COSC 301
Network Management

Lecture 19: Interior Routing

Zhiyi Huang

Computer Science, University of Otago

Today's Focus

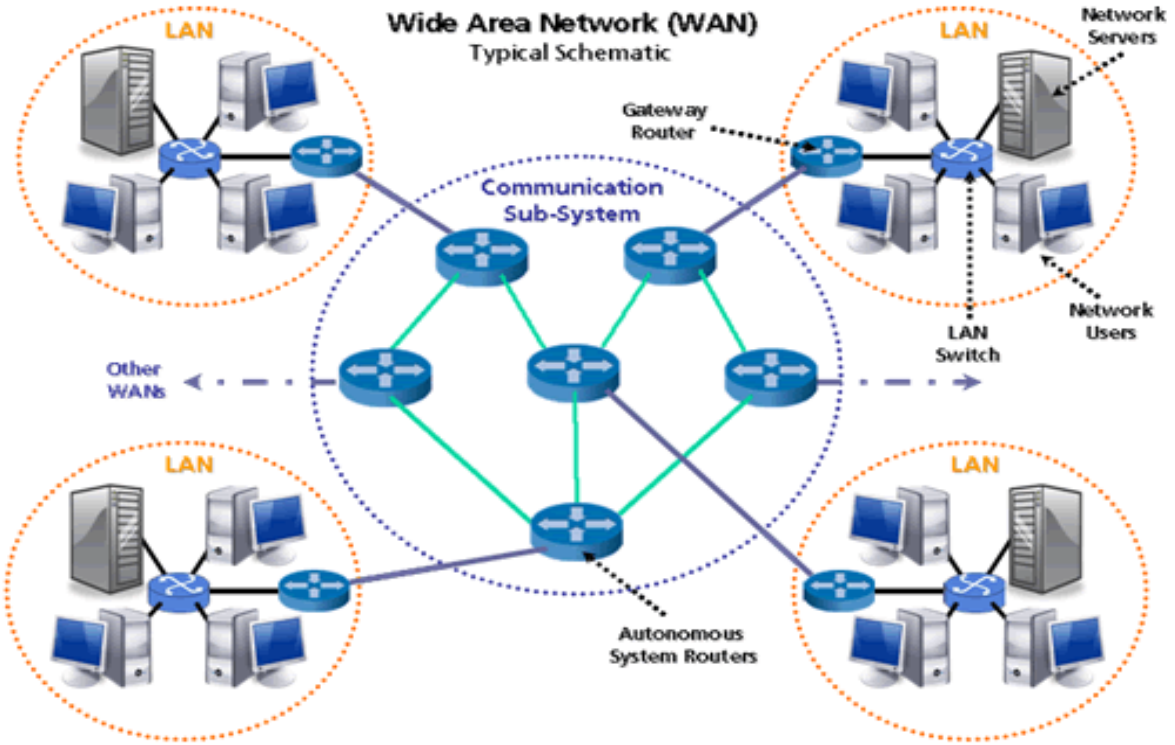

Interior Routing

-- Autonomous System

-- RIP

-- OSPF

Autonomous System

- The Internet does not consist of independent networks
  - Networks and routers are owned by organizations or individuals
  - Networks and routers owned by a given entity fall under a single administrative authority, which guarantees that internal routes remain consistent and viable.
- Autonomous System: a group of networks and routers controlled by a single administrative authority
  - Routers are free to choose their own mechanisms for discovering, validating, and checking consistency of routes
  - The Internet is divided into autonomous systems with each owned by a single administrative authority.
  - In the current Internet, each large ISP is an autonomous system.

Internet Routing

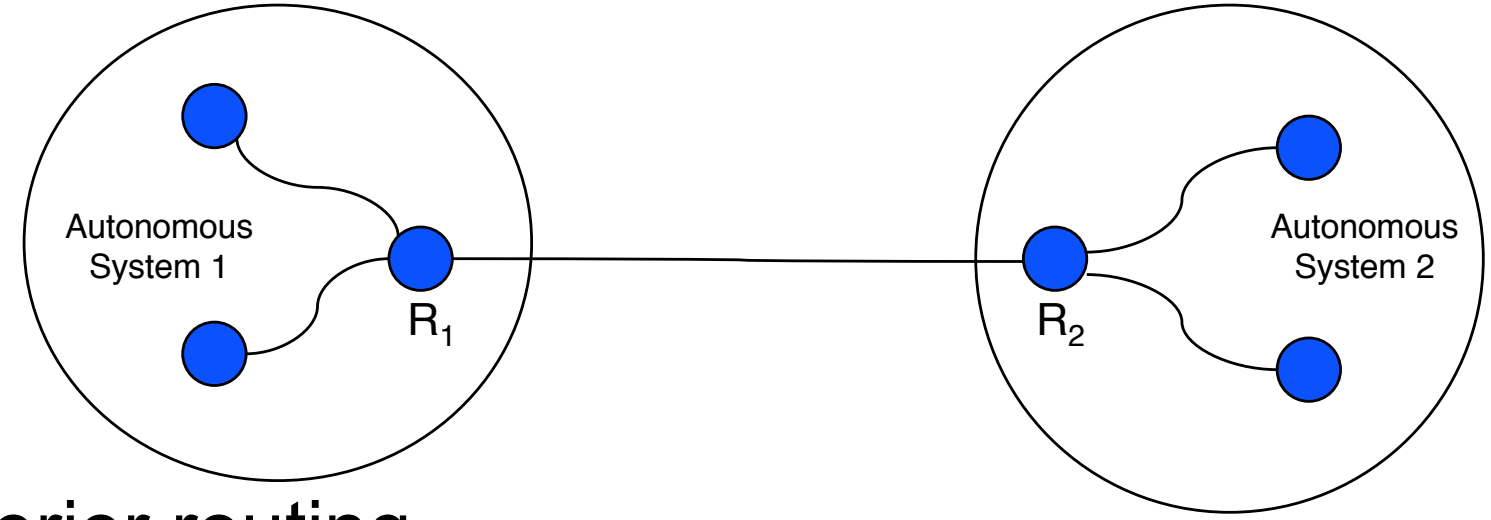

- Interior routing
  - Used within an autonomous system
  - Used within an area of administrative control
  - RIP, OSPF
- Exterior routing
  - Used between autonomous systems
  - Used to peer with networks without administrative control
  - BGP

Static vs. Dynamic

- Static routing
  - Advantage: Simplicity, less software that might fail
  - Disadvantage: Limits on scalability, less resilience and load-balancing, slow response to handle problems
- Dynamic routing
  - Advantage: Scalable, optimal routes, automatic fail-over and fail-back, load balancing
  - Disadvantage: route update consumes bandwidth, additional load on router CPU

Routing Information Protocol (RIP)

- Early, easy, and common
  - Still used today in SOHO gear
  - Set and forget; few knobs
- Straightforward implementation of Distance Vector algorithm
  - routing by rumour
  - Alternative routes are not kept
  - Metric is hop-count, and metric of 16 indicates unreachbility

RIP Metric

Network with Fast Ethernet links between A-B and B-C, and a 802.11g backup link A-C

What is RIP's next hop for A to C?

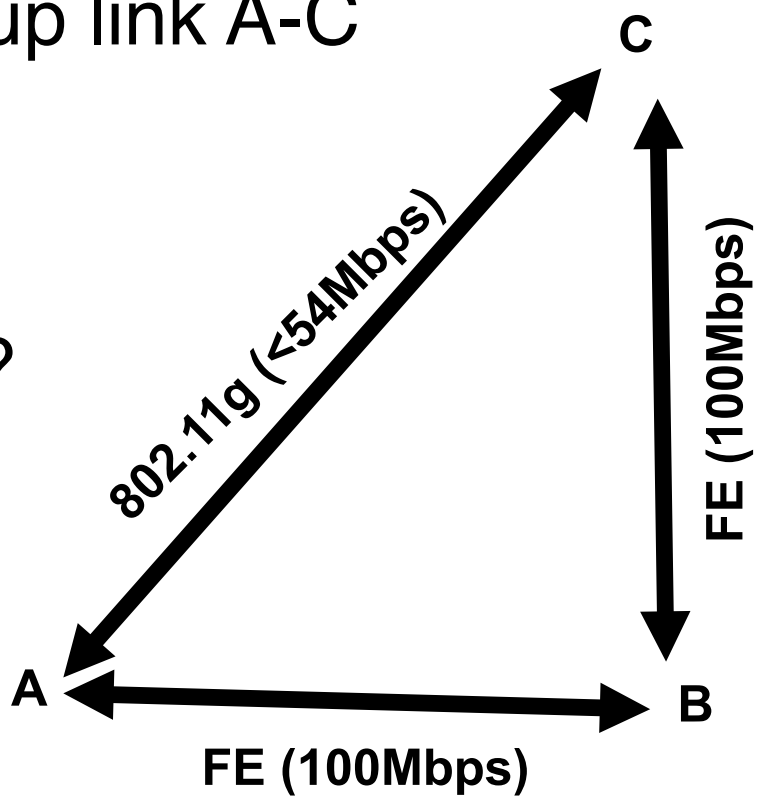

RIP Limitations (1)

- No Variable Length Subnet Masks support
  - e.g. /28, but is supported in RIPv2
- Slow convergence
  - Full advertisements broadcast every 30s
  - Count-to-Infinity

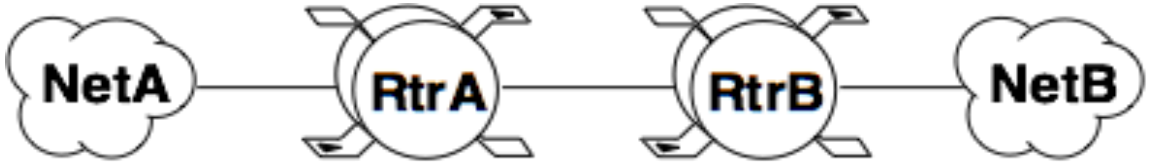

- **RtrA** loses connection to **NetA**
- Before **RtrA** sends update (of break), **RtrB** sends full update, advertising **NetA** at a cost of 1
- **RtrA** now thinks it can get to **NetA** via **RtrB** at a cost of 1+1=2 hops
- **RtrA** advertises this to **RtrB**, which sees an increased cost. **RtrB** advertises 2+1 = 3, continues to 16 (infinity/unreachability)

RIP Limitations (2)

- Limited diameter of network (15 hops)
- Only metric is hop-count
  - poor support for heterogeneous networks
- Advertisements not authenticated (v1)
  - attacking the network made very easy

RIP is okay when …

- Small network or minimal IT support
- Homogeneous network
  - With respect to link speed
- Convergence time is acceptable

RIP Versions

- RIP V1
  - No support for variable length subnet masks (VLSM)
  - No support for router authentication
- RIP V2
  - Support Classless Inter-Domain Routing (CIDR)
  - Support authentication (MD5)
- RIPng
  - RIP next generation: an extension of RIPv2 to support IPv6
  - No router authentication. IP routers were supposed to use IPsec for authentication

RIP Optimizations

- Split Horizon
  - prevent a router from advertising a route back onto the interface from which it was learned.
- Hold Down Timer
  - upon failure, hold the route as down until the network has converged, to expire old information
- Poison Reverse
  - Route can be deleted due to split horizon or route timeout.
  - A router actively advertises routes as unreachable over the interface over which they were learned
- Triggered Update
  - send updates as soon as something changes about a route's metric or state

Link State vs. Distance Vector

- Distance Vector
  - routing by rumour
  - unaware of topology
  - distributed Bellman-Ford algorithm

- Link State
  - multicast local link states to all routers
  - each router then assembles topology
  - Dijkstra's algorithm
  - higher memory and processor requirements

OSPF

- Open Shortest Path First
  - Link State protocol
- Break AS into smaller areas
  - simplifying calculation requirements
  - routes summarised crossing areas
- Areas attached to backbone
  - area 0 is backbone area

OSPF

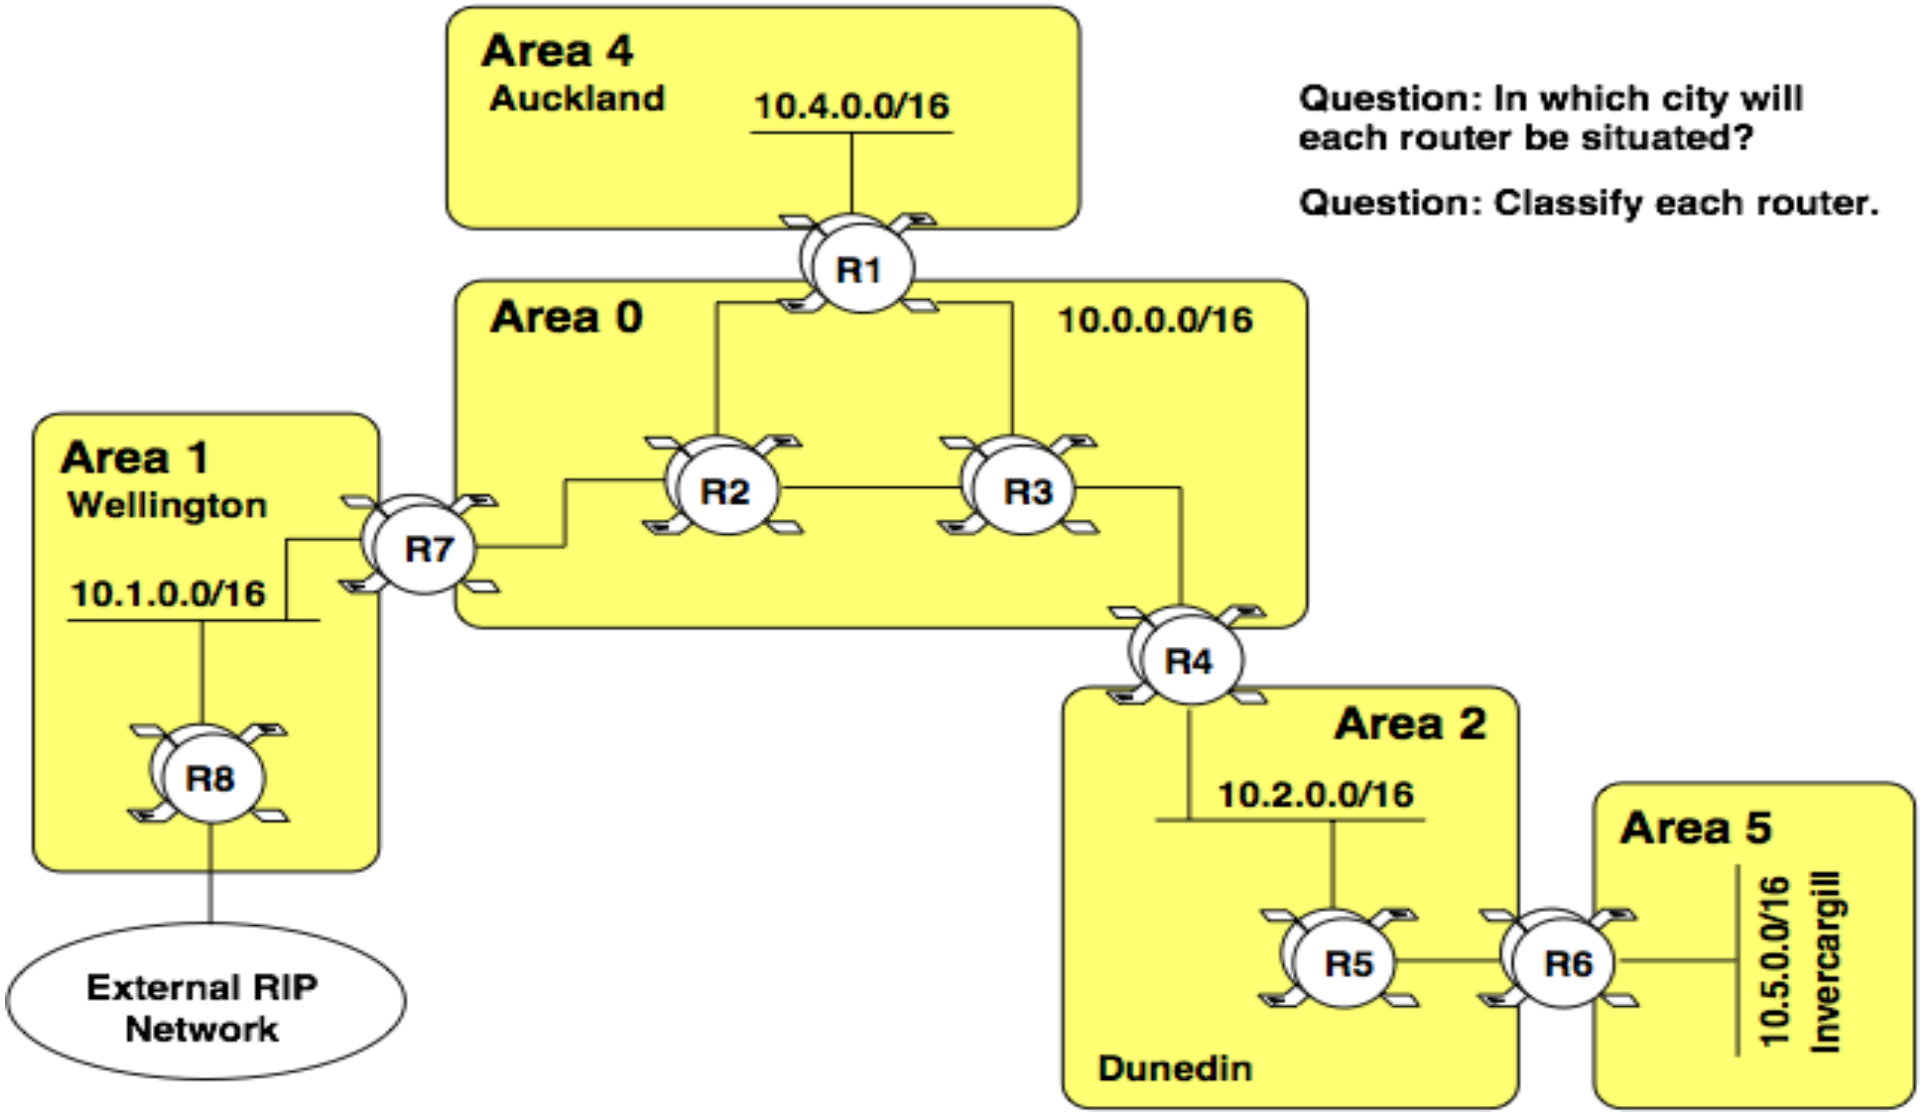
Area 4
Auckland
10.4.0.0/16
R1
Area 0
10.0.0.0/16
R2
R3
Area 1
Wellington
10.1.0.0/16
R7
R8
External RIP
Network
R4
Area 2
10.2.0.0/16
R5
R6
Dunedin
Area 5
10.5.0.0/16
Invercargill
Question: In which city will each router be situated?
Question: Classify each router.

Route Summarisation

- Reduce size of routing tables
  - supernetting (route aggregation)

- Reflects hierarchical nature of a network

  192.168.1.32/28 + 0010 0000
  192.168.1.48/28 = 0011 0000
  192.168.1.32/27 .001

  we own everything in 192.168.1.32/27

Route Summarisation

- Very large, complex, well-designed networks with skilled staff
- High memory and CPU requirements
- Fast convergence (but not fastest)

Virtual LANs

- A local area network configured by software, not by physical wiring
  - Virtually connecting devices in different physical LANs
- IEEE802.1Q
  - The protocol most commonly used today to configure VLANs

Broadcast Domains

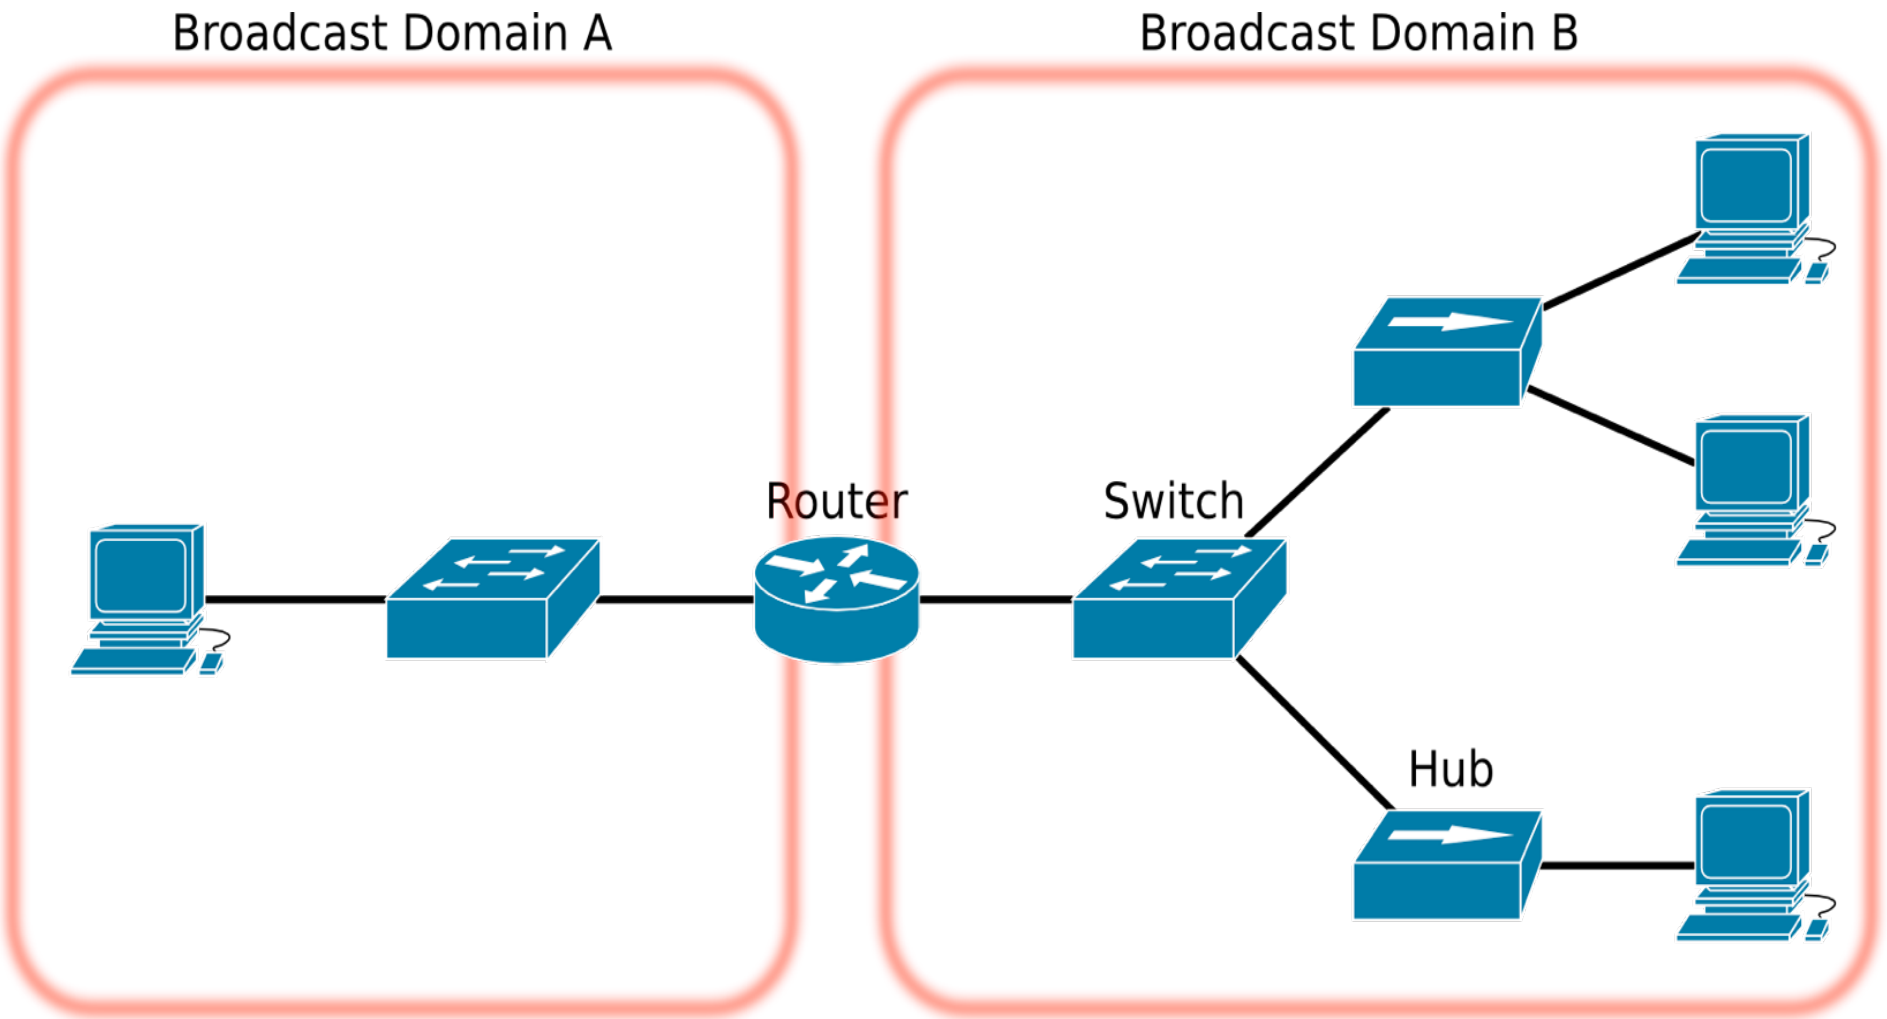

VLAN Motivations

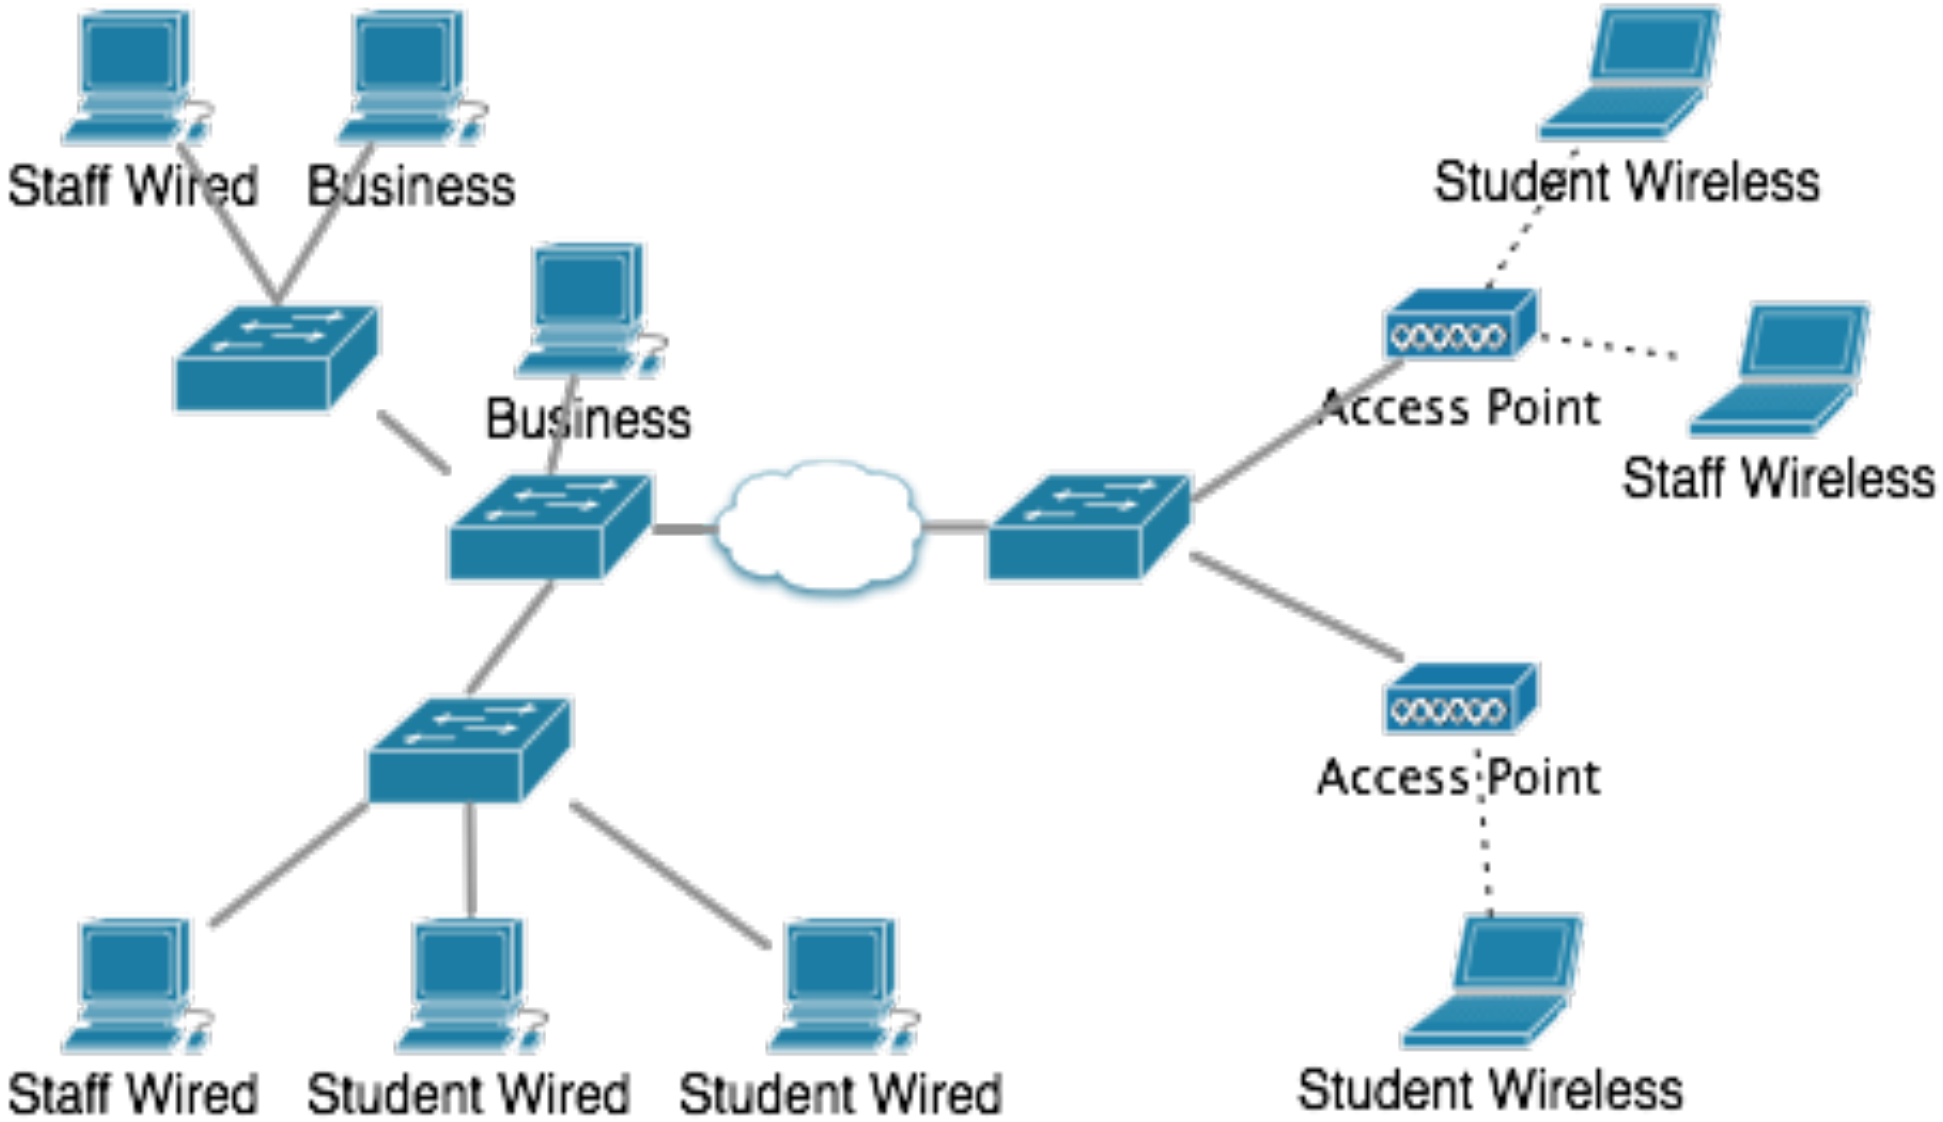

Virtual LAN Types

- Layer 1 VLAN: Membership by Port
  - Does not allow for user mobility
- Layer 2 VLAN: Membership by MAC Address
  - Support user mobility
  - VLAN membership must be assigned initially.
- Layer 3 VLAN: Membership by IP Subnet Address
  - Support user mobility
  - takes longer to forward packets using Layer 3 information than using MAC addresses.

Port Assignments

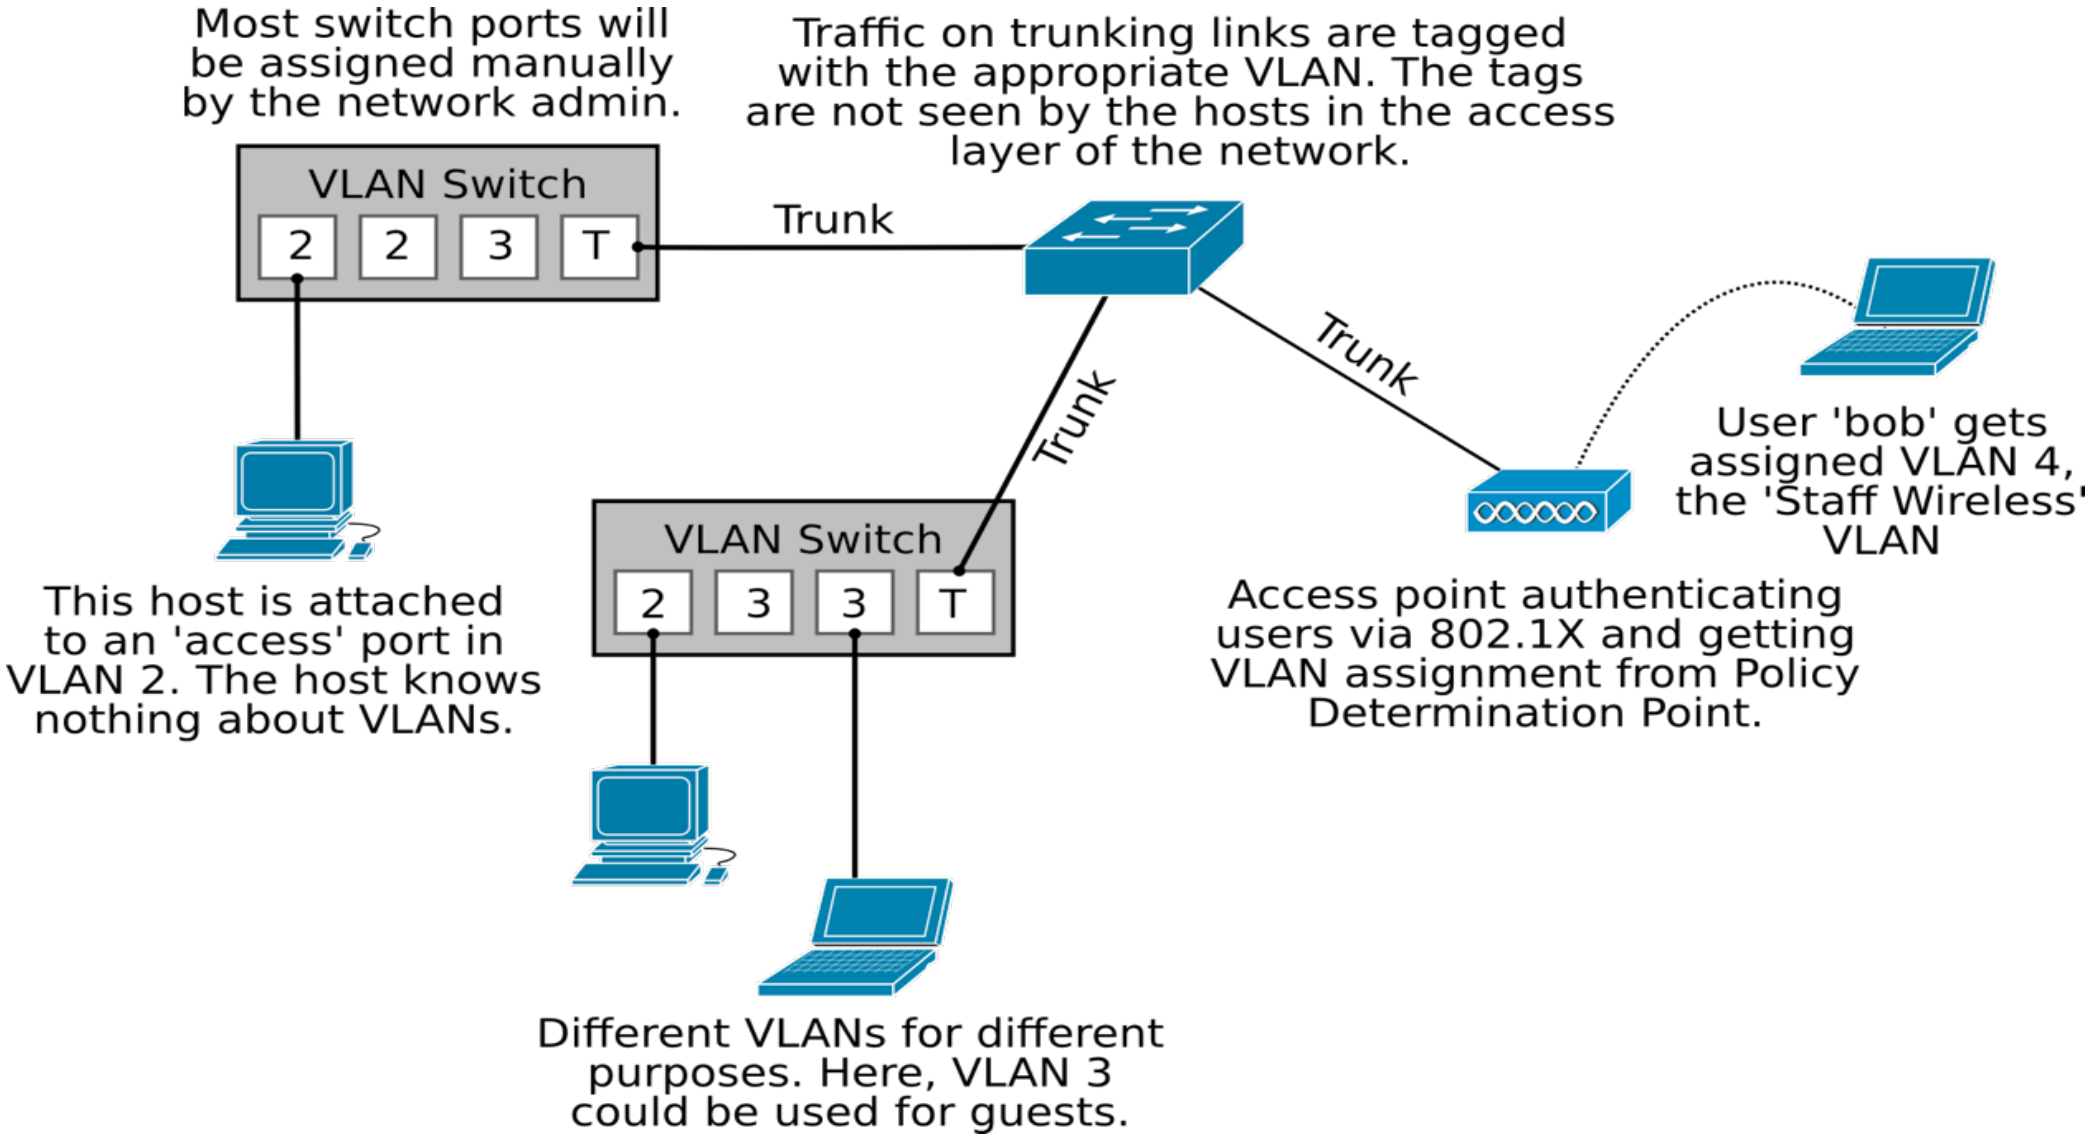

You will experiment with VLANs and internal routing in labs