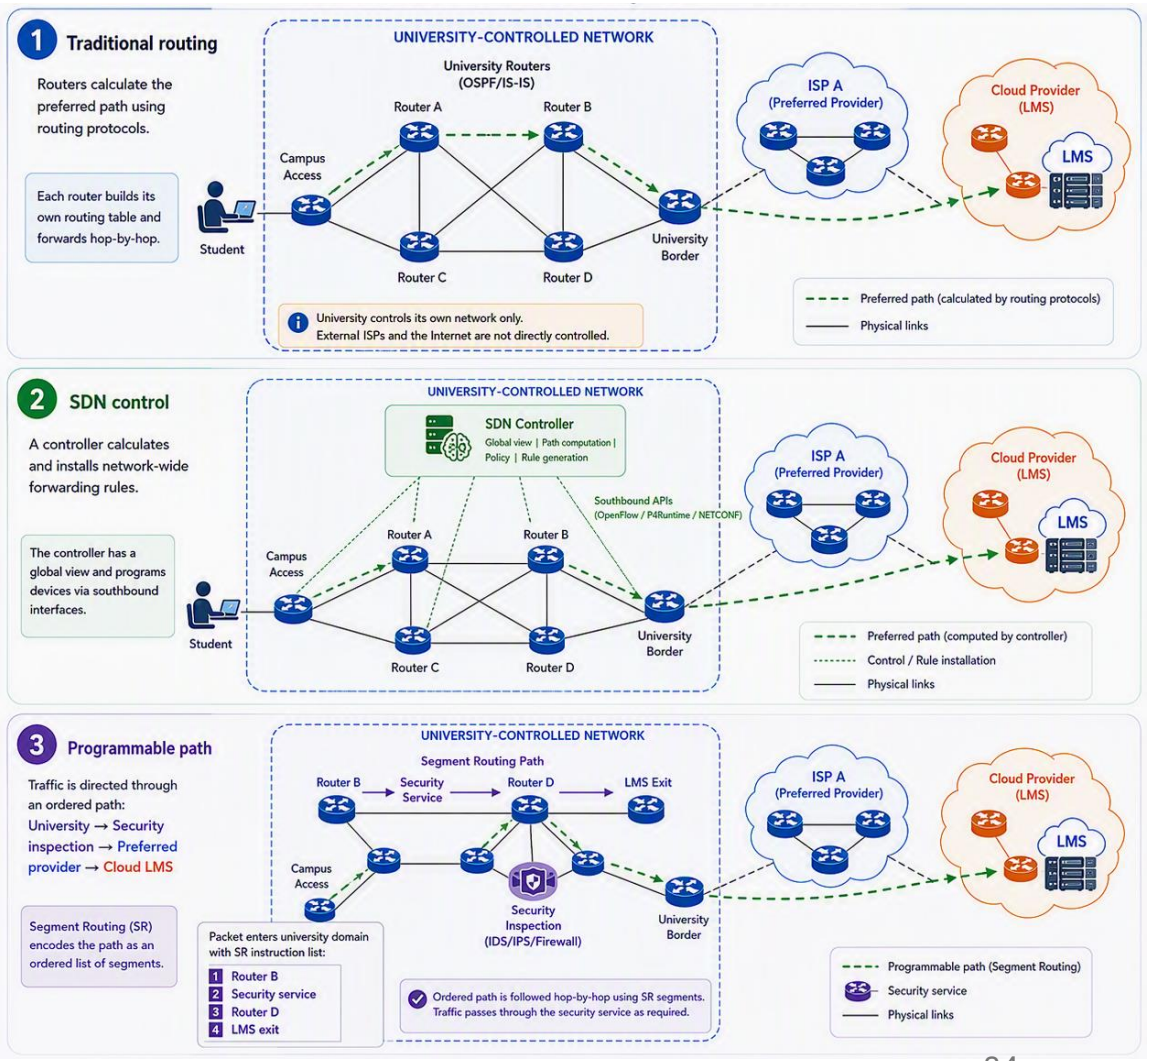

The LMS is healthy. The path is not.

- **Investigate**
  - Where does successful traffic differ from failed traffic?
- **Identify**
  - Which control-plane decision created the problem?
- **Correct**
  - What route, policy or validation control must change?
- **Verify**
  - Has reachability been restored for every affected network?

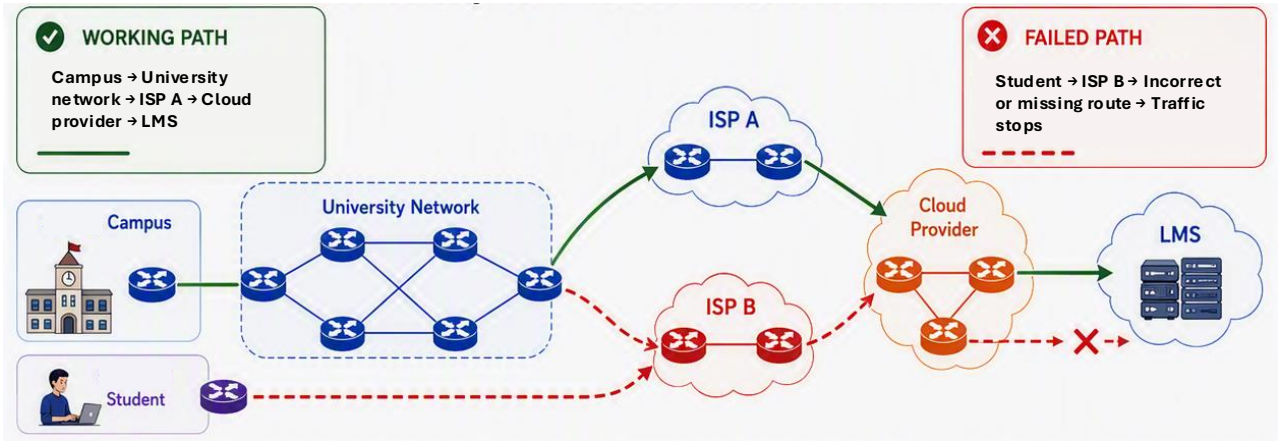

Do not just restore the path. Understand why it failed.

Questions?

- The destination may be healthy, but the path can still fail.

