Review Question

Examine the topology below

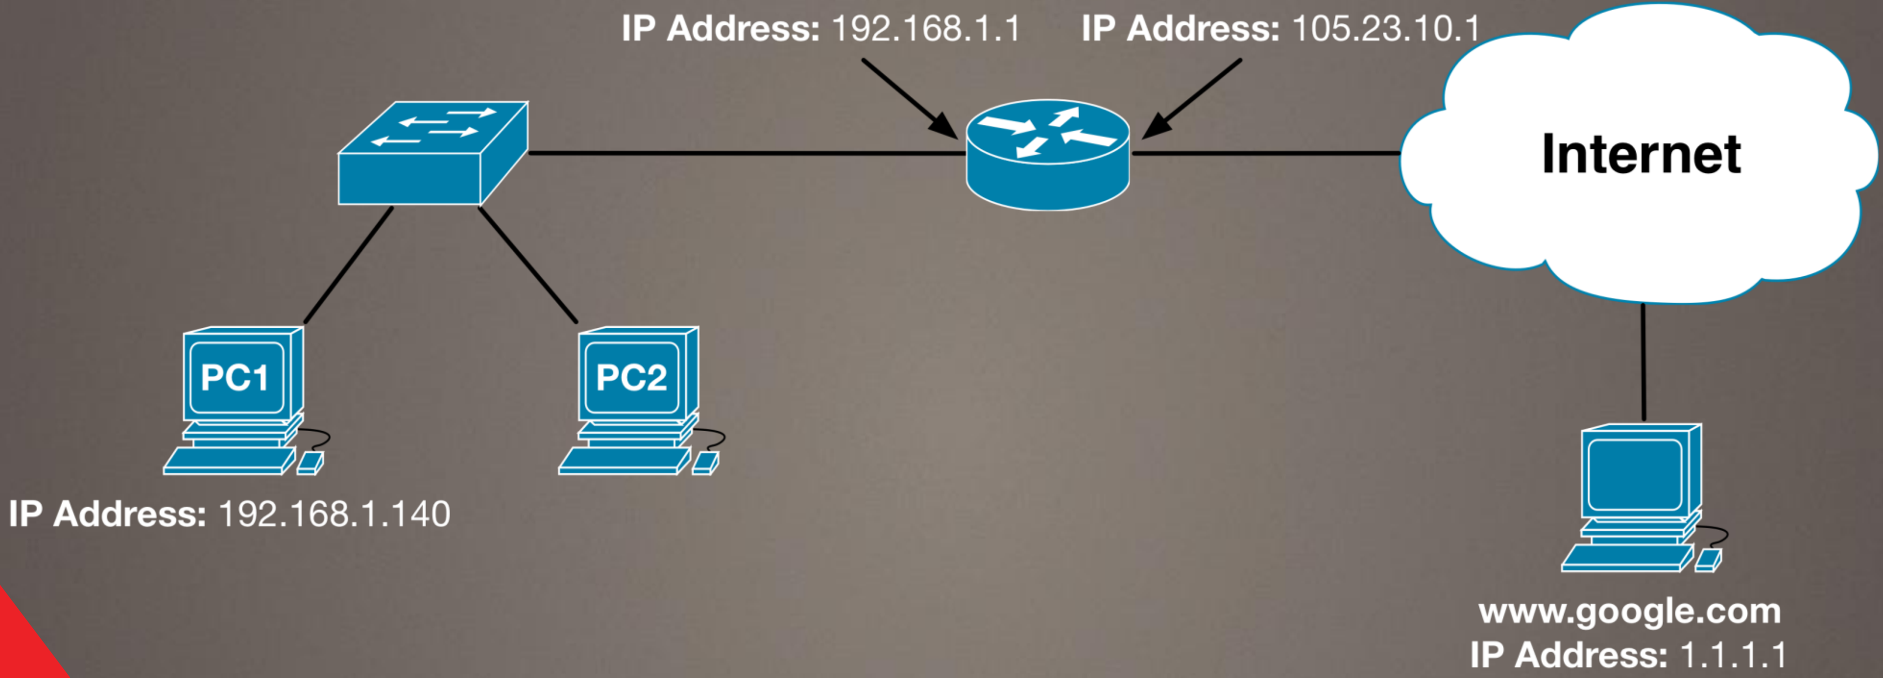

- What technology is required to allow PC1 to reach Google?
- What will the source IP address be for a packet sent by PC1 when received by Google's server?

Review Question

Examine the topology below

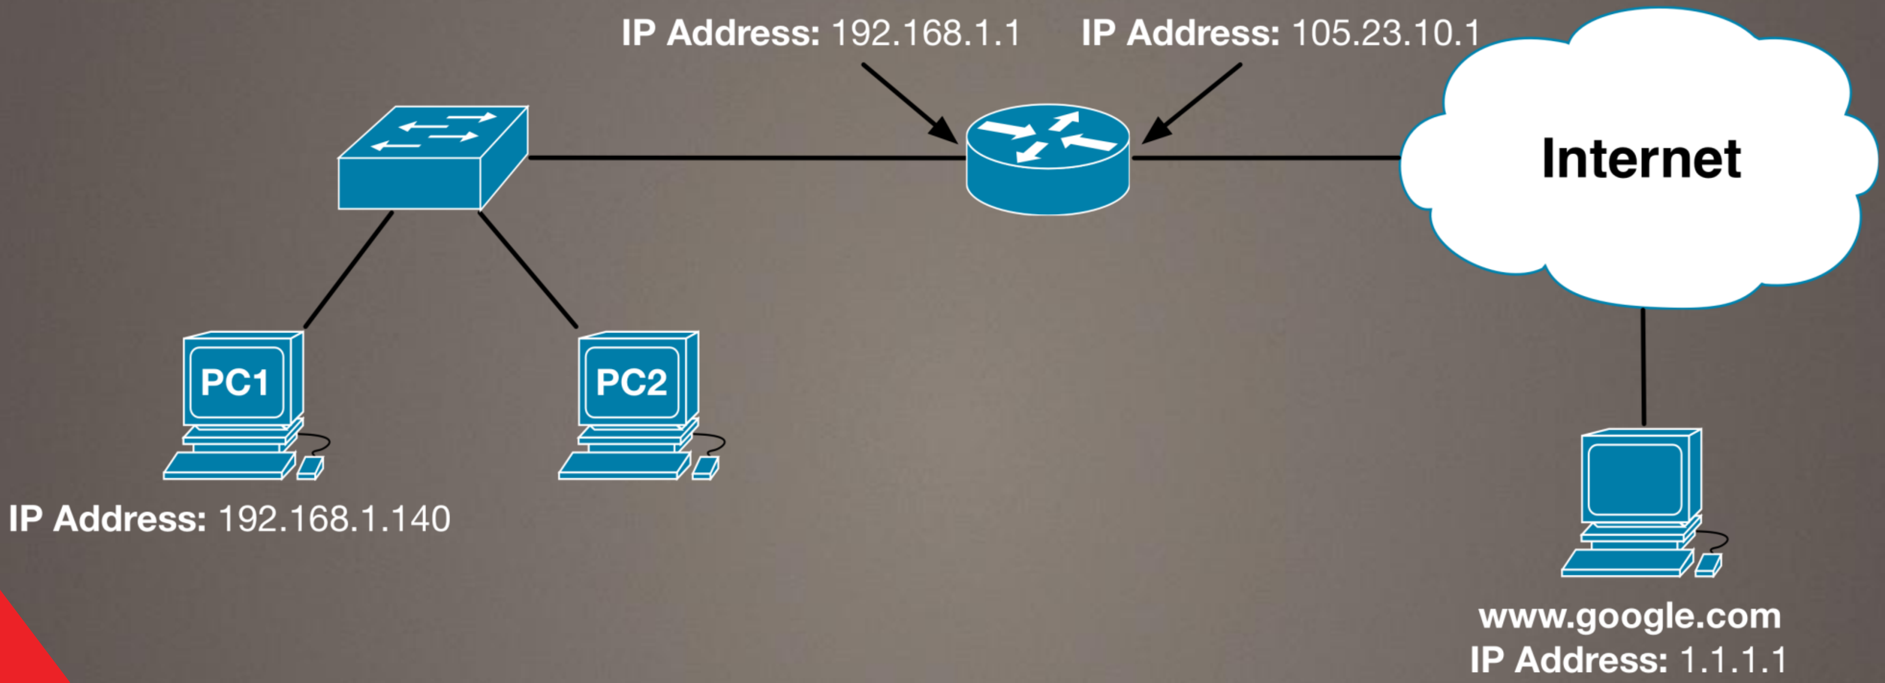

IP Address: 192.168.1.1

IP Address: 105.23.10.1

IP Address: 192.168.1.140

www.google.com
IP Address: 1.1.1.1

- What technology is required to allow PC1 to reach Google?

Network Address Translation

Review Question

Examine the topology below

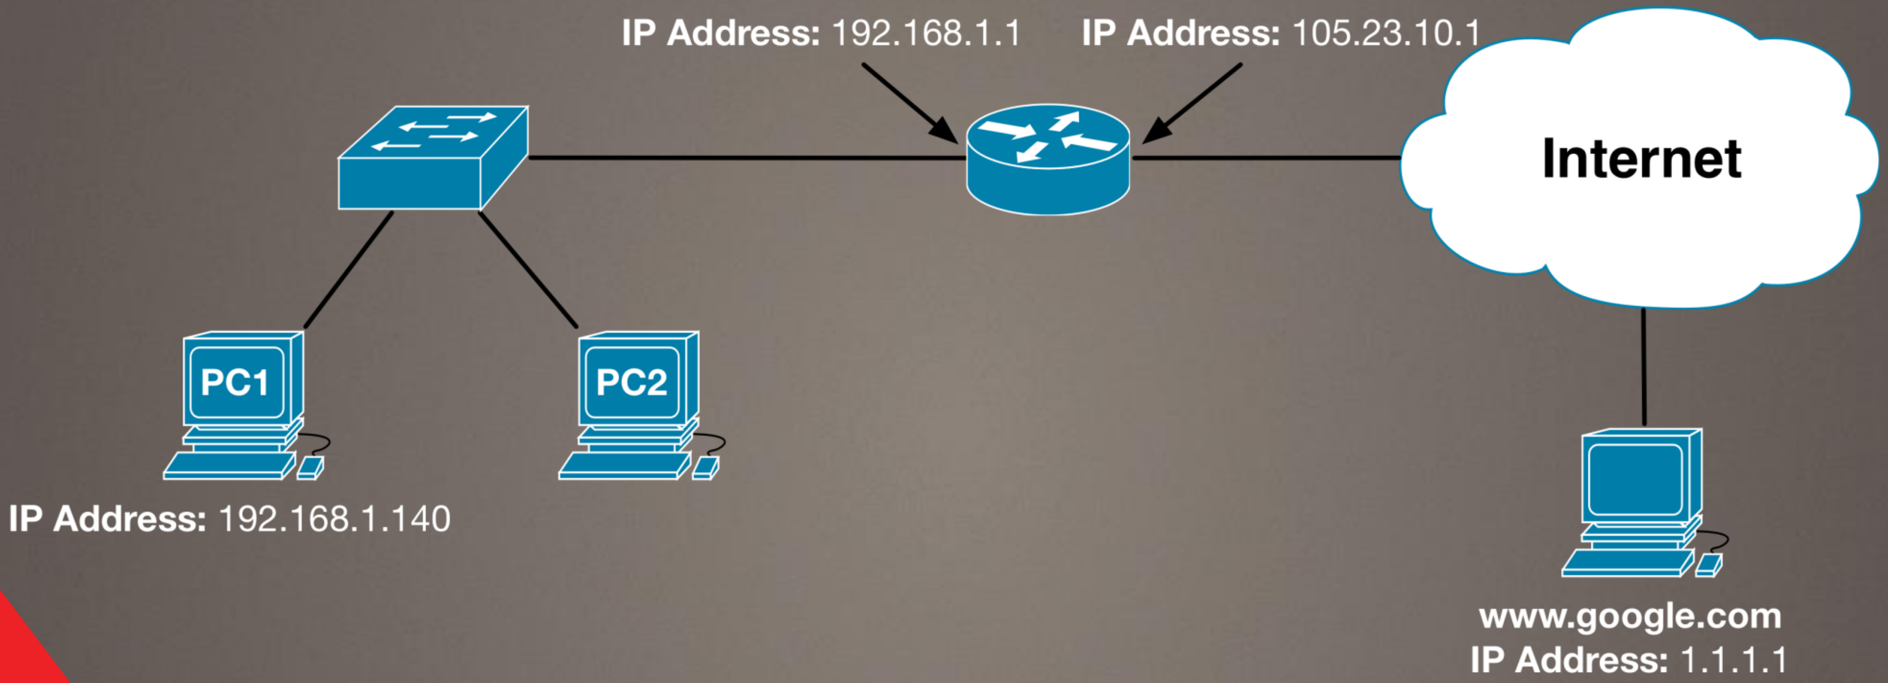

- What source IP address will be in the headers of the packet received by Google's server?

105.23.10.1

Review Question

Examine the topology below

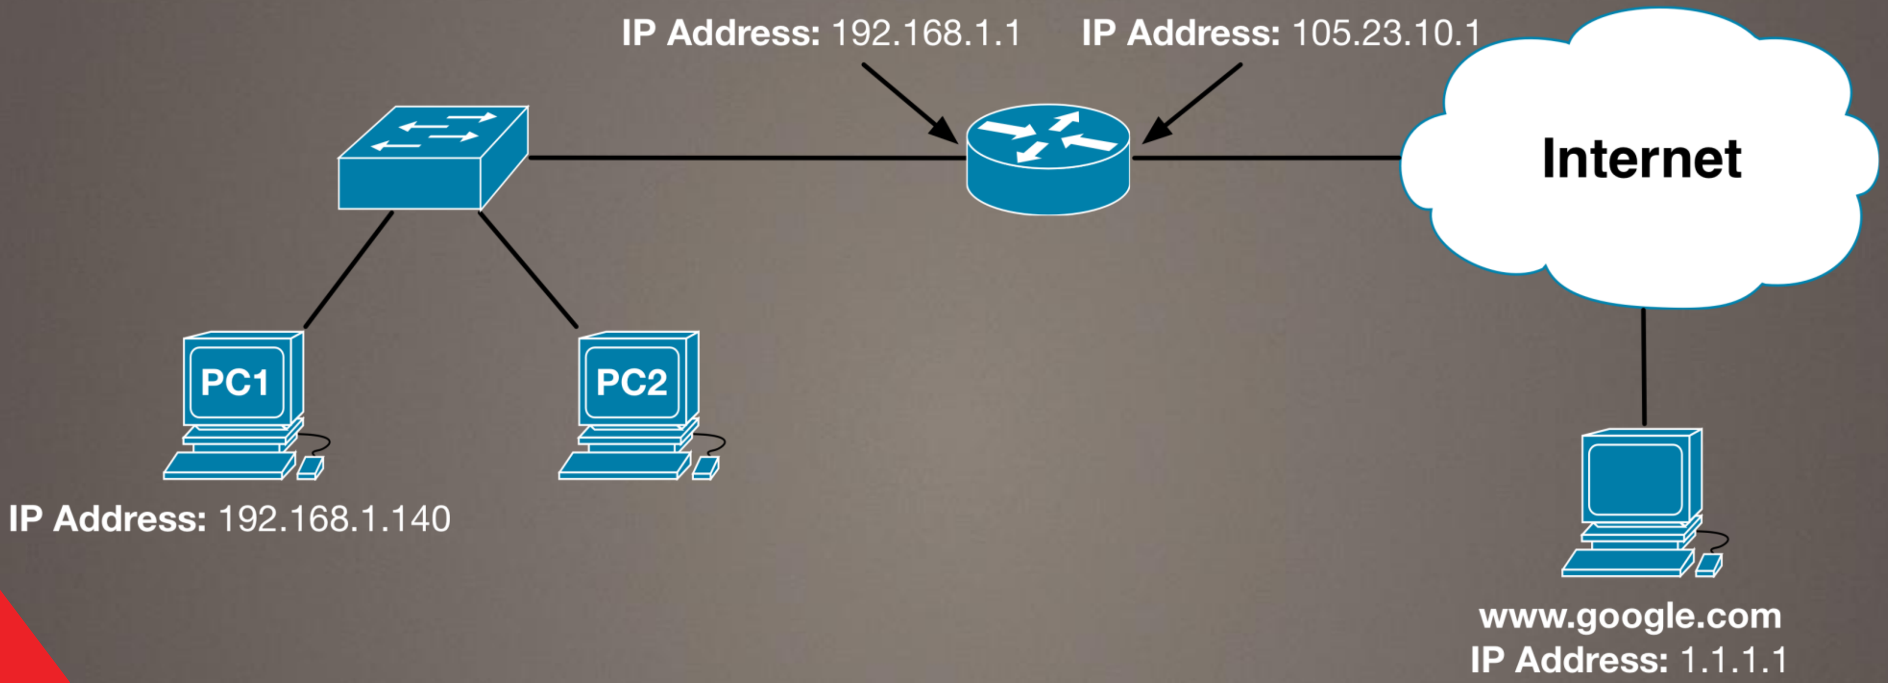

- What will happen if Google attempts to initiate a connection with PC1?

The packet will be dropped as the router is not expecting any communication from Google

Interior Routing

ICT169

Foundations of Data Communications

Admin

- Mid-Semester Test was last week
    - Contact me if you missed it due to illness or other extenuating circumstance
    - Results expected before next lecture
- Exam (currently) scheduled for Friday, 16 November
    - Two sessions beginning 9:30AM and 11:30AM
    - Allocation between sessions TBD
    - Ensure that you check the Exam Timetable as changes are possible until the morning of the exam

Last Week

- We discussed issues with IPv4 address depletion and some of the practices that allowed it to occur
- Also looked at two solutions:
  - Network Address Translation
  - IP version 6

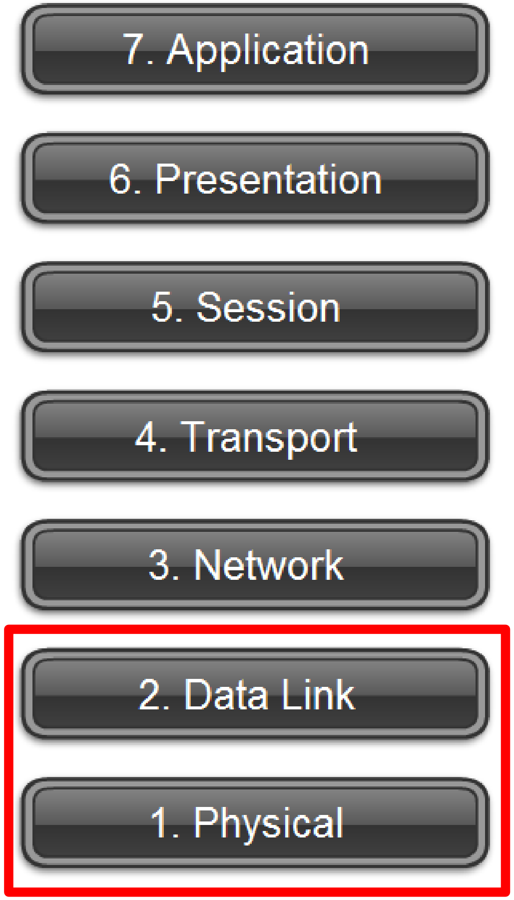

7. Application
6. Presentation
5. Session
4. Transport
3. Network
2. Data Link
1. Physical

Lecture Overview

- Examine how packets are forwarded between IP networks
- Different methods of generating routes:
    - Static routing
    - Dynamic routing
- Dynamic routing protocols
    - Interior and exterior gateway routing
    - Distance Vector and Link-State routing algorithms

Packet Forwarding Revisited

- When a device is readying a packet for submission, it first checks whether the destination is in its local subnet
- If the destination is in the local subnet, the devices will use Layer 2 addressing
- Otherwise, the packet will be forwarded to the default gateway

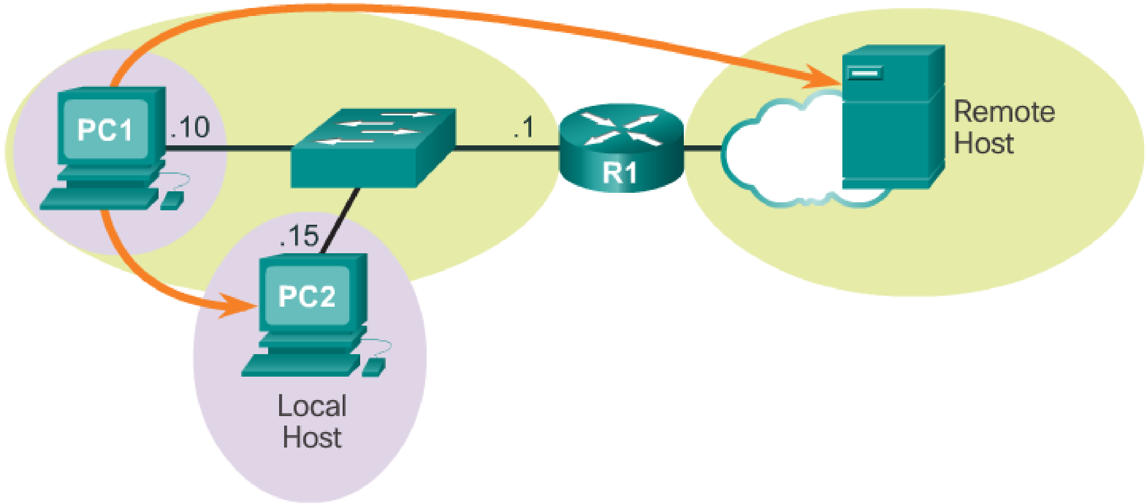

The Default Gateway Revisited

- The default gateway is the router responsible for forwarding traffic outside of the local subnet
- As the gateway exists within the local subnet, packets are forwarded to the gateway using its Layer 2 address

Gateways Enable Communications between Networks

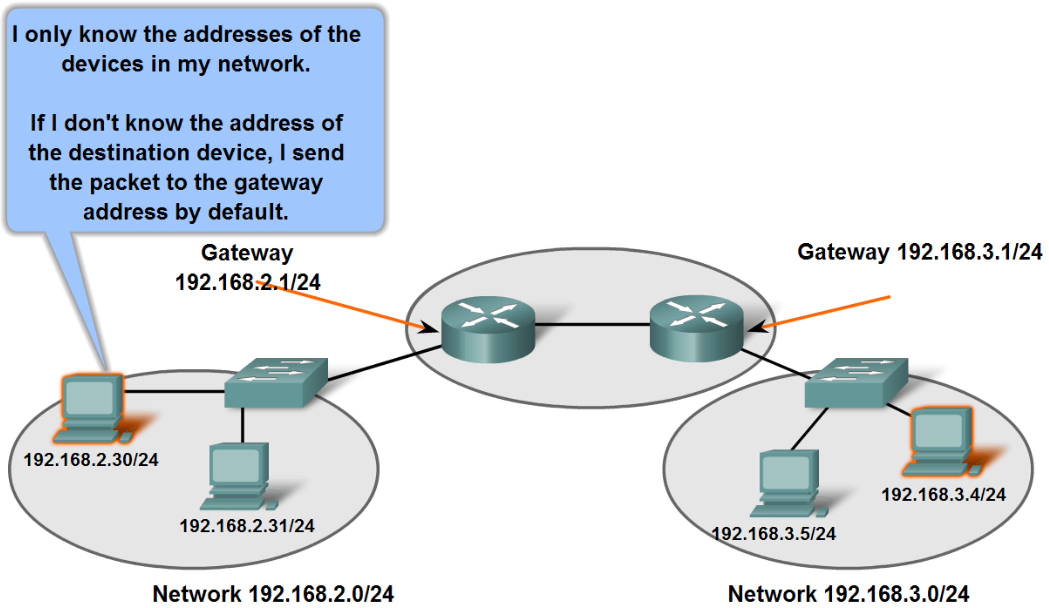

Routers and the Routing Process

- Routers are used to forward packets from one network to another
- Routers interfaces must belong to different IP networks
- Decision to forward the packet made based on the destination IP address

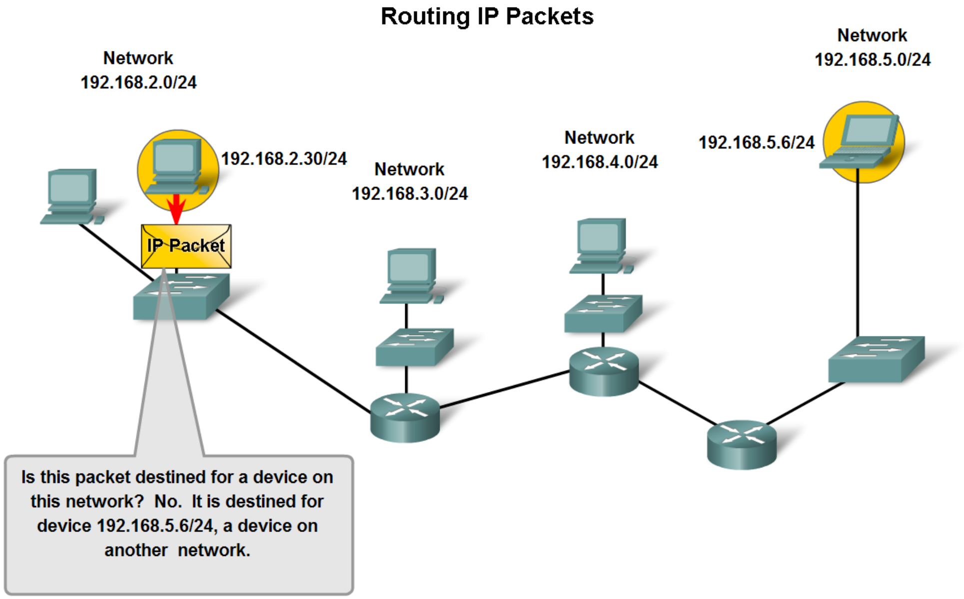

Is this packet destined for a device on this network? No. It is destined for device 192.168.5.6/24, a device on another network.

Routing Principles

Three key principles to routing:

1. Every router makes its decision alone, based on the information it has in its own routing table.
2. The fact that one router has certain information in its routing table does not mean that other routers have the same information.
3. Routing information about a path from one network to another does not provide routing information about the reverse, or return, path.

Murdoch UNIVERSITY

From Zinin, A. (2002). *Cisco IP routing: packet forwarding and intra-domain routing protocols*. Boston: Addison-Wesley.

The Routing Table

- Map of known possible destination networks
- Devices that operate at the network layer will have one
- Packets forwarded based on the longest match
- Packets for which no match exists will be sent to the default gateway or dropped

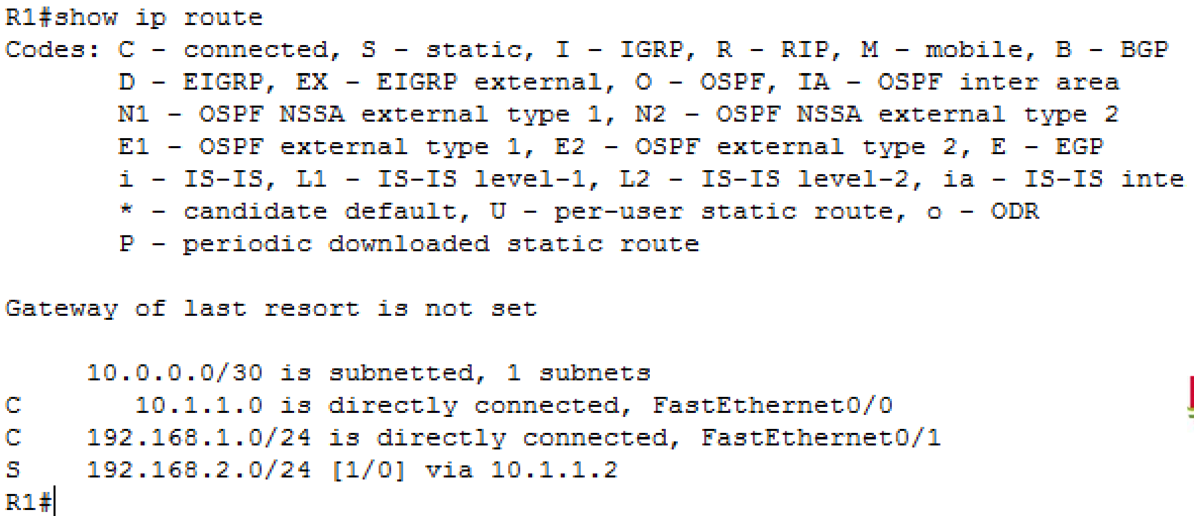

```
R1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inte
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     10.0.0.0/30 is subnetted, 1 subnets
C       10.1.1.0 is directly connected, FastEthernet0/0
C    192.168.1.0/24 is directly connected, FastEthernet0/1
S    192.168.2.0/24 [1/0] via 10.1.1.2
R1#
```

Matching Routes

- In the event that multiple routes match the destination address for a packet routes will select the longest match
- Longest match is based on **the number of bits that match**
- In the example, Route 3 is the longest match so the packet will be routed using it

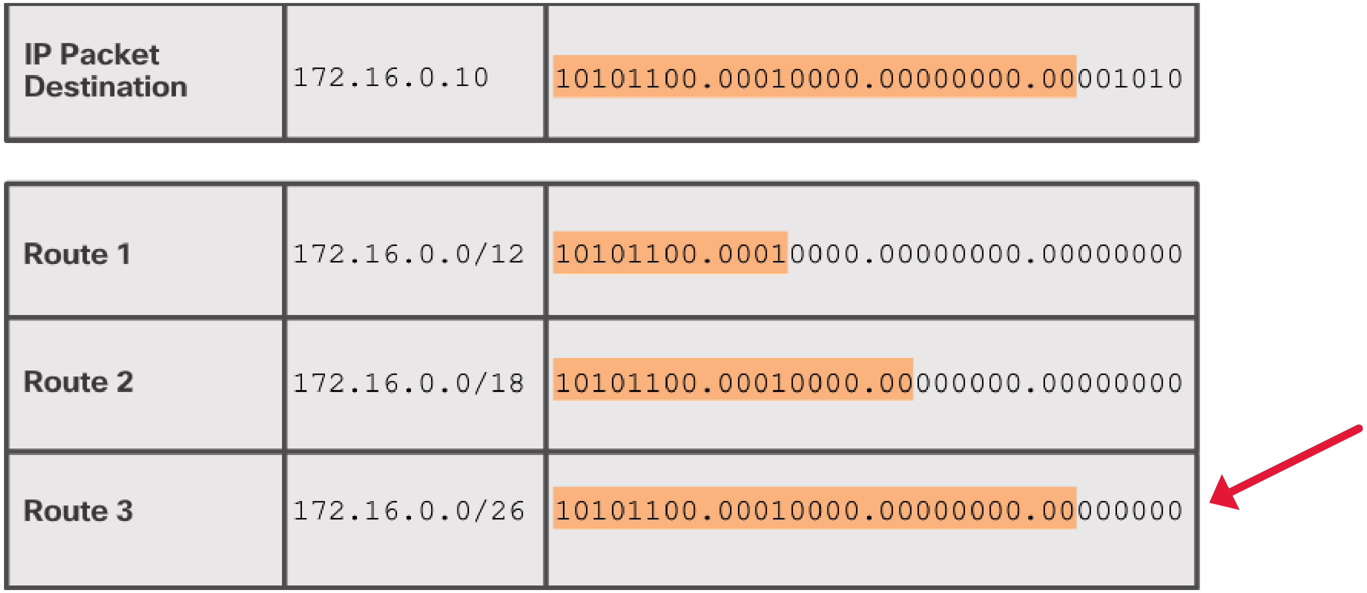

IP Packet Destination	172.16.0.10	10101100.00010000.00000000.00001010

Route 1	172.16.0.0/12	10101100.00010000.00000000.00000000
Route 2	172.16.0.0/18	10101100.00010000.00000000.00000000
Route 3	172.16.0.0/26	10101100.00010000.00000000.00000000

Longest match

The Routing Table (cont.)

- Exact format of routing table will depend on the device’s OS, but many elements in common
- Example of Cisco IOS routing table accessed using **`show ip route`**

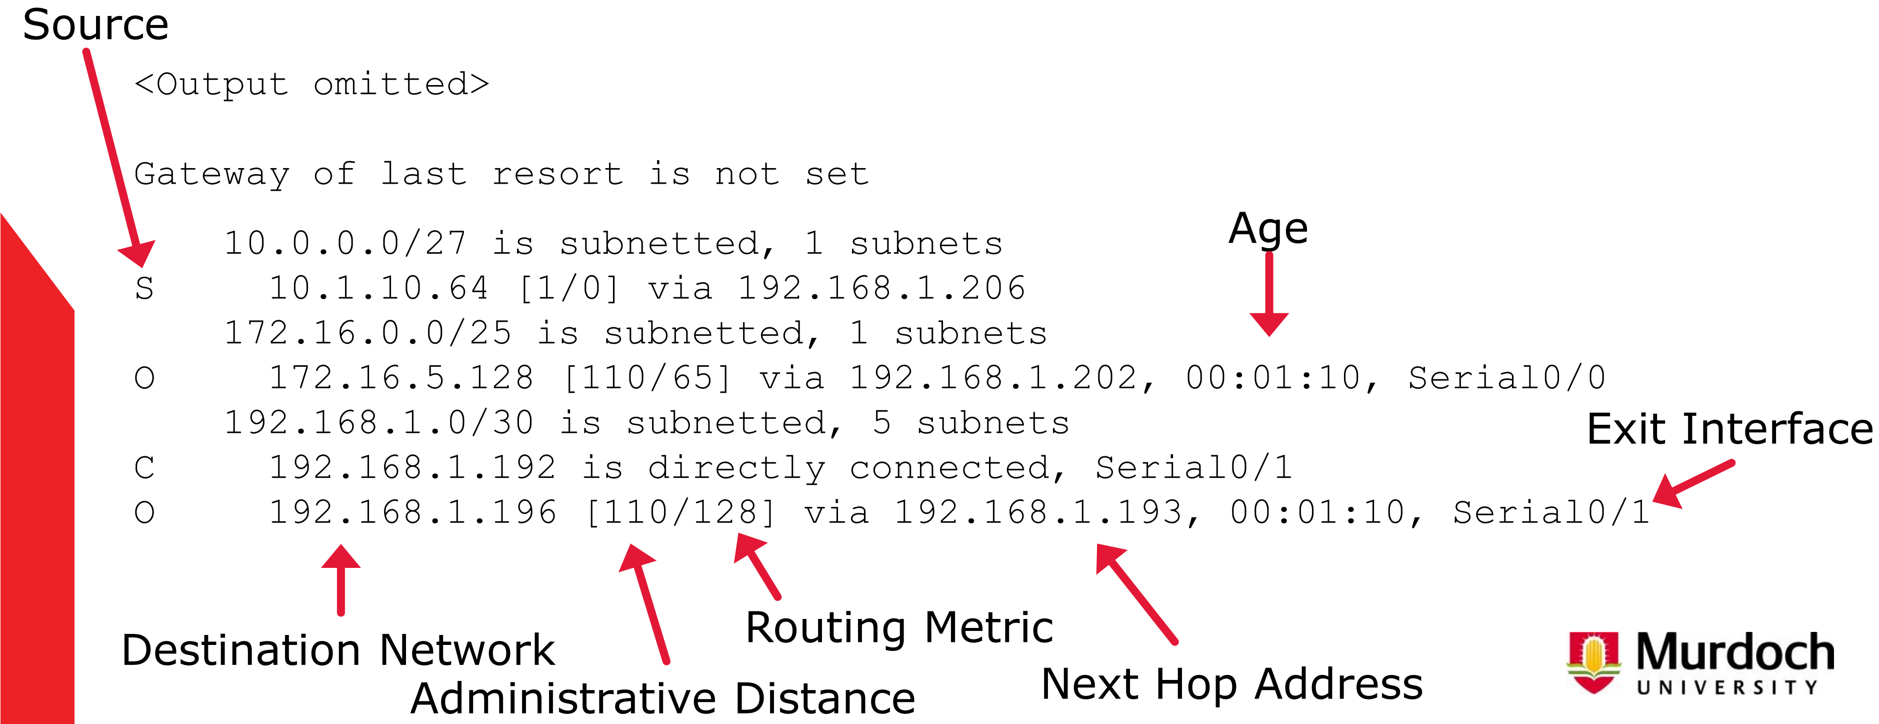

```
<Output omitted>

Gateway of last resort is not set

    10.0.0.0/27 is subnetted, 1 subnets
S      10.1.10.64 [1/0] via 192.168.1.206
    172.16.0.0/25 is subnetted, 1 subnets
O      172.16.5.128 [110/65] via 192.168.1.202, 00:01:10, Serial0/0
    192.168.1.0/30 is subnetted, 5 subnets
C      192.168.1.192 is directly connected, Serial0/1
O      192.168.1.196 [110/128] via 192.168.1.193, 00:01:10, Serial0/1
```

Source

Age

Exit Interface

Destination Network

Administrative Distance

Routing Metric

Next Hop Address

The Routing Table in Practice

- Assume that the router receives a packet destined for 192.168.2.10
- Should match the 192.168.2.0/24 route and forward via Gi0/1
- ARP will determine how the packet is forwarded

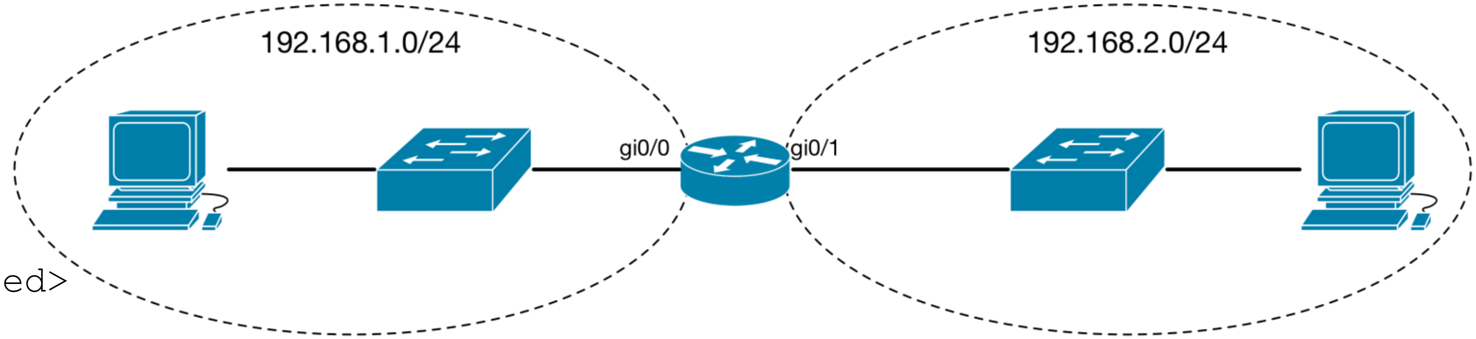

```
<Output omitted>

Gateway of last resort is not set

     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, GigabitEthernet0/0
L       192.168.1.1/32 is directly connected, GigabitEthernet0/0
     192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.2.0/24 is directly connected, GigabitEthernet0/1
L       192.168.2.1/32 is directly connected, GigabitEthernet0/1
```

The Routing Table in Practice (cont.)

- Assume R1 receives a packet destined for 192.168.2.10
- Should match the 192.168.2.0/24 route, but where is 10.1.1.2?

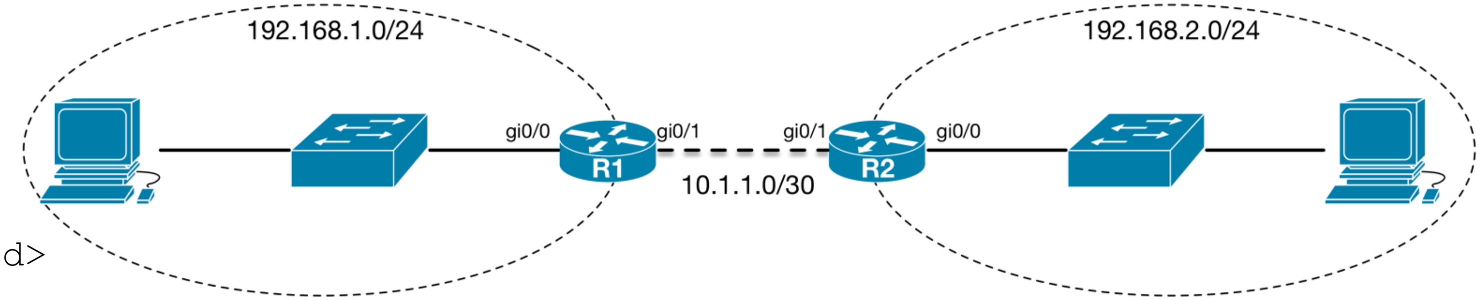

```
<Output omitted>

Gateway of last resort is not set

 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.1.1.0/30 is directly connected, GigabitEthernet0/1
L       10.1.1.1/32 is directly connected, GigabitEthernet0/1
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, GigabitEthernet0/0
L       192.168.1.1/32 is directly connected, GigabitEthernet0/0
S    192.168.2.0/24 [1/0] via 10.1.1.2
```

The Routing Table in Practice (cont.)

- R1 now needs to do a second lookup to determine how to reach 10.1.1.2
- Should now match against 10.1.1.0/30 and forward the packet out of gi0/1 to R2
- R2 is now responsible for forwarding the packet

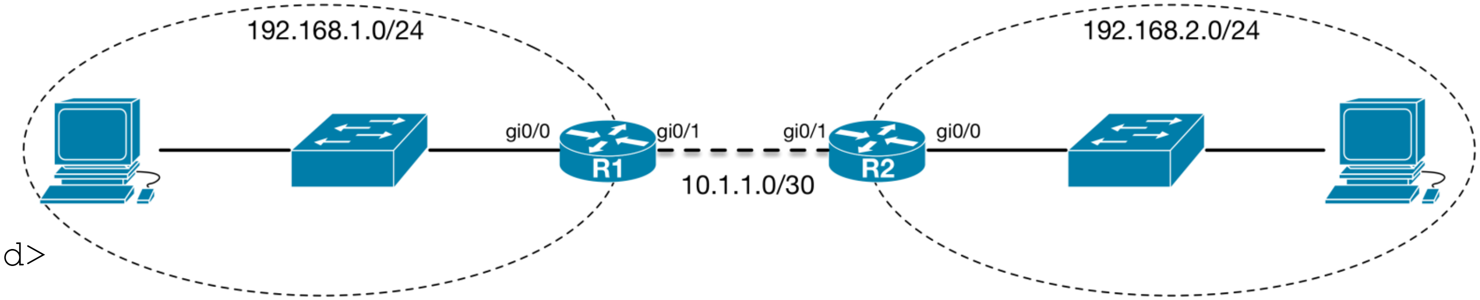

```
<Output omitted>

Gateway of last resort is not set

 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.1.1.0/30 is directly connected, GigabitEthernet0/1
L       10.1.1.1/32 is directly connected, GigabitEthernet0/1
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, GigabitEthernet0/0
L       192.168.1.1/32 is directly connected, GigabitEthernet0/0
S    192.168.2.0/24 [1/0] via 10.1.1.2
```

Network Discovery

- When a router first boots, routing information about **directly connected networks** is installed in the routing table
- In this example, R1 would only have routes to 192.168.1.0/24 and 10.1.1.0/30
- How do we tell R1 about 192.168.2.0/24?

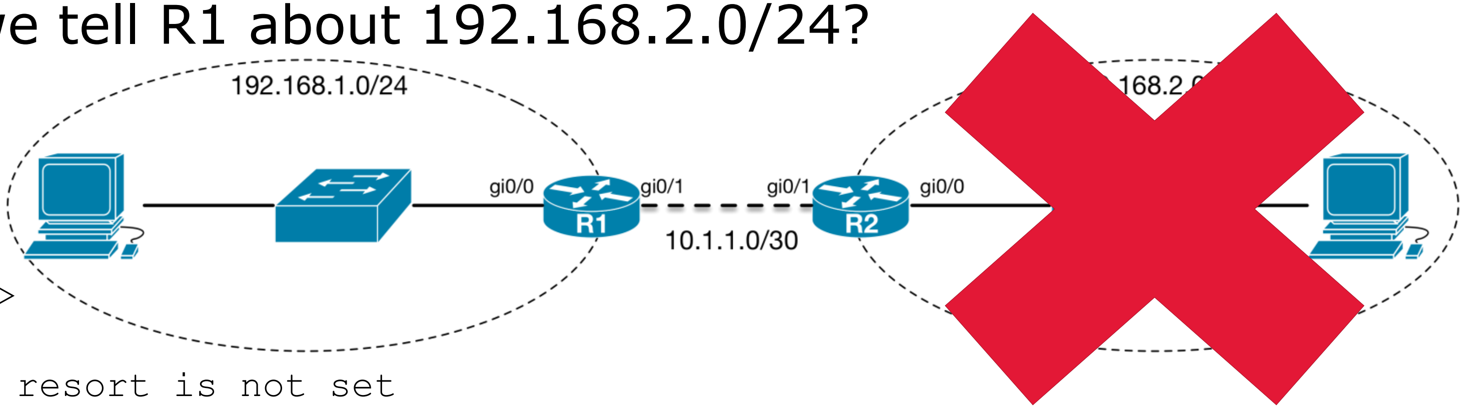

```
<Output omitted>

Gateway of last resort is not set

 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.1.1.0/30 is directly connected, GigabitEthernet0/1
L       10.1.1.1/32 is directly connected, GigabitEthernet0/1
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, GigabitEthernet0/0
L       192.168.1.1/32 is directly connected, GigabitEthernet0/0
```

Establishing Routes

- Routers are able to acquire knowledge of remote networks using two techniques:
  - A network engineer may manually enter a **static route**
  - **Routing protocols** may be used to propagate information between routers

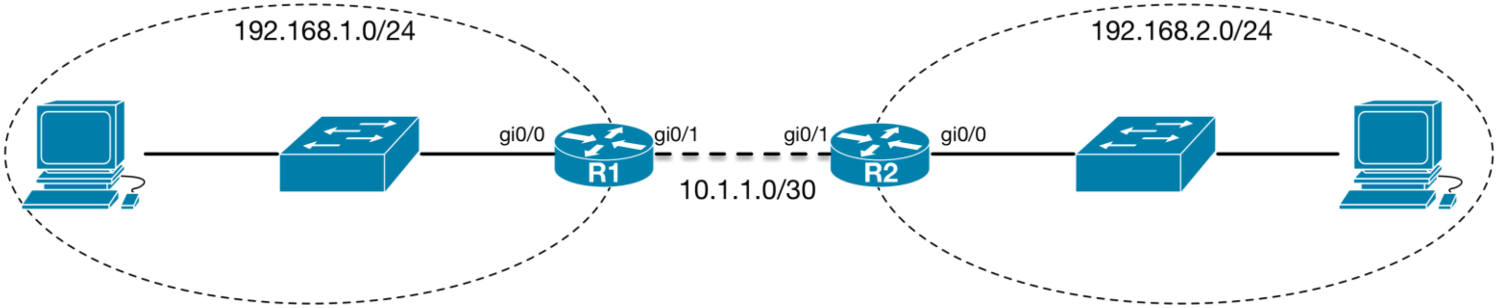

Static Routing

- Static routes are **manually configured** by the network engineer
- Static routing is ideal for small networks with limited redundancy
- Also used when control over routing must be maximised, or the overheads of routing protocols are to be avoided

Configuring Static Routing

- Static routes are usually configured by specifying:
  - Destination network
  - Subnet mask (of destination network)
  - Next hop IP address or;
  - Exit interface (only when using point-to-point links)

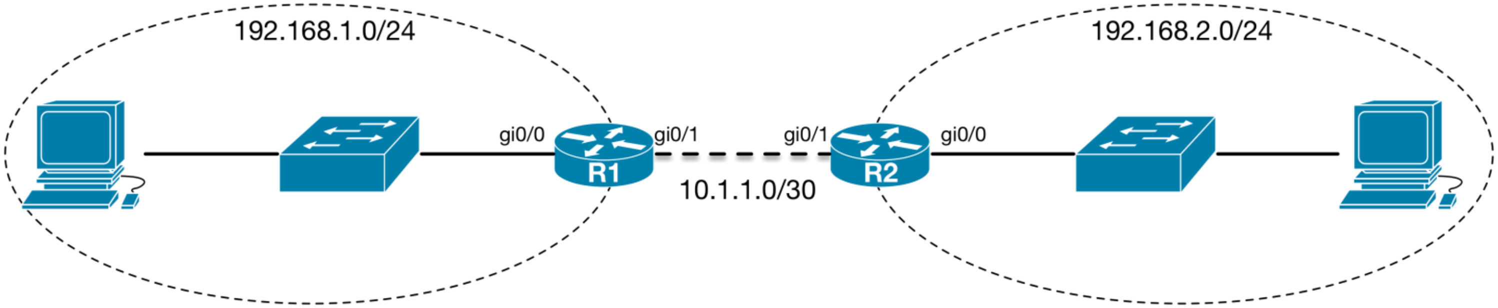

```
S     192.168.2.0/24 [1/0] via 10.1.1.2
```

The Default Route

- A ‘catch all’ static route that specifies how packets should be forwarded when no specific route is available
    - The default gateway is a default route
- Can also be used when a network has only one exit point (eg. residential networks)
- If no default route is configured, **packets for which no route exists will be dropped**

Limitations of Static Routing

- Static routing is not scalable; labour intensive as network grows
- Difficult to predict consequences of installing a new route
- Not adaptable to changes in the network; routes must be manually updated

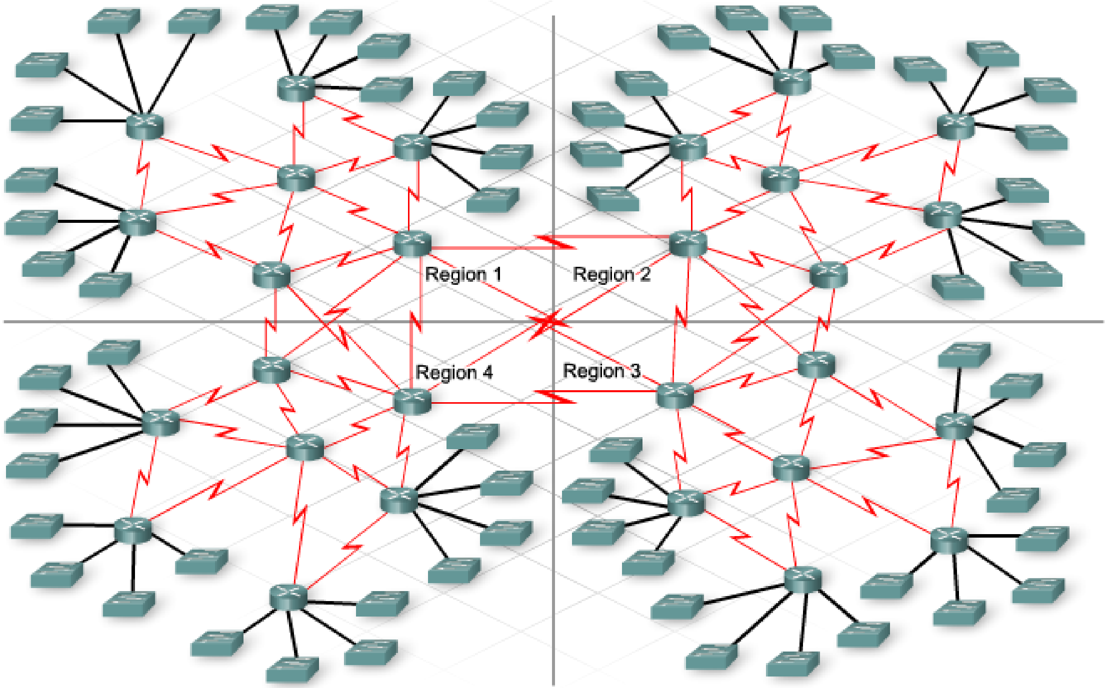

Dynamic Routing

- More commonly, dynamic routing protocols are used to propagate routing information throughout the network
  - Advertise networks the router knows how to reach, rather than specify networks you want to get to
- Determines best path to a given destination without human intervention
- Routers **automatically update routing table** when topology changes

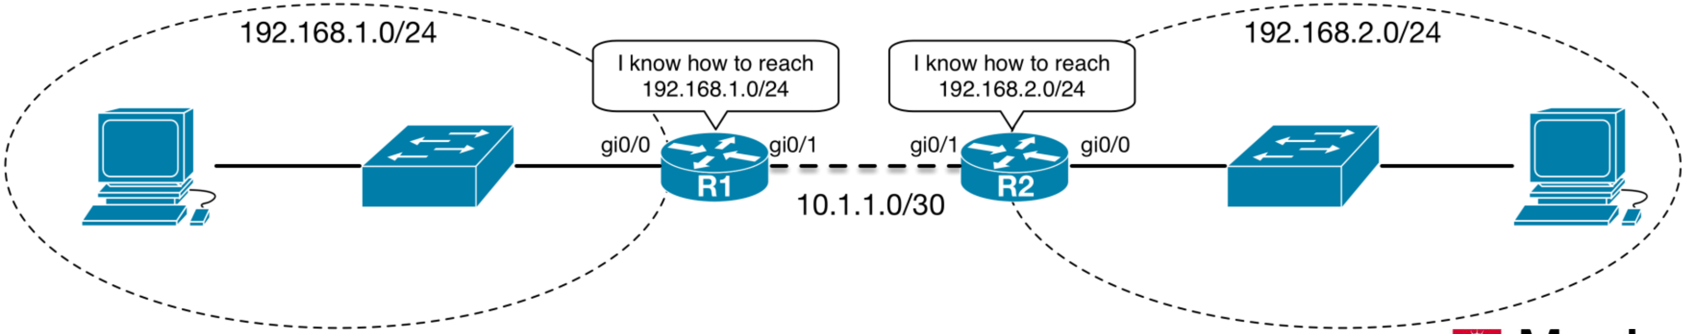

Dynamic Routing (cont.)

- To achieve this functions, dynamic routing protocols must:
    - Discover remote networks
    - Update routing information
    - Identify the 'best' path to destination networks and install this information in the routing table
    - Identify alternate path(s) to destination networks in the event of link failure
- Routing protocols don't make forwarding decisions

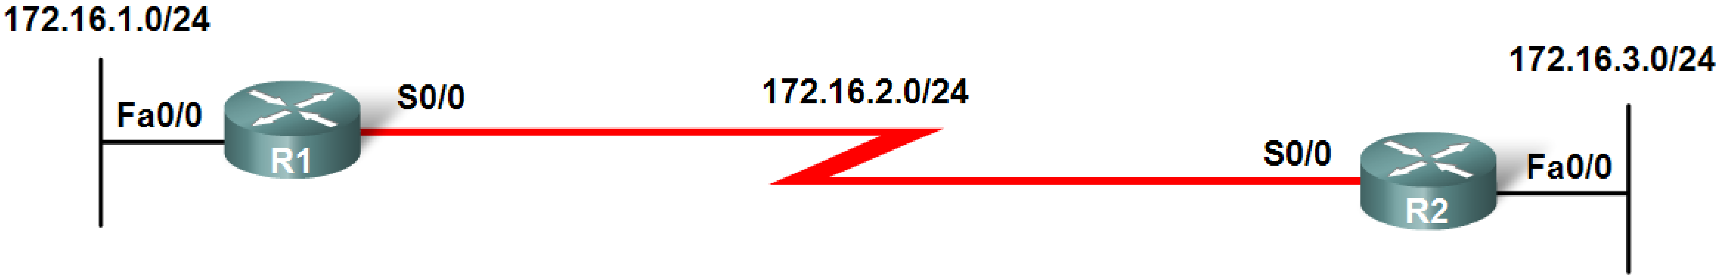

Dynamic Routing Protocols

- Routing protocols can be divided into two categories
- **Interior Gateway Protocols (IGPs)** are used to route within an administrative domain (eg. a single organisation)
  - Some examples include: Router Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), and Open Shortest Path First (OSPF)
- **Exterior Gateway Protocols (EGPs)** are used for routing packets between autonomous systems (eg. between organisations)
  - Only one EGP in common use: Border Gateway Protocol (BGP)
  - Used to route Internet traffic

Routing Protocol Terminology

- **Convergence** describes a state where the routing table for all routers is consistent
- **Metrics** are used by routing protocols to determine the best route
- **Administrative Distance** specifies the priority level of a route; usually determined by the information source
- **Load balancing** describes the capability to split load across multiple equivalent paths

```
<Output omitted>
O     192.168.1.196 [110/128] via 192.168.1.193, 00:01:10, Serial0/1
```

Routing Metric

Administrative Distance

Routing Metrics

- Metric used varies by routing protocol
- Some common examples:
    - Hop count
    - Bandwidth
    - Delay
    - Load
    - Reliability
- Lower costs are ‘better’

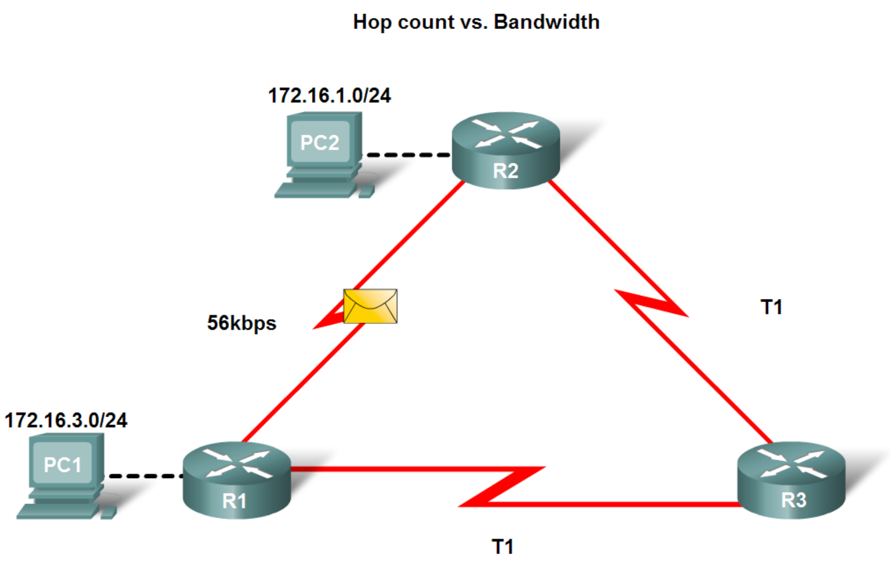

```
<Output omitted>
O      192.168.1.196 [110/128] via 192.168.1.193, 00:01:10, Serial0/1
```

Routing Metric

Administrative Distance

- Each source of routing information is assigned an administrative distance
- Lower administrative distance indicates higher priority

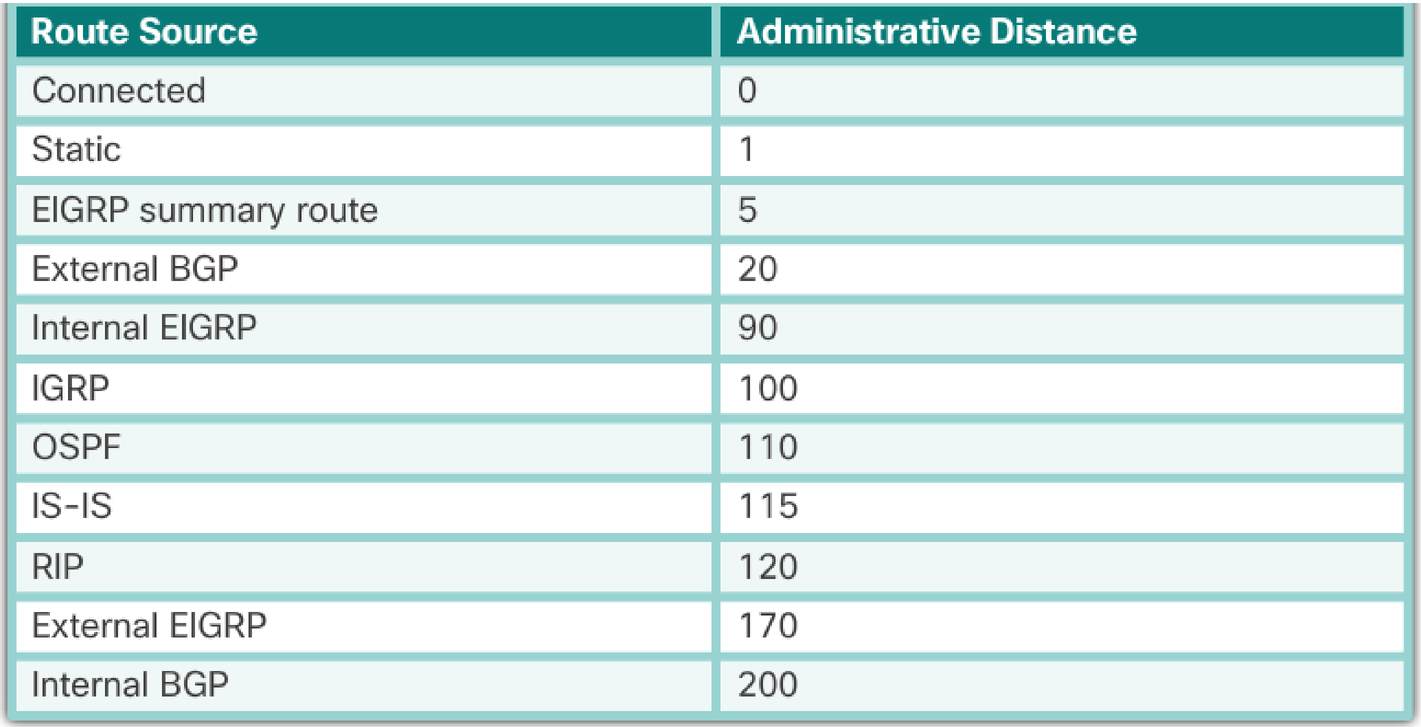

Route Source	Administrative Distance
Connected	0
Static	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200

Interior Gateway Protocols

- Two classes of IGPs also exist
- Distance Vector Routing Protocols
    - Routers advertise all knowledge about the network to neighbours
    - Routes are advertised as vectors of distance and direction
    - Advertisements are transmitted periodically
    - Each router has a limited or incomplete view of the network
    - Example: RIP
- Link-State Routing Protocols
    - Routers advertise directly connected networks to all others
    - Each router constructs its own shortest path tree (topology diagram)
    - Updates are sent when network topology changes
    - Each router forms a complete view of the network
    - Example: OSPF

Break

When we return: A closer look at Distance Vector and Link-State routing

Distance Vector Routing

- In traditional distance vector routing protocols, the entire routing table is periodically broadcasted to all neighbours
- Routers must construct routing table based only on what they receive from neighbours
    - No knowledge of the full network topology
    - Sometimes referred to as ‘routing by rumour’

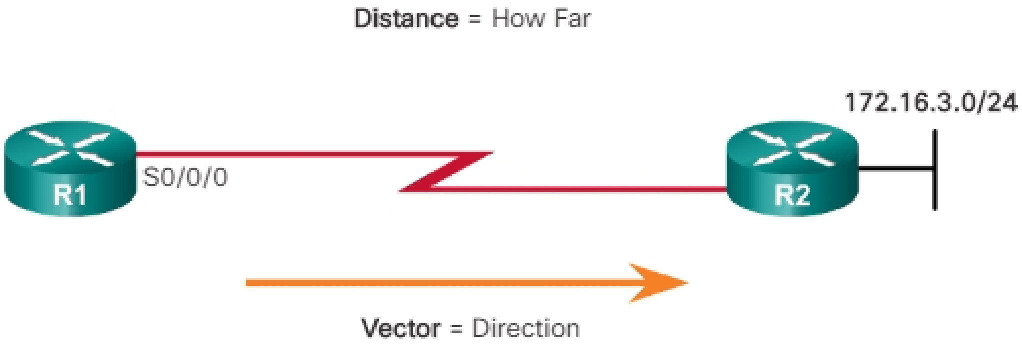

Distance Vector Algorithm

- Iterative algorithm that runs asynchronously
- Distributed; router notifies
- Iteration triggered by:
    - Local link cost change
    - Neighbour updates
- Bellman-Ford algorithm

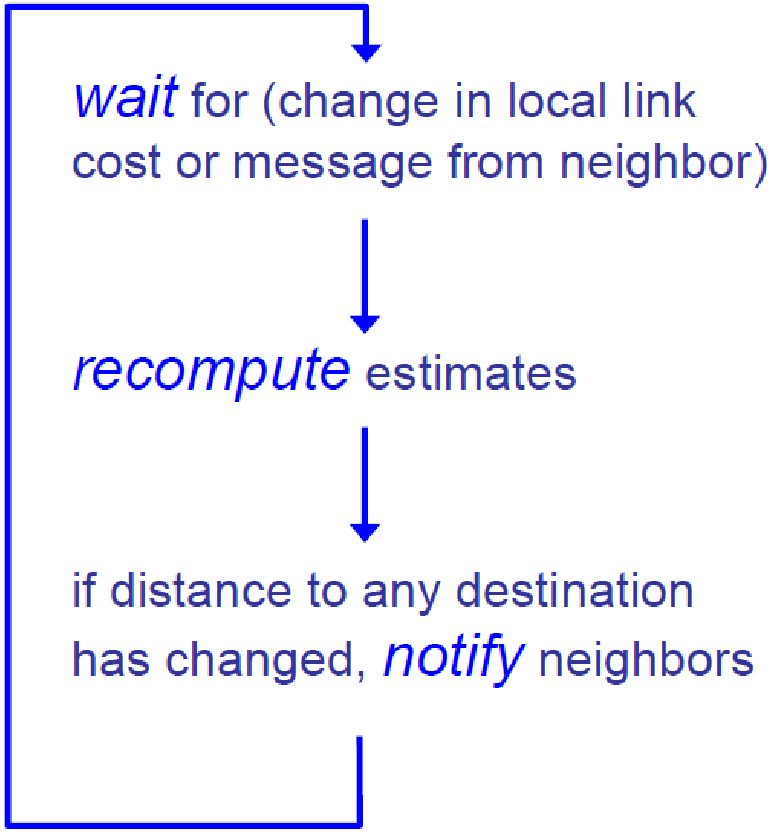

wait for (change in local link cost or message from neighbor)

recompute estimates

if distance to any destination has changed, *notify* neighbors

Distance Vector Algorithm – An Example

- Initial network state: all routers have knowledge of all other directly connected routers
- At this stage, all non-connected routers have costs of ∞

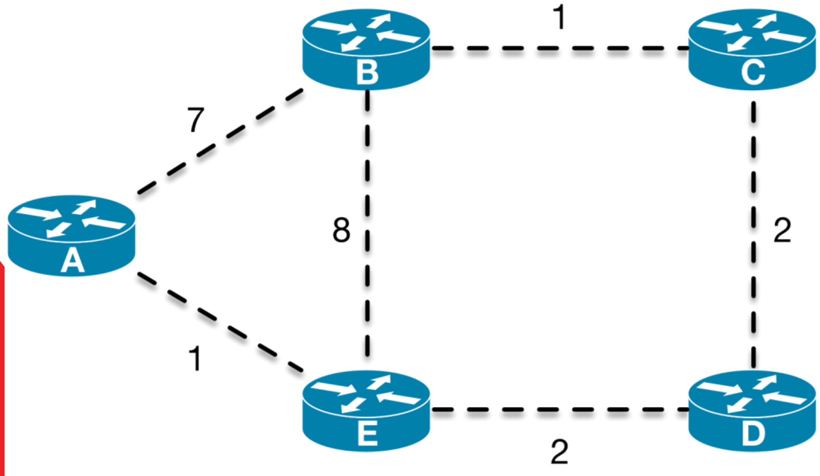

	Distance to node				
Info at node	A	B	C	D	E
A	0	7	∞	∞	1
B	7	0	1	∞	8
C	∞	1	0	2	∞
D	∞	∞	2	0	2
E	1	8	∞	2	0

Distance Vector Algorithm – An Example (cont.)

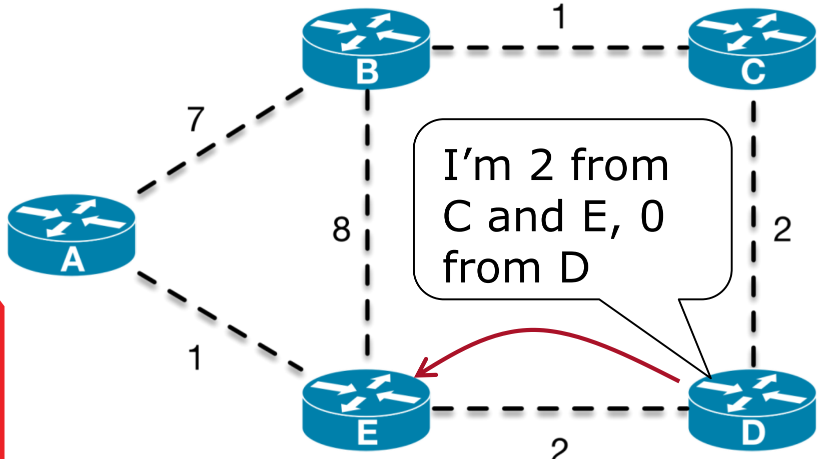

I’m 2 from C and E, 0 from D

D is 2 away, and 2+2 < ∞ so C is 4 away

Distance to node

Info at node	A	B	C	D	E
A	0	7	∞	∞	1
B	7	0	1	∞	8
C	∞	1	0	2	∞
D	∞	∞	2	0	2
E	1	8	**4**	2	0

Distance Vector Algorithm – An Example (cont.)

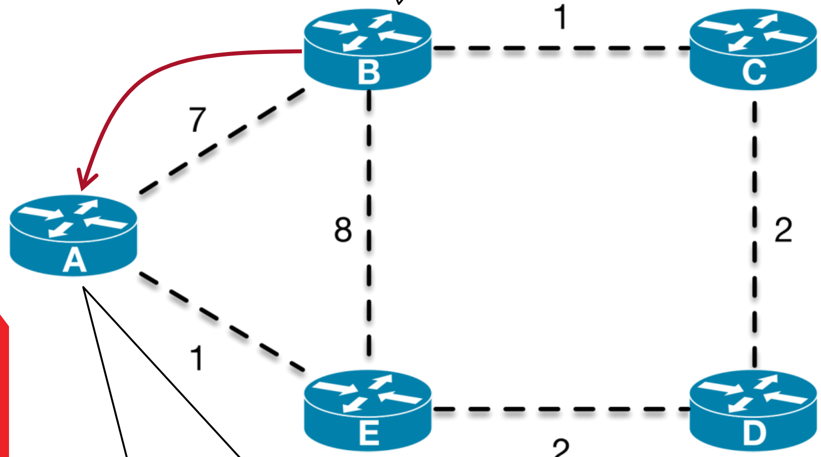

I'm 7 from A, 0 from B, 1 from C, 8 from E

B is 7 away, and 7+1 < ∞ so C is 8 away

Distance to node

Info at node	A	B	C	D	E
A	0	7	**8**	∞	1
B	7	0	1	∞	8
C	∞	1	0	2	∞
D	∞	∞	2	0	2
E	1	8	4	2	0

Distance Vector Algorithm – An Example (cont.)

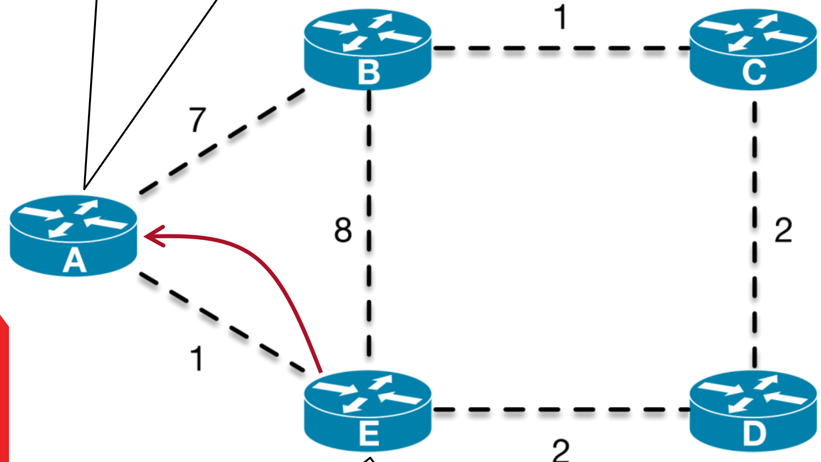

E is 1 away, and
1+4 < 8 so C is 5 away
1+2 < ∞ so D is 3 away

I'm 1 from A, 8 from B, 4 from C, 2 from D, 0 from E

	Distance to node				
Info at node	A	B	C	D	E
A	0	7	**5**	**3**	1
B	7	0	1	∞	8
C	∞	1	0	2	∞
D	∞	∞	2	0	2
E	1	8	4	2	0

Distance Vector Algorithm – An Example (cont.)

- Eventually, we reach a converged state
- All routers now have a path to reach all others

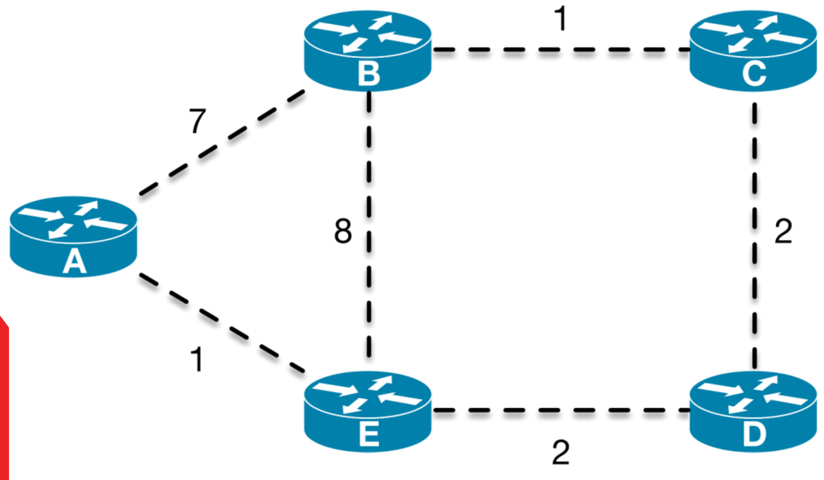

Distance to node

Info at node	A	B	C	D	E
A	0	6	5	3	1
B	6	0	1	3	5
C	5	1	0	2	4
D	3	3	2	0	2
E	1	5	4	2	0

Distance Vector Algorithm – Recovering from Link Failure

- Assume the link between Routers A and E fails
- Both routers must now mark the cost of their link to ∞
- Inform their neighbouring routers who must then recompute costs

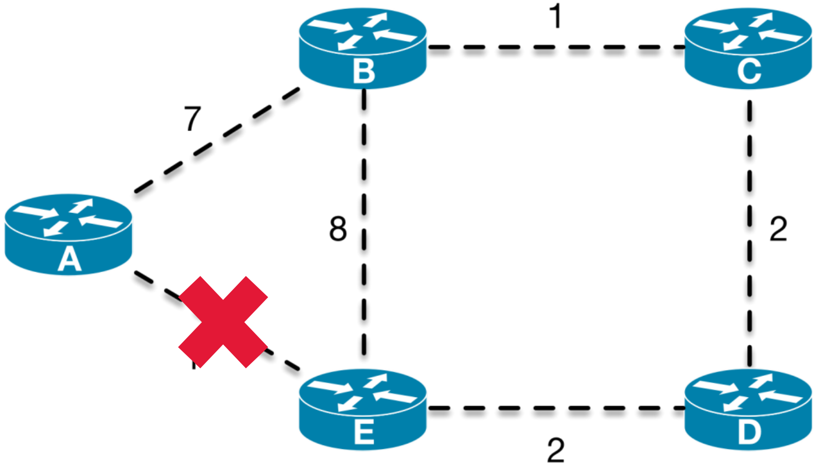

Distance to node

Info at node	A	B	C	D	E
A	0	7	8	10	12
B	7	0	1	3	5
C	8	1	0	2	4
D	10	3	2	0	2
E	12	5	4	2	0

Distance Vector Routing Loops

- Distance Vector routing is referred to as ‘route by rumour’
- Loops can form due to mistiming of periodic updates
- Referred to as **count to infinity** loops

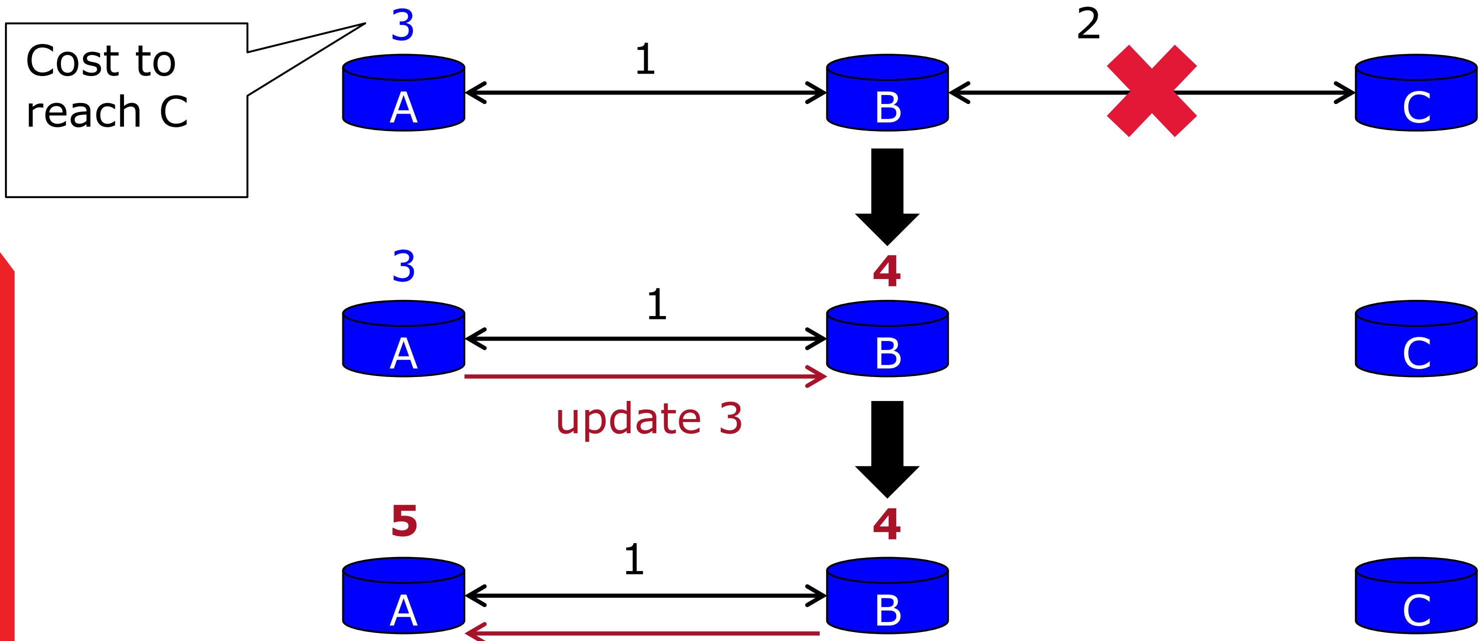

Murdoch UNIVERSITY

Count to Infinity

- Occurs due to lack of information (and some bad timing)
- In the example, Router B accepts a route to C via A (which is actually through B)
- Three techniques exist to mitigate this issue:
    - Triggered updates – update neighbours immediately when topology change occurs
    - Split horizon – don't re-advertise network through the interface it was learned through (Router A should not advertise Network C to Router B)
    - Poison reverse – update links that go down to have a cost of ∞ (Router B advertises Network C as having cost = ∞)

Link-State Routing

- More complex than distance vector protocols
- Establishes and maintains relationships with neighbouring routers
- Provides Link State Advertisements on all interfaces:
    - LSAs contain information about networks the router is connected to
    - Neighbouring routers propagate LSAs
- Each router computes shortest paths to all destinations
- Uses Dijkstra's algorithm

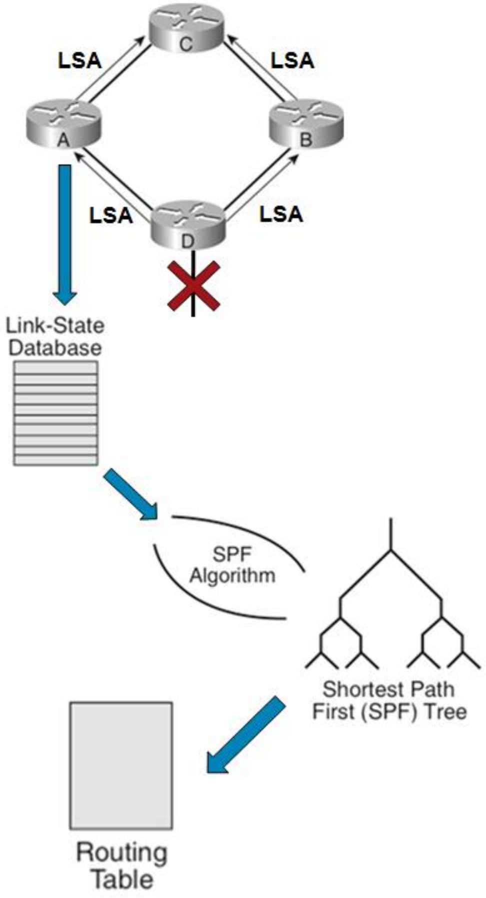

Dijkstra’s Algorithm – An Example

- Let’s build the tree for Router A
- Start by calculating the cumulative cost to reach the other routers in the network
- Routers (and costs) marked in green are confirmed
- Routers (and costs) marked in orange are tentative

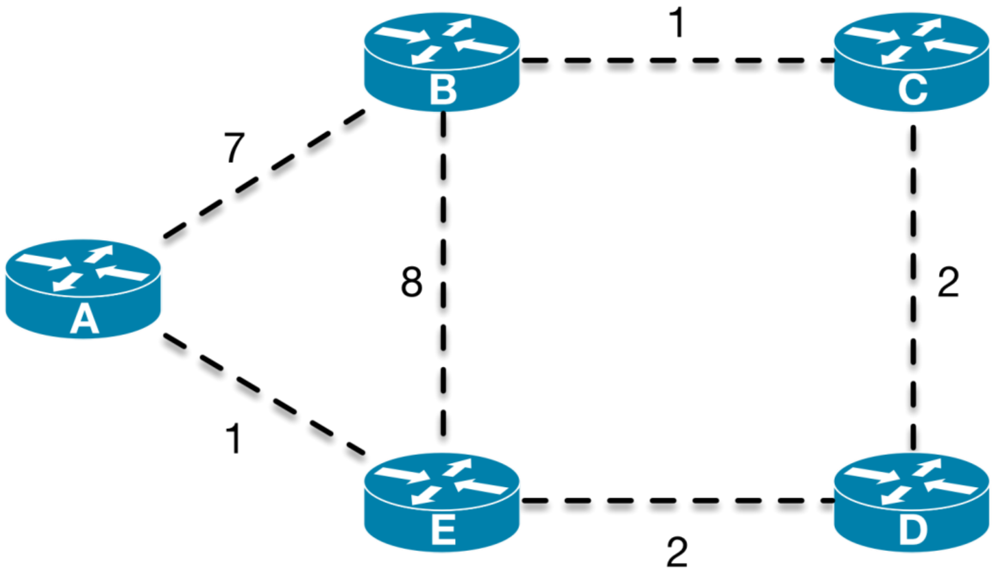

Dijkstra’s Algorithm – An Example

- Router A can immediately be confirmed; cost to reach yourself should always be 0
- Cost to reach Routers B and E can also be confirmed
- Cost to reach Routers C and D tentative

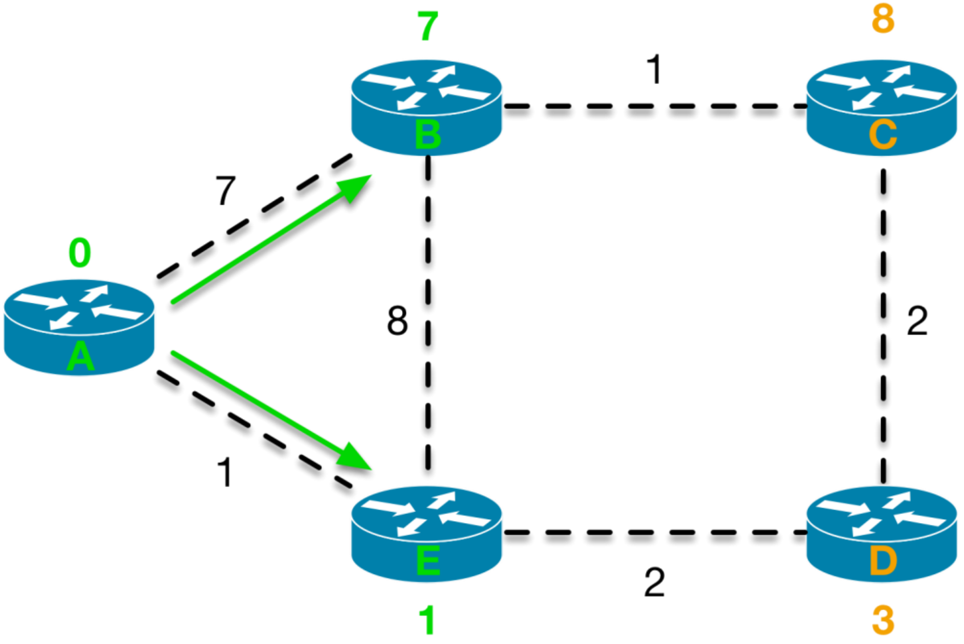

Dijkstra’s Algorithm – An Example

- Cost to Router D can now be confirmed too
- Cost to Router C is updated; a lower cost path was found

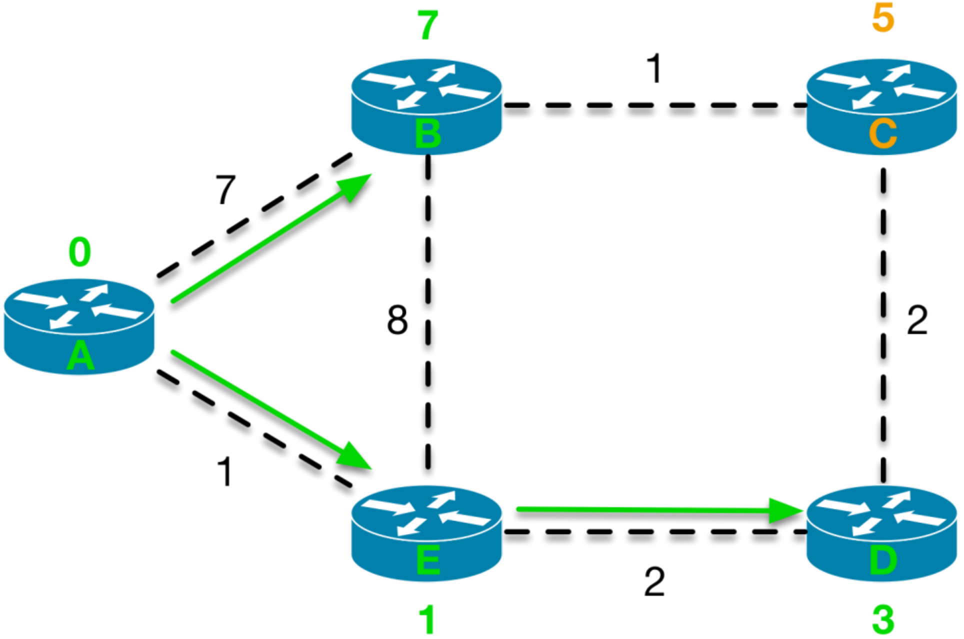

Dijkstra’s Algorithm – An Example

- Cost to Router C can now be confirmed
- Network (at least as far as Router A is concerned) is now converged

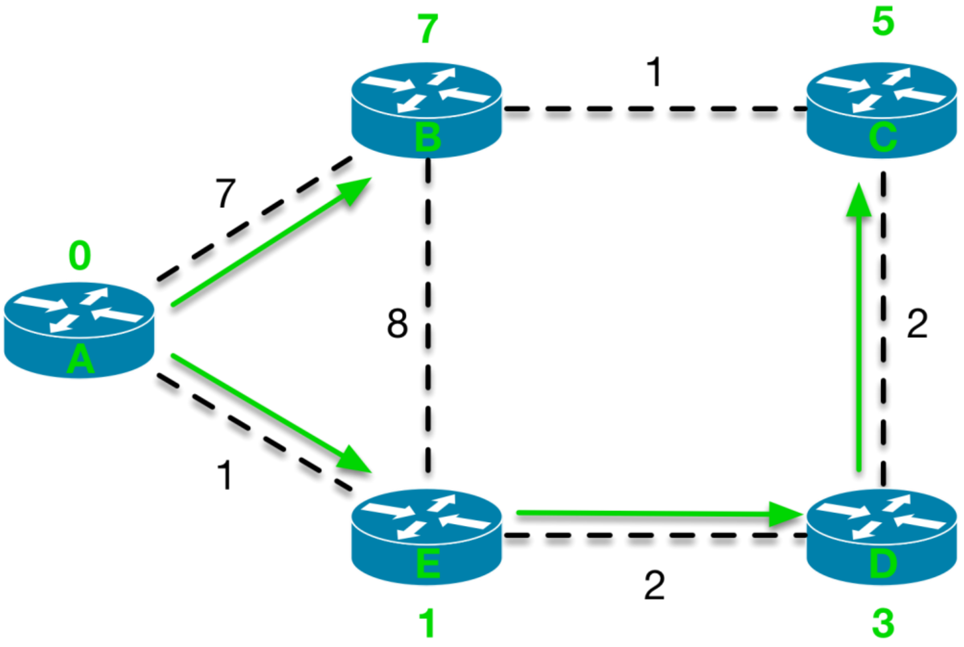

Link State Routing Loops

- Link-State routing should remain loop-free as long as the link state databases are consistent
- Transient (temporary) loops can form after link failure, but should resolve themselves once the network re-converges

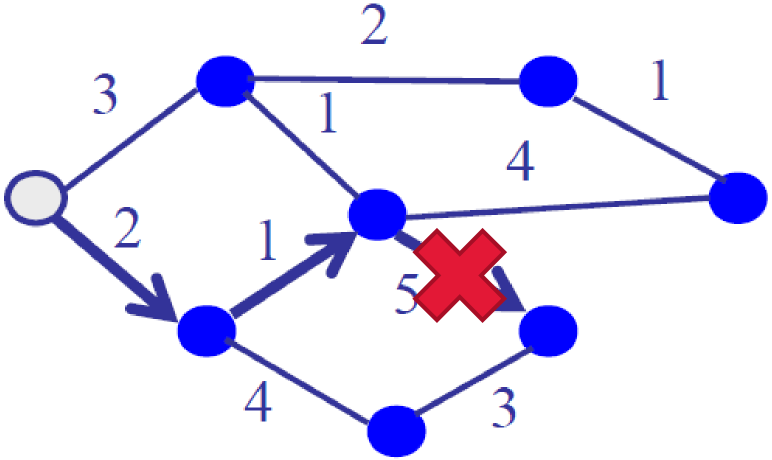

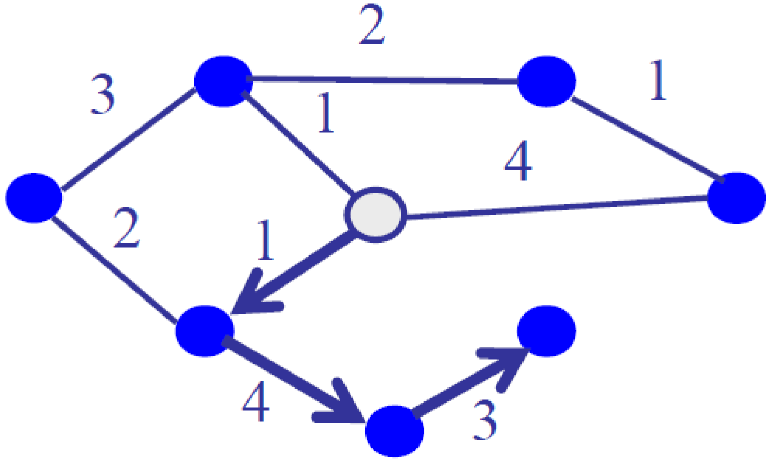

Distance Vector Routing Protocols

- Routing Information Protocol (RIP)
  - Simple routing protocol that uses hop count as metric
  - Periodic updates every 30 seconds, but also includes support for triggered updates
  - Uses split horizon and poison reverse to prevent loops
- Enhanced Interior Gateway Routing Protocol (EIGRP)
  - Developed by Cisco (converted to open standard in 2016)
  - Uses DUAL algorithm
  - Includes enhancements usually associated with link-state protocols (eg. neighbour table, remove periodic updates)

Link-State Routing Protocols

- Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS)
    - Widely used protocols for routing within organisations
    - IS-IS is more commonly used inside ISP networks
    - Support for authentication
    - Can subdivide network into areas for better performance

Comparing Routing Protocols

	Distance Vector	Link-State
Message complexity	Exchange entire routing table between neighbours	Flood link-state advertisements throughout network
Robustness / Accuracy of path cost	Errors by one router can propagate throughout network	Each router computes its own table
Resource requirements	Low (just use advertised routing tables)	High memory and CPU requirements (compute own routing table)
Convergence	Slow (unless triggered updates available	Fast

Comparing the Convergence Process – Distance Vector

- Compare the process of achieving a converged state after a link failure
- Distance vector protocols must wait for periodic updates to occur before a change is propagated
- In the example topology, three periodic updates must take place

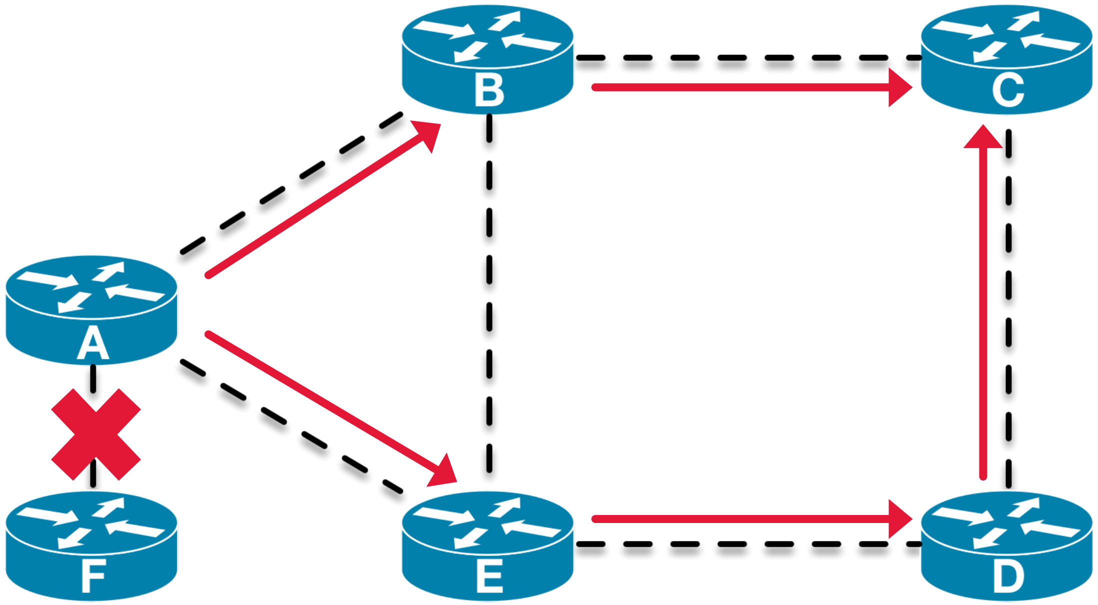

Comparing the Convergence Process – Link-State

- Compare the process of achieving a converged state after a link failure
- Link-state protocols will flood LSAs through the network as soon as Router A notices its link with F is down

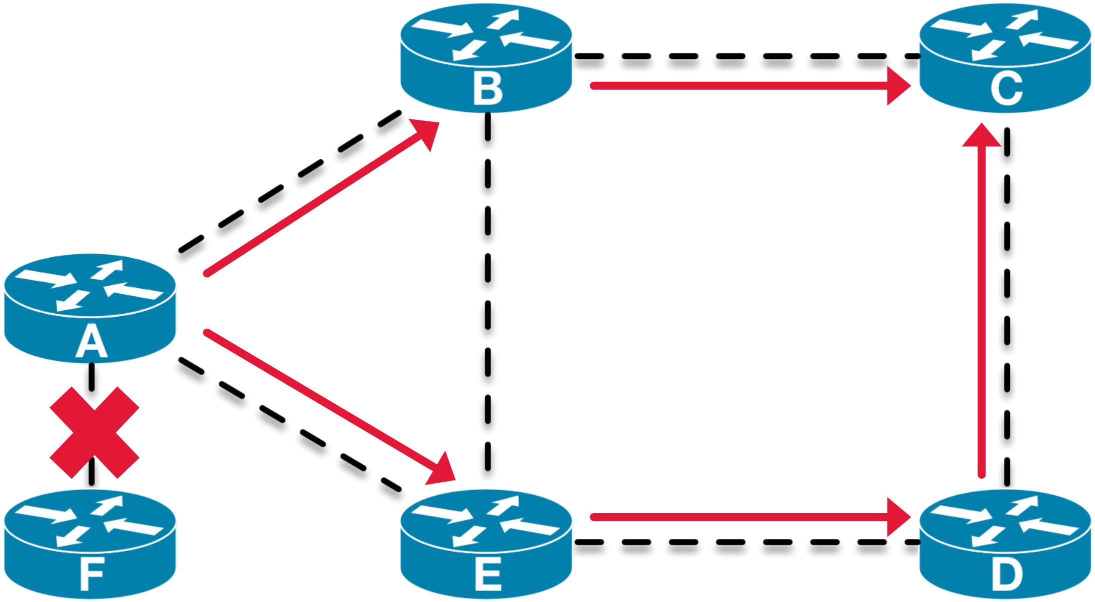

Lecture Objectives

You should now be able to:

- Describe the packet forwarding process
- Describe the purpose of the default gateway
- Describe the process of routing packets between networks
- Describe the role of the routing table in routing packets between networks
- Identify key attributes of the routing table
- Differentiate between static and dynamic routing
- Describe the purpose of default routes
- Describe the characteristics of distance vector routing protocols
- Describe the characteristics of link-state routing protocols
- Differentiate between distance vector and link-state routing protocols

Lecture Summary and the Week Ahead

- We've taken a look at the process of packet forwarding and routing between networks
- Discussed two ways in which routes can be determined
- Identified different categories of dynamic routing protocols (and briefly described their operation)
- Readings for this week are Routing and Switching Essentials Chapters 4, 6, and 7
- In the labs: configuring static routes

Next Week

- Continue our look at interior gateway routing with a closer look at OSPF
- Exterior routing using Border Gateway Protocol
- Structure of the Internet