Packet Propagation and Switching Within a Router

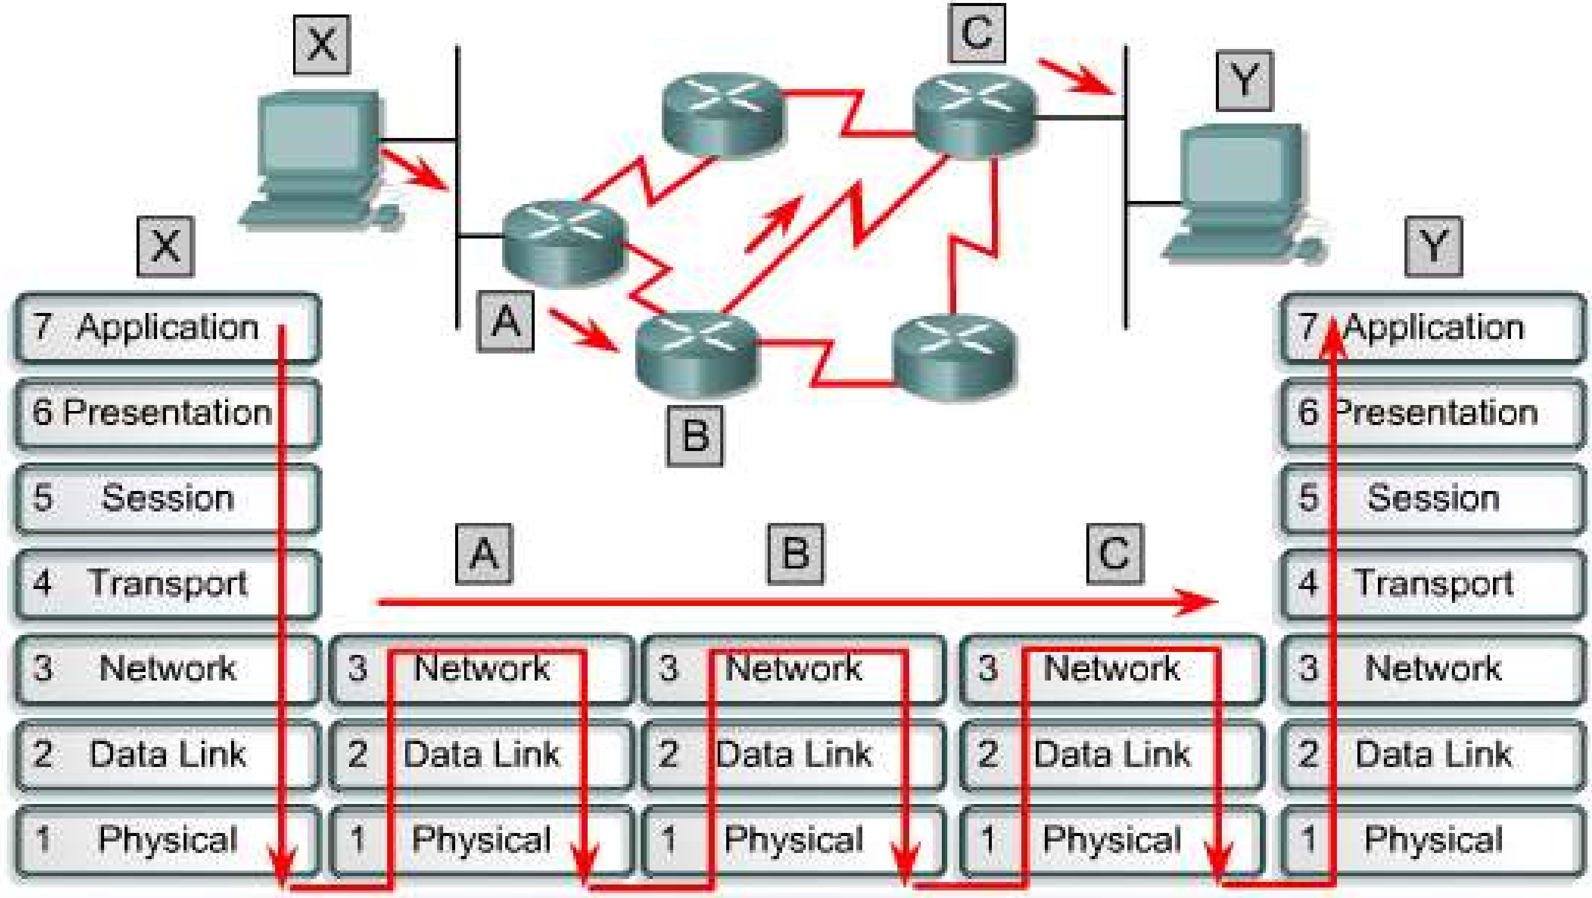

Each router provides its services to support upper-layer functions.

Packet Propagation and Switching Within a Router

- As a frame is received at a router interface.
- The MAC address is checked to see if the frame is directly addressed to the router interface, or a broadcast.
- The frame header and trailer are removed and the packet is passed up to Layer 3.
- The destination IP address is compared to the routing table to find a match.
- The packet is switched to the outgoing interface and given the proper frame header.
- The frame is then transmitted.

Routing Overview

- A router is a network layer device that uses one or more routing metrics to determine the optimal path.
- Routing metrics are values used in determining the advantage of one route over another.
- Routing protocols use various combinations of metrics for determining the best path for data.

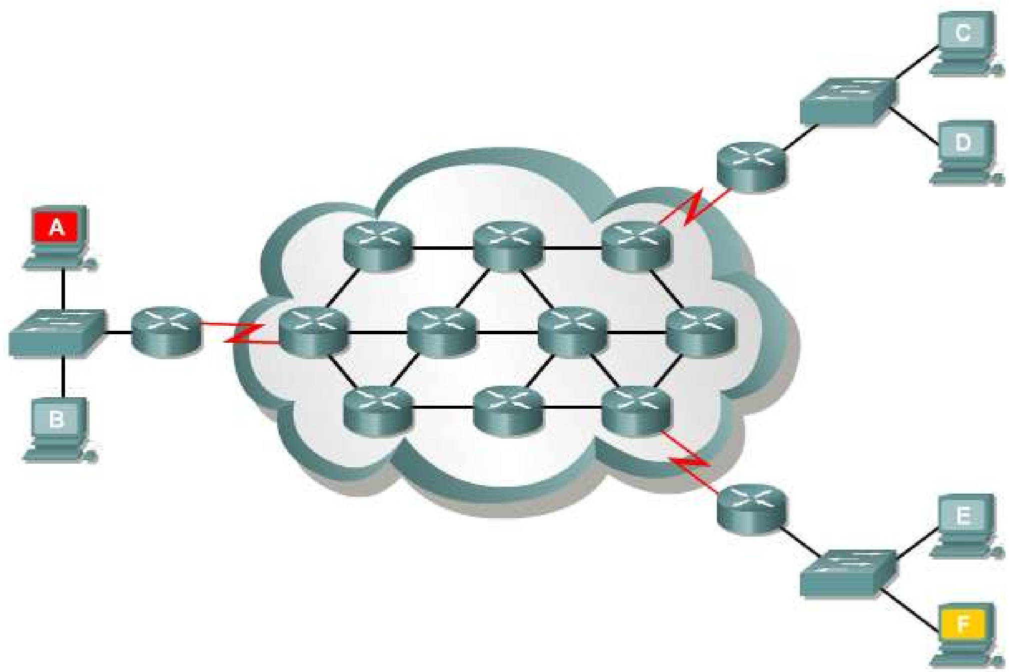

Routing Versus Switching

- This distinction is routing and switching use different information in the process of moving data from source to destination.

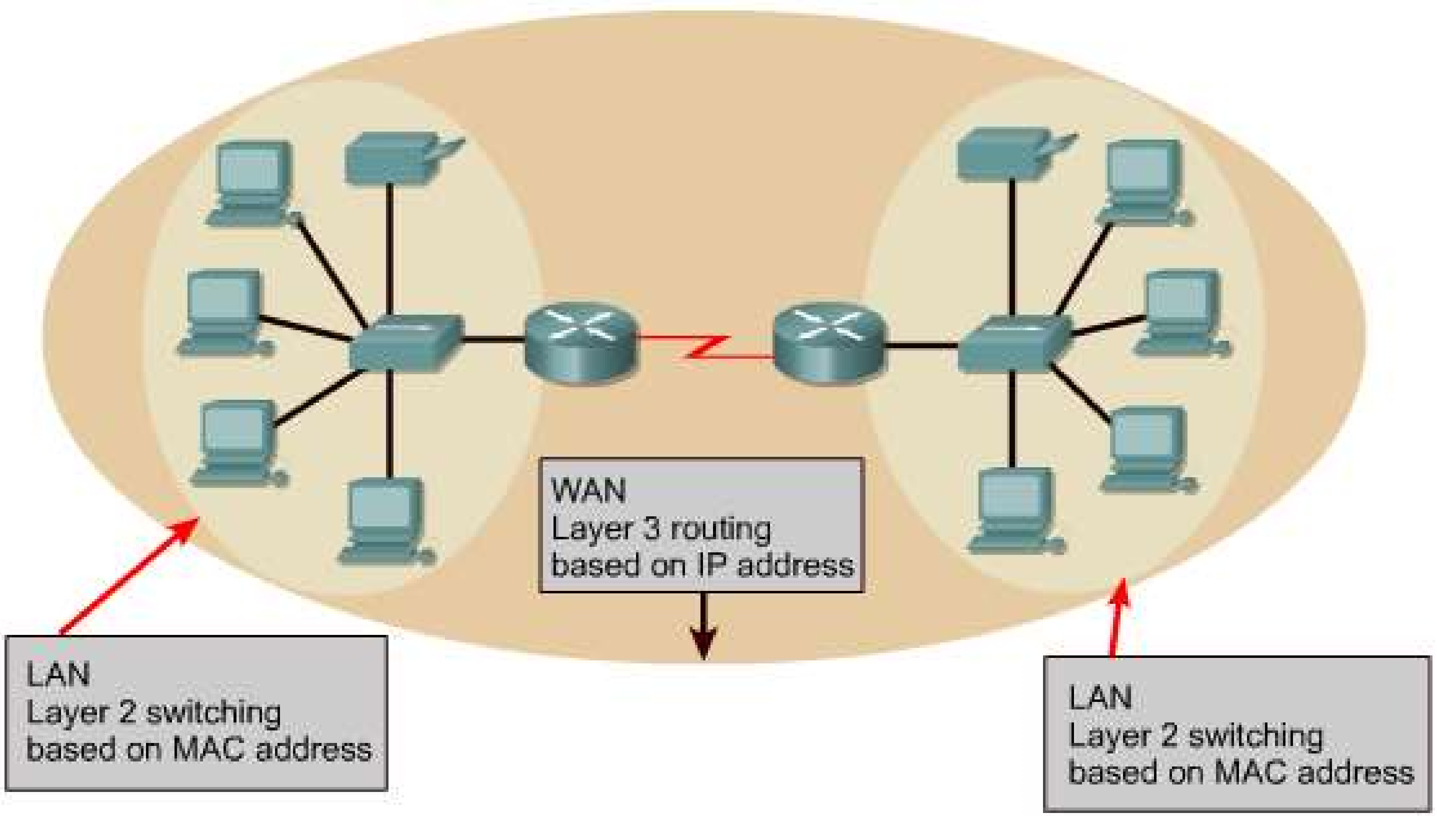

Routing Versus Switching

Features	Router	Switch
Speed	Slower	Faster
OSI Layer	Layer 3	Layer 2
Addressing used	IP	MAC
Broadcasts	Blocks	Forwards
Security	Higher	Lower

Routed Versus Routing

- A routed protocol:
  - Includes any network protocol suite that provides enough information in its network layer address to allow a router to forward it to the next device and ultimately to its destination.
  - Defines the format and use of the fields within a packet.
- A routing protocol:
  - Provides processes for sharing route information.
  - Allows routers to communicate with other routers to update and maintain the routing tables.

Routed Versus Routing Protocol

Routed protocol used between routers to direct user traffic
Examples: IP and IPX

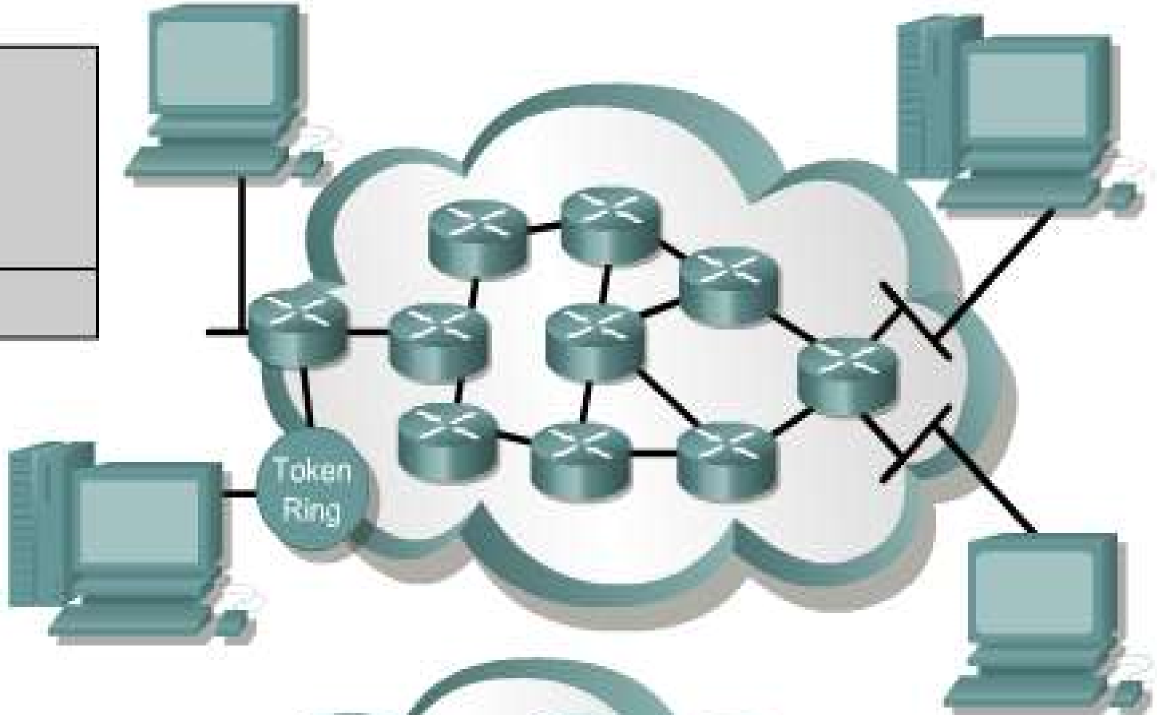

Routing protocol used between routers to maintain tables
Examples: RIP, IGRP, OSPF

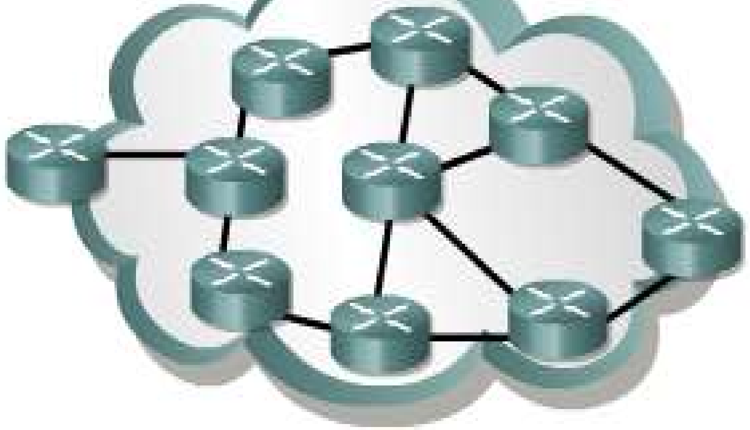

Path Determination

- Path determination enables a router to compare the destination address to the available routes in its routing table, and to select the best path.

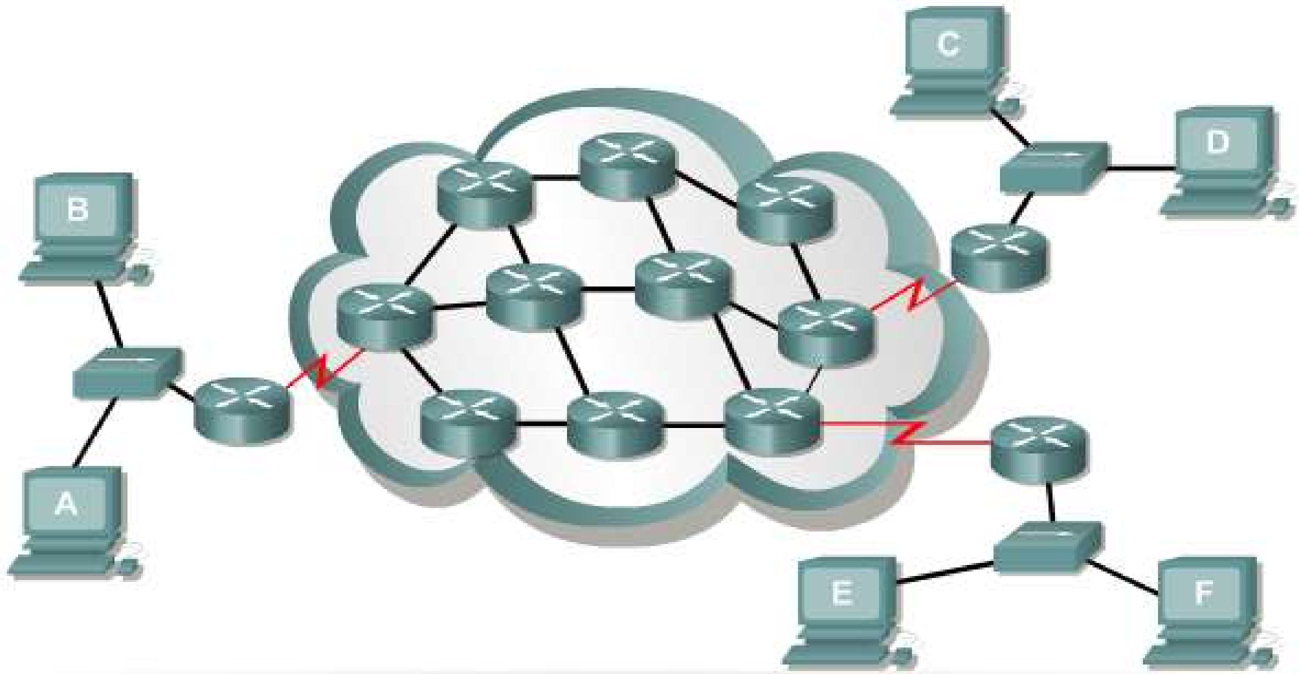

If computer A was sending data to computer F, what path would the data take? That is determined by the information in the routing table.

Routing Tables

- Routers keep track of the following:
  - Protocol type
  - Destination/next-hop associations
  - Routing metric
  - Outbound interfaces

Routing Algorithms and Metrics

- Routing protocols have one or more of the following design goals:
  - Optimization
  - Simplicity and low overhead
  - Robustness and stability
  - Flexibility
  - Rapid convergence

Protocol	Metric	Maximum number of routers	Origins
RIP	Hop count	15	Xerox
IGRP	• Bandwidth • Load • Delay • Reliability	255	Cisco

Routing Metric Components

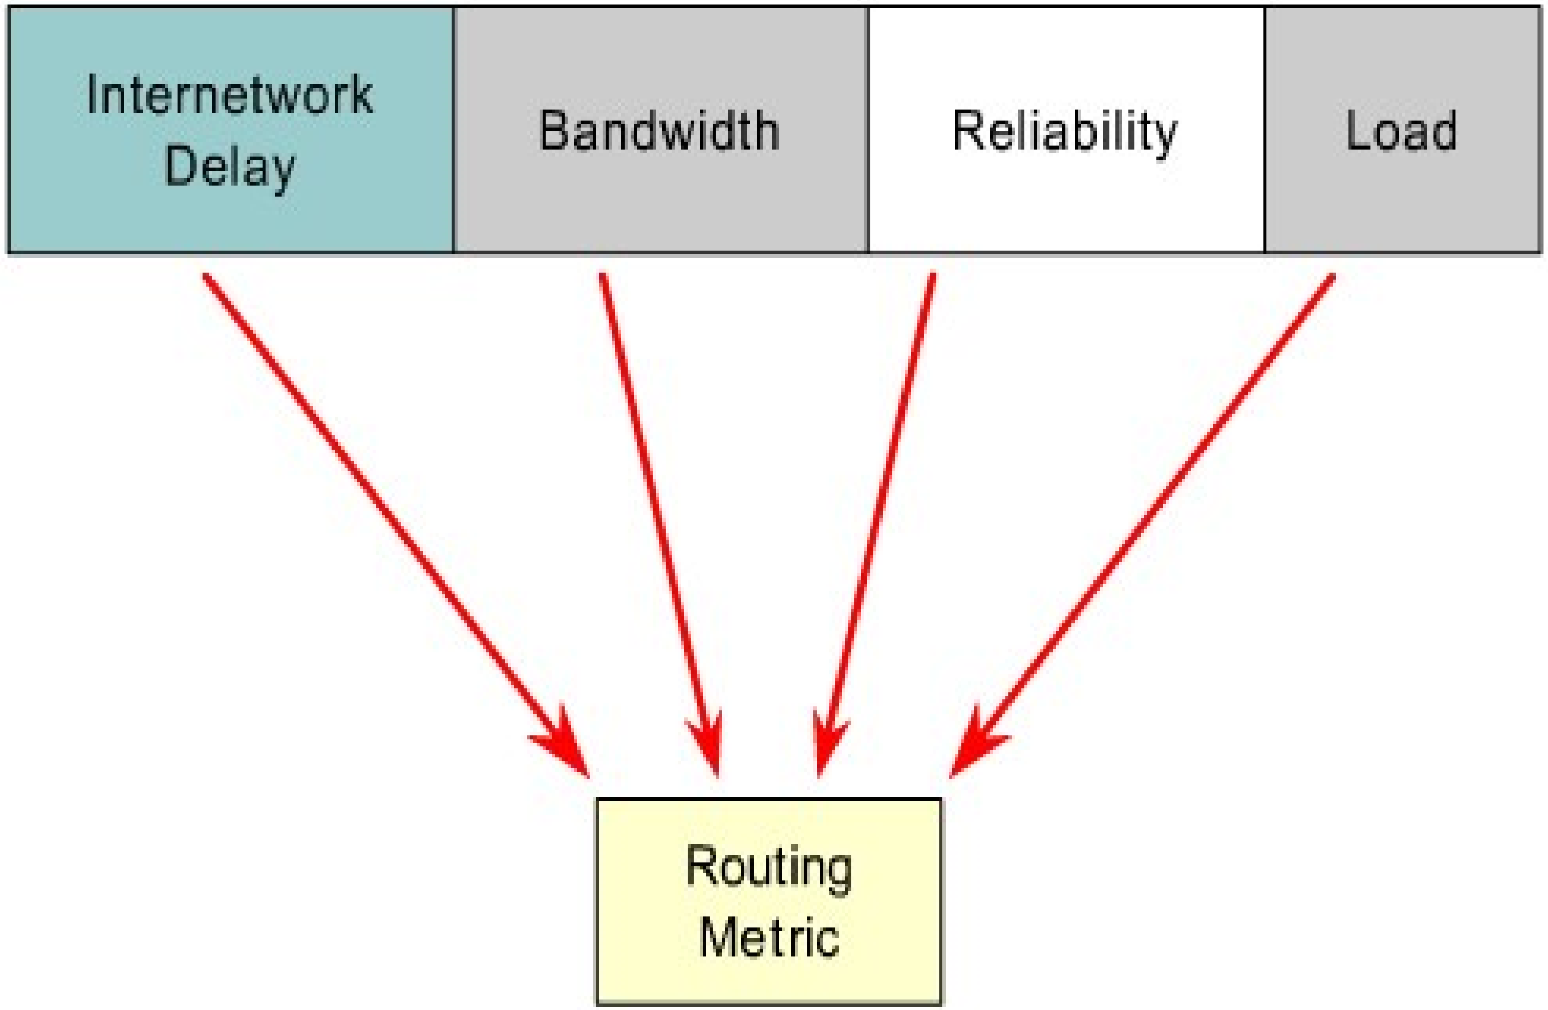

Determining Route Source and Destination

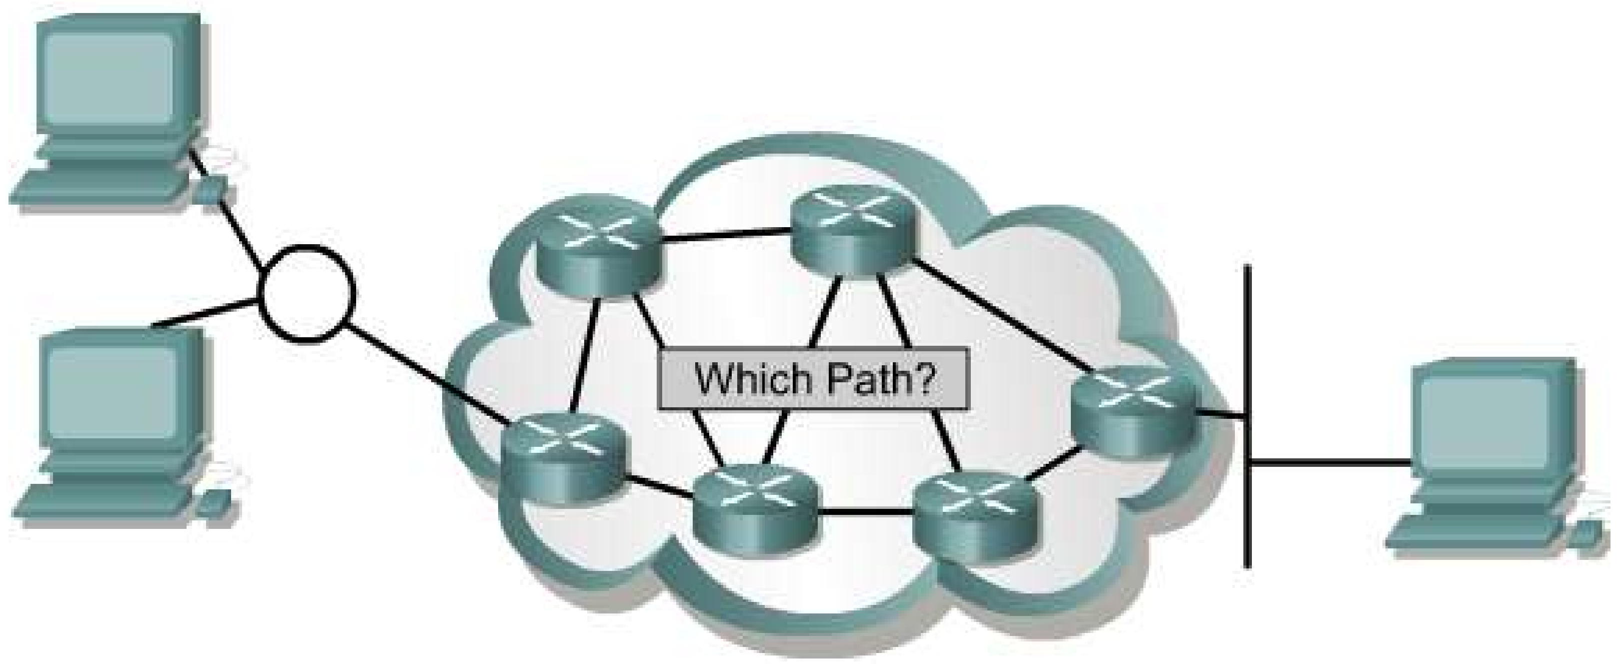

Layer 3 functions to find the best path through the internetwork.

Determining L2 and L3 Addresses

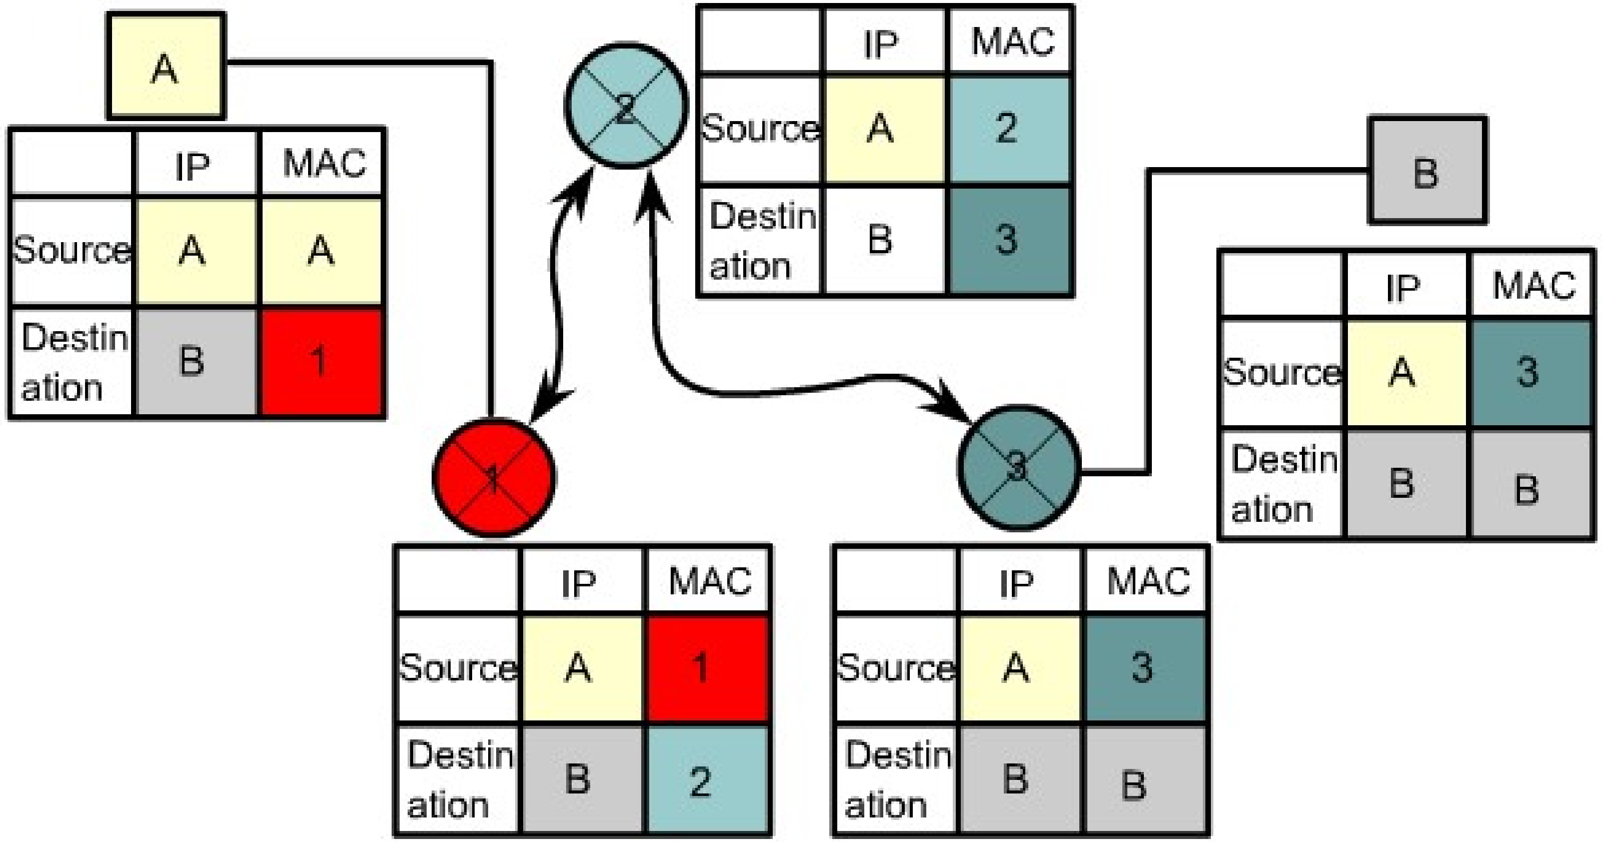

At each interface, as the packet moves across the network, the routing table is examined and the router determines the next hop. The packet is then forwarded using the MAC address of the next hop. The IP source and destination heads do not change, at any time.

Determining the Route Metric

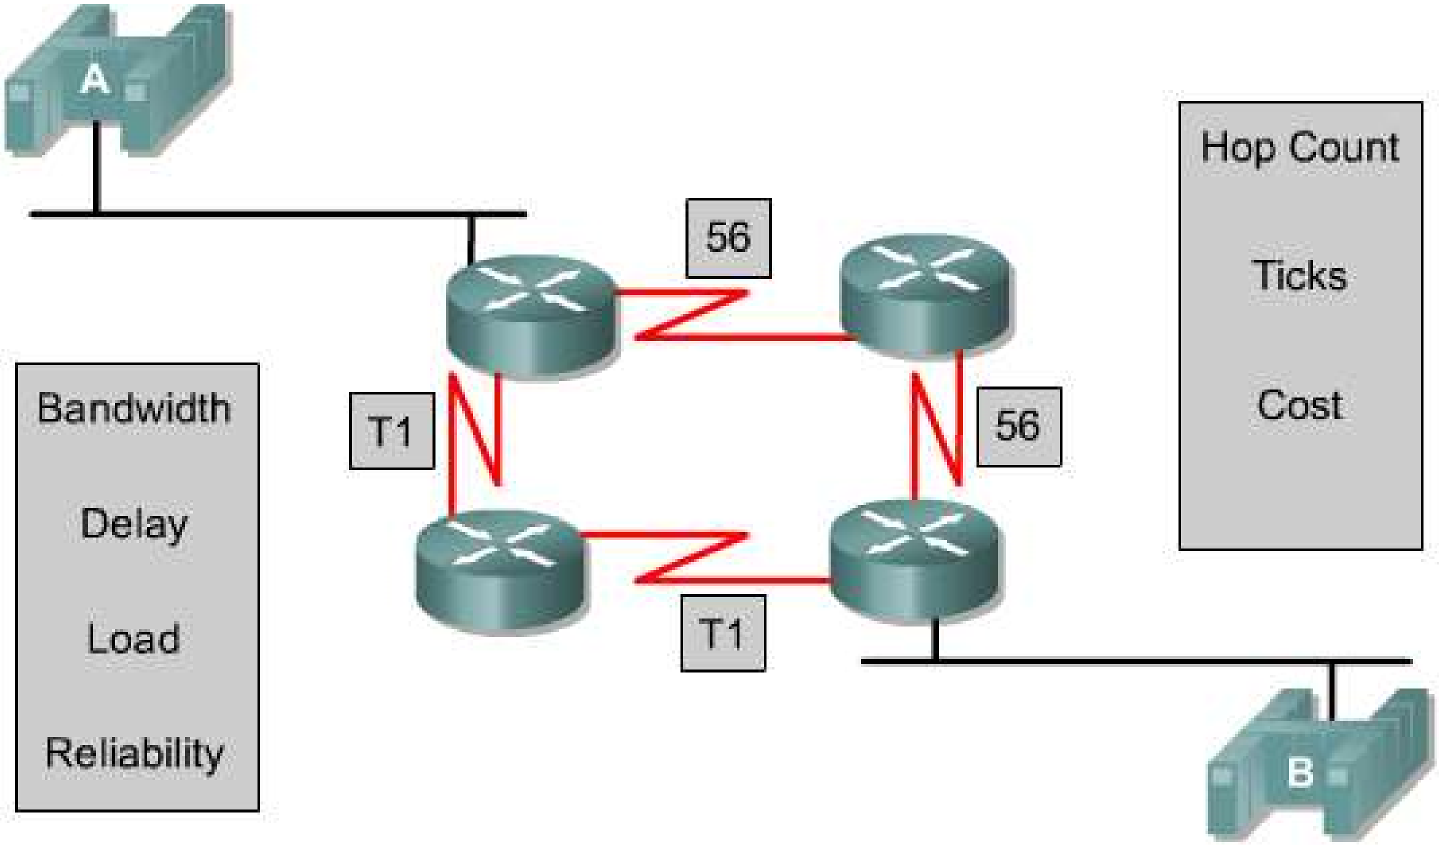

Routing protocols use metrics to determine the best route to a destination.

Determining the Route Next Hop

- Destination/next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router.

Observing Multiple Paths to a Destination

- Some routing protocols support multiple paths to the same destination.
- Unlike single path algorithms, these multi-path algorithms permit traffic over multiple lines, provide better throughput, and are more reliable.

Route Types

Static

Uses a programmed route that a network administrator enters into the router

Dynamic

Uses a route that a routing protocol adjusts automatically for topology or traffic changes

Static Routing

Static Routing Advantages	Static Routing Disadvantages
Low processor overhead. Routers do not spend valuable CPU cycles calculating the best path. This requires less processing power and less memory (and therefore, a less expensive router).	*High maintenance configuration*. Administrators must configure all static routes manually. Complex networks may require constant reconfiguration.
No bandwidth utilization. Routers do not take up bandwidth updating each other about static routes.	*No adaptability*. Statically configured routes can not adapt to changes in link status.
Secure operation. Routers that do not send updates will not inadvertently advertise network information to an untrusted source. Routers that do not accept routing updates are less vulnerable to attack.	
Predictability. Static routes enable an administrator to precisely control a router's path selection. Dynamic routing sometimes yields unexpected results, even in small networks.	

Configuring Static Routes

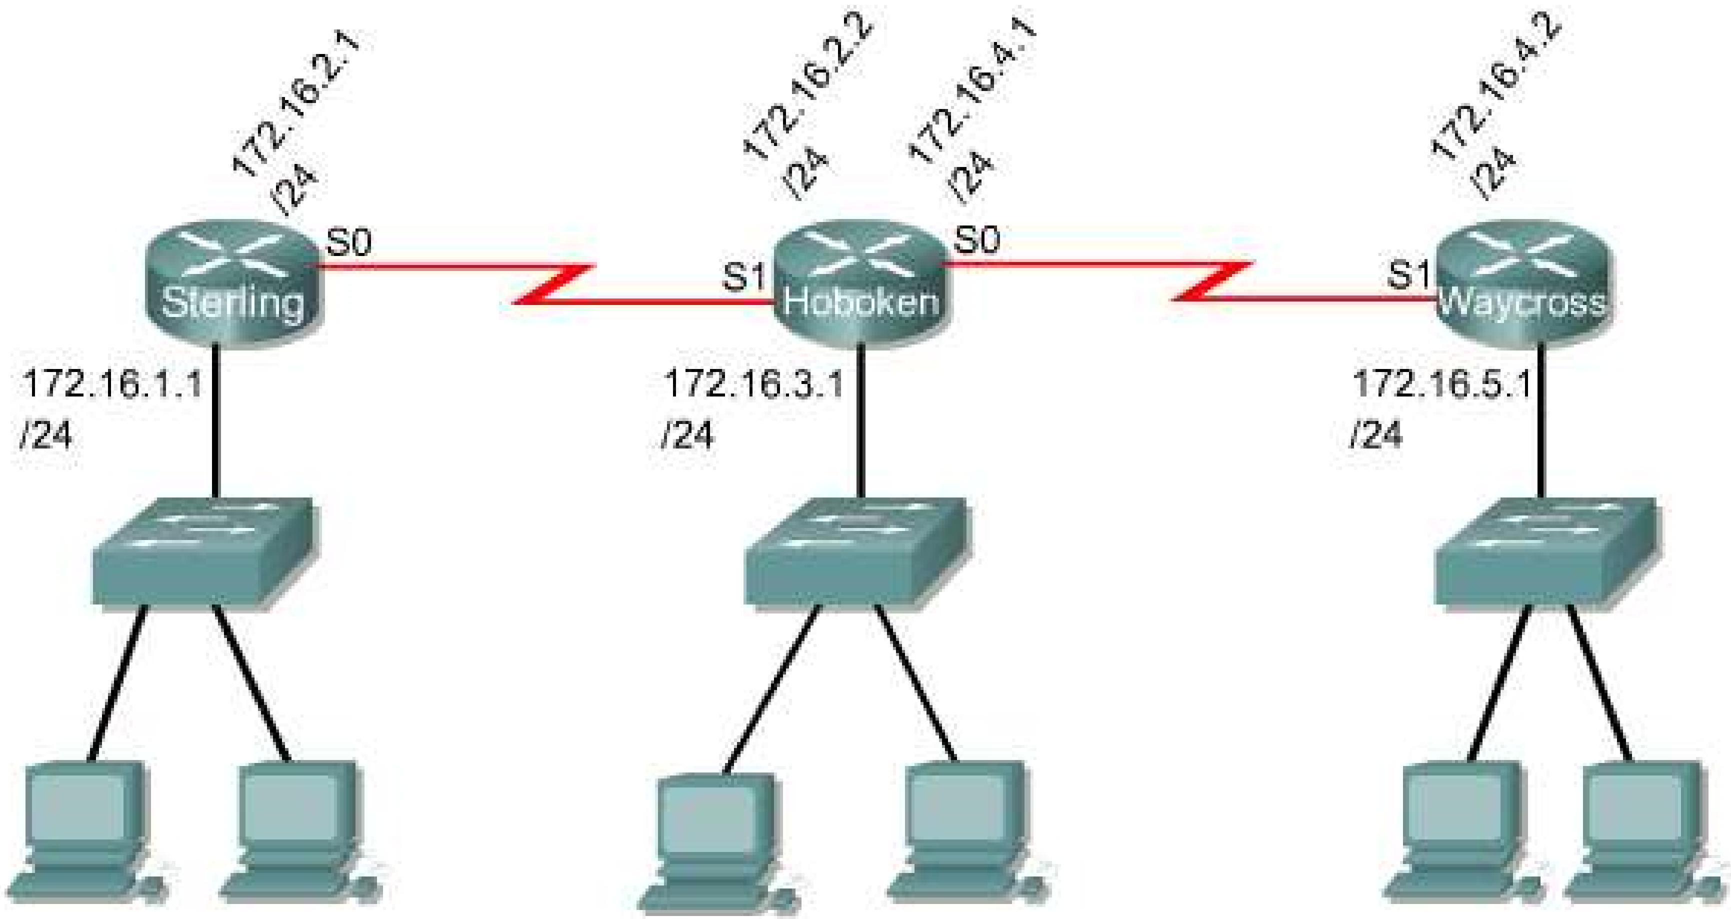

Non-directly Connected Networks

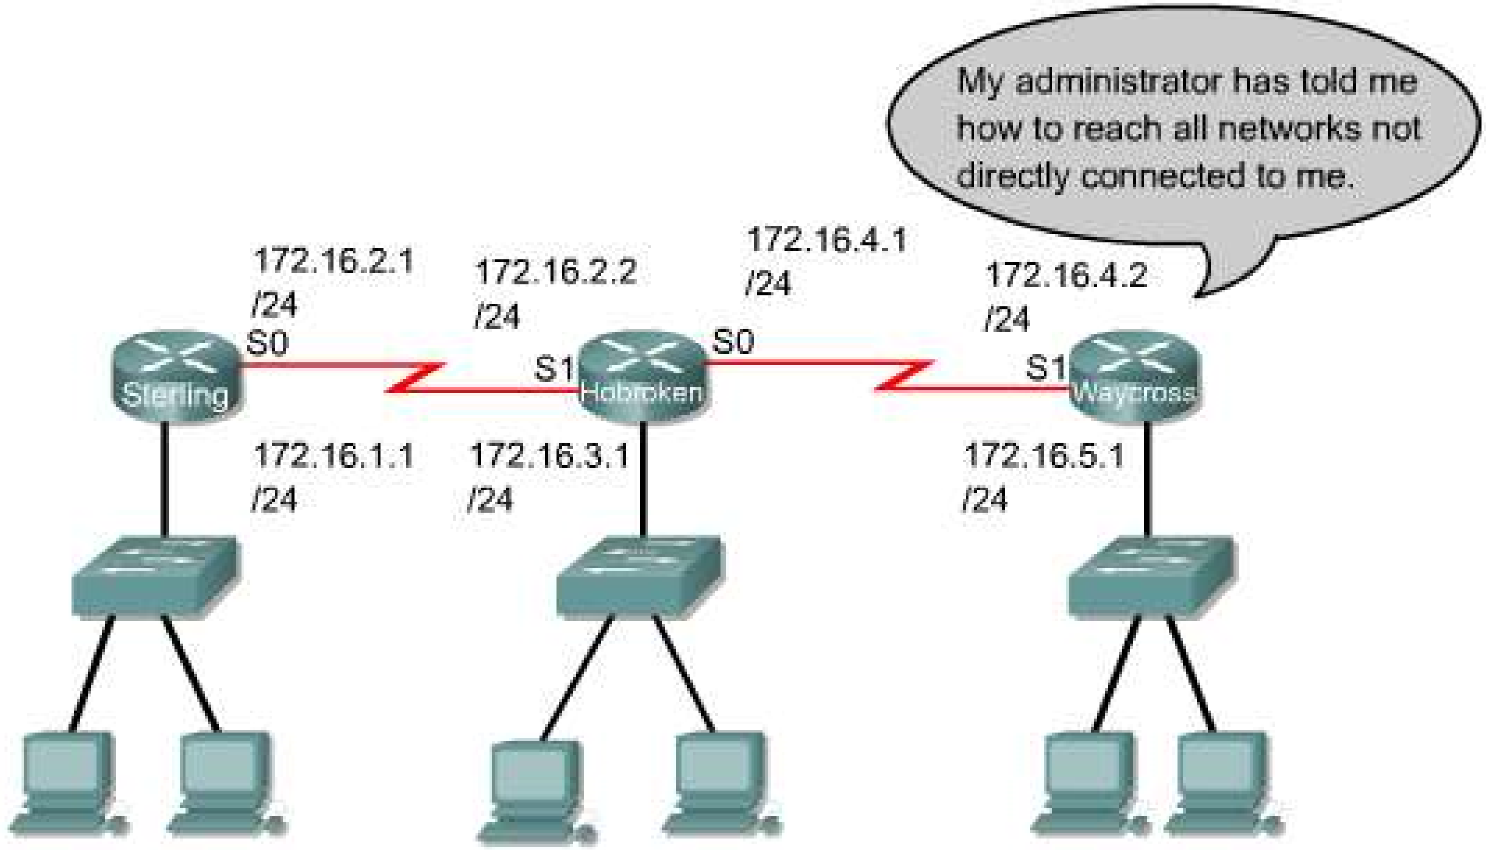


```
Waycross(config)#ip route 0.0.0.0  0.0.0.0 S1
```

This command points to all non-directly-connected networks

Dynamic Routing

Dynamic Routing Advantages	Dynamic Routing Disadvantages
High degree of adaptability. Routers can alert each other about links that are down or about newly discovered path. Routers automatically "learn" a network's topology and select optimum paths.	*Increased processor overhead and memory utilization*. Dynamic routing processes can require a significant amount of CPU time and system memory.
Low maintenance configuration. After the basic parameters for a routing protocol are set correctly, administrative intervention is not required.	*High bandwidth utilization*. Routers use bandwidth to send and receive routing updates, which can detrimentally affect perfomance on slow WAN links.

Dynamic Routing Operations

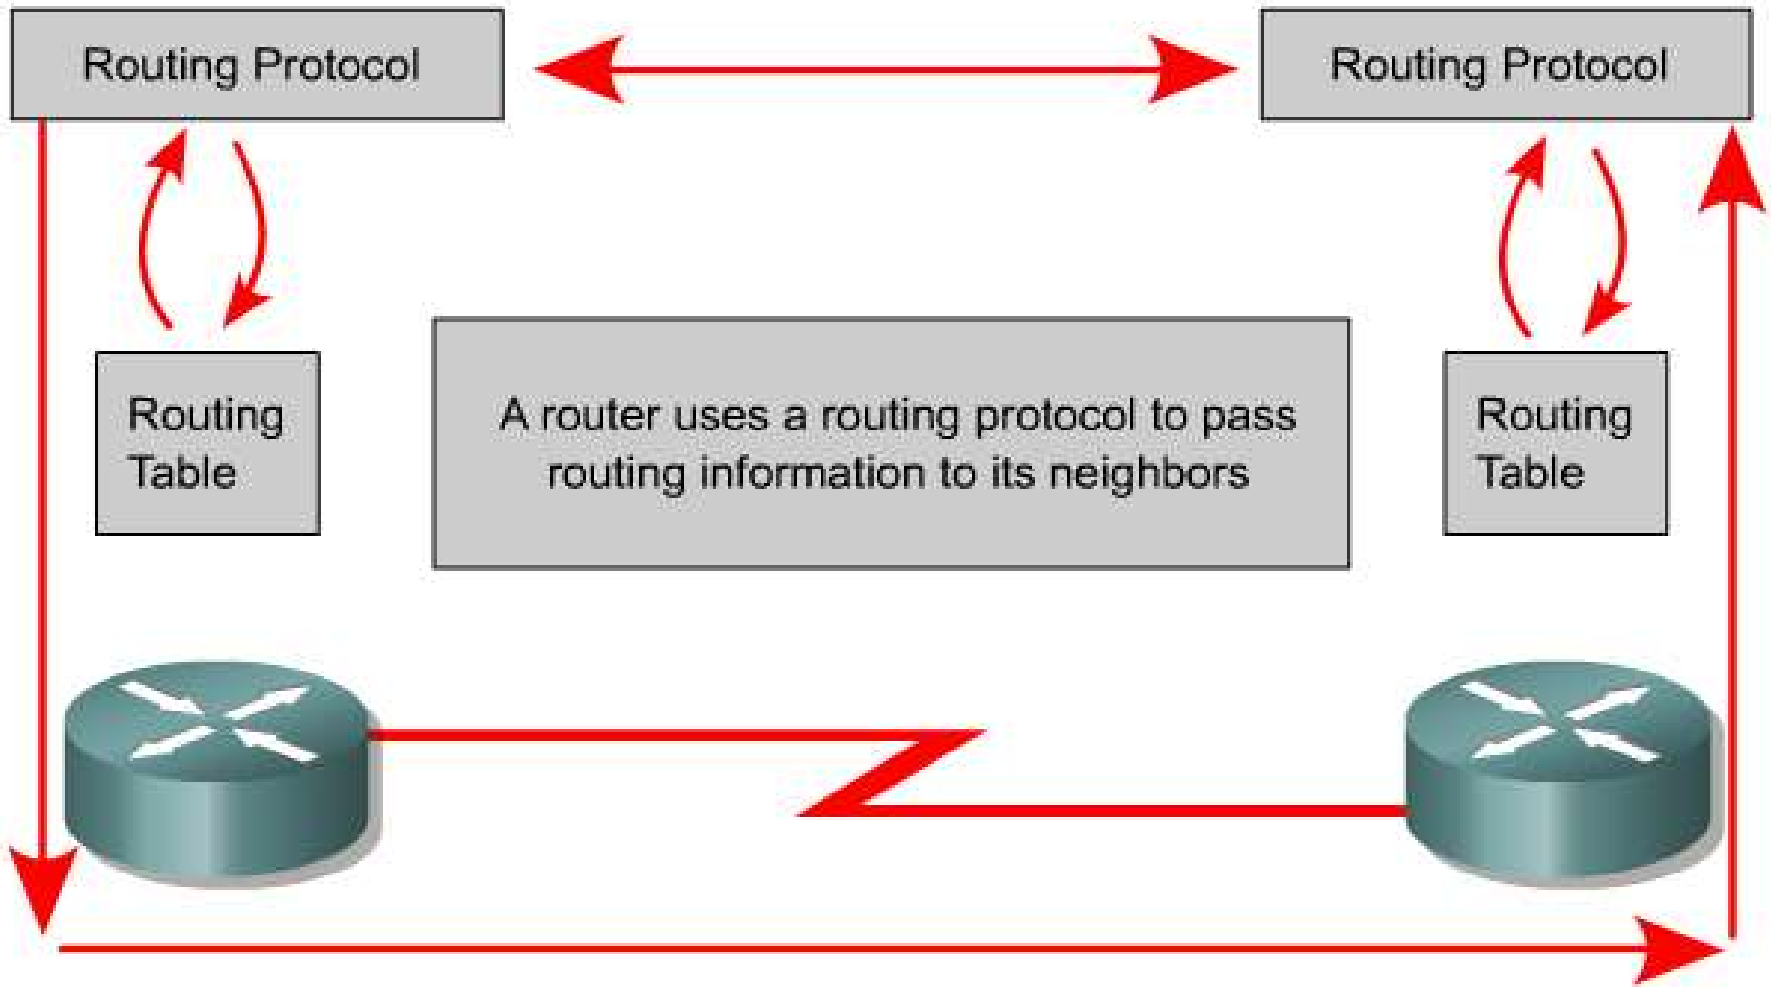

Classes of Routing Protocols

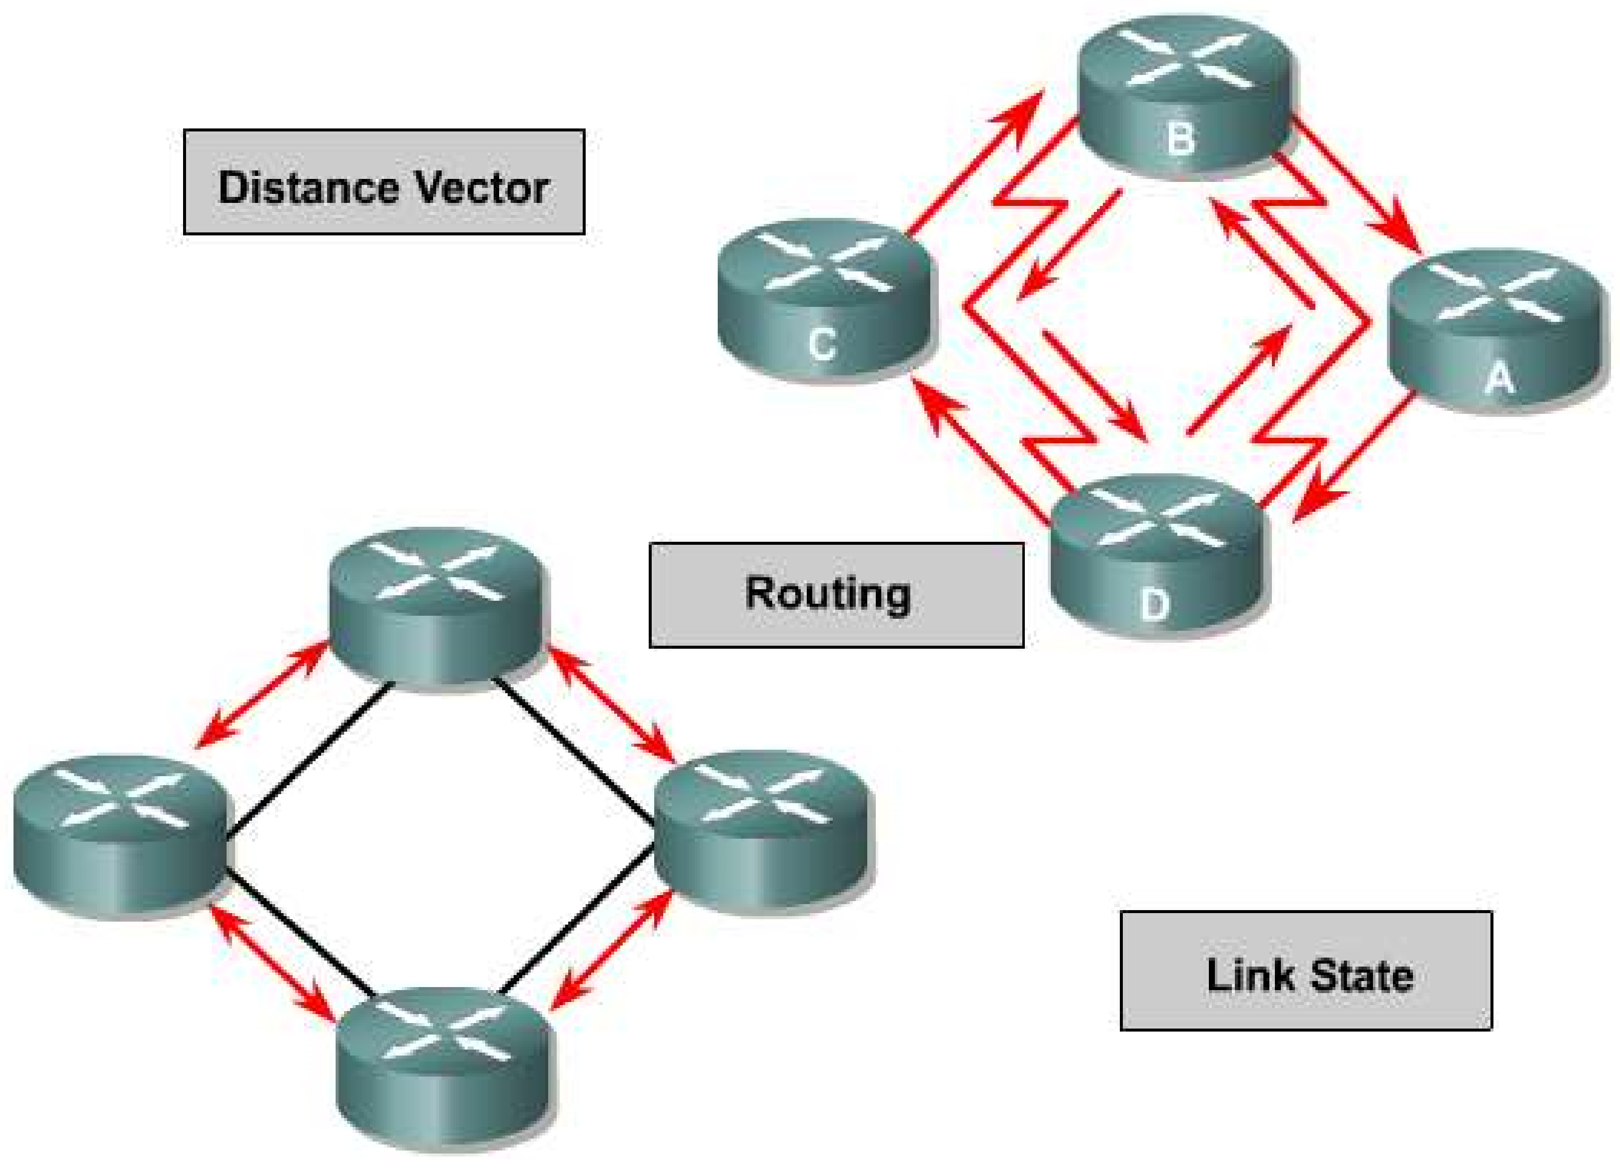

Link State and Distance Vector

- Examples of distance-vector protocols:
  - Routing Information Protocol (RIP)
  - Interior Gateway Routing Protocol (IGRP)
- Examples of link-state protocols:
  - Open Shortest Path First (OSPF)
  - Intermediate System-to-Intermediate System (IS-IS)

Distance Vector Concepts

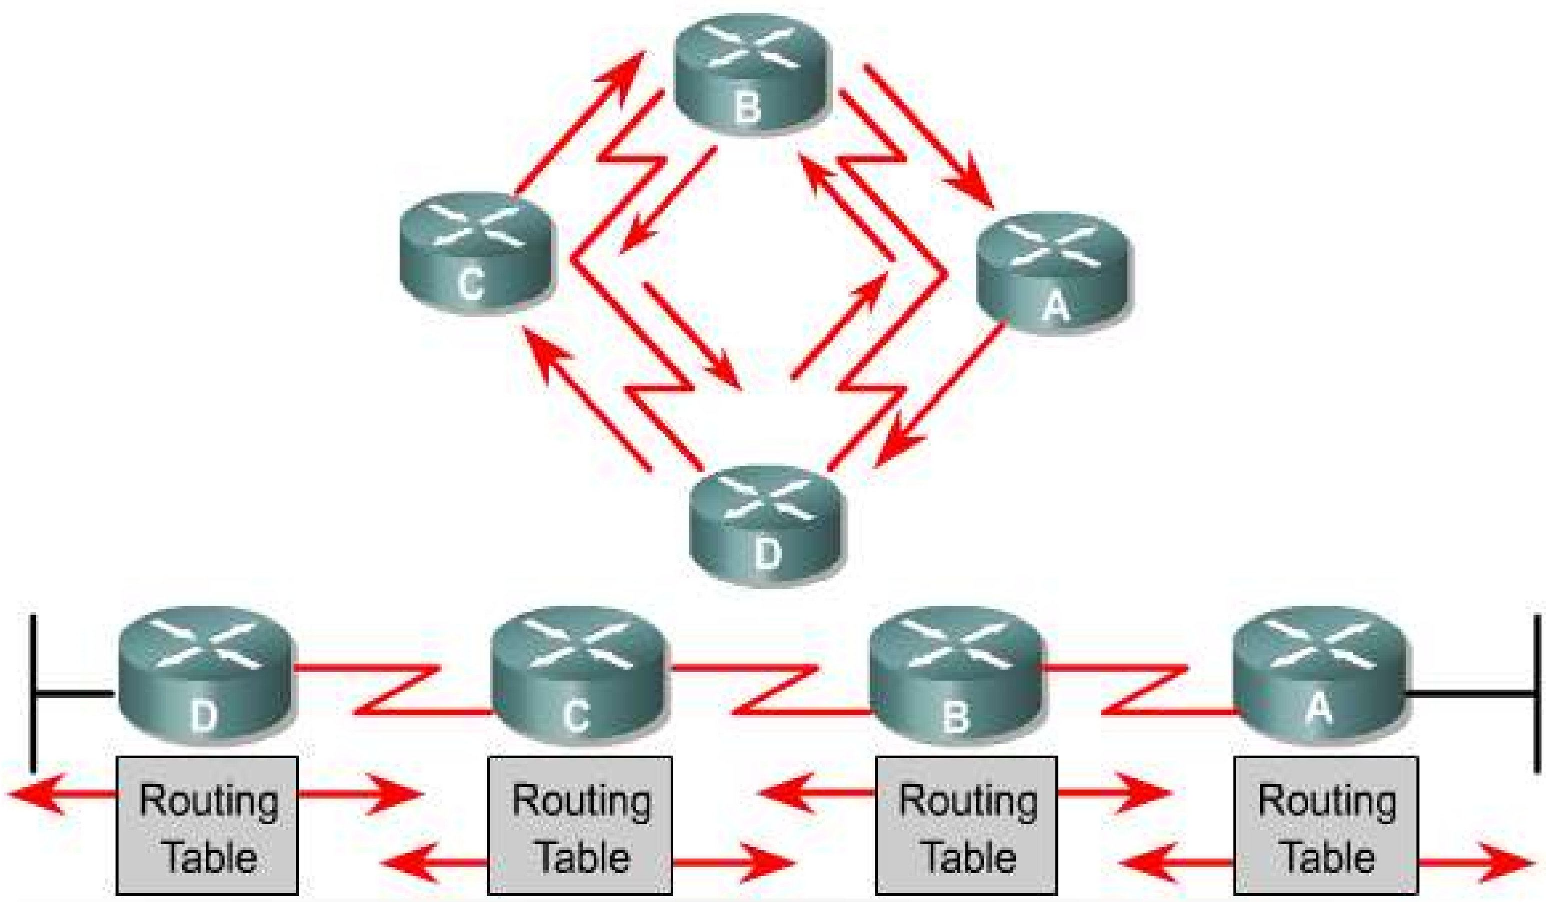

Pass periodic copies of a routing table to neighbor routers and accumulate distance vectors.

Distance Vector Network Discovery

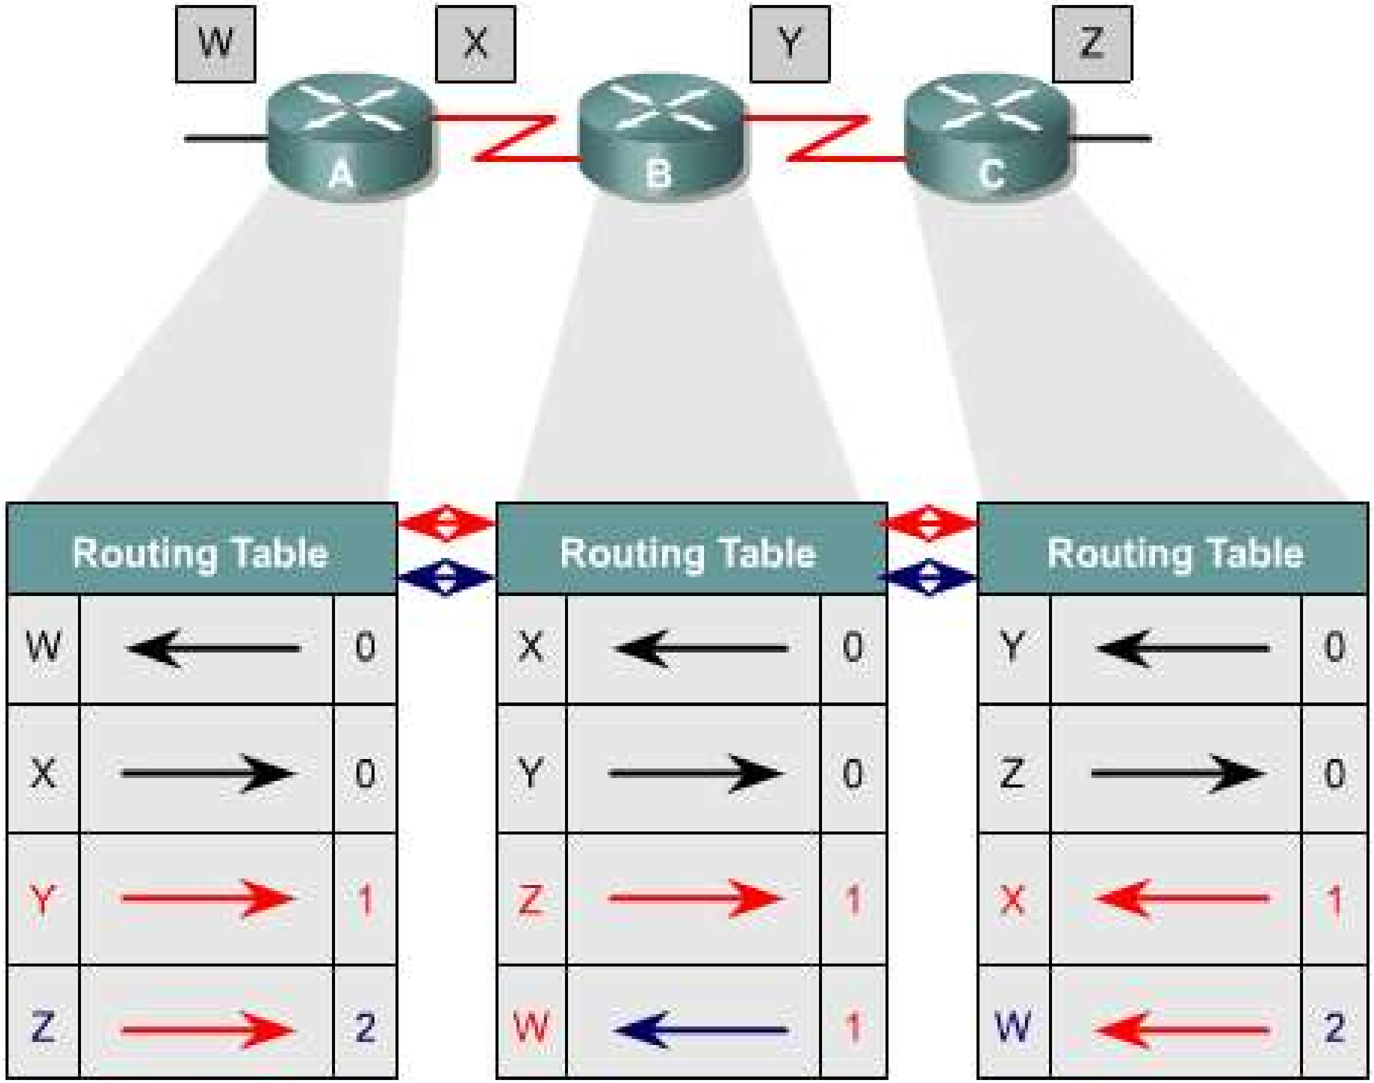

Link-State Concepts

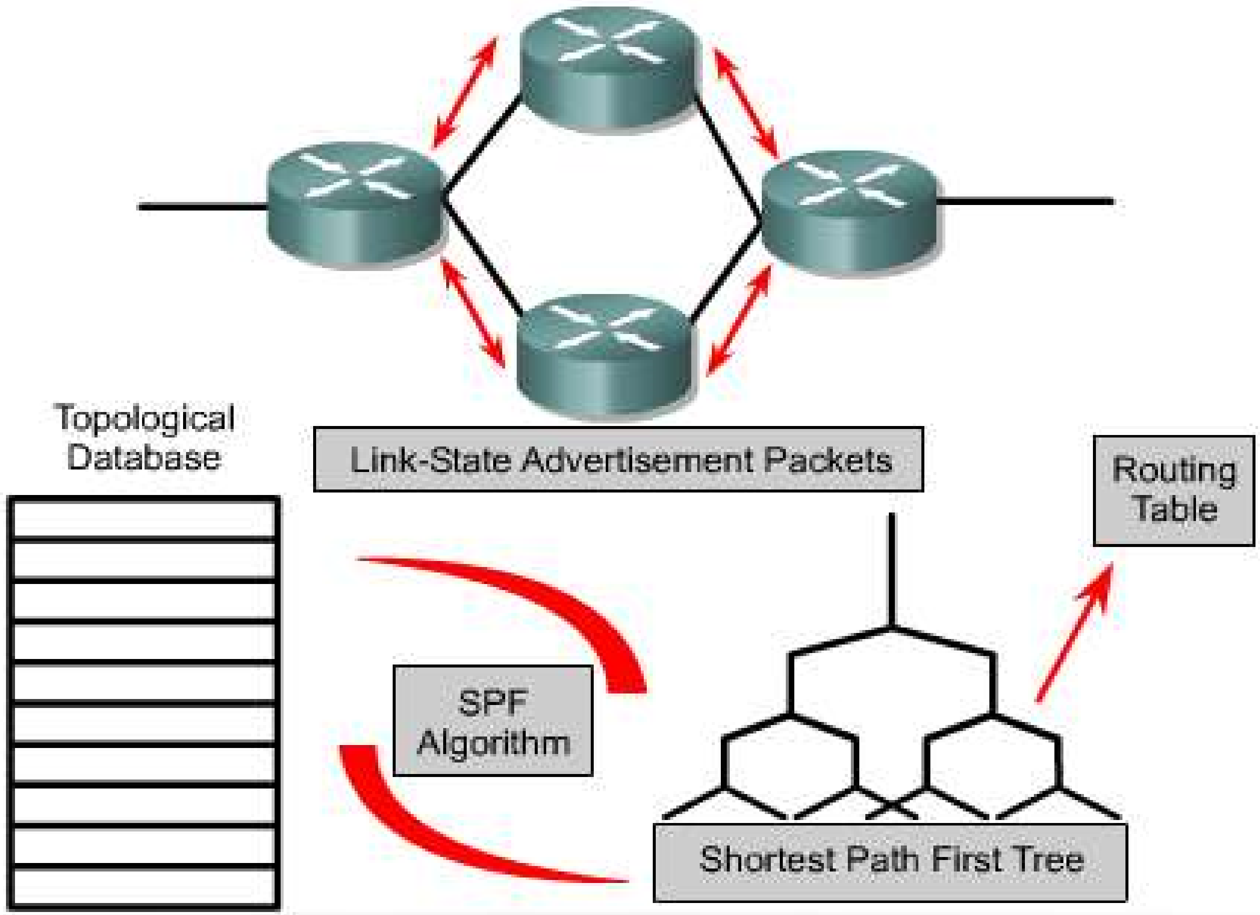

Routers send LSAs to their neighbors. The LSAs are used to build a topological database. The SPF algorithm is used to calculate the shortest path first tree in which the root is the individual router and then a routing table is created.

Link-State Network Discovery

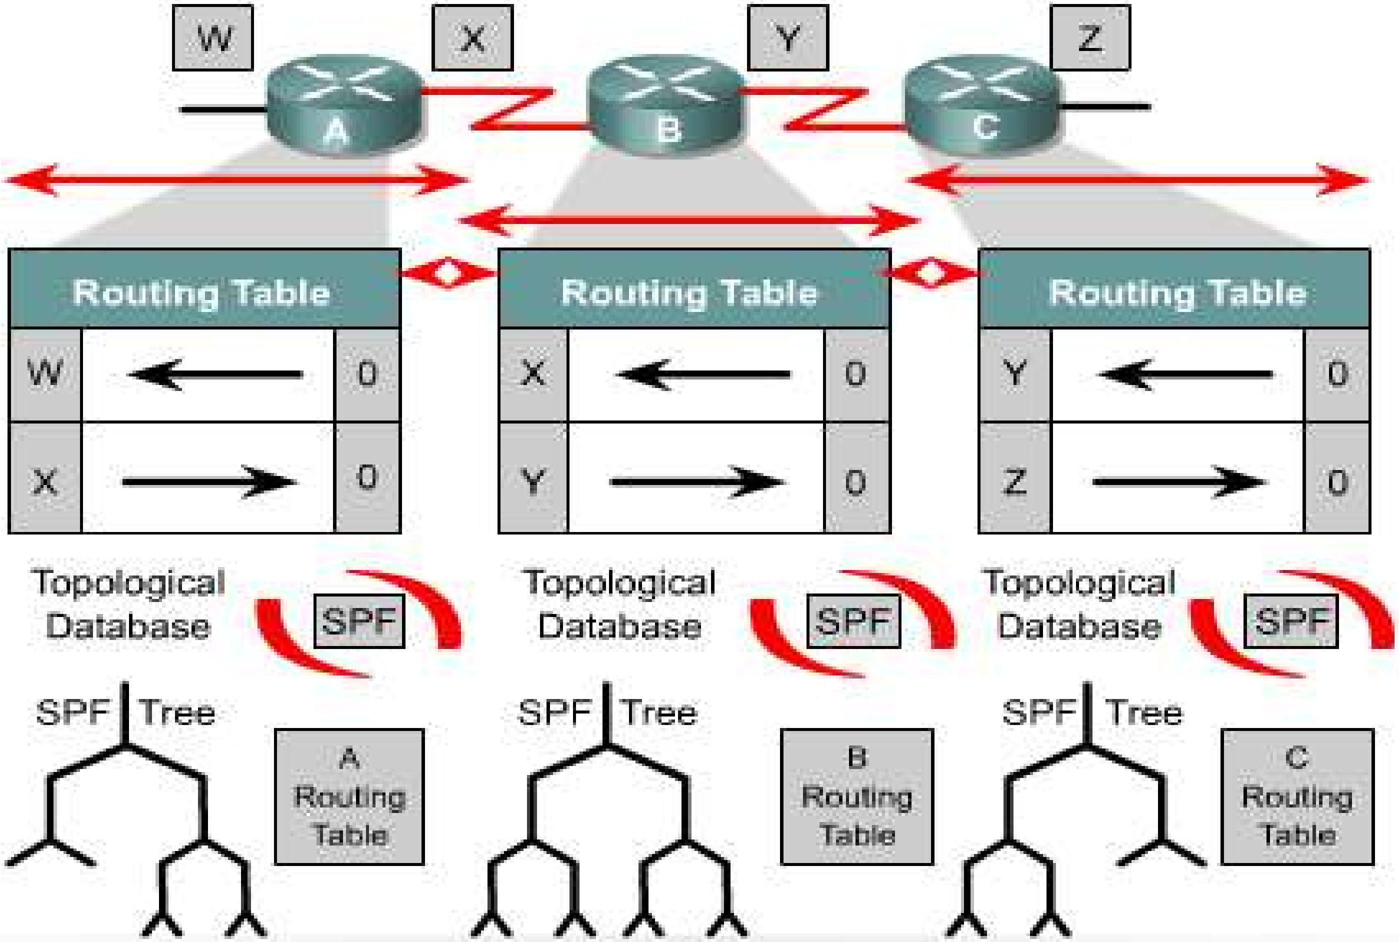

Each router has its own topological database on which the SPF algorithm is run.

Link-State Topology Changes

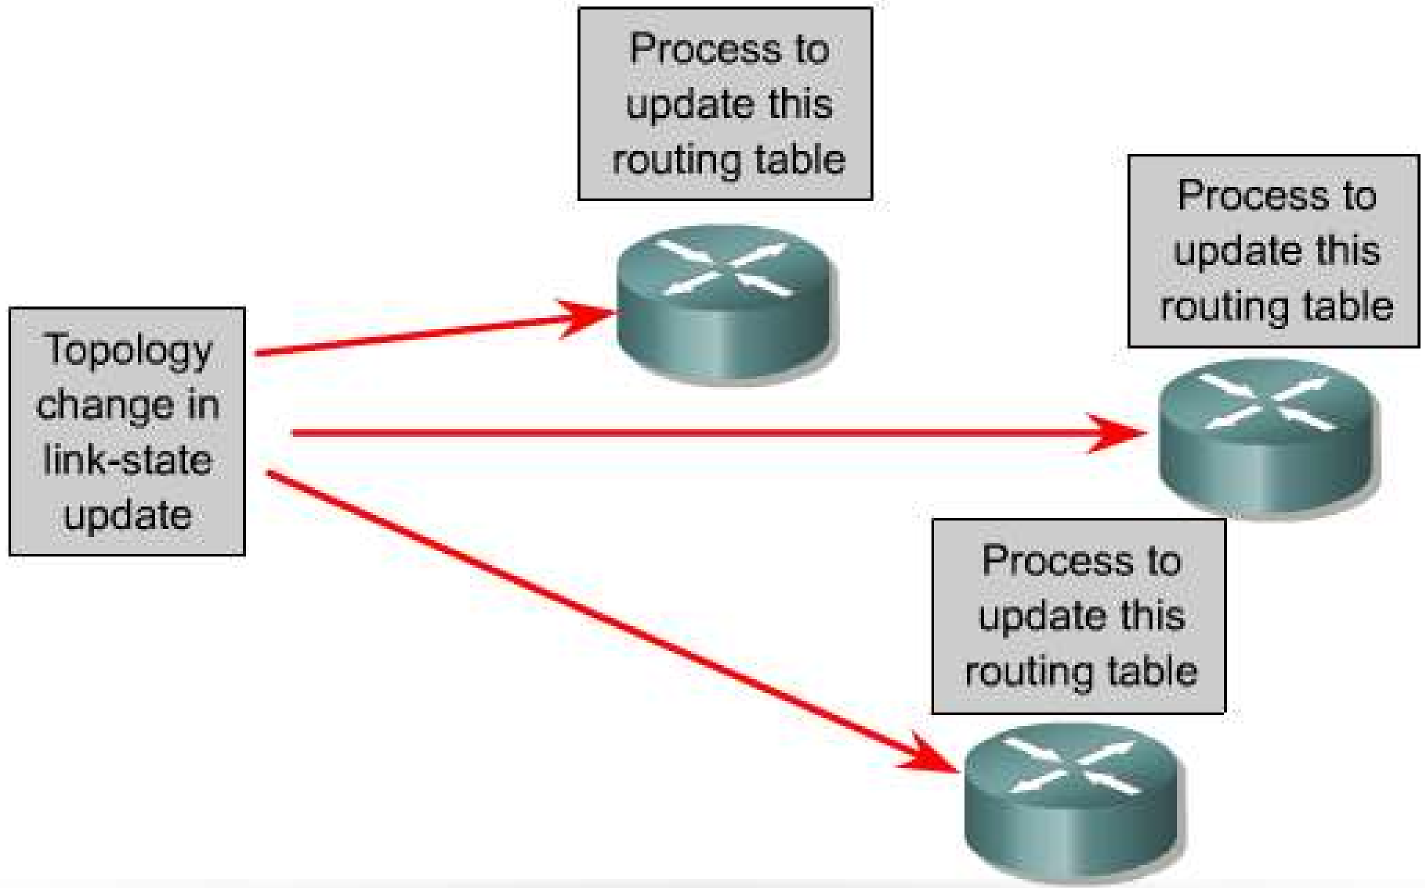

Each router has its own topological database on which the SPF algorithm is run.

Link-State Concerns

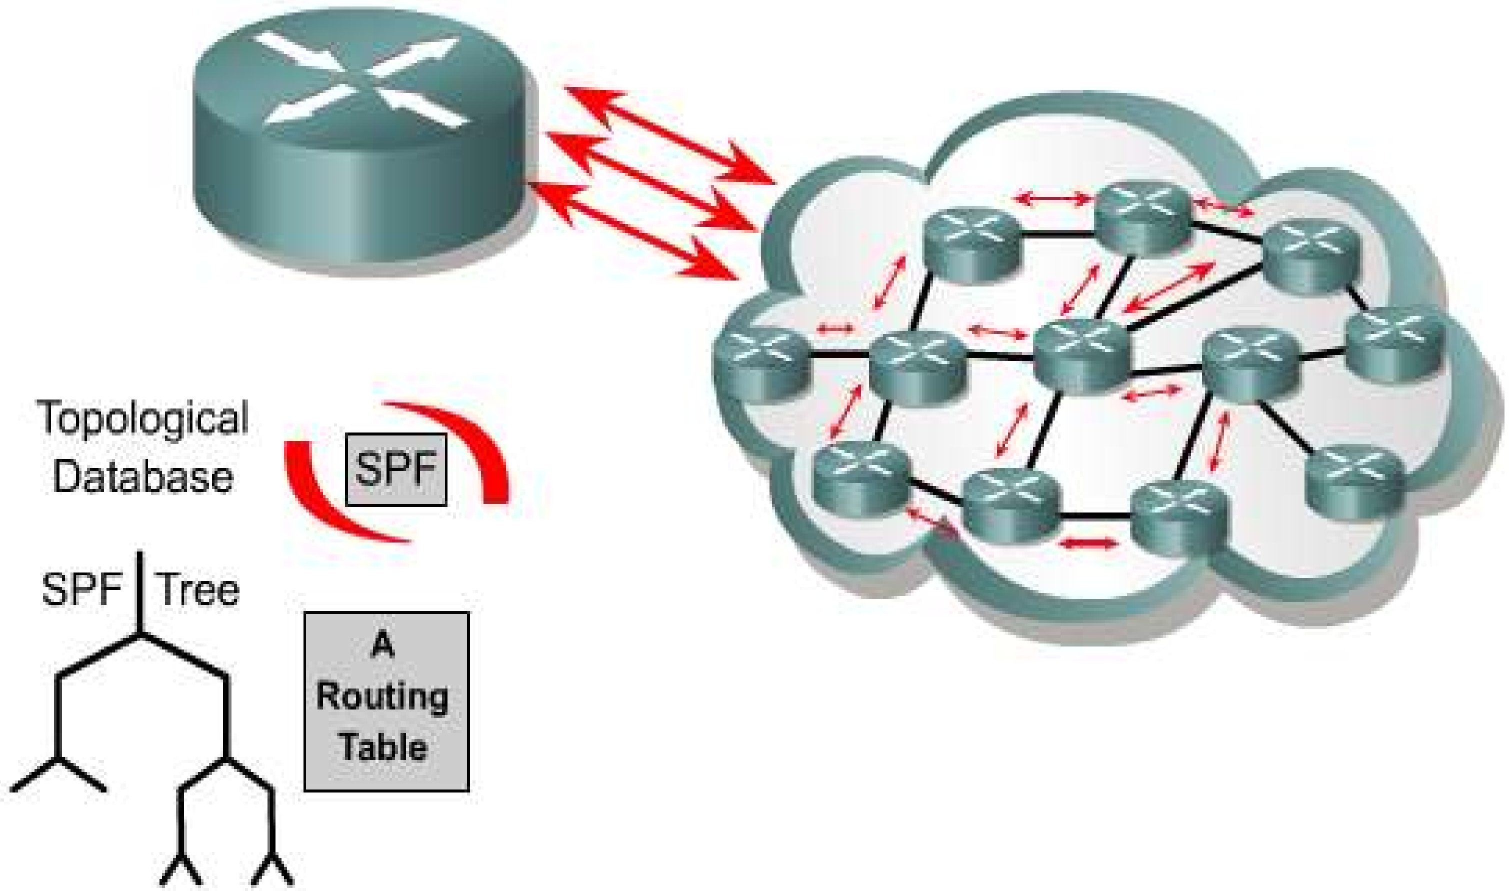

- Processing and memory requirements are increased for link-state routing.
- Bandwidth is consumed during the initial link-state flooding of LSAs.

Autonomous Systems

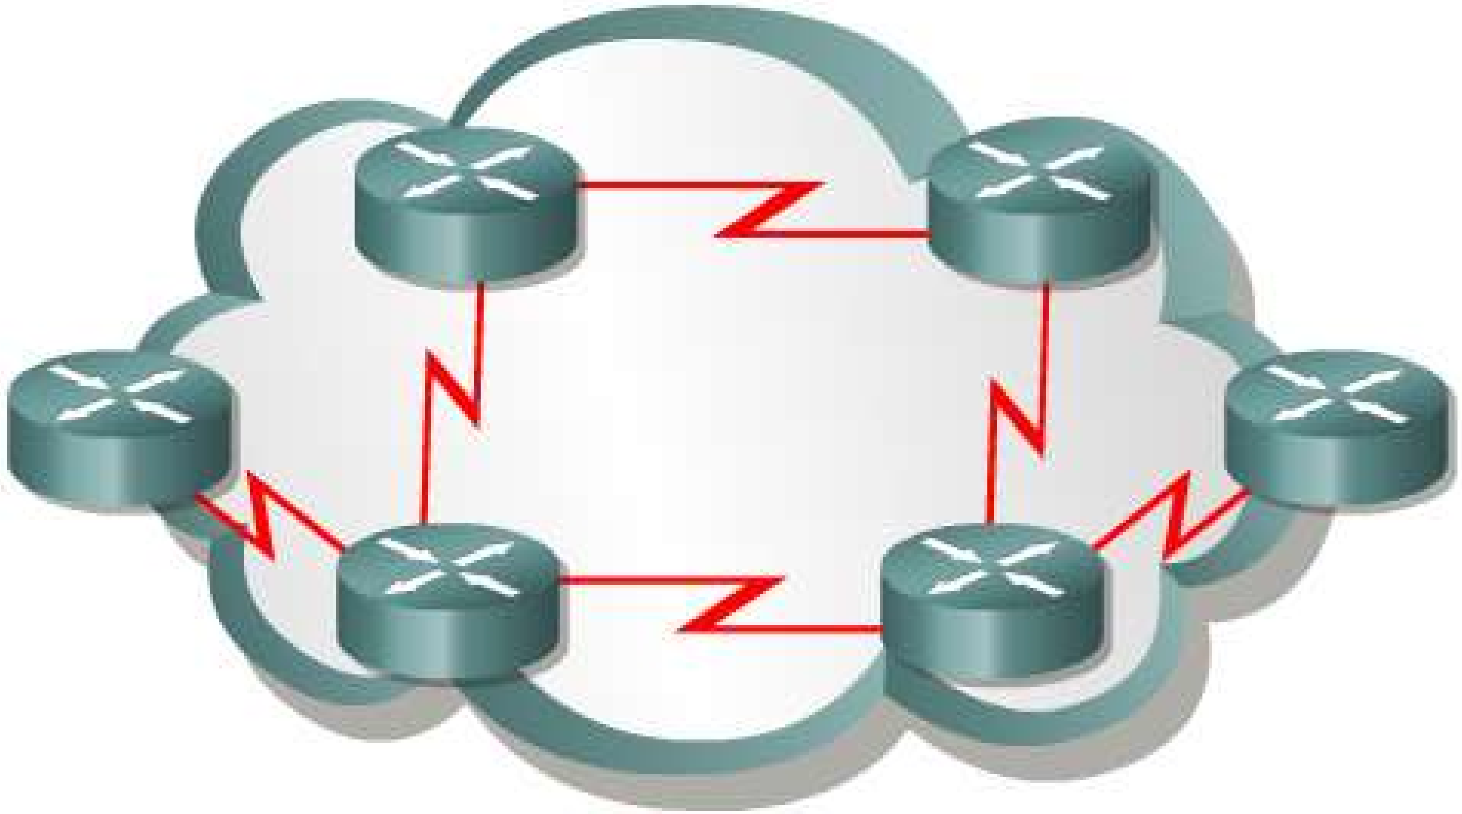

Routers under a common administration

Routing Protocols

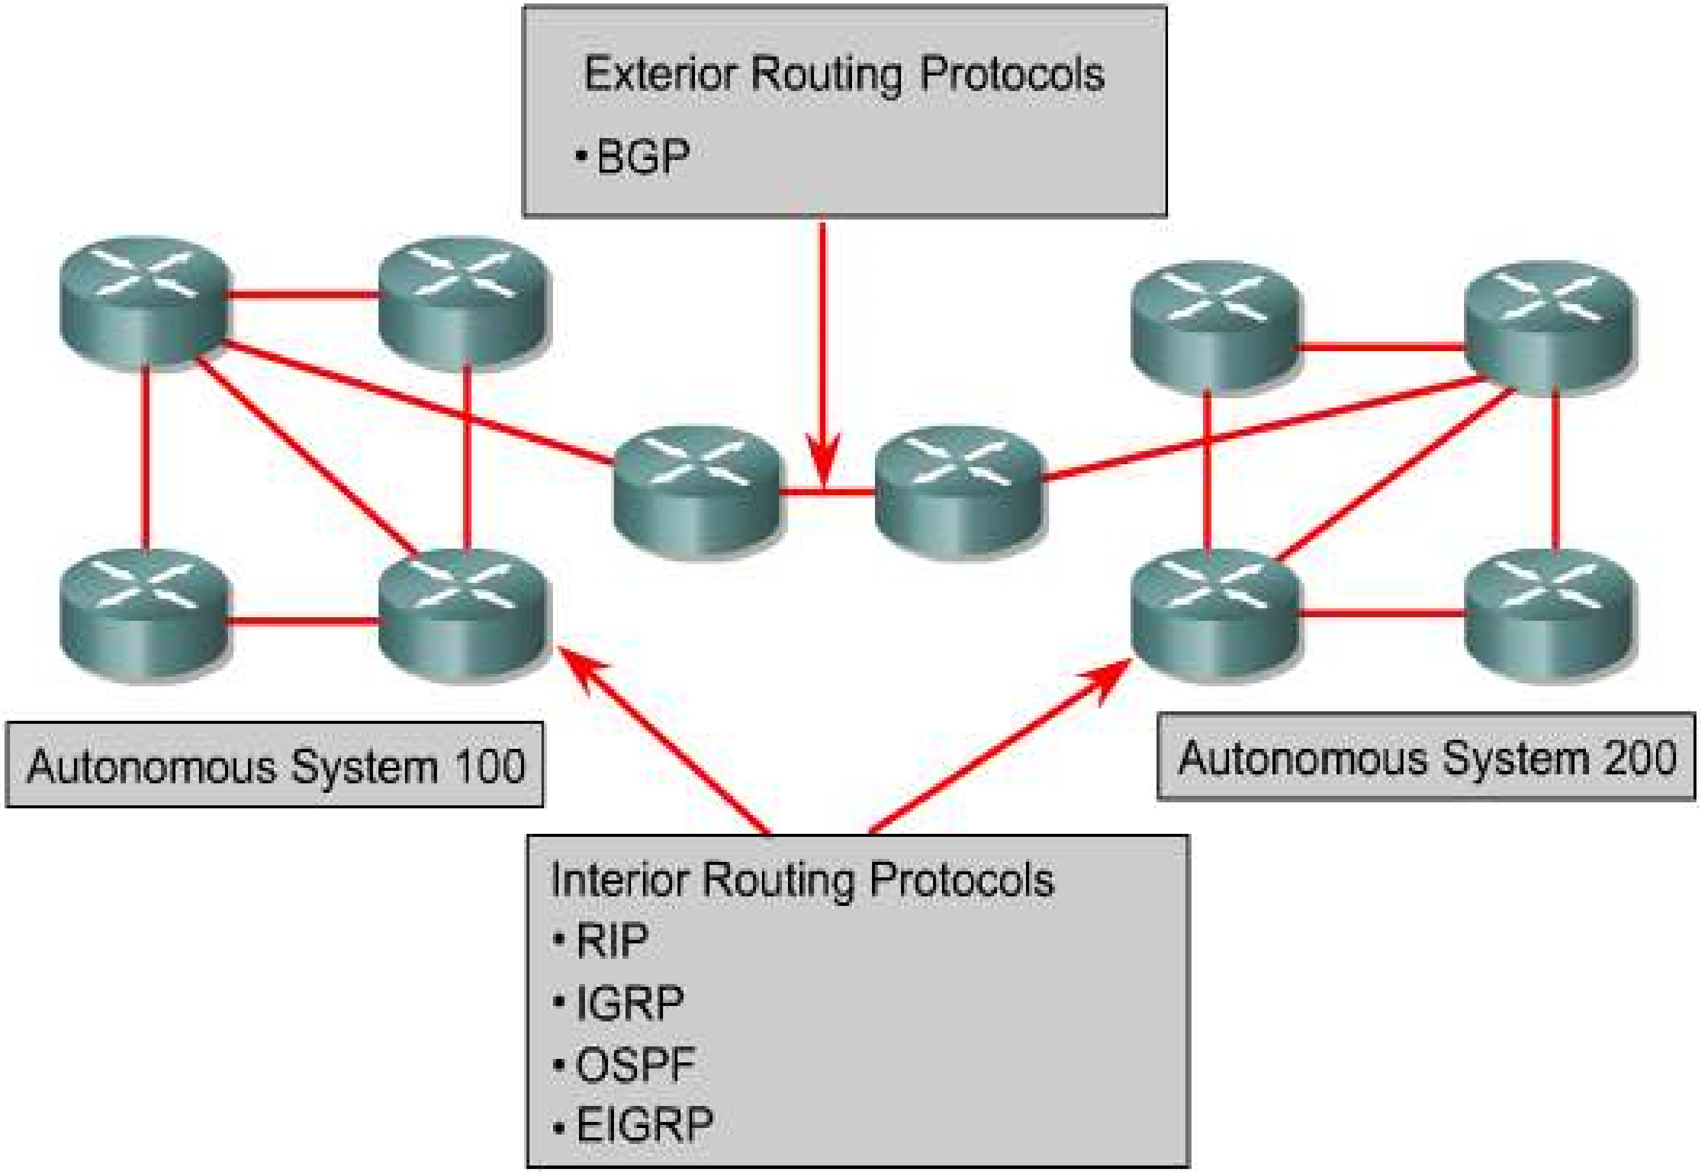

Interior/Exterior Routing Protocols

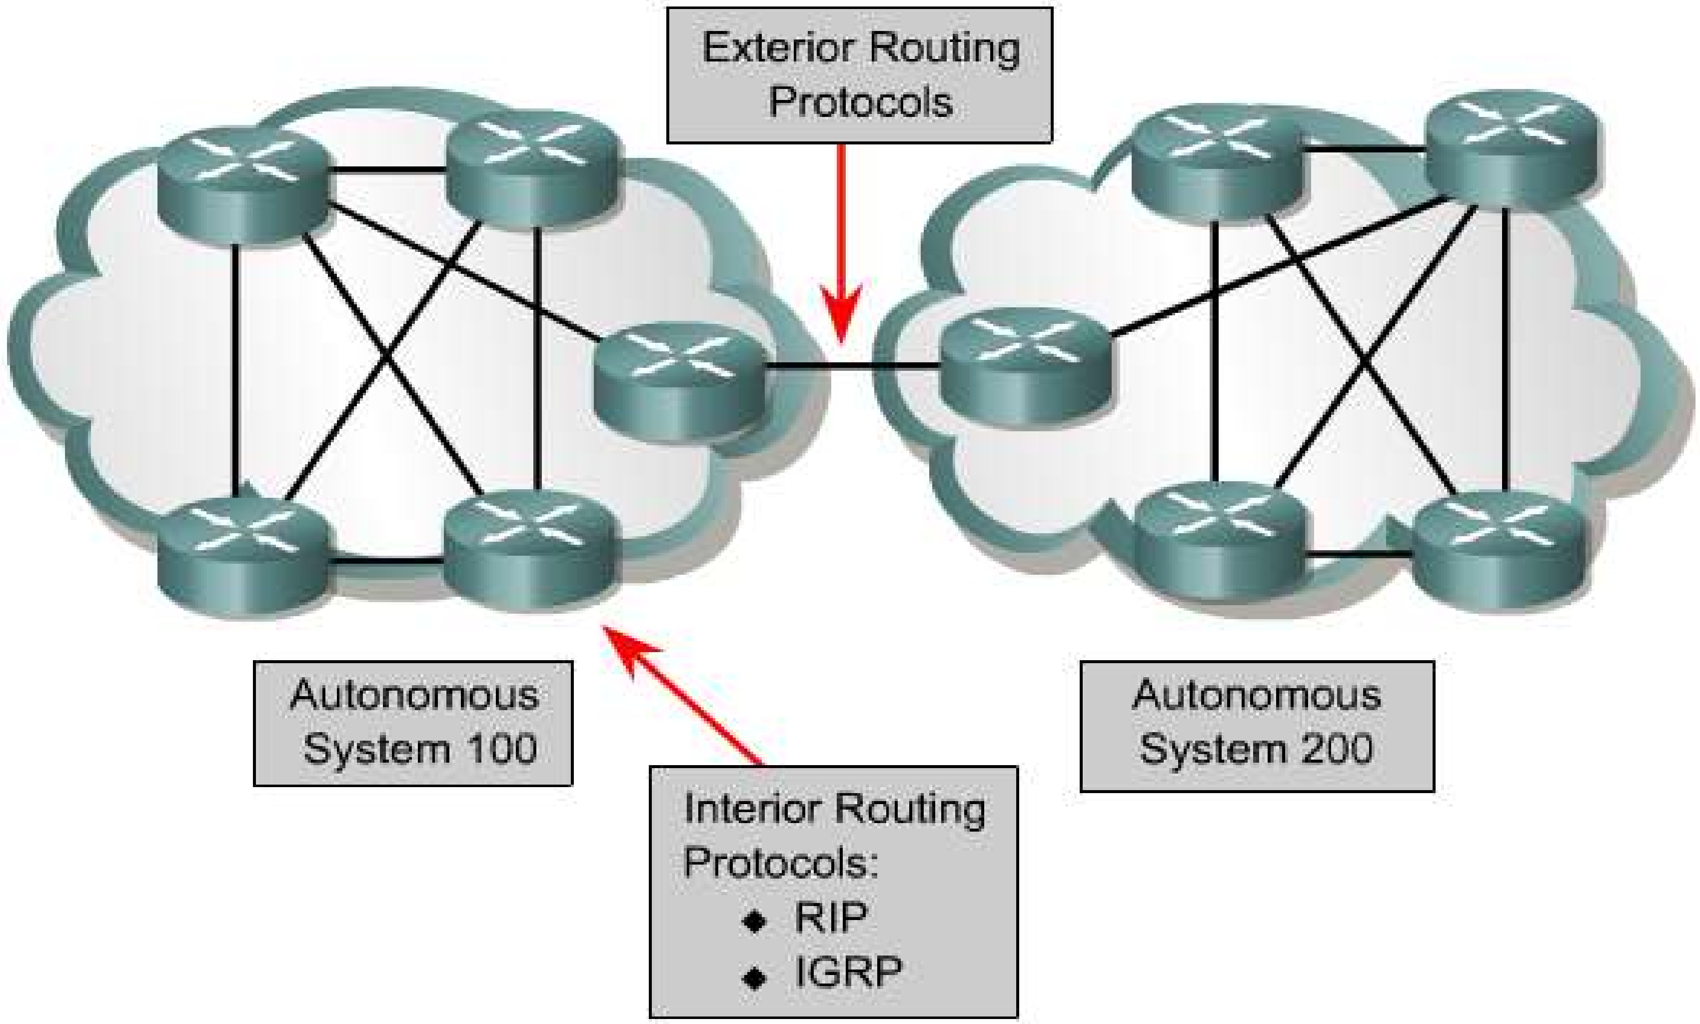

IGP and EGP

- IGPs route data within an autonomous system.
  - RIP, RIPv2, OSPF, IS-IS
- EGPs route data between autonomous systems
  - Border Gateway Protocol (BGP)

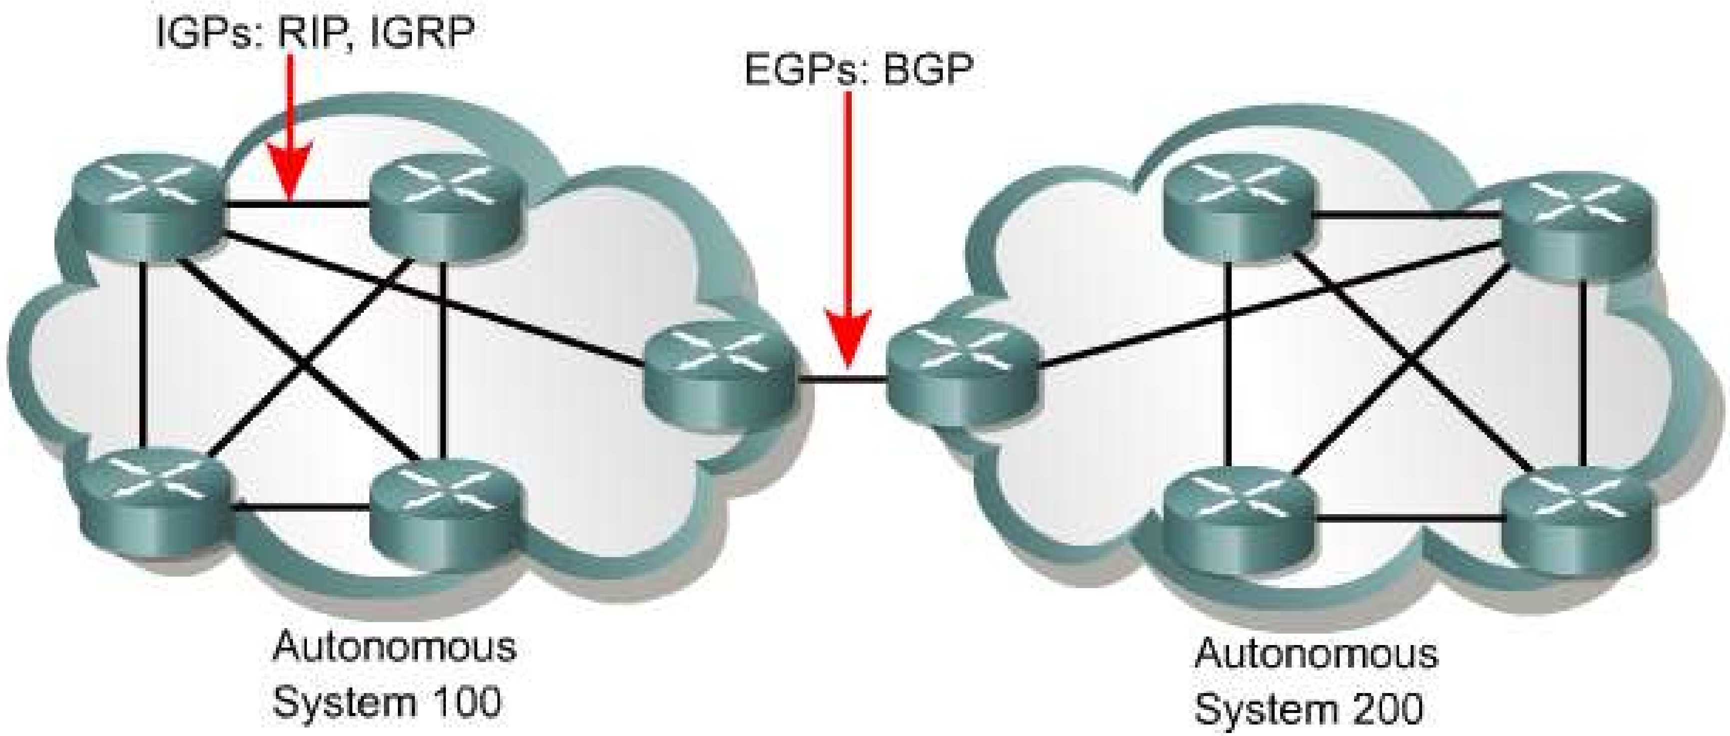

Routing Protocols

- RIP
- RIP v2 (VLSM)
- EIGRP
- OSPF
- IS-IS
- BGP
- ....

Classifying Routing Protocols

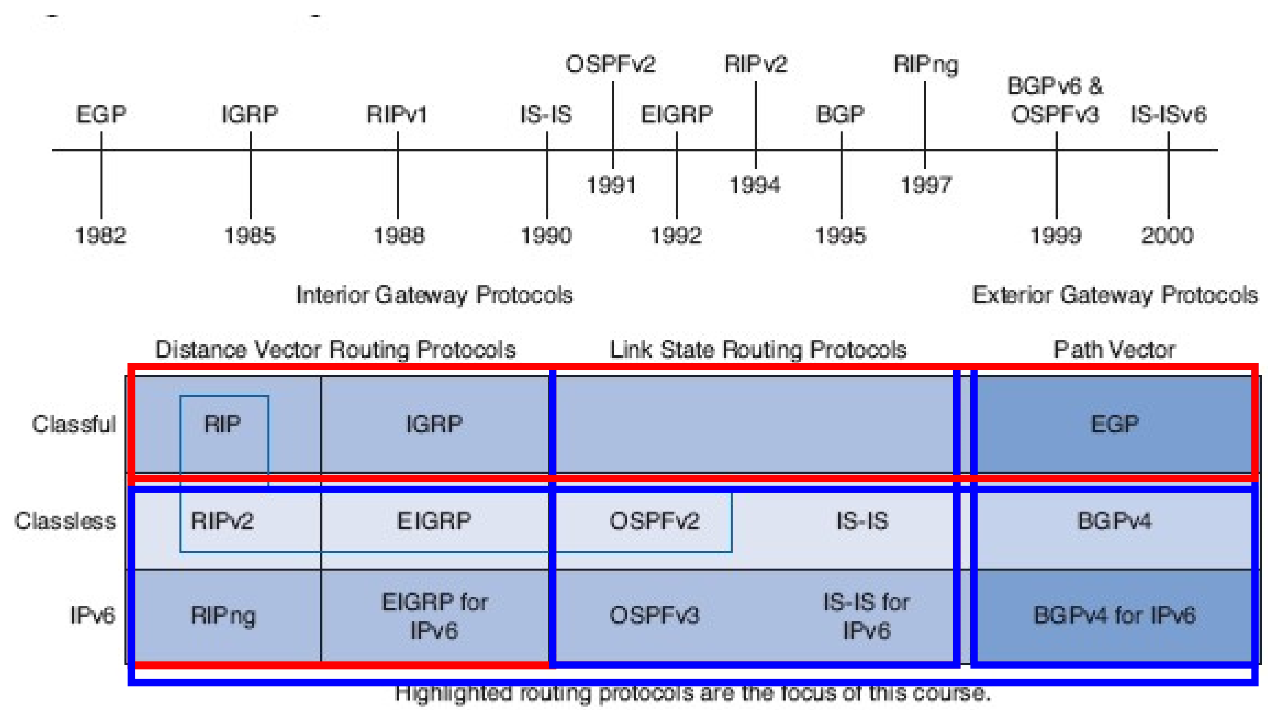

- Routing Protocols can be classified by:
  - IGP or EGP
  - Distance vector or link-state
  - Classful or classless

Distance Vector and Link-State Routing Protocols

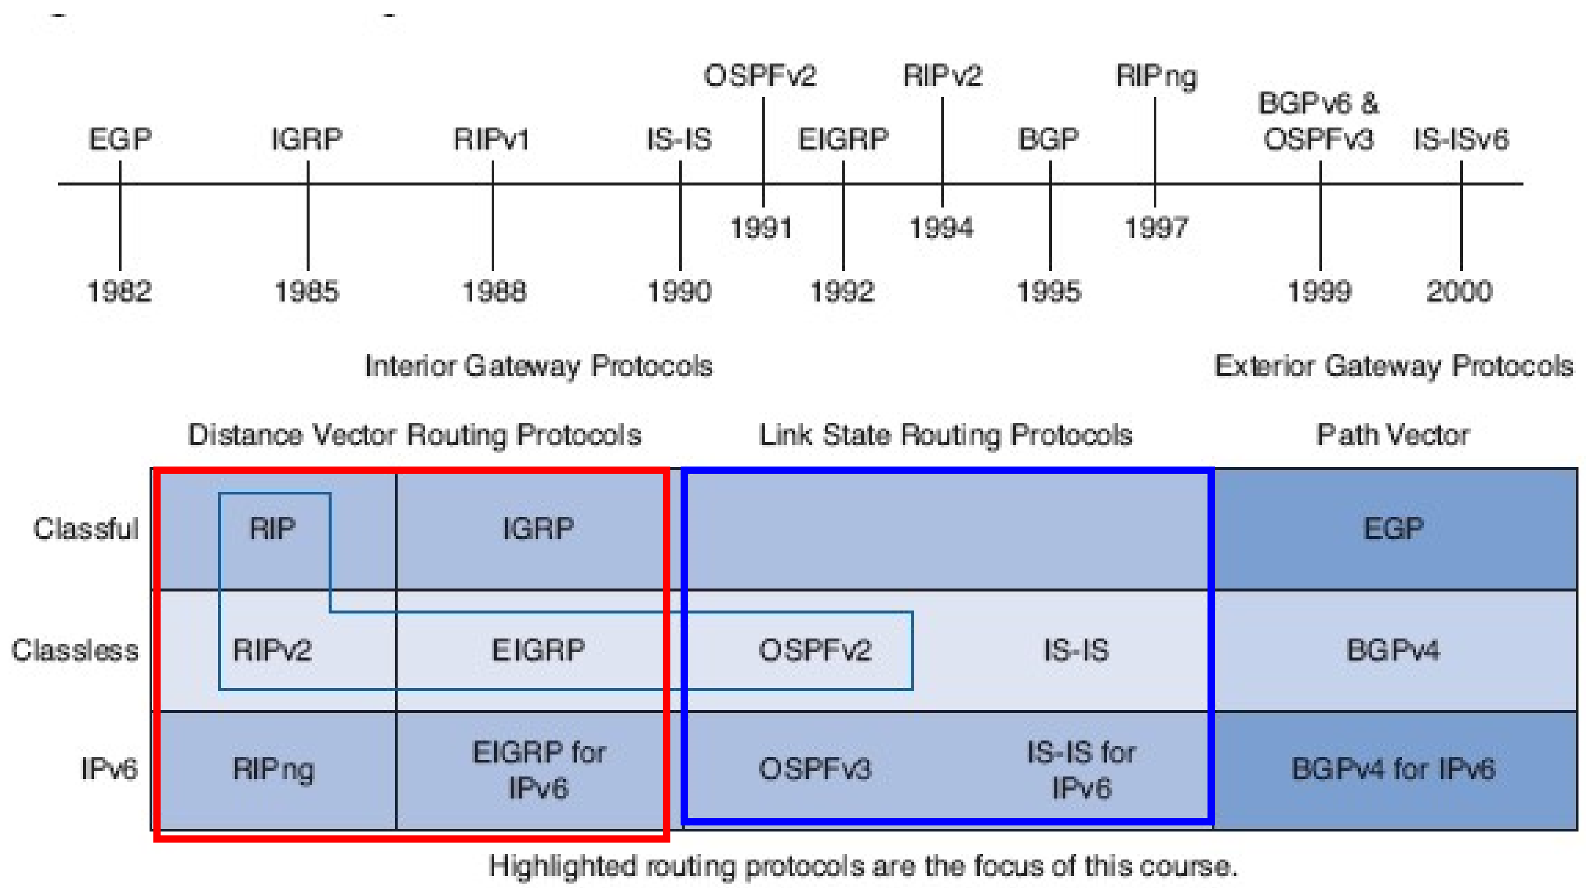

- **Interior gateway protocols (IGP)** can be classified as two types:
  - Distance vector routing protocols
  - Link-state routing protocols

Distance Vector Routing Protocol Operation

- **What does a street sign like this tell you?**
  - **How far (distance)**
  - **Which way (direction)**
- ***Distance vector***
  - Routes are advertised as ***vectors*** of distance and direction.
- **Distance** is defined in terms of a metric
  - Such as hop count
- **Direction** is simply the:
  - Nexthop router or
  - Exit interface
- Typically use the **Bellman-Ford algorithm** for the best-path (shortest) route determination

Link-State Protocol Operation

- ***Link-state*** routing protocol can create a "complete view," or topology, of the network.
- Link-state protocols are associated with Shortest Path First (SPF) calculations.
- A ***link-state router*** uses the link-state information to:
  - Create a topology map
  - Select the best path to all destination networks in the topology.
  - Each router makes the decision!

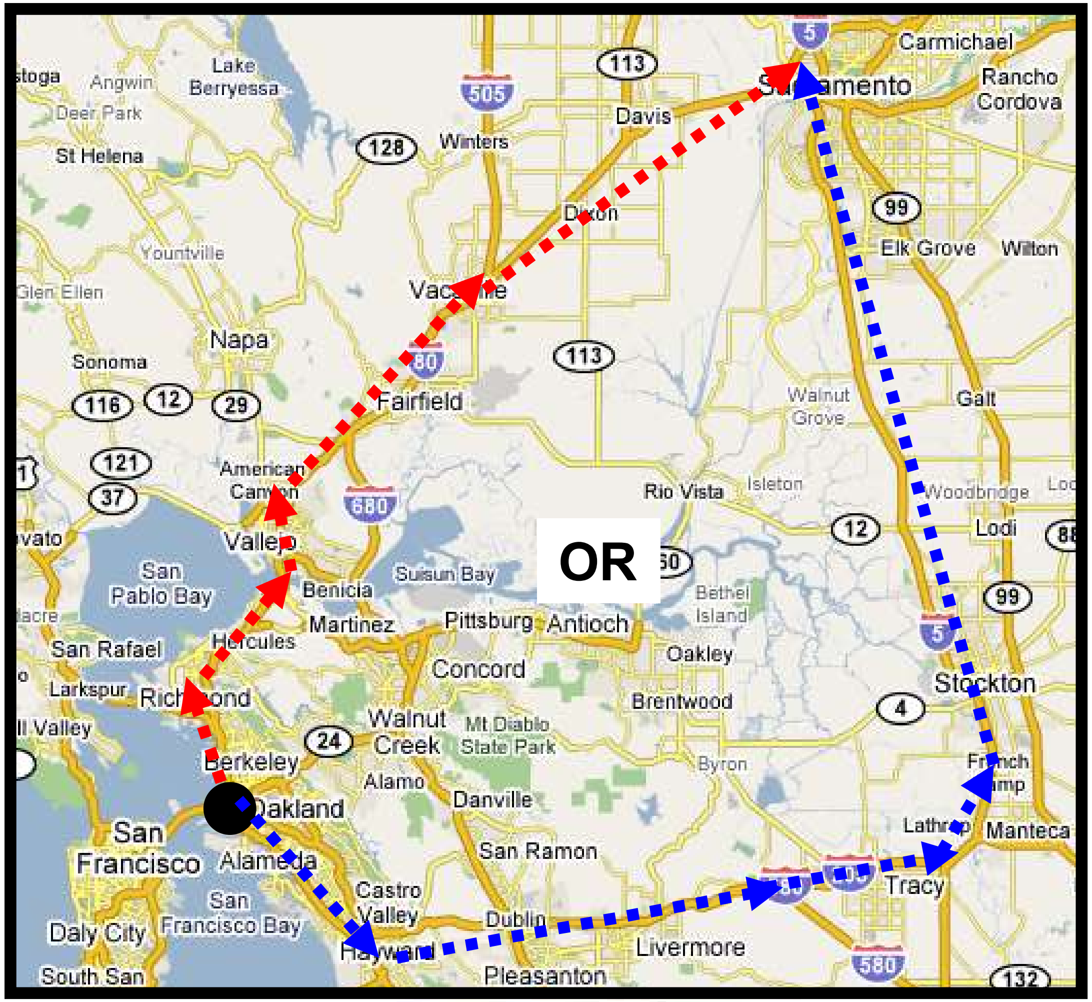

Link State routing protocols is like having a complete map of the network topology

Classful and Classless Routing Protocols

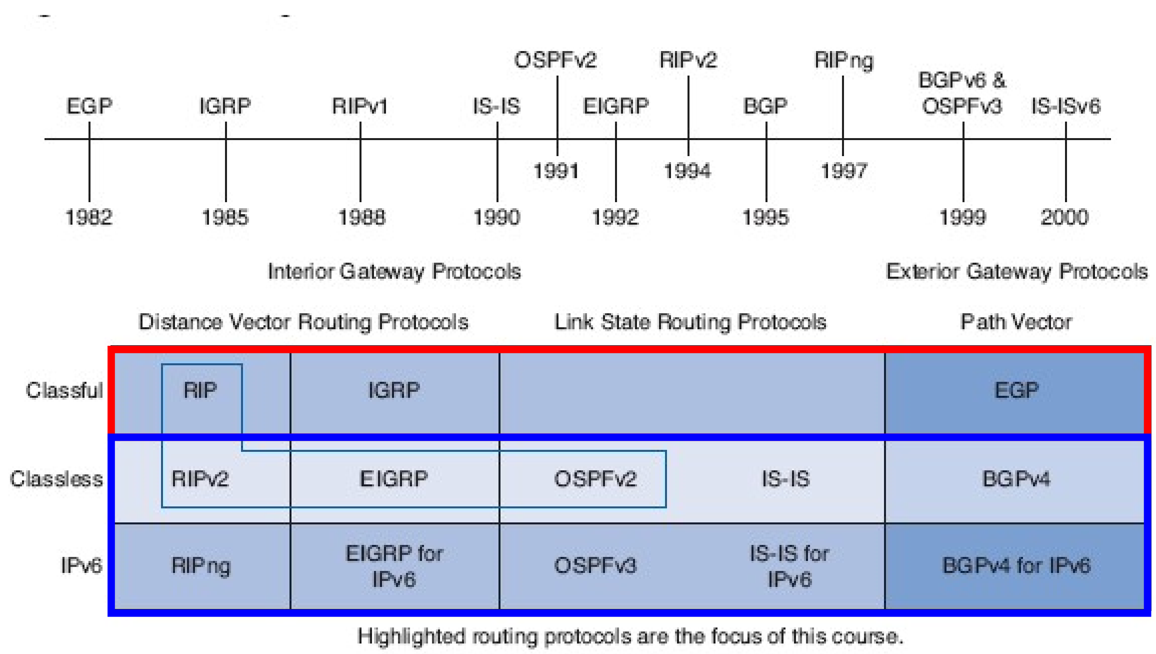

- All routing protocols can also be classified as either
  - **Classful routing protocols**
  - **Classless routing protocols**
    - IPv6 routing protocols are classless

Classful Routing Protocols

	1st Octet	2nd Octet	3rd Octet	4th Octet	Subnet Mask
Class A	Network	Host	Host	Host	255.0.0.0 or /8
Class B	Network	Network	Host	Host	255.255.0.0 or /16
Class C	Network	Network	Network	Host	255.255.255.0 or /24

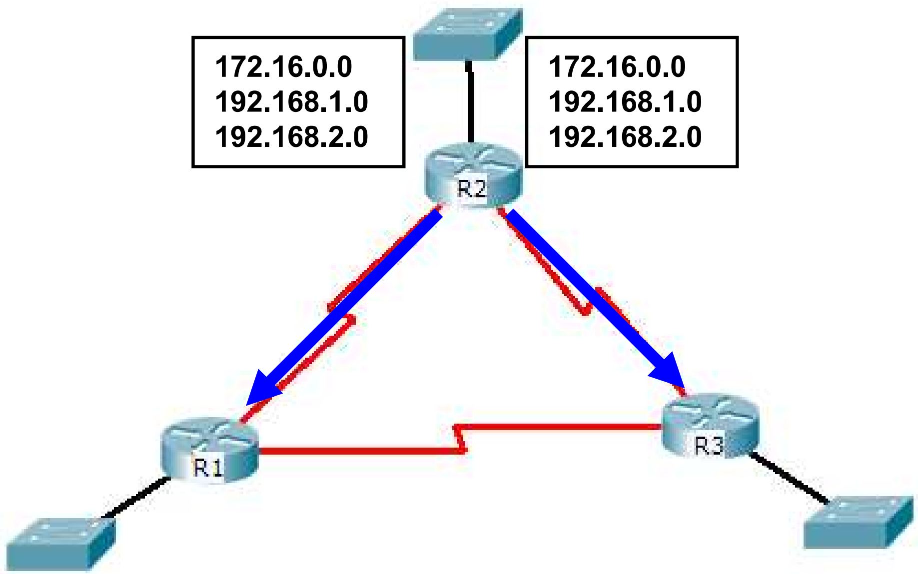

- **Classful routing protocols** do NOT send subnet mask information in routing updates.
  - The first routing protocols, such as RIP
  - When network addresses were allocated based on classes.
    - Class A, B, or C.
    - Network mask determined based on value of first octet of the network address.

Classless routing Protocols

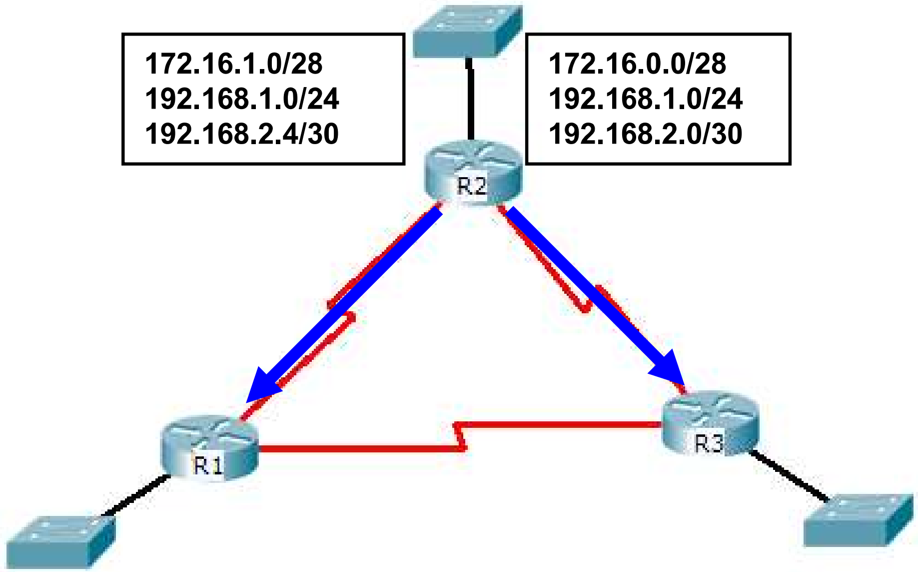

- **Today's networks** are no longer allocated based on classes
  - *Subnet mask cannot be determined by the value of the first octet*.
- **Classless routing protocols** include the subnet mask with the network address in routing updates.

Classless routing Protocols

172.16.0.0/16 Major Classful Network

/27 and /30 subnets

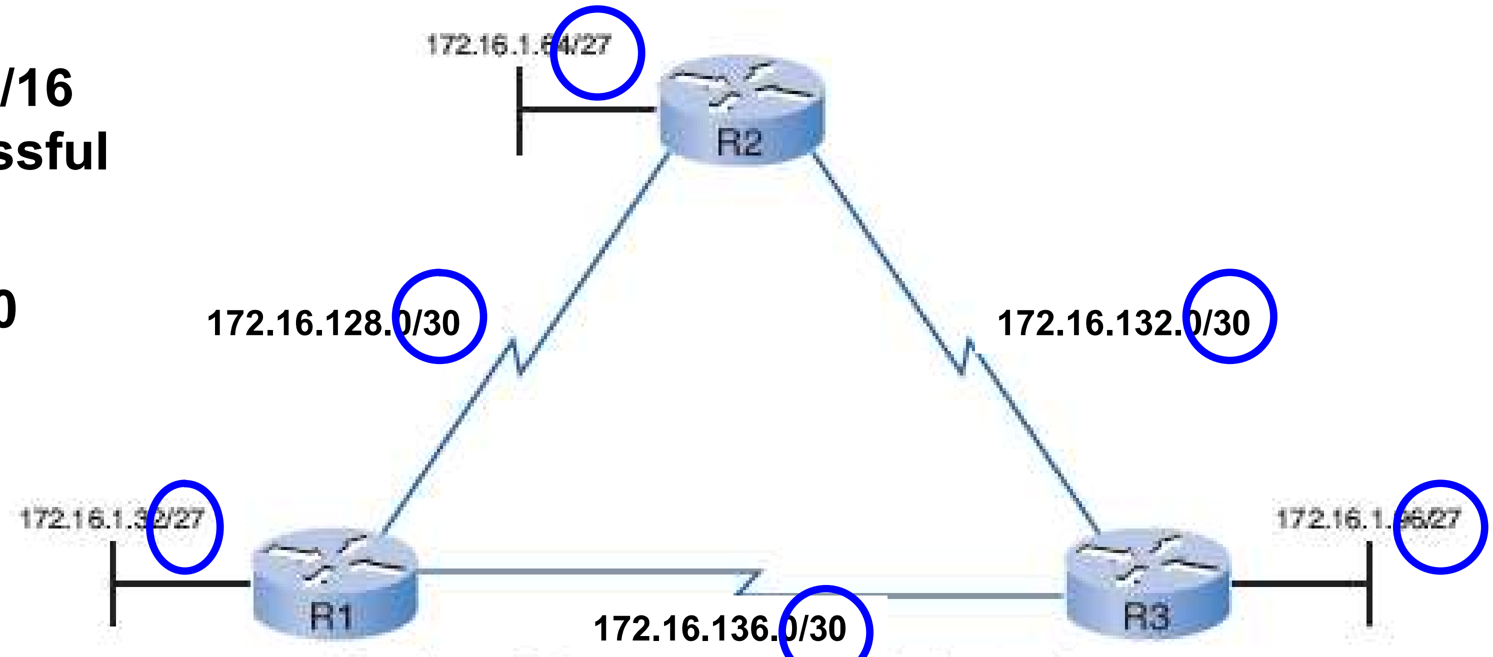

Classless: Subnet mask can vary in the topology.

- Classless routing protocols are required in most networks today because of their **support** for:
  - VLSM
  - CIDR
  - Discontiguous networks.

Dynamic Routing Protocols and Convergence

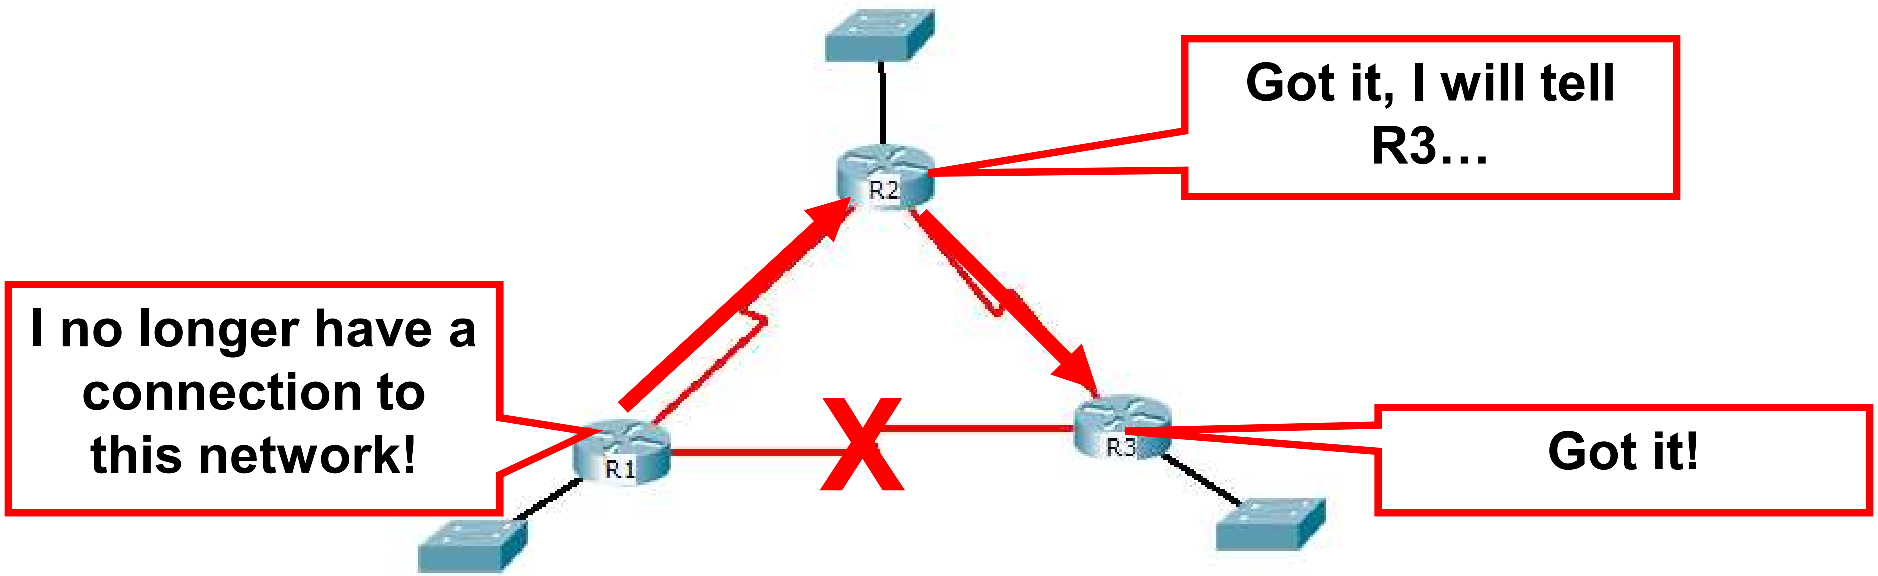

- ***Convergence*** is when *<u>the routing tables of all routers are at a state of consistency.</u>*
  - **Network has converged:** When <u>all routers have complete and accurate information</u> about the network.
- **Convergence time** is the <u>time it takes routers to:</u>
  - <u>Share information</u>
  - <u>Calculate best paths</u>
  - <u>Update their routing tables</u>
- A <u>network is not completely operable</u> <u>until the network has converged</u>; therefore, most networks require short convergence times.

Dynamic Routing Protocols and Convergence

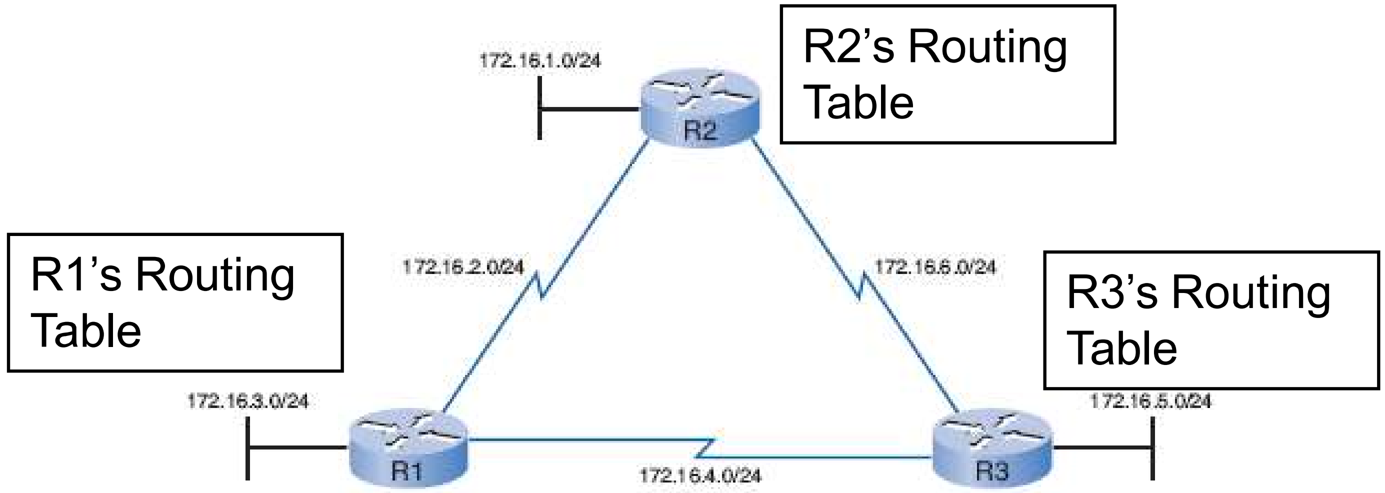

- Generally, convergence time:
  - Slow: RIP and IGRP
  - Faster: EIGRP, OSPF, and IS-IS

Purpose of a Metric

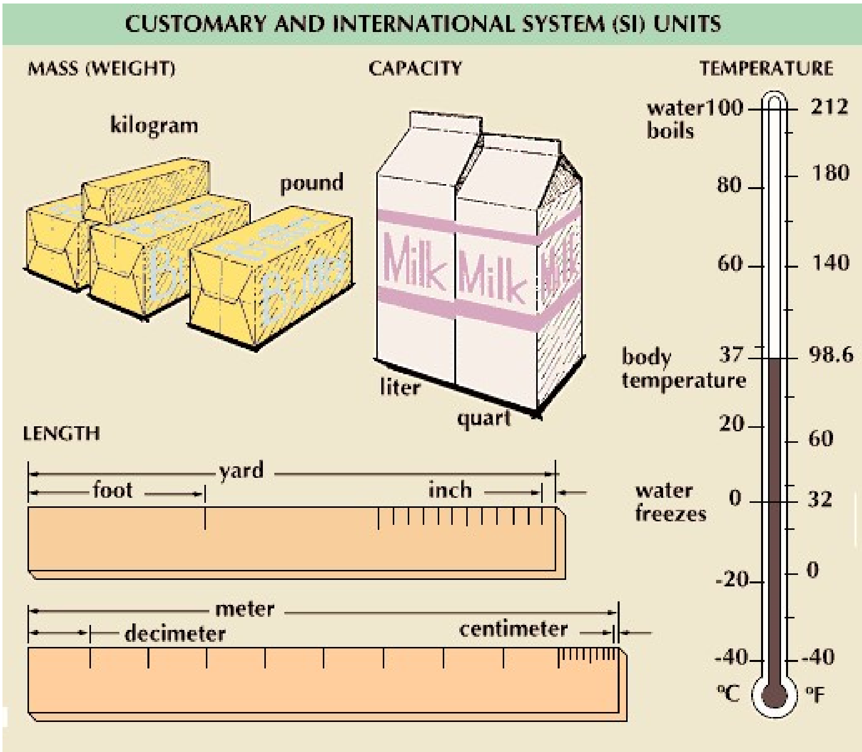

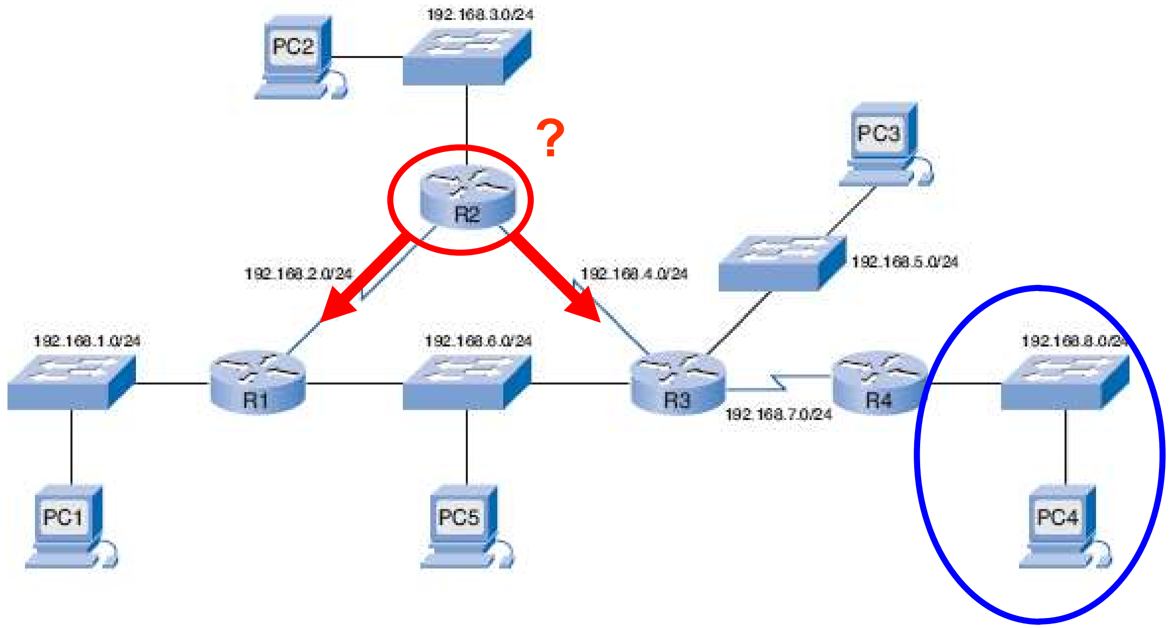

- **Metrics** are a way to measure or compare.
  - Determine the best path.
- **Routing protocol** learns multiple routes to the same destination.
  - Metric is used to determine which path is most preferable

Purpose of a Metric

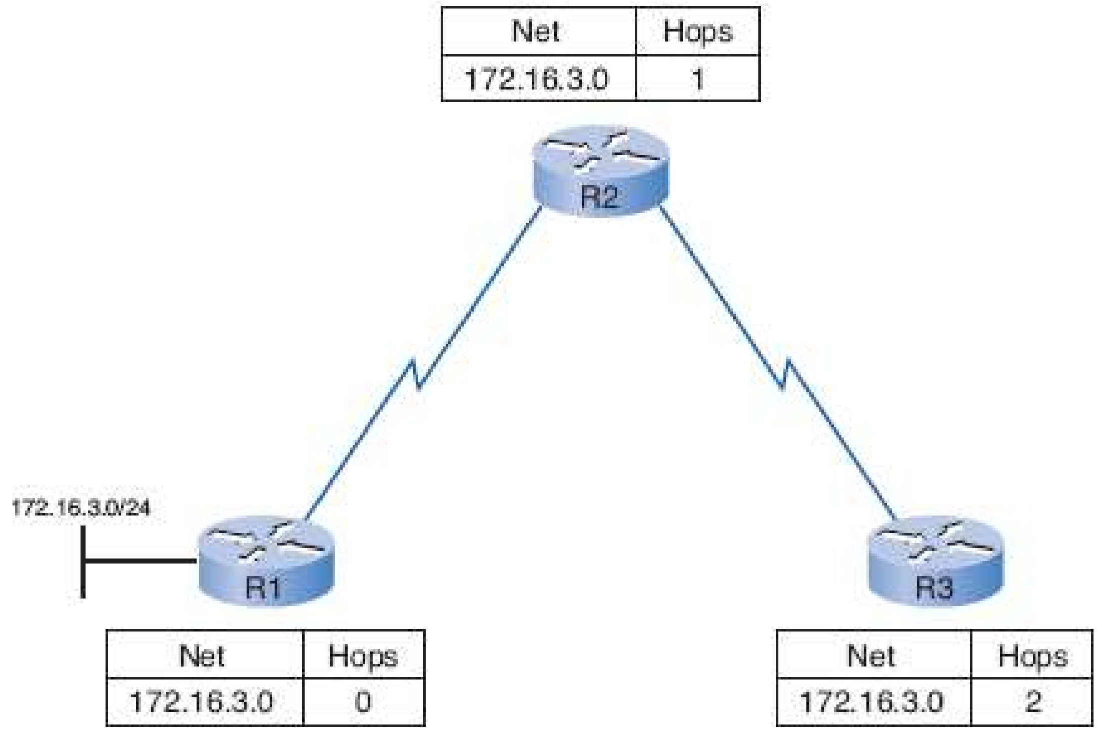

- What might be some ways (metrics) that routing protocols might use to determine the “best path?
- Routing protocol metrics:
  - **RIP**: Hop count
  - **IGRP and EIGRP**: Bandwidth, delay, reliability and load
  - **OSPF** (Cisco’s version): Bandwidth
  - **IS-IS**: Four values
  - **BGP**: Attributes

Metric Parameters

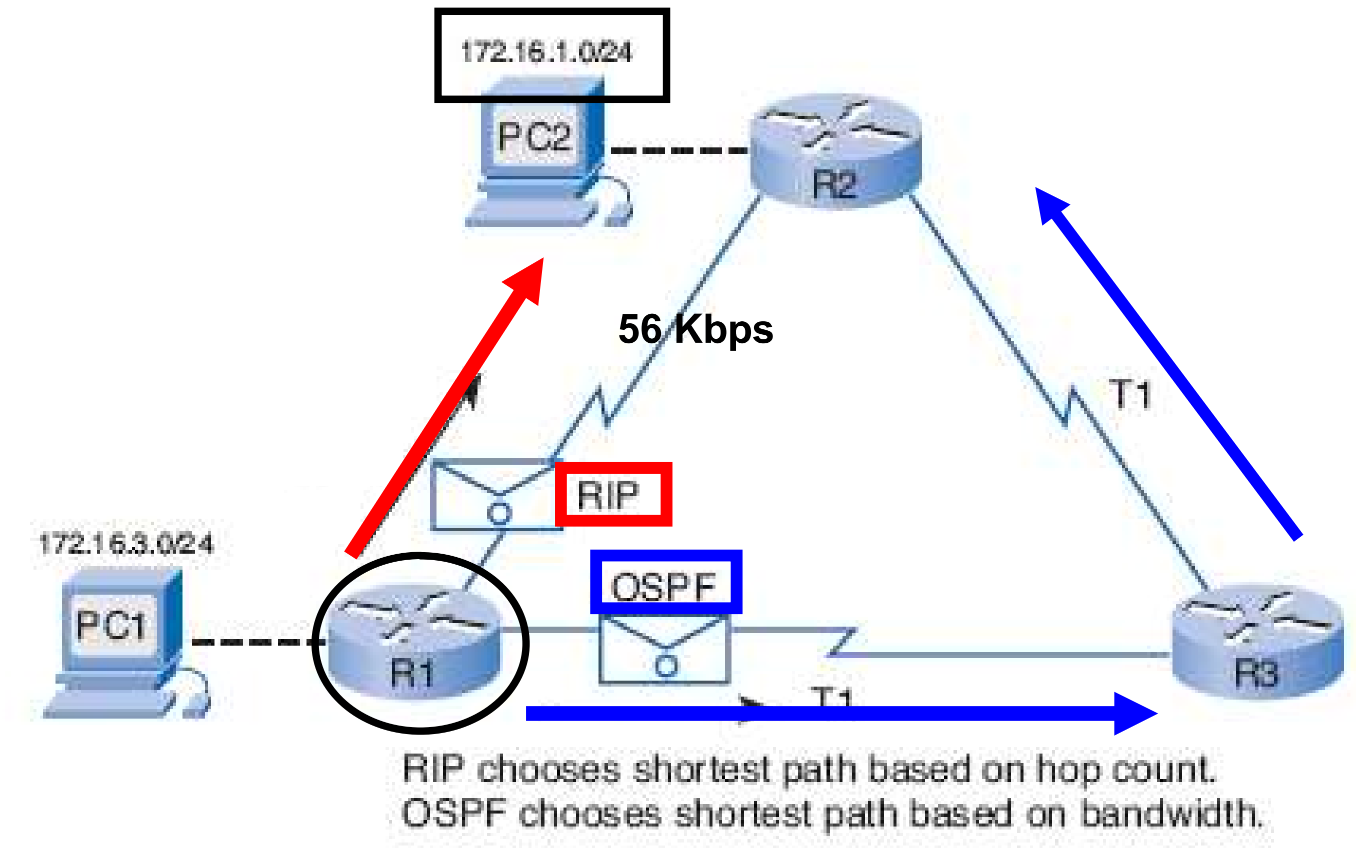

- **R1** to reach the **172.16.1.0/24** network.
- **RIP**: Fewest number of hops via R2.
- **OSPF**: Path with the highest cumulative bandwidth through R3.
  - This results in faster packet delivery.

Metric Field in the Routing Table

```
R2# show ip route

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.1.0 is directly connected, FastEthernet0/0
C       172.16.2.0 is directly connected, Serial0/0/0
C    192.168.1.0/24 is directly connected, Serial0/0/1
S    192.168.2.0/24 [1/0] via 192.168.1.1
R    192.168.7.0/24 [120/1] via 192.168.4.1, Serial0/0/1
R    192.168.8.0/24 [120/2] via 192.168.4.1, Serial0/0/1
```

- The **routing table** displays the metric for each dynamic and static route.
  - **Static routes** always have a metric of 0.
- Routing protocols install route in routing table with the lowest metric.

```
R2# show ip route
<output omitted>
Gateway of last resort is not set
R     192.168.1.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
C     192.168.2.0/24 is directly connected, Serial0/0/0
C     192.168.3.0/24 is directly connected, FastEthernet0/0
C     192.168.4.0/24 is directly connected, Serial0/0/1
R     192.168.5.0/24 [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
R     192.168.6.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
                     [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
R     192.168.7.0/24 [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
R     192.168.8.0/24 [120/2] via 192.168.4.1, 00:00:26, Serial0/0/1
```

- All routers running RIP
- What is the metric (distance) for R2 to reach the 192.168.8.0 network?
  - **2 (hops away)**
- What is the direction (vector) for R2 to reach the 192.168.8.0 network?
  - **Serial 0/0/1 (via R3)**

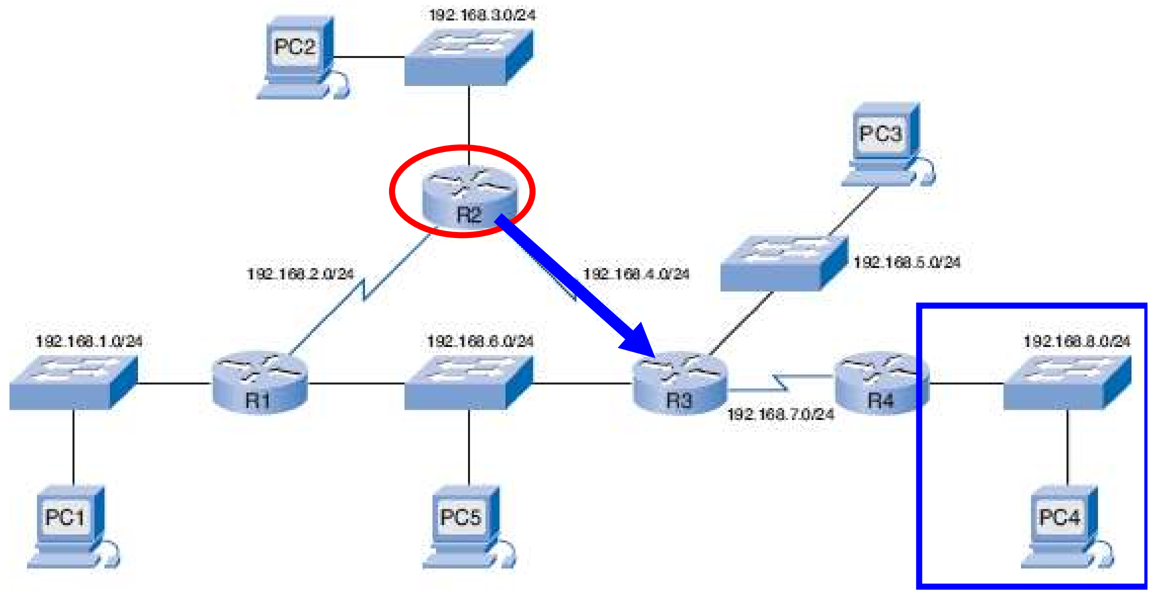

Load Balancing

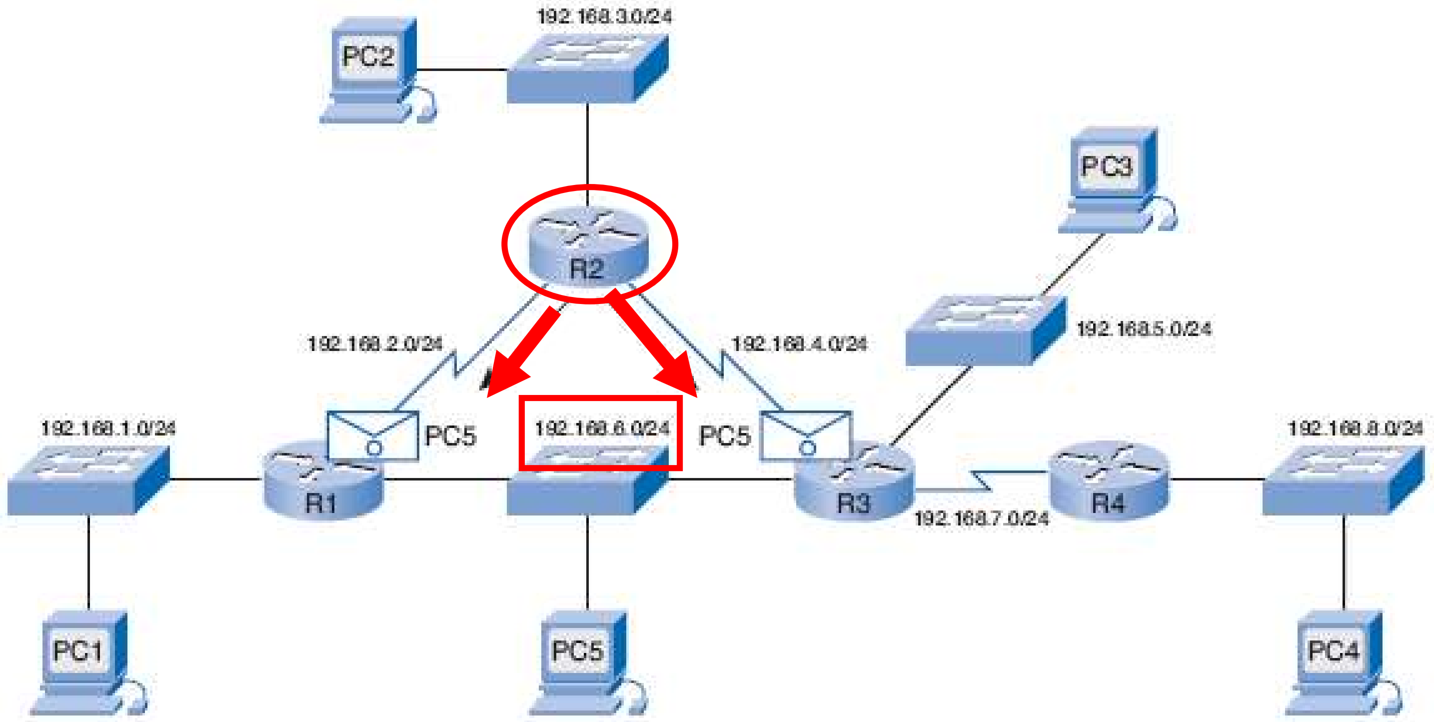

R2 load balances traffic destined for the 192.168.6.0/24 network.

```
R2# show ip route
<output omitted>
R     192.168.6.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
                     [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
```

- What happens when two or more routes to the same destination have identical metric values?
  - The router **load balances** between these equal-cost paths.
  - All routing protocols do equal cost load balancing.
  - EIGRP also does unequal cost load balancing.

Purpose of Administrative Distance

- What if a router learns about a remote network from two different routing sources.
- What if RIP advertises the network as 10 hops away but OSPF advertises it as a cumulative bandwidth of 100,000.
- Which is better RIP or OSPF?
  - Can't tell
  - Can't compare apples and oranges.
  - Note: This is not common.
- ***Administrative distance (AD)*** is:
  - Used to determine which routing source takes precedence.
  - Used when there are multiple routing sources for the same destination network address.
- **Lower the AD** the more preferred the routing source.

Route Source	AD
Connected	0
Static	1
EIGRP summary route	5
E[illegible]	20
I[illegible]	90
IGRP	100
OSPF	110
IS-IS	115
[illegible]	120
[illegible]	170
[illegible]	200

So, which one would be preferred RIP or OSPF? OSPF

Which route would be preferred, OSPF or a Static Route to the same network? Static Route

Purpose of Administrative Distance

- Cisco uses Administrative distance (AD) to define the preference of a routing source.
- Routing sources:
  - **Directly connected networks**
  - **Static routes**
  - **Specific routing protocols**
- AD for static and dynamic can be modifed

Route Source	AD
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

Note

- The term ***trustworthiness*** is commonly used when defining administrative distance.
- *The lower the administrative distance value, the more trustworthy the route.*

Purpose of Administrative Distance

- AD: 0 to 255.
- The lower the value, the more preferred the route source.
- **AD of 0** is the most preferred.
  - *Only a directly connected network has an administrative distance of 0, which cannot be changed.*
  - No better route to a network than being directly connected to that network.
- **AD of 255** means the router will not believe the source of that route
  - Route will not be installed in the routing table.

Route Source	AD
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

Verifying AD: `show ip route`

```
R2# show ip route

D    192.168.6.0/24 [90/2172416] via 192.168.2.1, 00:00:24, Serial0/0/0
```

What is the AD of this route? 90

Route Source	AD
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

Verifying AD: **show ip protocols**

```
R2# show ip protocols
Routing Protocol is "eigrp 100 "
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  Automatic network summarization is in effect
  Automatic address summarization:
  Maximum path: 4
  Routing for Networks:
     192.168.2.0
     192.168.3.0
     192.168.4.0
Routing Information Sources:
  Gateway         Distance     Last Update
  192.168.2.1     90           2366569
Distance: internal 90 external 170
<continued next slide?
```

Static Routes and Administrative Distance

```
R2# show ip route

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.1.0 is directly connected, FastEthernet0/0
C       172.16.2.0 is directly connected, Serial0/0/0
S       172.16.3.0 is directly connected, Serial0/0/0
C    192.168.1.0/24 is directly connected, Serial0/0/1
S    192.168.2.0/24 [1/0] via 192.168.1.1
```

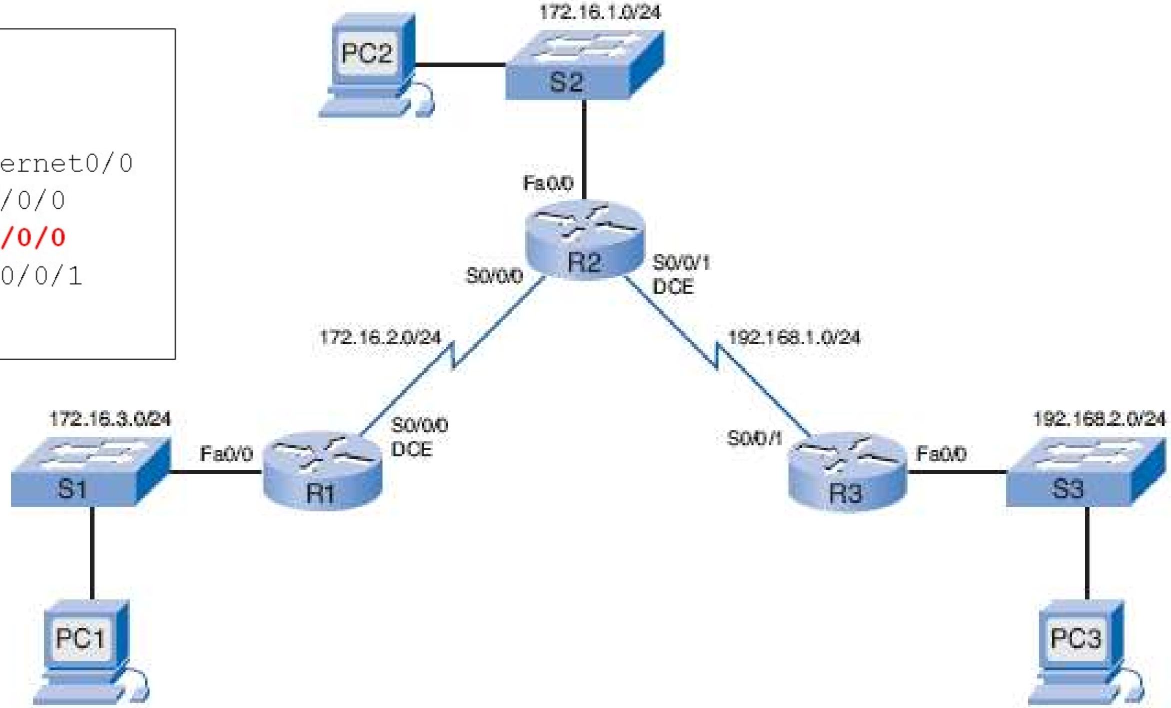

- **Static routes**
  - Default AD = 1
- <u>After directly connected networks (AD = 0), static routes are the most preferred route source.</u>

Static Routes and Administrative Distance

Exit Interface: `ip route 172.16.3.0 255.255.255.0 serial 0/0/0`

```
R2# show ip route

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.1.0 is directly connected, FastEthernet0/0
C       172.16.2.0 is directly connected, Serial0/0/0
S       172.16.3.0 is directly connected, Serial0/0/0
C    192.168.1.0/24 is directly connected, Serial0/0/1
S    192.168.2.0/24 [1/0] via 192.168.1.1
```

Next-hop: `ip route 192.168.2.0 255.255.255.0 192.168.1.1`

- **What is the AD of a Static Route?**
- **Static route**: default AD = 1 (never 0)
  - Exit-interface: AD = 1
  - Next-hop IP address: AD = 1
- After directly connected networks (AD = 0), static routes are the most preferred route source.

Floating Static Route

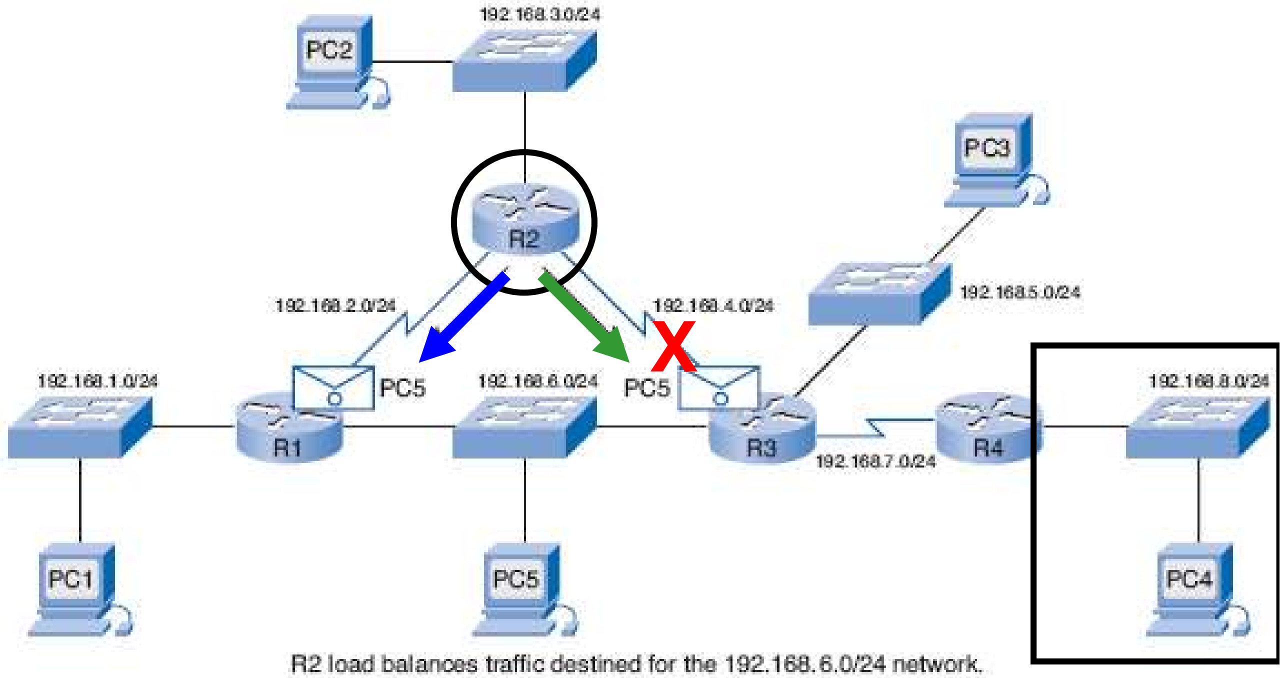


```
R2: ip route 192.168.8.0 255.255.255.0 192.168.4.1
R2: ip route 192.168.8.0 255.255.255.0 192.168.2.1 5
```

- There are situations when an administrator will configure a static route to the same destination that is learned using a dynamic routing protocol, but using a different path.
- The static route will be configured with an AD greater than that of the routing protocol.
- If there is a link failure in the path used by the dynamic routing protocol, the route entered by the routing protocol is removed from the routing table.
- The static route will then become the only source and will automatically be added to the routing table.
- This is known as a ***floating static route***.

What Is VLSM and Why Is It Used?

What is VLSM and why use it?

- Short term extensions to IPv4
- Subnetting 1985
- Variable length subnetting 1987
- Classless Interdomain Routing 1993
- Private IP addresses
- Network Address Translation (NAT)

VLSM is supported by:

- OSPF
- IS-IS
- EIGRP
- RIP v2
- Static routing

What is VLSM and why use it?

- Ultimate solution: IPv6 128 bit address space
- Allows for: 340,283, 366, 920, 938, 463, 374, 607, 431, 768, 211, 456 possibilities

A Waste of Space

Subnet Number	Subnet Address	
Subnet 0	192.168.187.0	/27
Subnet 1	192.168.187.32	/27
Subnet 2	192.168.187.64	/27
Subnet 3	192.168.187.96	/27
Subnet 4	192.168.187.128	/27
Subnet 5	192.168.187.160	/27
Subnet 6	192.168.187.192	/27
Subnet 7	192.168.187.224	/27

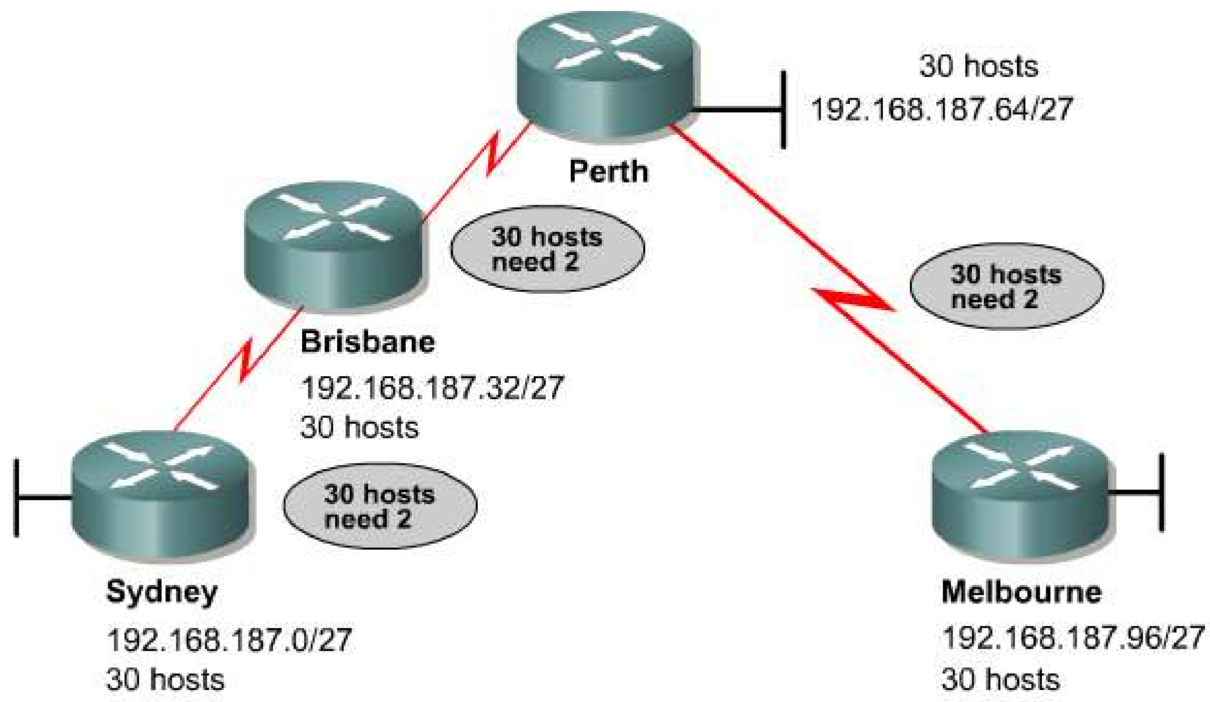

When to Use VLSM?

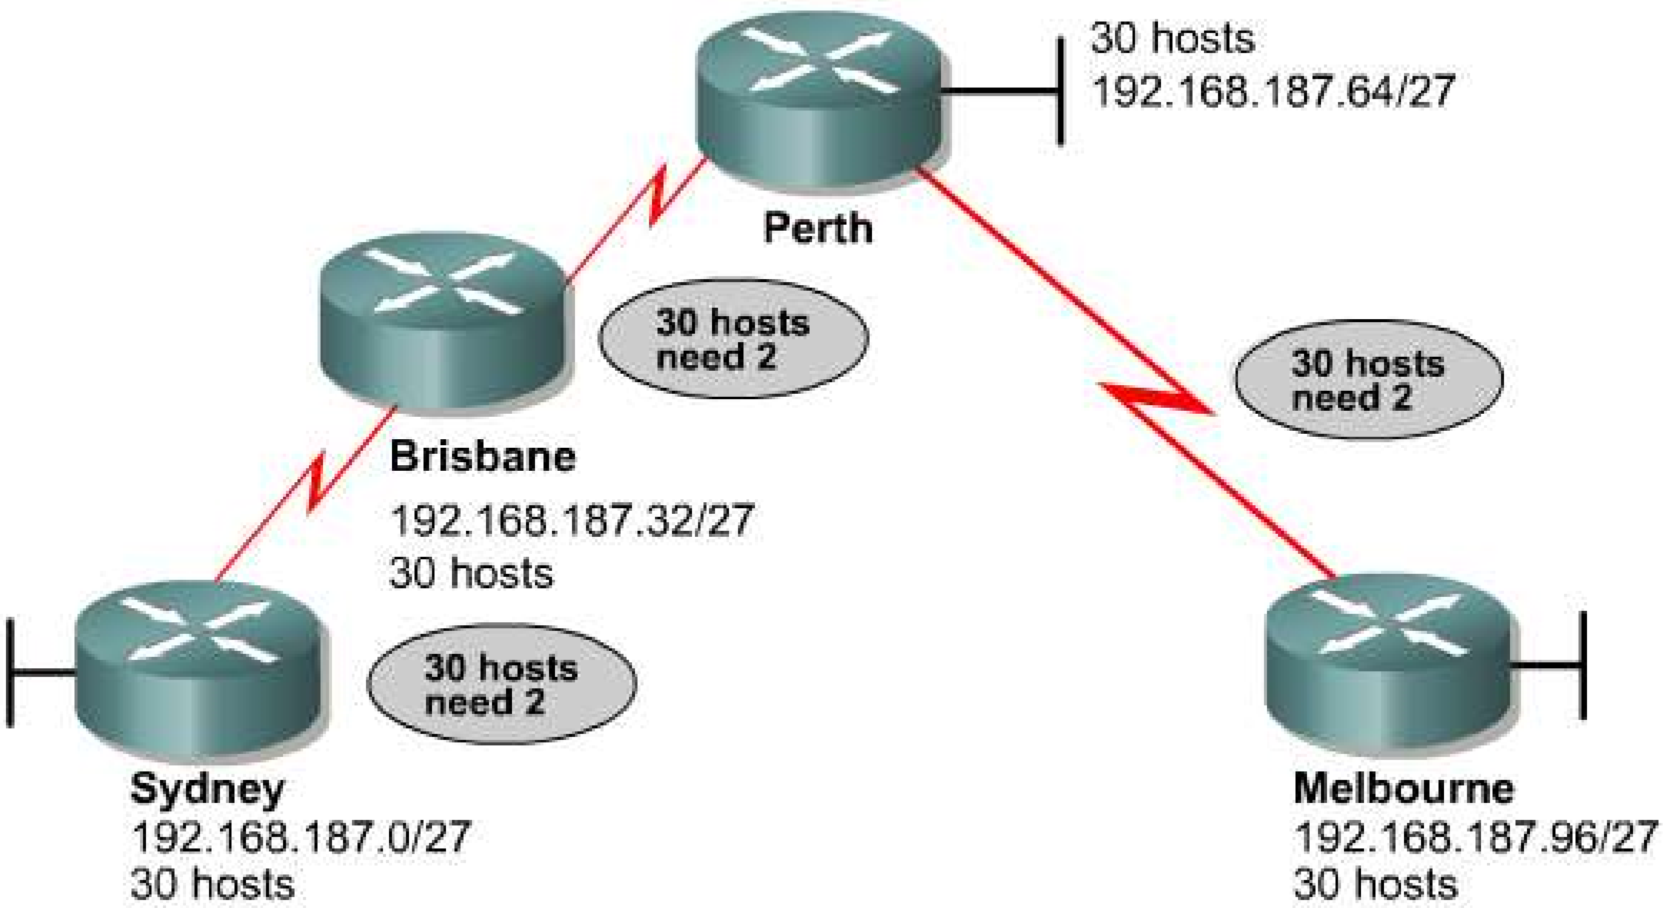

Use VLSM on the point to point links to use only 2 valid host addresses instead of wasting 30.

Calculating Subnets with VLSM

Subnetted Address: 172.16.32.0/20

In Binary 101011100.00010000.00100000.00000000

VLSM Address: 172.16.32.0/26

In Binary 101011100.00010000.0010 | 0000. 00 | 000000

	Network			Subnet	VLSM Subnet	Host	
1st subnet:	172	•	16	.0010	0000.00	000000	= 172.16.32.0/26
2nd subnet:	172	•	16	.0010	0000.01	000000	= 172.16.32.64/26
3rd subnet:	172	•	16	.0010	0000.10	000000	= 172.16.32.128/26
4th subnet:	172	•	16	.0010	0000.11	000000	= 172.16.32.192/26
5th subnet:	172	•	16	.0010	0001.00	000000	= 172.16.33.0/26

A Working VLSM Example

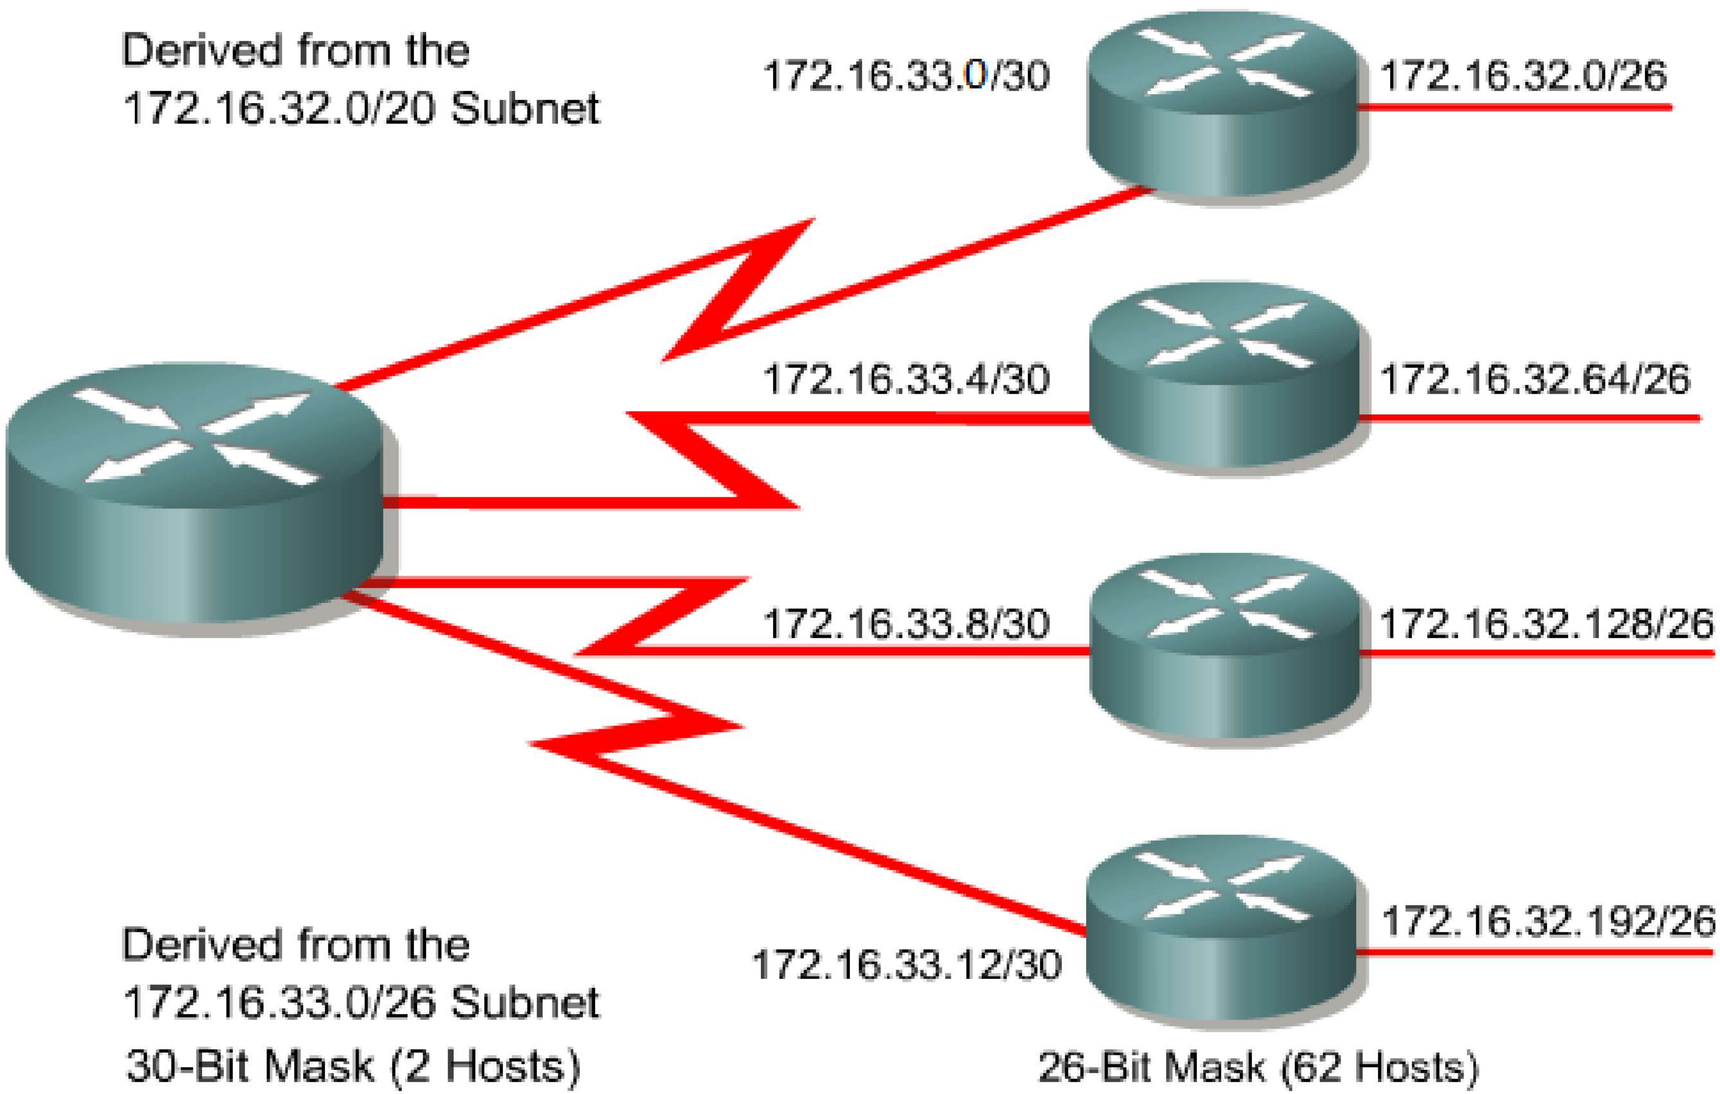

Route Aggregation

- Keeping the networks close to one another saves routing table space
- Every network needs a separate entry in the routing table
- Each subnet needs a separate entry.
- Allow for aggregation can reduce the size of the routing table.

Route Summarization

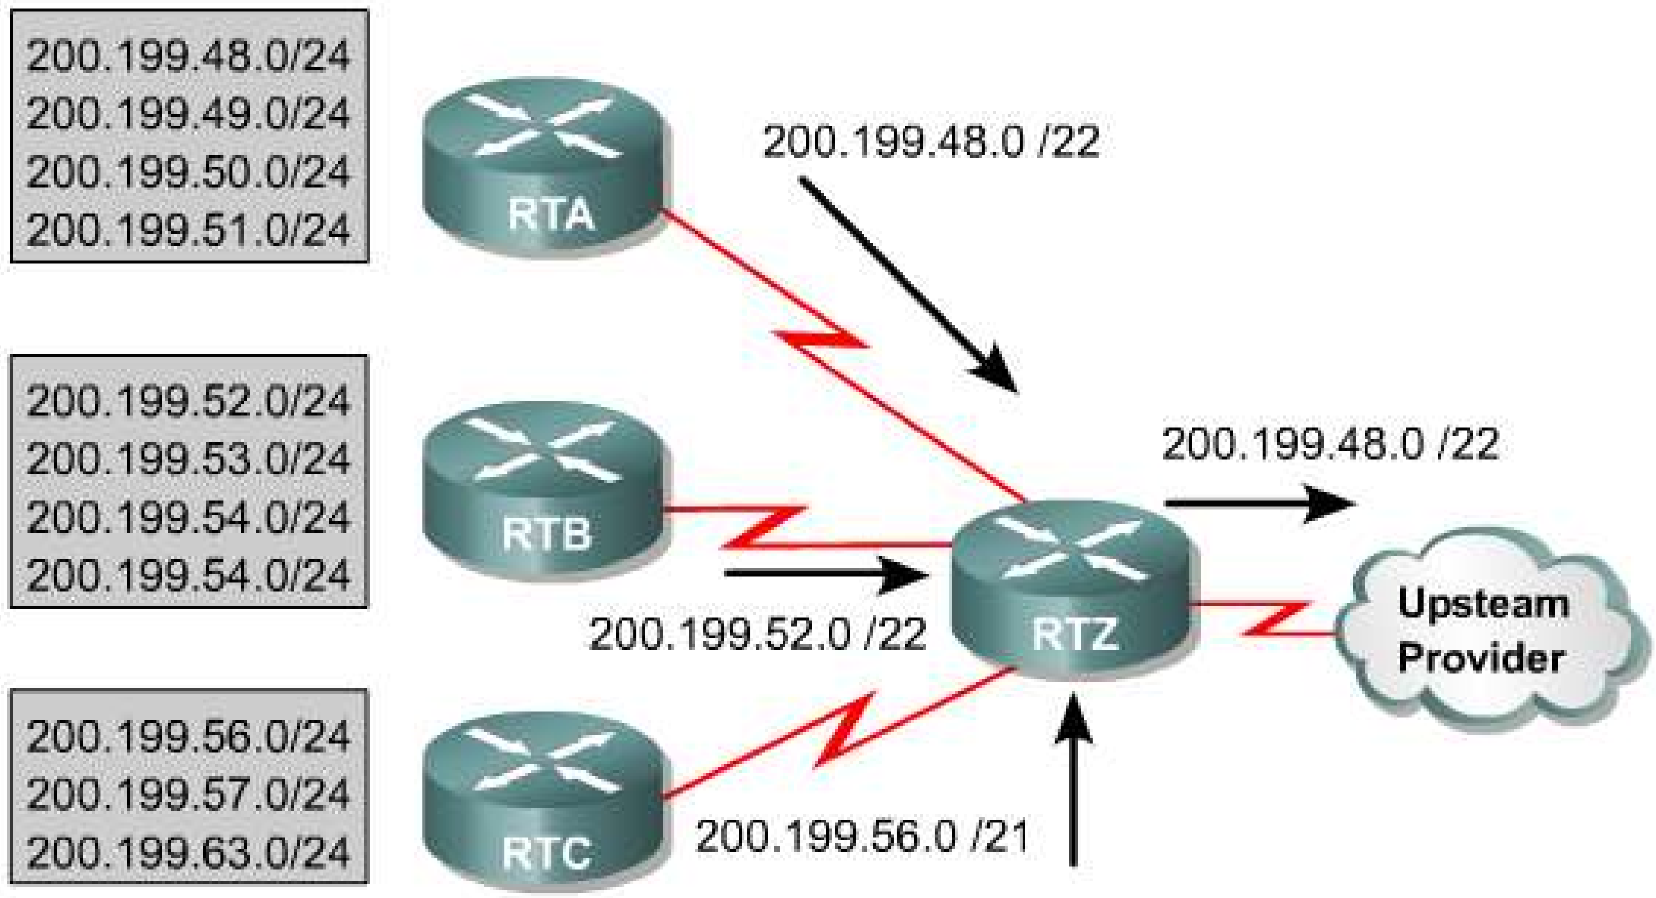

Route summarization reduces routing table size by aggregating routes to multiple networks into one supernet.

RIP History

- RIPv1 has the following limitations:
  - It does not send subnet mask information in its updates.
  - It sends updates as broadcasts on 255.255.255.255.
  - It does not support authentication.
  - It is not able to support VLSM or classless interdomain routing (CIDR).

RIPv2 Features

Feature	Description
Transmits subnet mask with route	Enables VLSM by passing the mask along with each route so that the subnet is exactly defined.
Provides authentication	Both clear text and/or MD5
Includes a next-hop route IP address in its routing update	A router can advertise a route and direct any listeners to a different router on the same subnet (if the other router has a better route).
Uses external route tags	RIP can pass information about routes learned from an external source and redistributed into RIP. This is used to separate RIP routes from externally learned routes.
Provides multicast routing updates	Instead of sending updates to 255.255.255.255, the destination IP address is 224.0.0.9. This reduces the amount of processing required on non-RIP speaking hosts on a common subnet.

Comparing RIPv1 and RIPv2

RIP v1	RIP v2
Easy to configure	Easy to configure
Only supports classful routing protocol	Supports use of classless routing
No subnet information with the routing update	Sends subnet mask information with the routing updates
Does not support prefix routing - all the devices in the same network must use the same subnet mask.	Supports prefix routing - different subnets within the same network can have different subnet masks (VLSM)
No authentication in updates	Provides for authentication in its updates
Broadcasts over 255.255.255.255	Multicasts routing updates over the Class D address 224.0.0.9 - makes it more efficient

OSPF
(Open Shortest Path First)

OSPF

- **OSPF is a link state protocol that is based on open standards.**
- **Among the other non-proprietary routing protocols such as RIP, OSPF is preferred due to its scalability.**
- **OSPF's superior scalability is due to its hierarchical design which is achieved through the use of multiple OSPF areas.**

OSPF Issues

- **OSPF offers many advantages over RIP, such as faster convergence, no hop count, routing updates which do not impact bandwidth through the use of multicast, support of VLSM, scalability, grouping of routers in areas versus a flat topology.**
- **The greatest disadvantage with OSPF is the degree of planning required for implementation, in other words, the configuration is more complex.**

Multicast addresses

- If a DR or BDR needs to send a link-state update, it will send it to all OSPF routers at 224.0.0.5.
- The DR and BDR have their own multicast address, 224.0.0.6. Non-DR/BDR routers send their LSUs to 224.0.0.6, or "all DR/BDR routers.

OSPF Operation

- **Being a link state protocol, OSPF merely advises neighbor routers of the status of their neighbor routers, or links. Thus, the entire network topology is known by each router via this link status information.**
- **Using the link status updates, each router may not have a complete routing table of the network, rather an entire link state database is built which is essentially a picture of the interconnection of all routers.**

OSPF Operation

- **All routers in a given area have identical link state databases. An OSPF area is a collection of routers and networks that have the same identification.**
- **Independently, each router runs the SPF algorithm which is merely a measure of cost, a value that is generally measured by bandwidth.**
- **Because all routers in given area have identical databases, DRs and BDRs are used to simplify the exchange of information between neighboring routers.**

OSPF States

- **Down State – no information has been exchanged between routers.**
- **Init State – The first Hello packet, generated approx every 10 seconds, has been received.**
- **Two-way State – Occurs when a router sees acknowledgement of itself in its neighbor’s hello packet.**

OSPF States

- **ExStart State – Characterized as an adjacency, but the routers have not established full adjacency yet, this state is where they determine which router is primary, master, and which is secondary, slave.**
- **Exchange State – Now that the routers have established the master/slave relationship, they describe their link-state databases to each other.**
- **Loading State – LSRs are now used to request the complete database information from neighbor routers. This information is sent via LSUs which contain the LSAs. LSUs are acknowledged with LSAcks**

OSPF States

- **<u>Full Adjacency</u> – Loading state is now complete and the routers now have established full adjacency. Each router keeps track of its adjacencies with its adjacencies database.**

Discovering Routes and reaching Full State

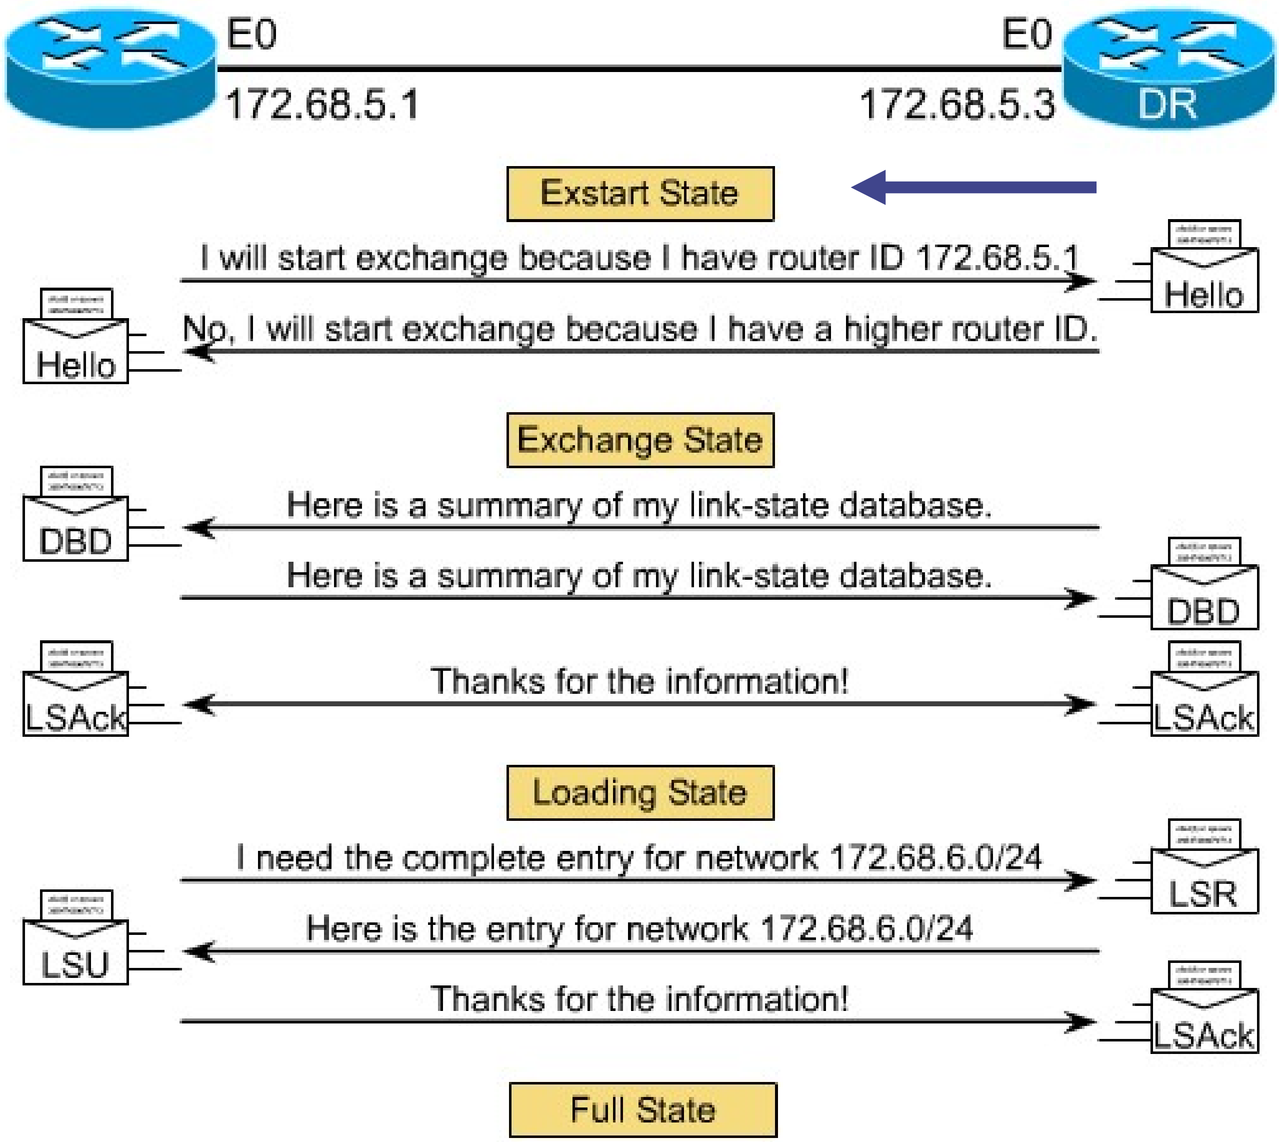

"adjacent"

OSPF Type-2 (DBD)

OSPF Type-2 (DBD)

OSPF Type-2 (DBD)

OSPF Type-2 (DBD)

OSPF Type-5 (LSAck)

OSPF Type-3 (LSR)

OSPF Type-4 (LSU)

OSPF Type-5 (LSAck)

"full adjacency"

OSPF Databases

- **Adjacency database:**

  List of all the neighbor routers with which two-way communication and adjacency have been established. This database is unique to each router.

- **Link-state database:**

  Identical between all routers within an area, this database shows the network topology.

- **Forwarding database:**

  This is the routing table that each router uses to determine best path.

OSPF Network Types

Broadcast multi-access:

This type of network has many routers connected to it, creating the need for the OSPF DR and BDR. Used with Token Ring, Ethernet and FDDI topologies.

Point-to-point:

With this type of network, only two routers exist, there is no need for a DR and BDR.

Non-broadcast multi-access (NBMA):

This type of network is found with Frame Relay, SMDS, X.25 and also needs the OSPF DR and BDR. These networks support many routers but have no broadcast capability.

Point-to-multipoint:

This type of network is generally created by a network administrator and does not need a DR and BDR.

DR and BDR

- **Found on every multi-access OSPF network, the Designated Router serves two functions:**
  1) **establish adjacency with every other router on the network**
  2) **to act as spokesperson for the network the DR sends and receives the LSAs.**
- **Because the DR can represent a single point of failure, the Backup Designated Router functions primarily for redundancy. It also must establish adjacencies with all other routers on the network and is the second focal point for LSAs.**

DR and BDR selection

- **The DR and BDR election process is handled via the OSPF Hello packet.**
- **Hello packets serve as the ballots in the DR/BDR elections. Each packet contains the router ID and a priority value.**
- **The router ID is the numerically highest IP address of all the interfaces on a router that operate within an OSPF area, including serial, Ethernet and loopbacks.**
- **A priority value can be applied to a router that will take precedence over any routing ID.**

DR and BDR selection

- **In a multi-access network, the router with the highest priority, if one exists – or the highest router ID if there is no priority value or if the priority values are tied - becomes the DR. The router with the second highest priority value (or router ID, if there is no priority value assigned) becomes the BDR. This is assuming there is no current DR and BDR.**
- **In the event the DR fails, the BDR becomes the DR and a new election is held to select a new BDR.**
- **There is no change in the DR and BDR status until a DR fails.**

Hello and Dead Intervals

- **Hellos are sent every 10 seconds by default on multi-access and point-to-point networks.**
- **On interfaces that connect to NBMA networks, such as Frame Relay, hellos are sent every 30 seconds, by default.**
- **The dead interval is four times the hello interval by default.**
- **All intervals are configurable.**

OSPF Cost

- Formula used to calculate OSPF cost is:

  100,000,000 / bandwidth value
- The lower the cost, the greater the desirability of the link.
- You can manually enter a cost value for a link. The cost can be between 1 and 65,535

*Note: the default bandwidth of a serial link with OSPF is 1.544 Mbps

NBMA networks

- **Two basic configurations are the full mesh and partial mesh. Used with frame relay, routers in a full-mesh topology have a PVC to every other router.**
- **Because full mesh is so expensive, a popular alternative is the partial mesh or hub and spoke configuration. However this creates a single point of failure, the hub router.**

OSPF Multi-Area Components

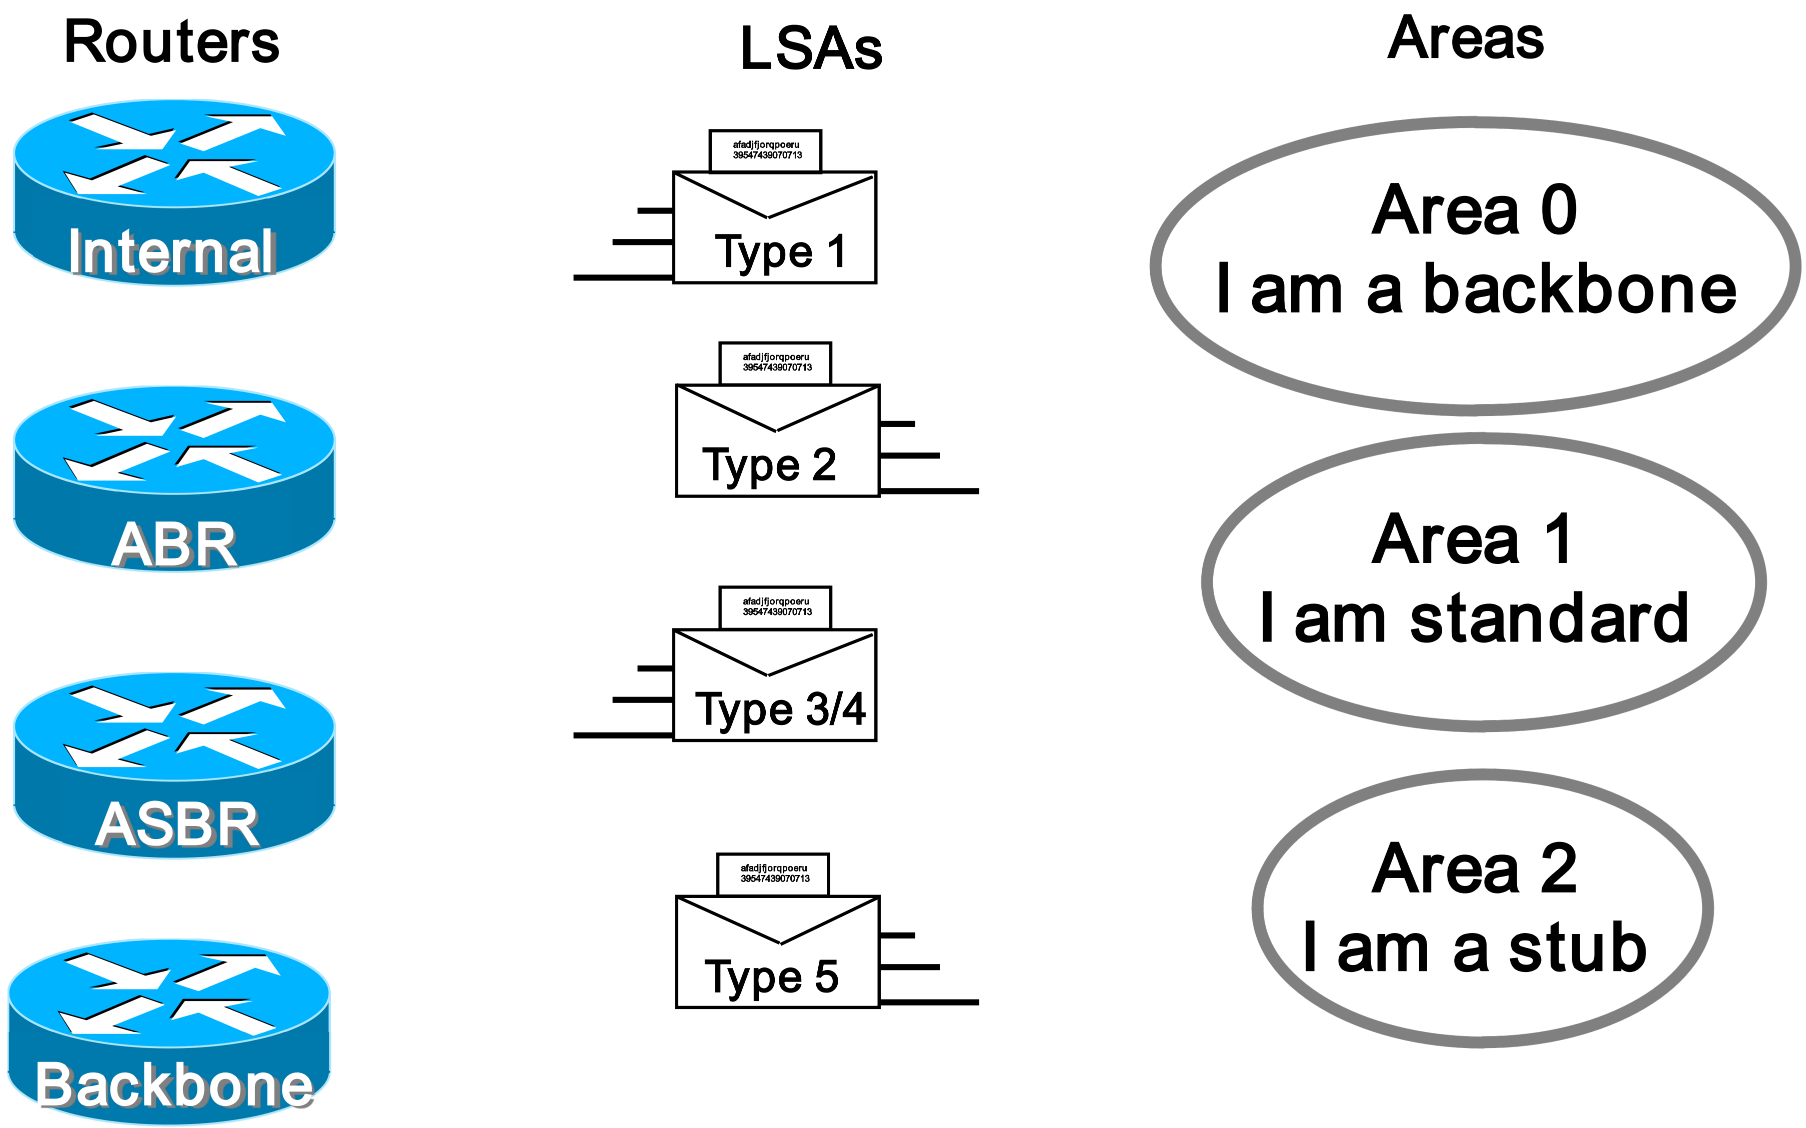

Types of OSPF Routers

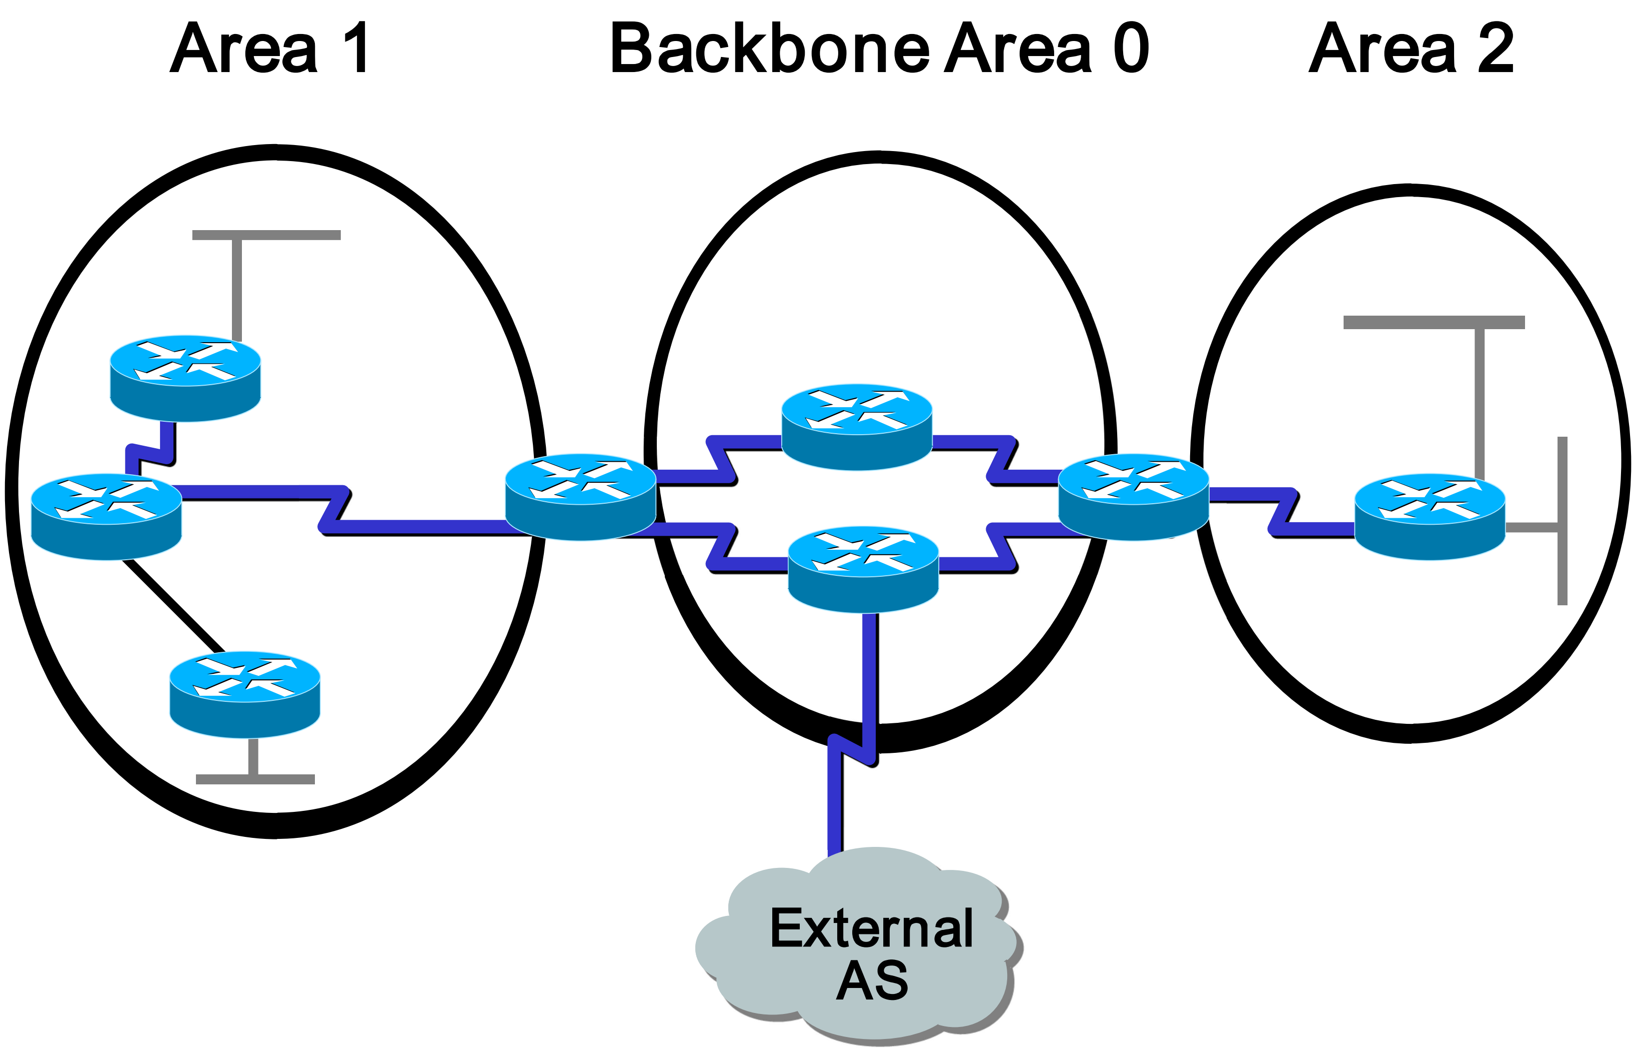

Types of OSPF Router - Internal router

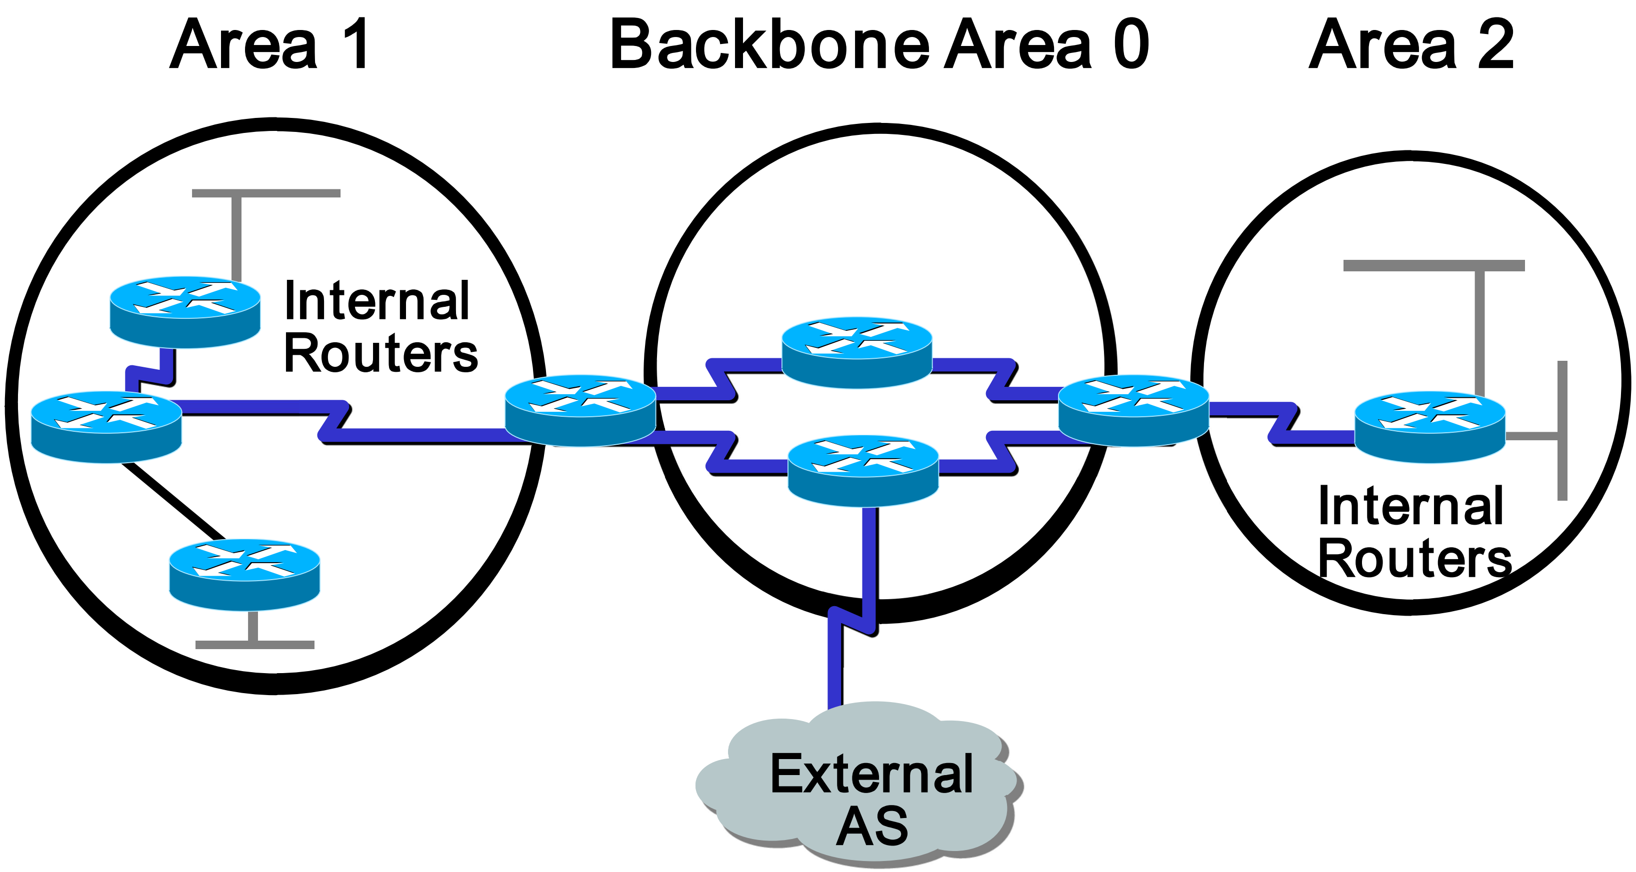

- All interfaces in the same area.
- Internal routers within the same area have:
  - Identical link-state databases
  - Run a single copy of the routing algorithm.

Types of OSPF Router – Backbone router

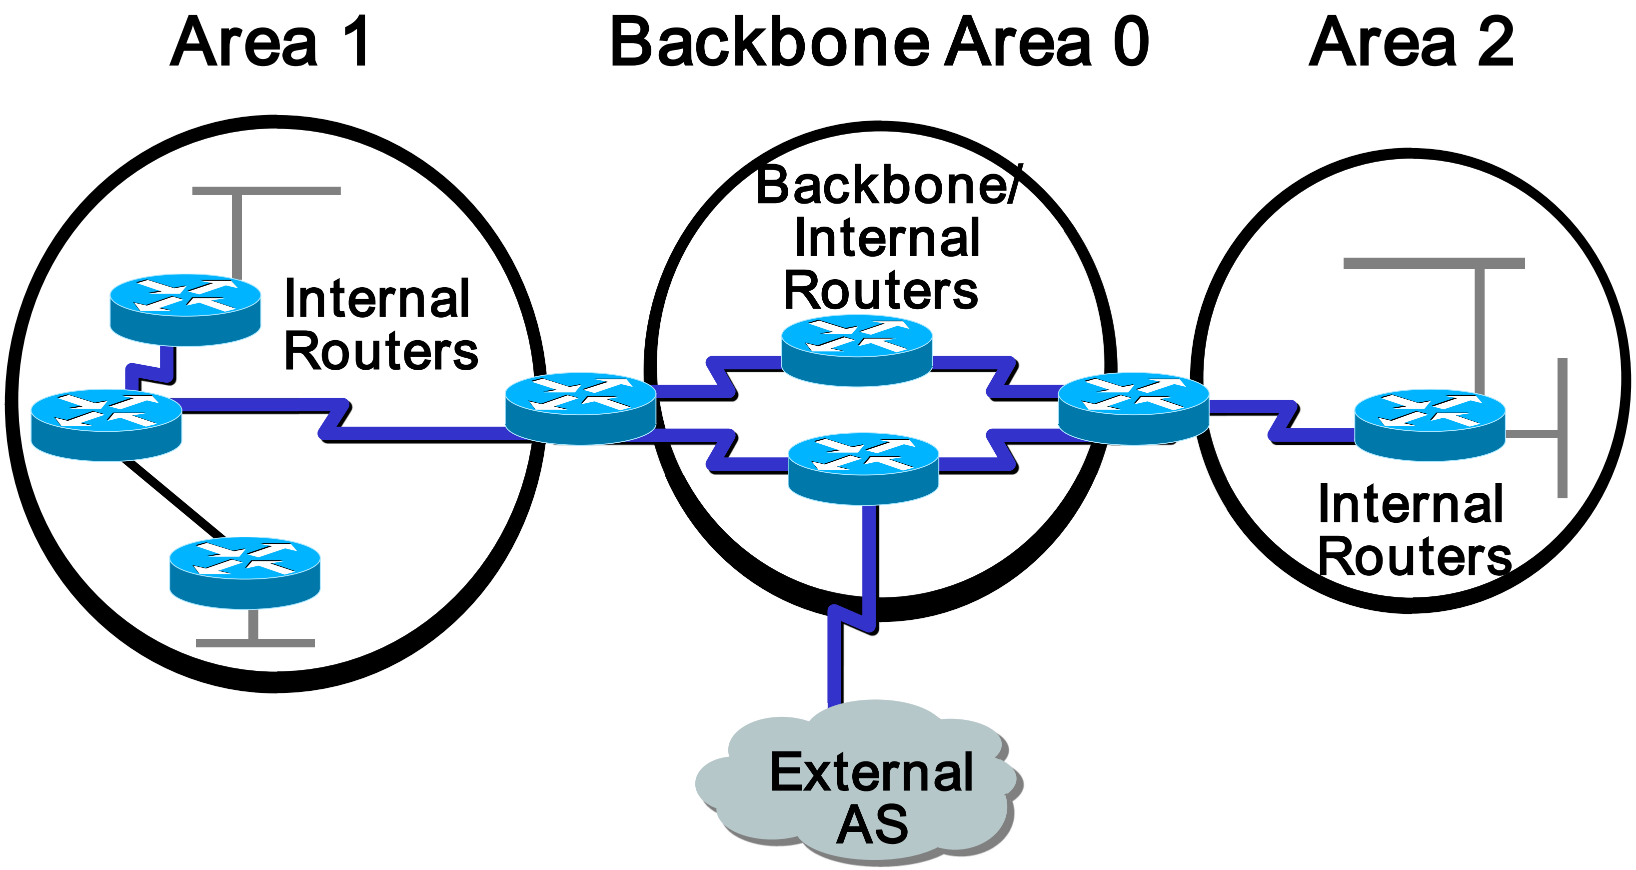

- Attached to the backbone area.
- Have at least one interface connected to Area 0

Types of OSPF Router - Area Border Router (ABR)

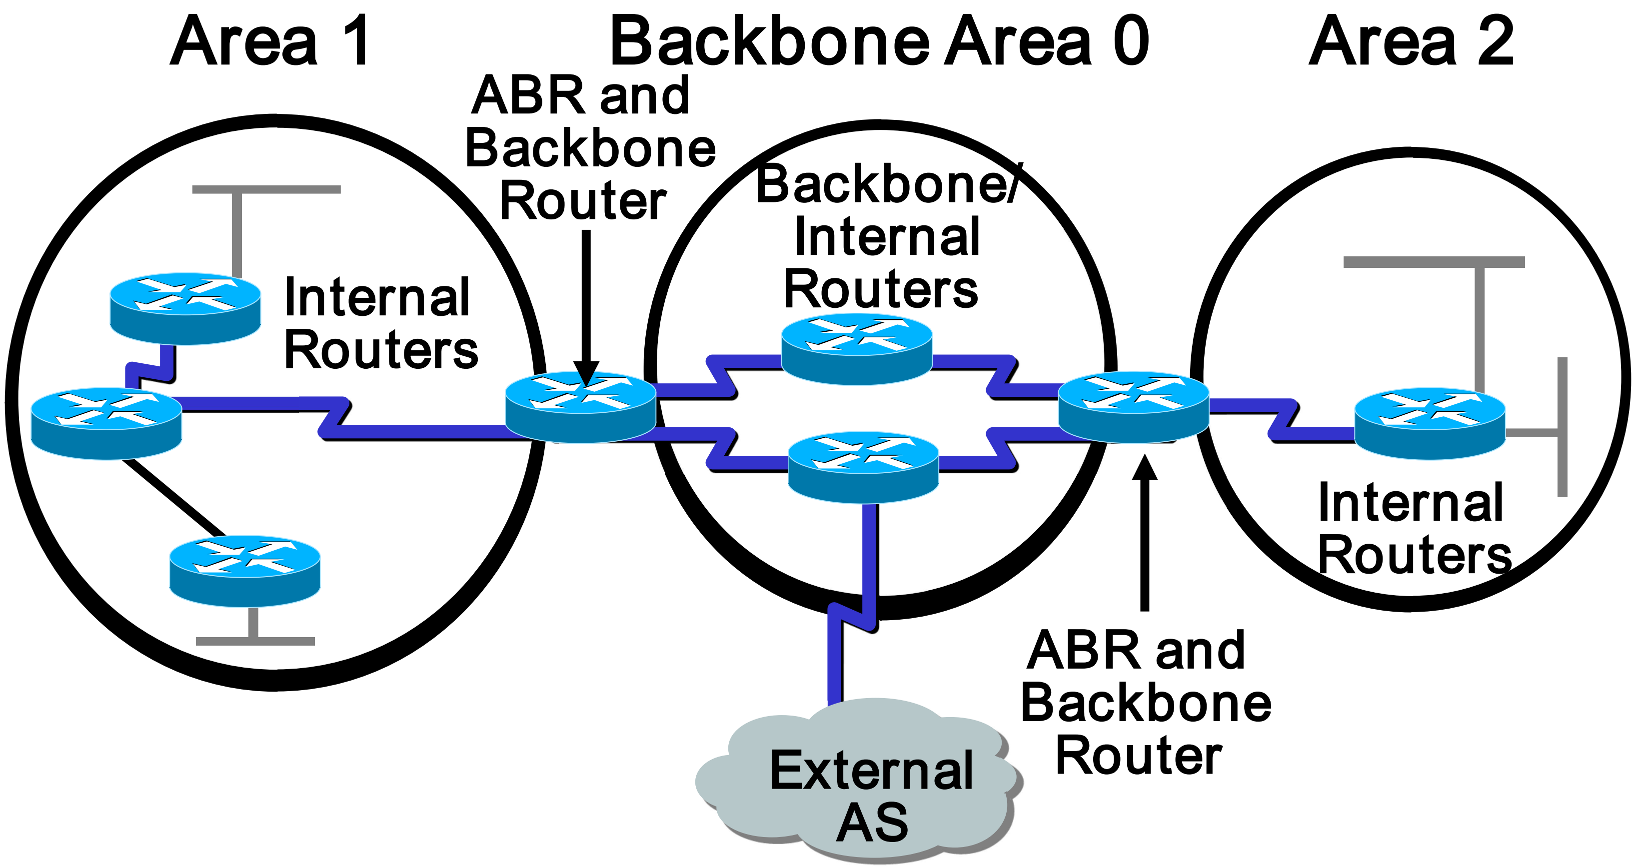

- Attached to multiple areas → area's exit point:
  - Summarize, Distribute & Routing information to/from another area via backbone area.
- Maintain separate LSDB for each connected area

Types of OSPF Router - Autonomous System Boundary Router (ASBR)

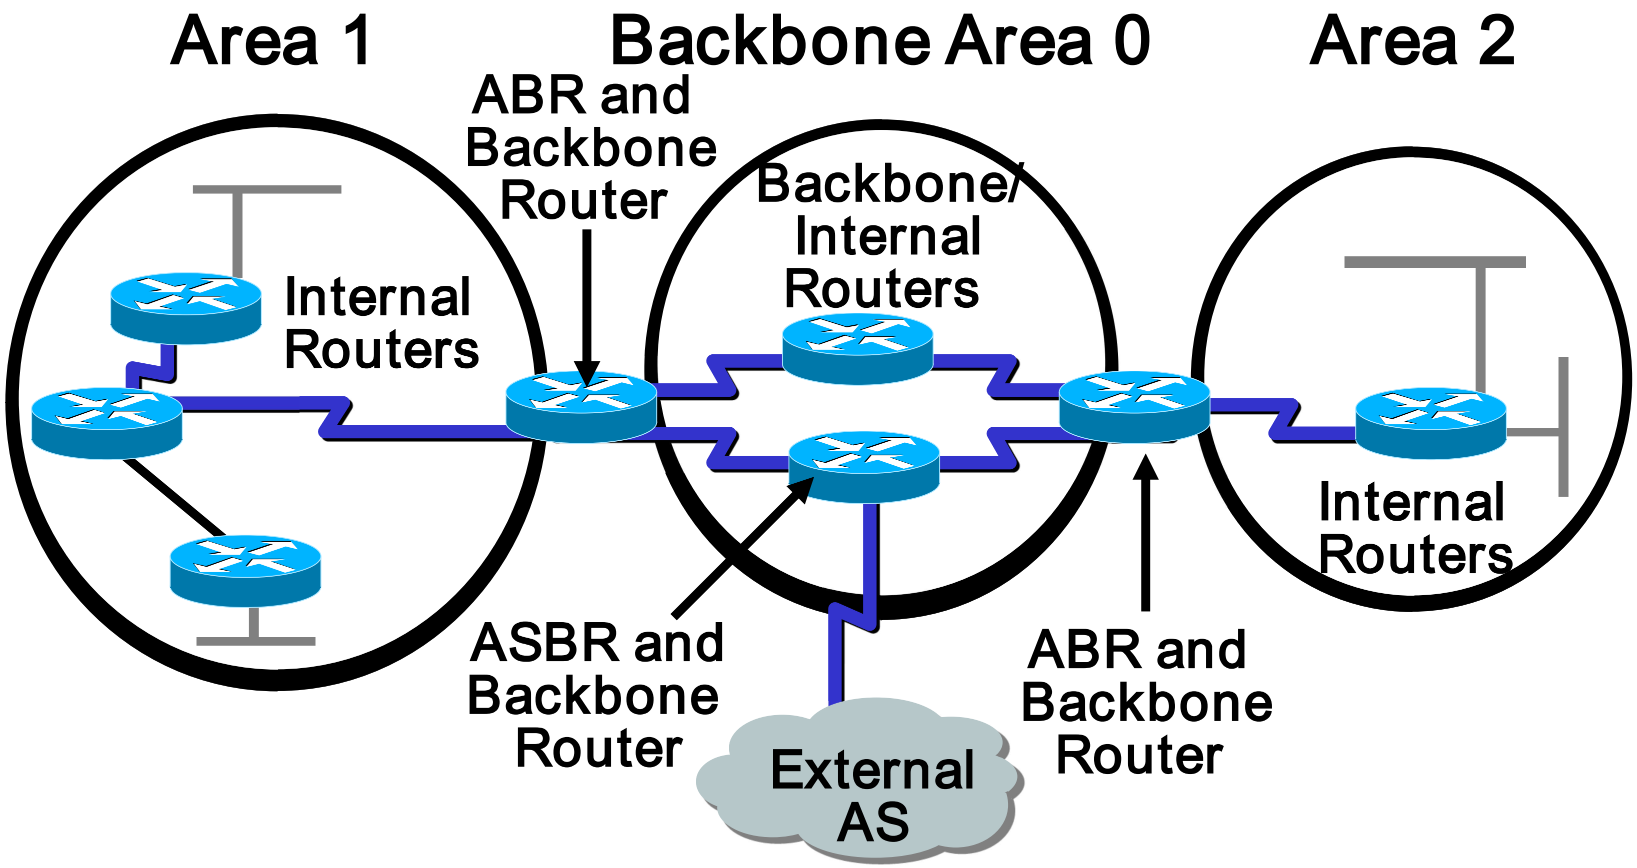

- Have at least one interface connected to another AS (external internetwork).
- Can import non-OSPF network information to the OSPF network, and vice versa (redistribution)

OSPF LSA All Types – RFC - 2328

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group membership
7	NSSA

OSPF LSA – Type 1

- Generated by each Router for each area it belongs to.
- Describes the states of the router's link to the area
- Flooded only within a particular area
- Descriptors:
  - Link status
  - Cost

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group Membership
7	NSSA

OSPF LSA – Type 2

- Generated by DR in multi-access networks
- Describes the set of routes attached to a particular network
- Flooded only within the area contain the network.

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group Membership
7	NSSA

OSPF LSA – Type 3

- Originated by ABRs
- Describes the links between the ABR and the internal routes of local area.
- Flooded throughout the backbone area to the other ABRs.

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group Membership
7	NSSA

OSPF LSA – Type 4

- Originated by ABR
- Describes a route to a destination outside the area still inside the AS.
- Flooded throughout the LSA's throughout the backbone area to the other ABRs except for totally stubby areas

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group Membership
7	NSSA

OSPF LSA – Type 5

- Originated by ASBRs
- Describe routes to destinations external to the AS.
- Flooded throughout an OSPF AS except for stub and totally stubby areas
- Default route for the AS can also be described by an AS-external-LSA
- Denoted : E1 or E2

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group Membership
7	NSSA

OSPF LSA – Type 5 - External LSA

External Type 1 - total cost

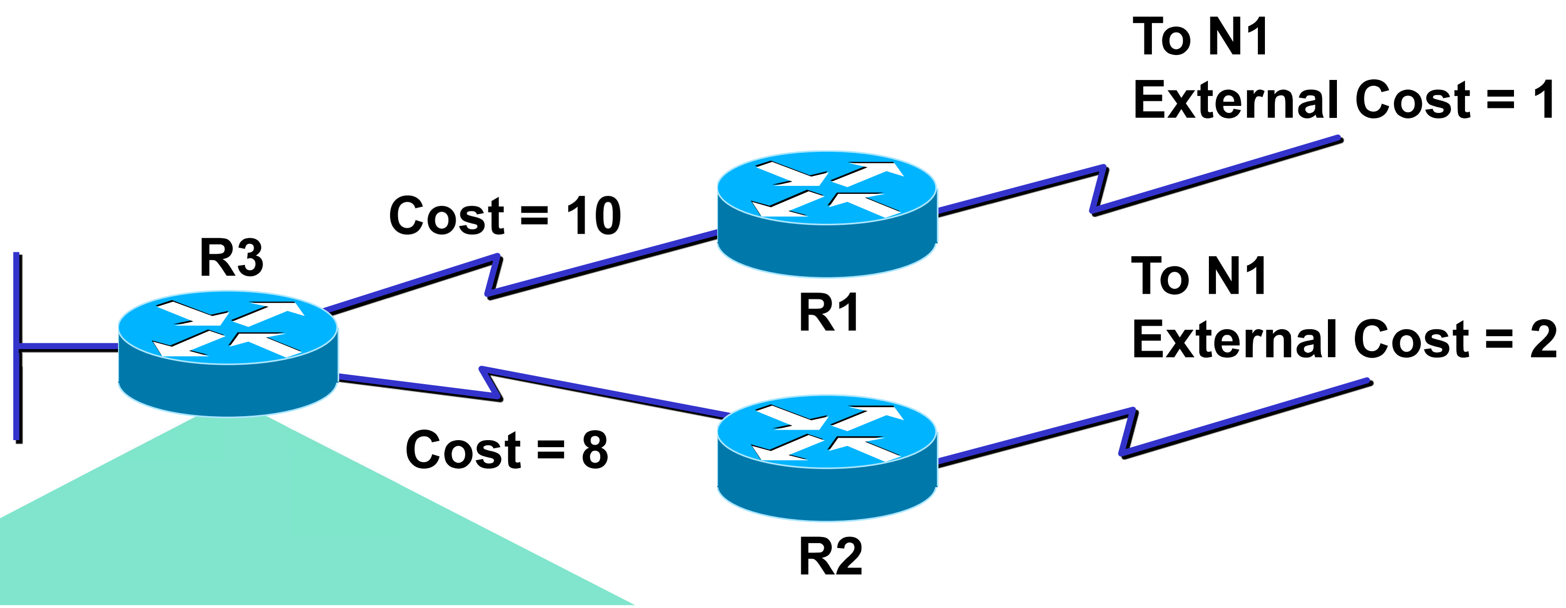

Network	Metric	Next Hop
N1	11	R1
N1	10	R2

← Selected Route

OSPF LSA – Type 5 - External LSA

External Type 2 - cost of the outgoing interface

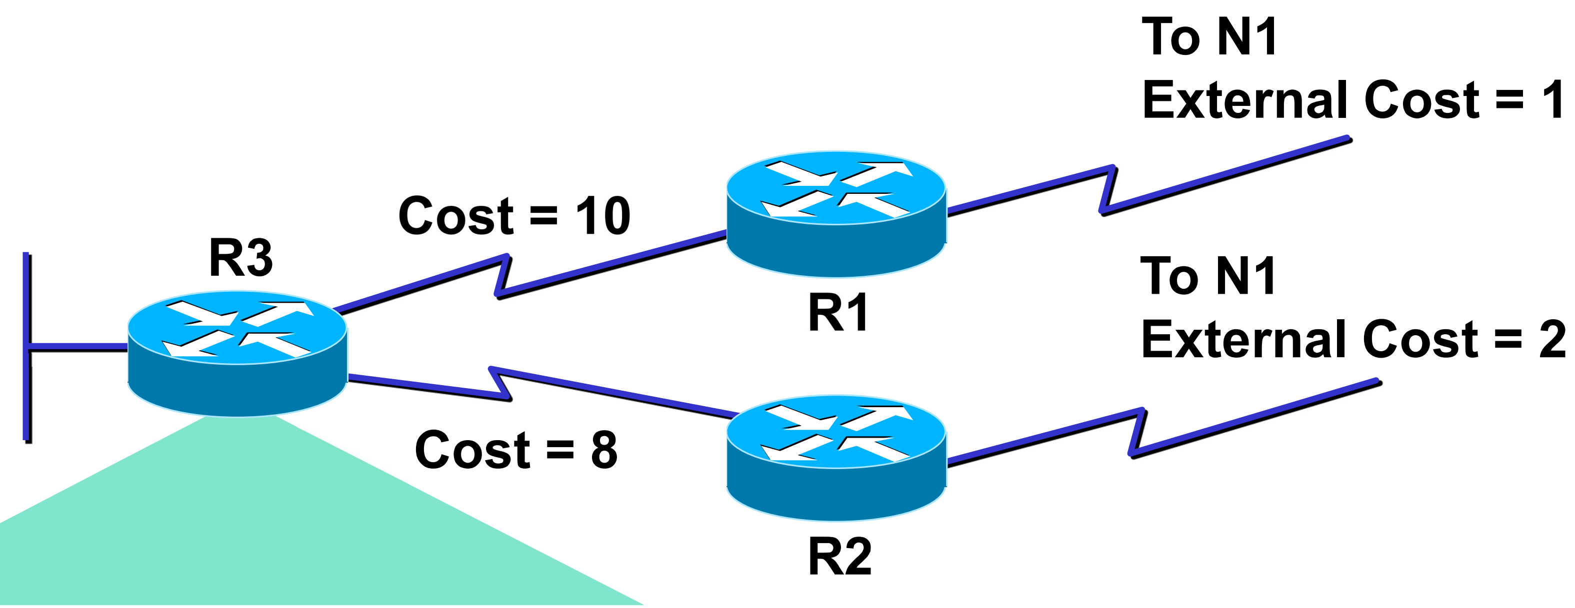

Network	Metric	Next Hop
N1	1	R1
N1	2	R2

OSPF LSA – Type 7

- Originated by ASBR connected to an NSSA (Not So Stubby Area).
- Provide for carrying external route information within an NSSA.
- Flooded throughout NSSAs and translated into Type 5 messages by ABRs.
- RFC-1587

Type	LSA Name
1	Router Link
2	Network Link
3	Summary Network
4	Summary ASBR
5	External
6	Group Membership
7	NSSA

LSA Type 1 – Router Link

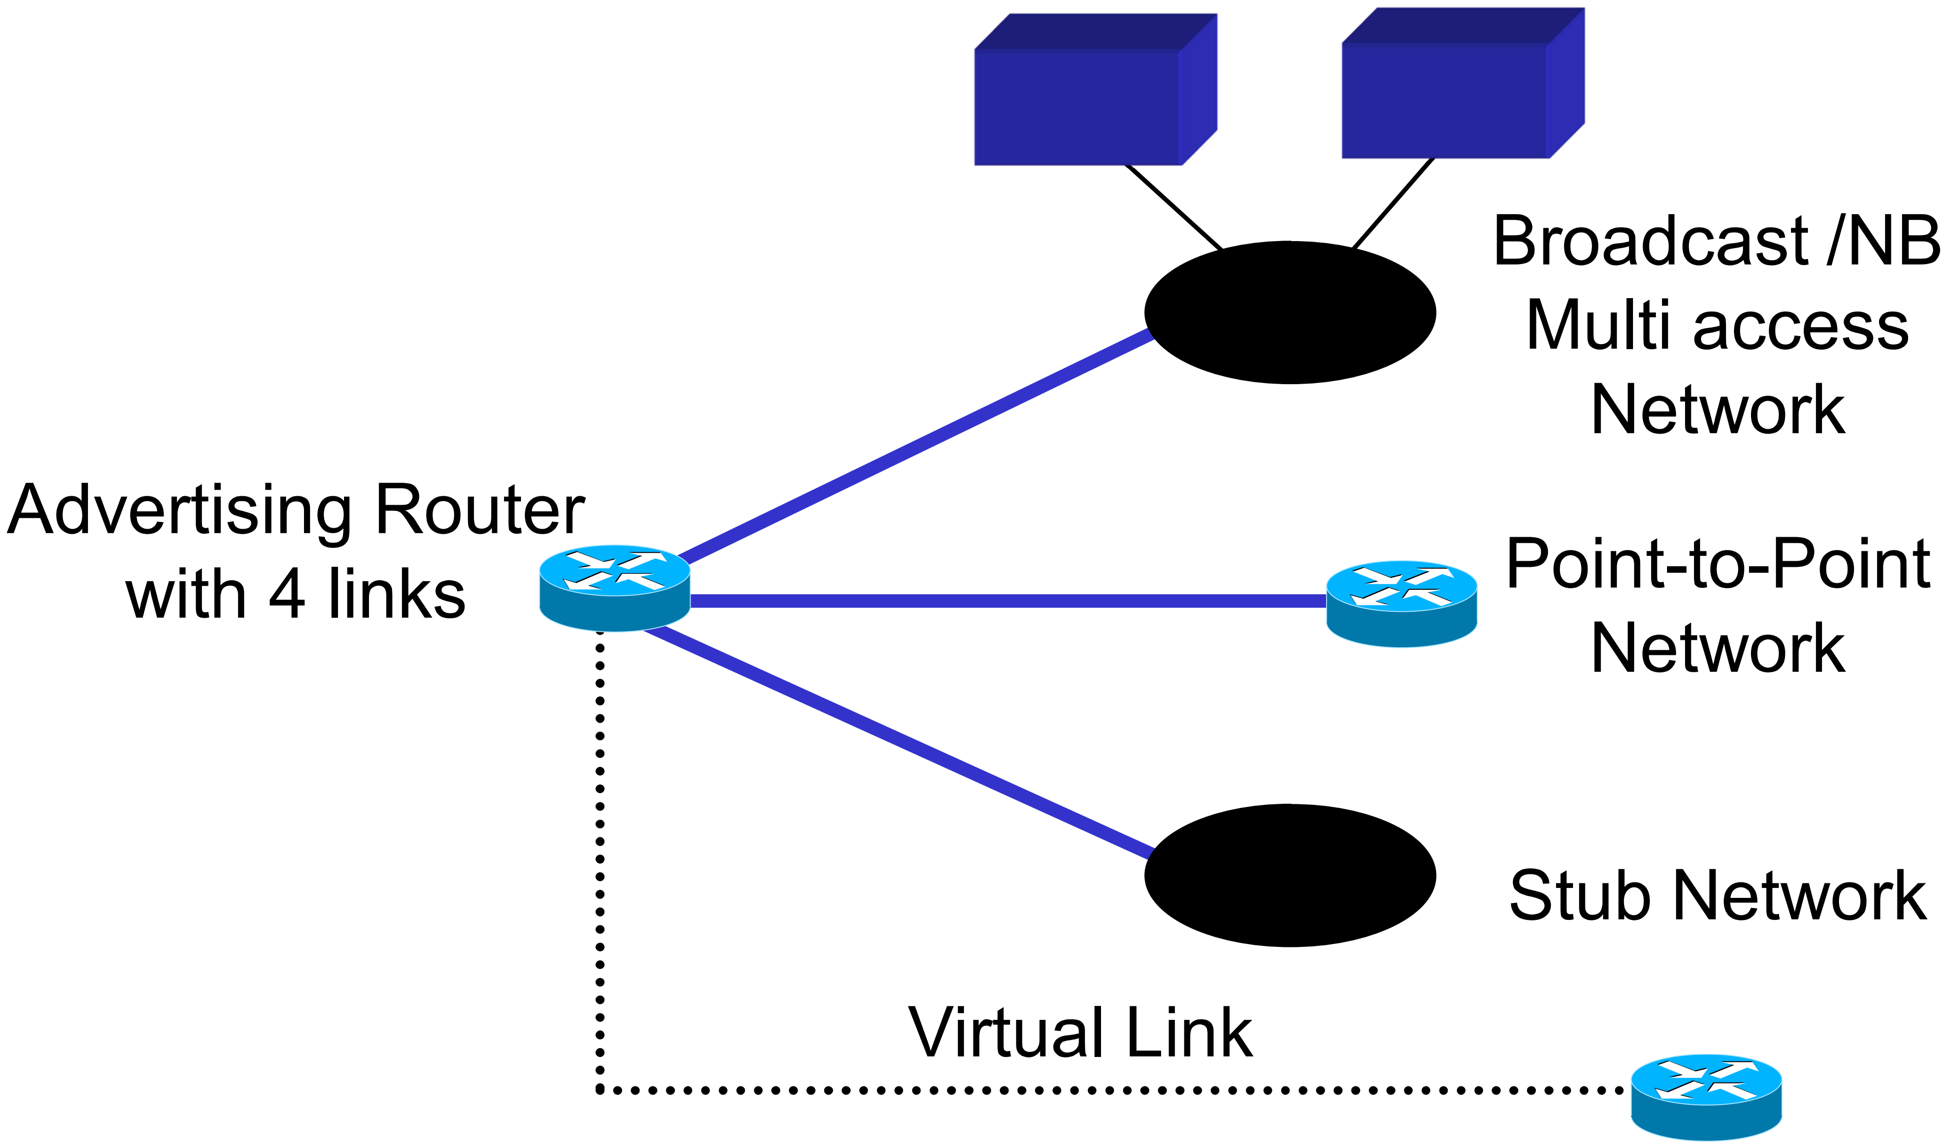

LSA Type 2 – Network Link

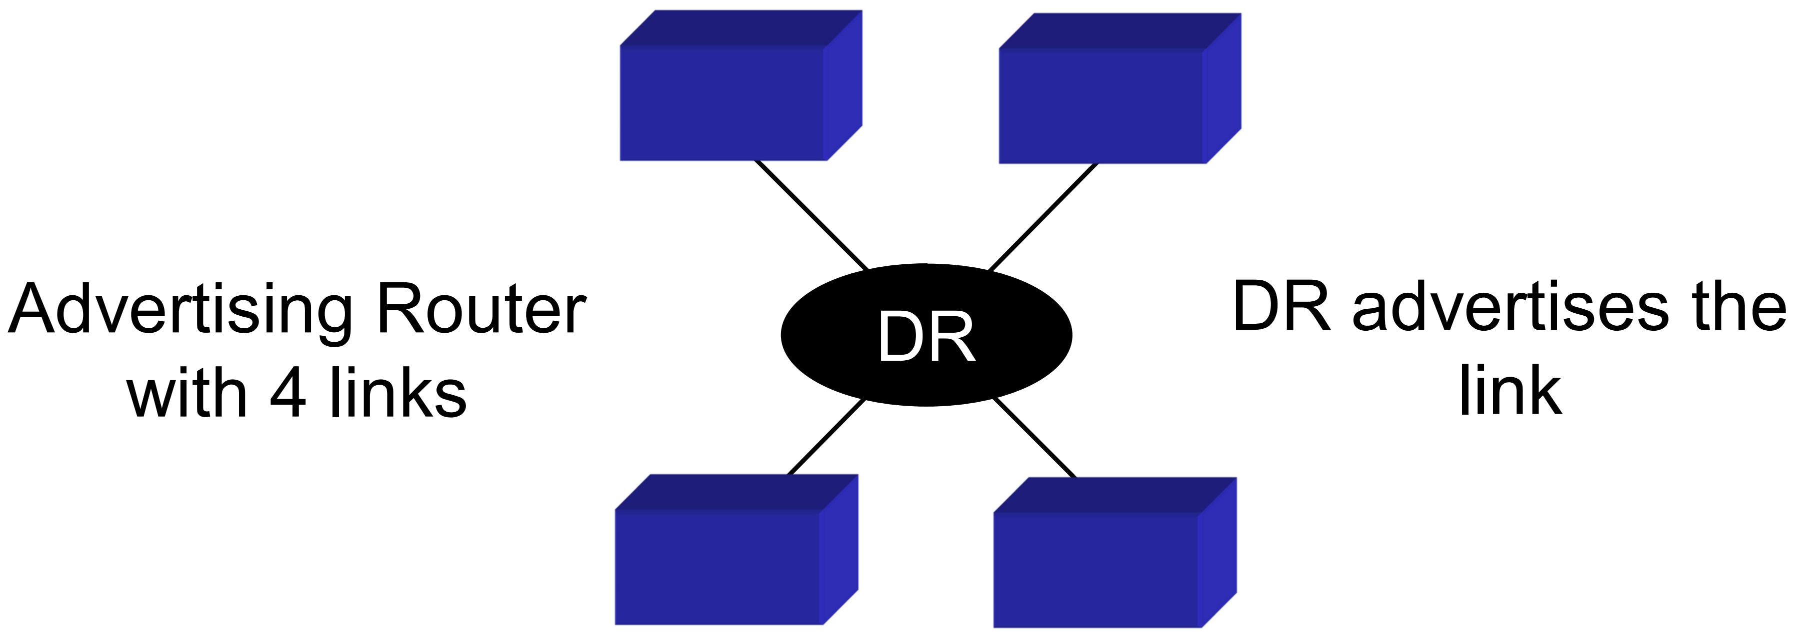

LSA Type 3 – Summary Link to Network

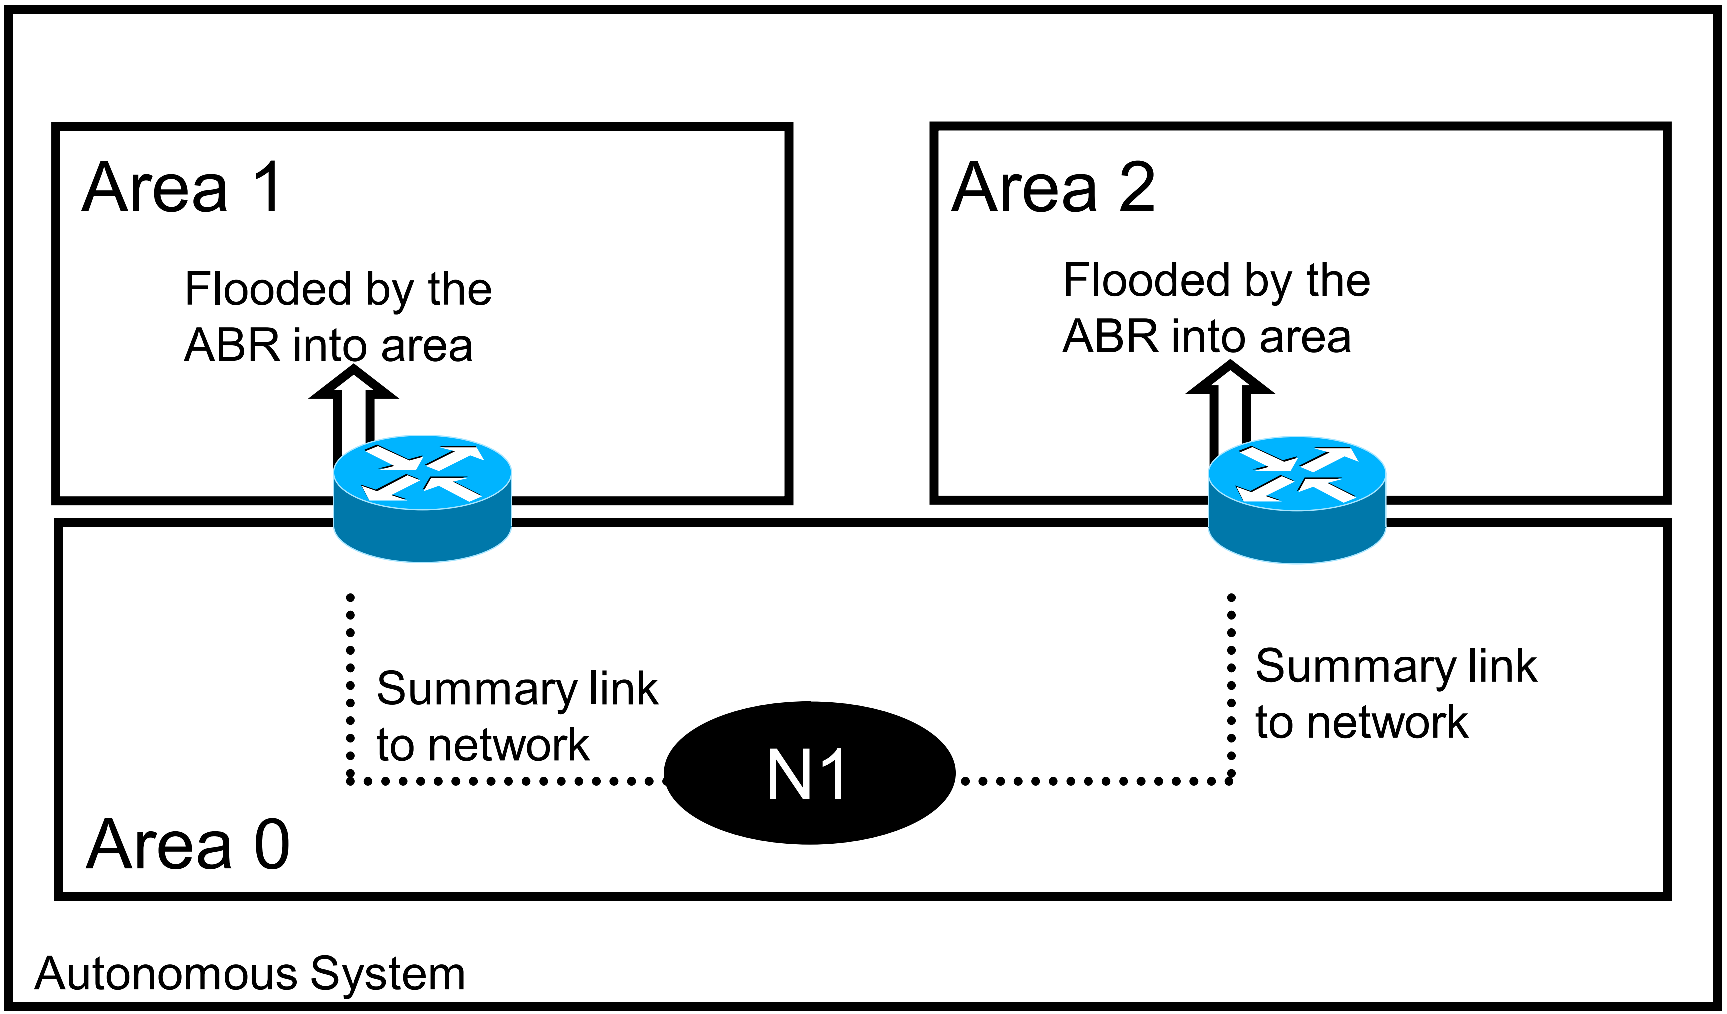

LSA Type 4 – Summary Link to ASBR

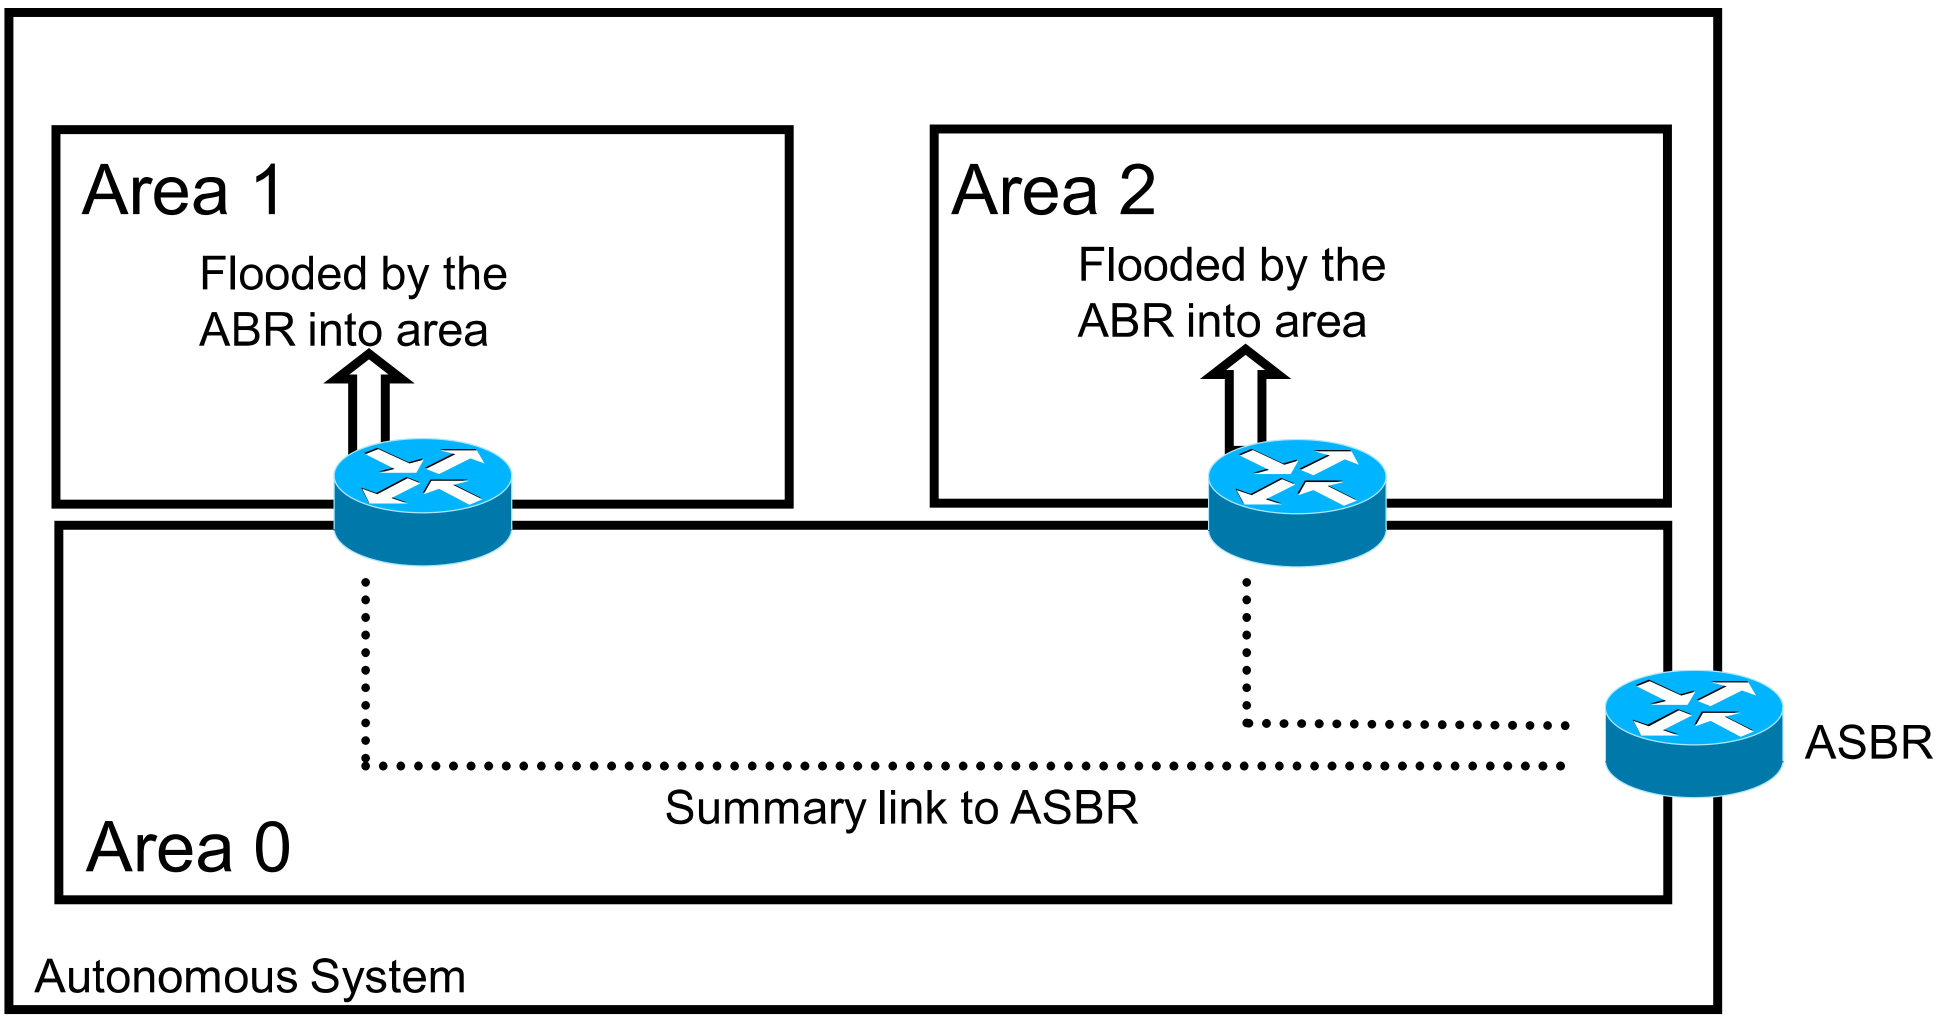

LSA Type 5 – AS External link

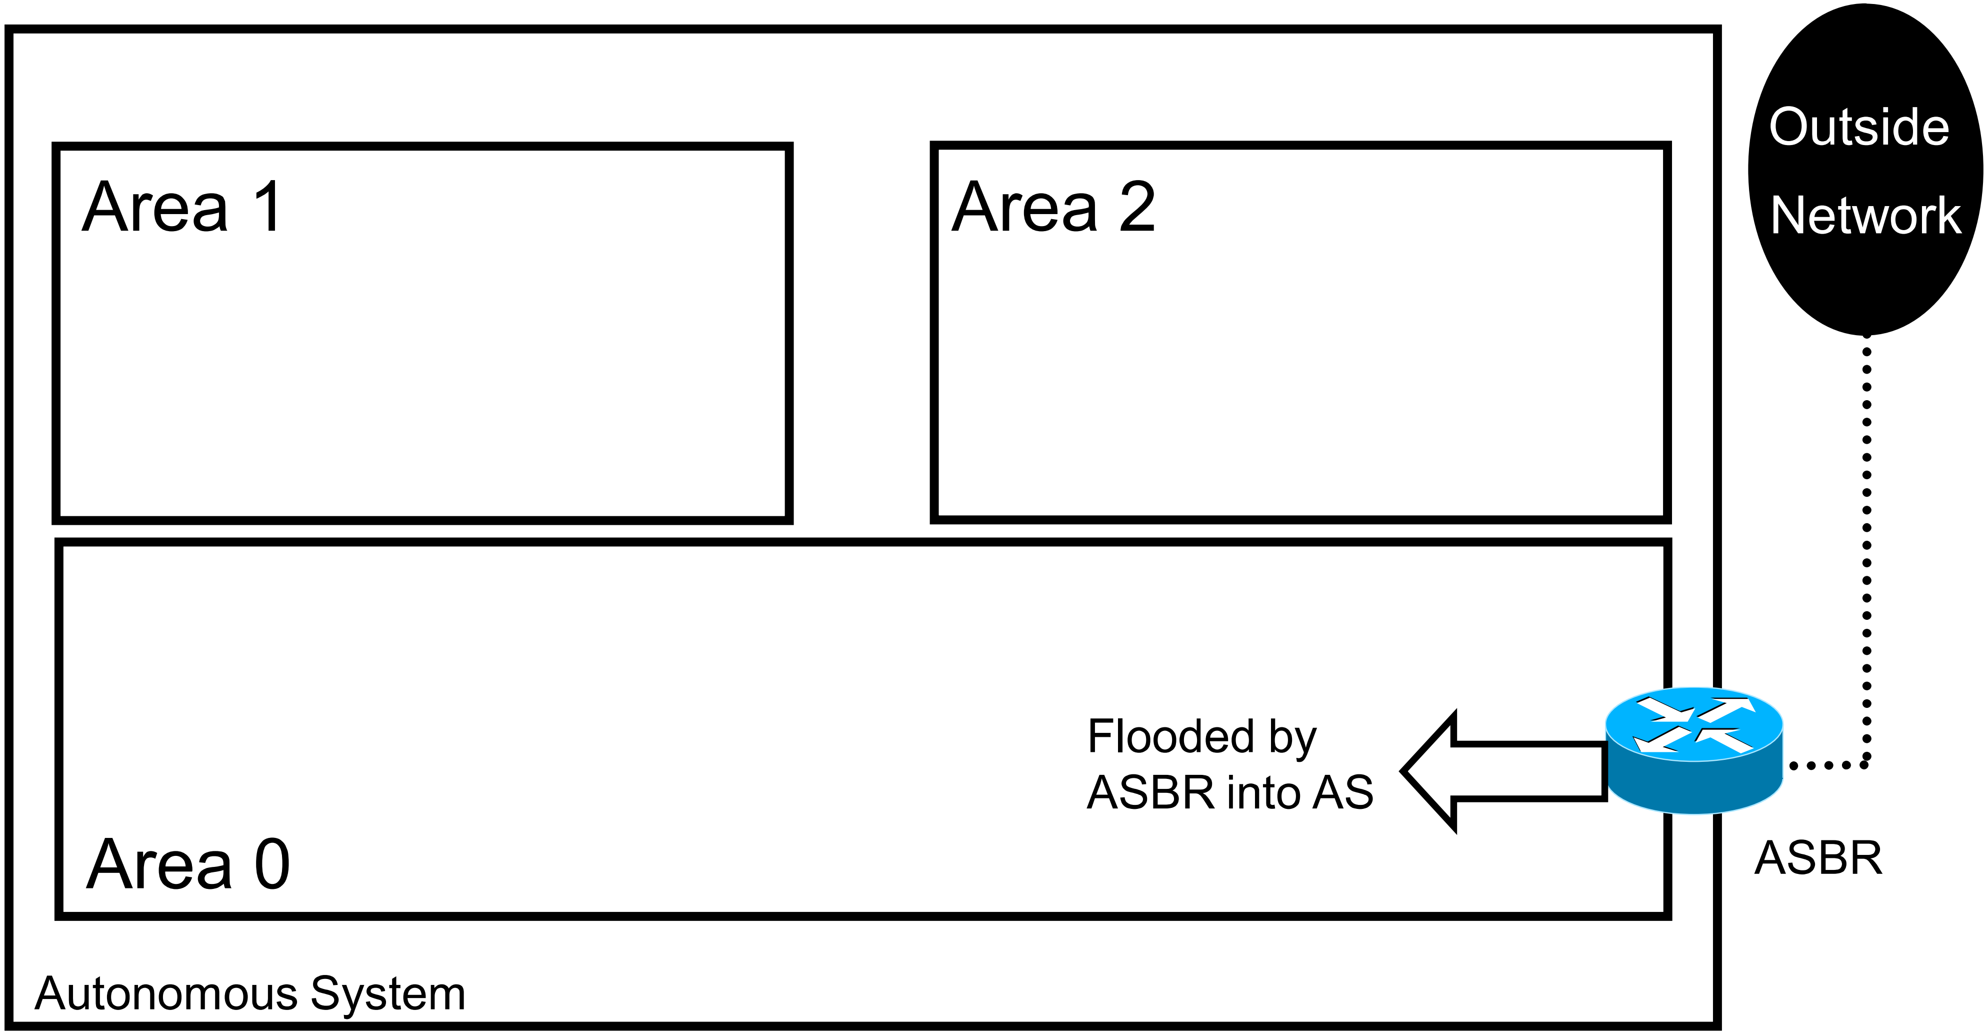

OSPF Area Types

- Standard area:
  - A standard area can accept link updates and route summaries.
- Backbone area (transit area):
  - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect.
  - The backbone area is always Area 0.
  - All other areas must connect to this area to exchange route information.
  - The OSPF backbone has all the properties of a standard OSPF area.
- Stub area:
  - A stub area is an area that does not accept information about routes external to the autonomous system (that is, the OSPF internetwork), such as routes from non-OSPF sources.
  - If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.

OSPF Area Types

- Totally stubby area (Cisco):
  - A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system.
  - Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route.
- Not-so-stubby area (NSSA):
  - An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

OSPF Area Types

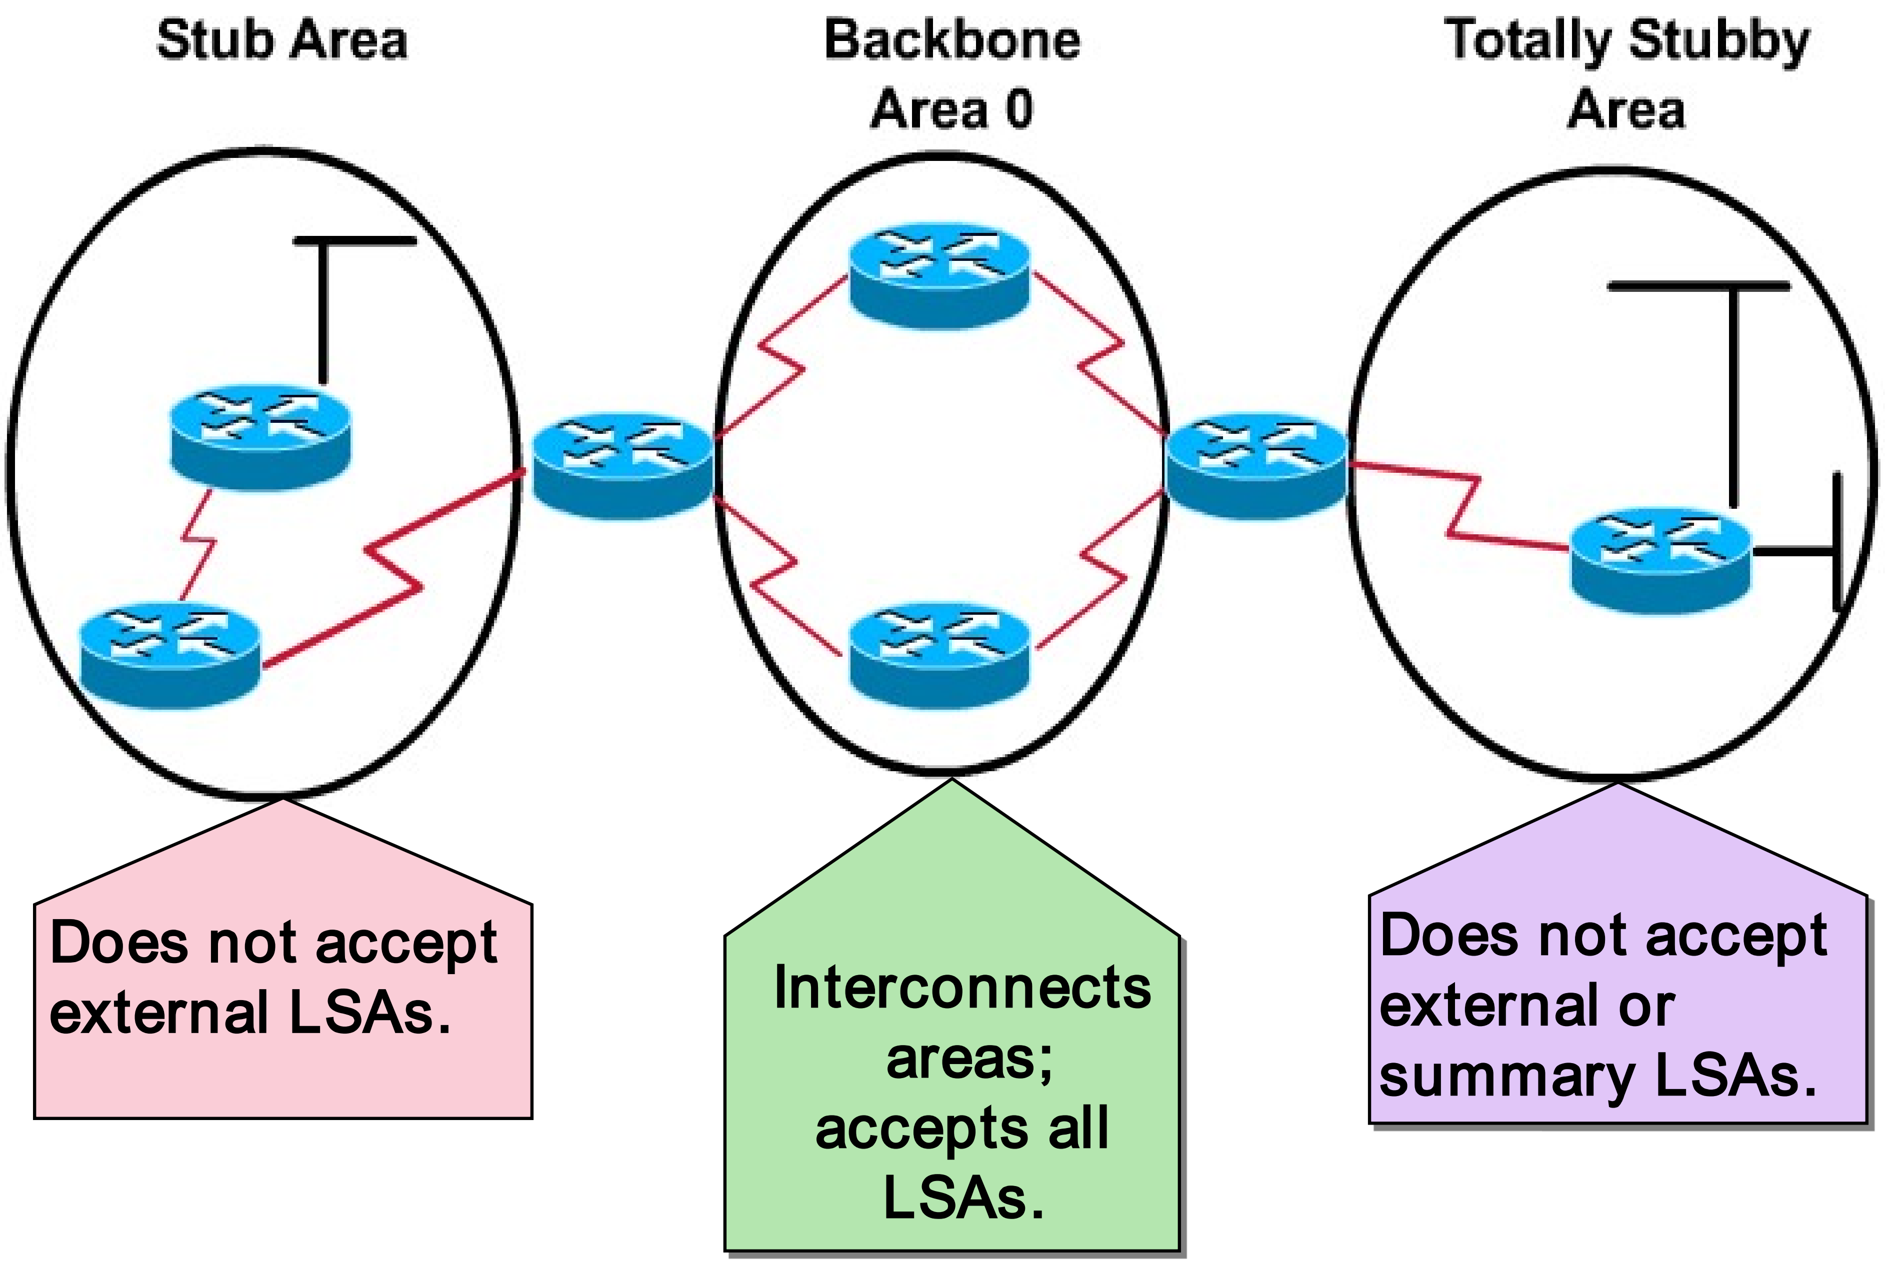

Location of Different LSAs

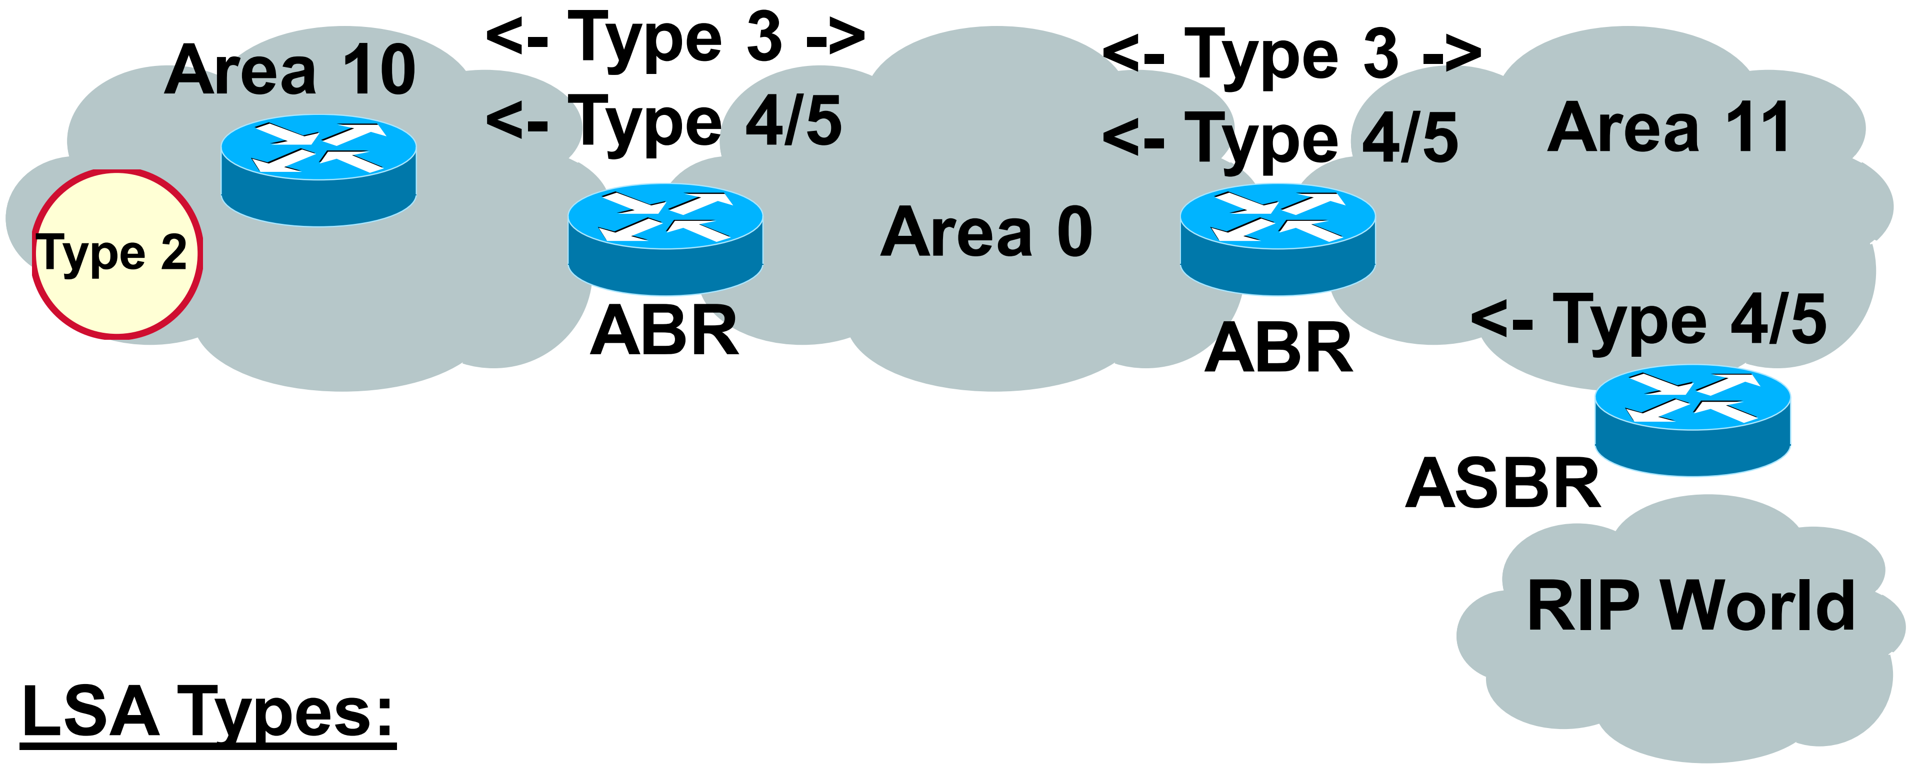

LSA Types:

1) Router
2) Network
3) Summary
4) ASBR Summary
5) External

Supporting Route Summarization

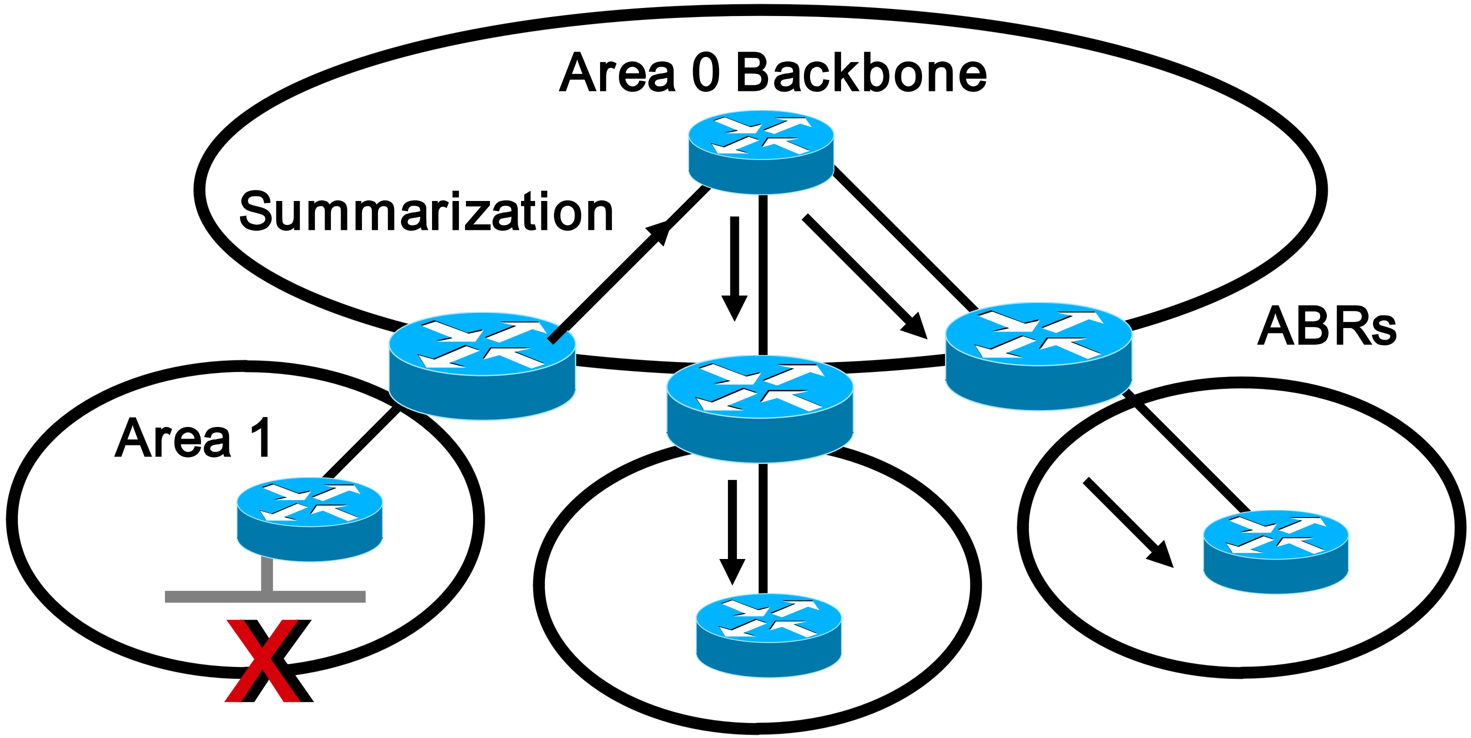

- Minimizes number of routing table entries
- Localizes impact of a topology change
- Reduces LSAs and saves CPU

Using Route Summarization

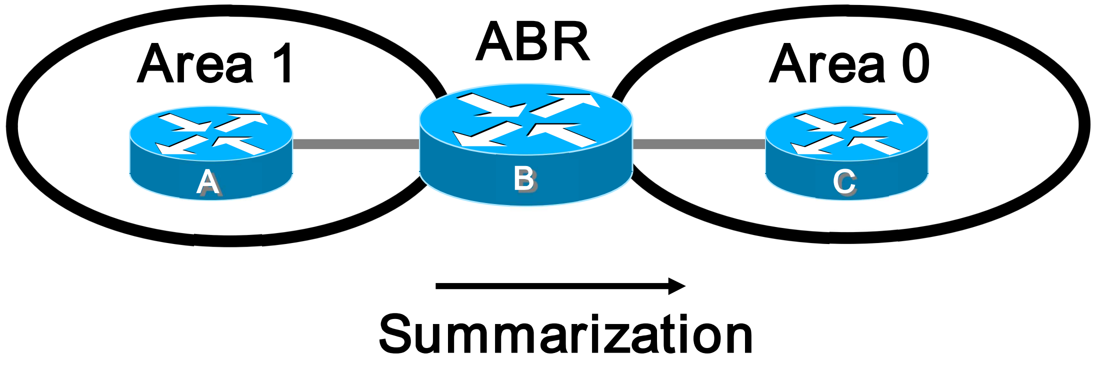

Routing Table for B

O 172.16.8.0	255.255.252.0
O 172.16.12.0	255.255.252.0
O 172.16.16.0	255.255.252.0
O 172.16.20.0	255.255.252.0
O 172.16.24.0	255.255.252.0
O 172.16.28.0	255.255.252.0

LSAs Sent to Router C

- IA 172.16.8.0 255.255.248.0
- IA 172.16.16.0 255.255.240.0

- InterArea (IA) summary link carries mask
- One entry can represent several subnets

Update Routing table

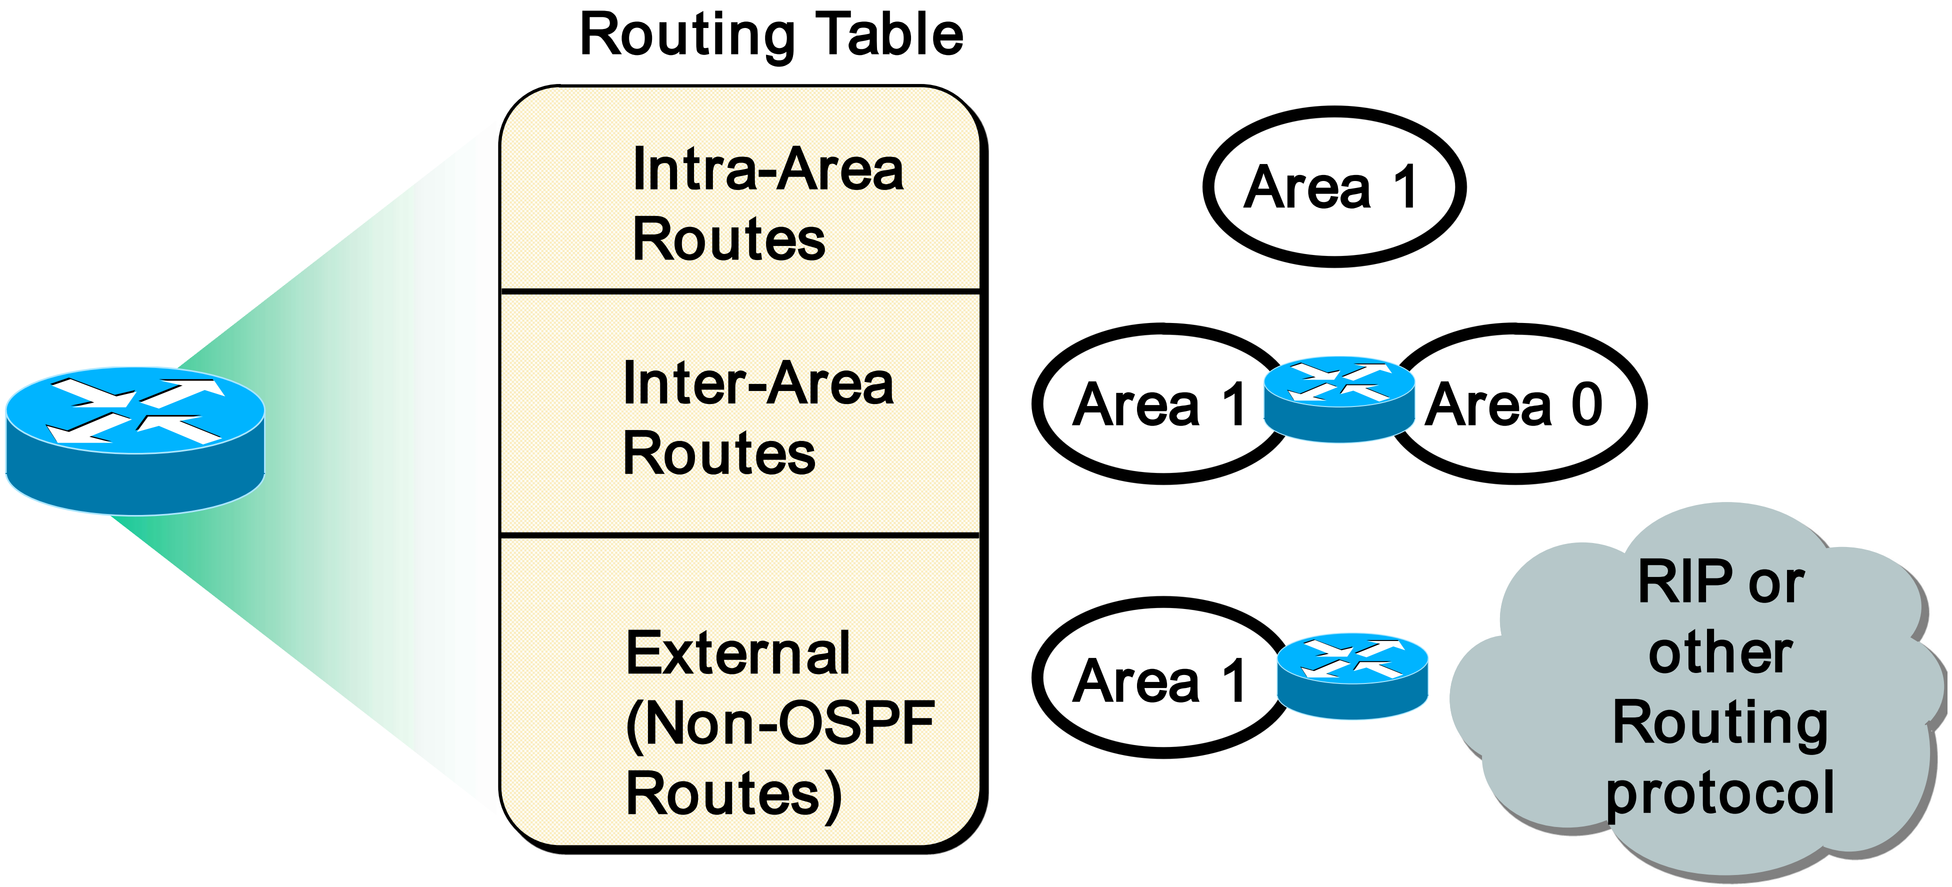

- Intra-area routes
- Inter-area routes
- External destination

Maintaining Routing Information

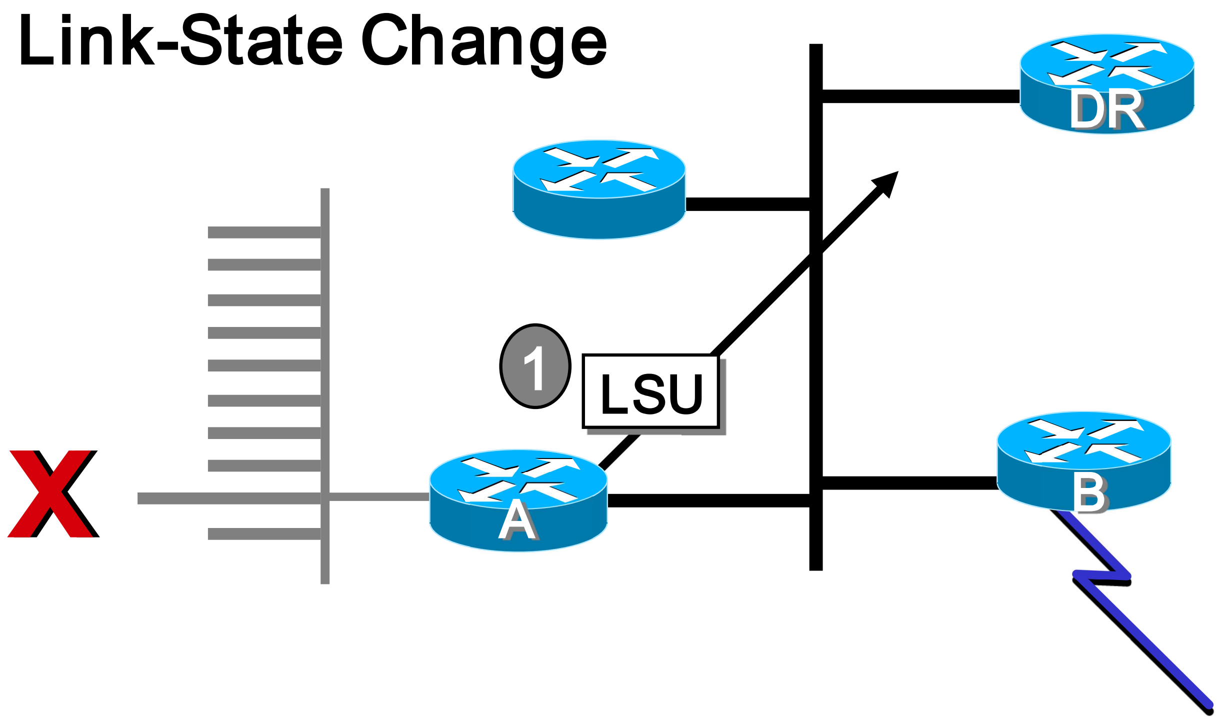

- Router A tells all OSPF **DRs** on **224.0.0.6**

Maintaining Routing Information

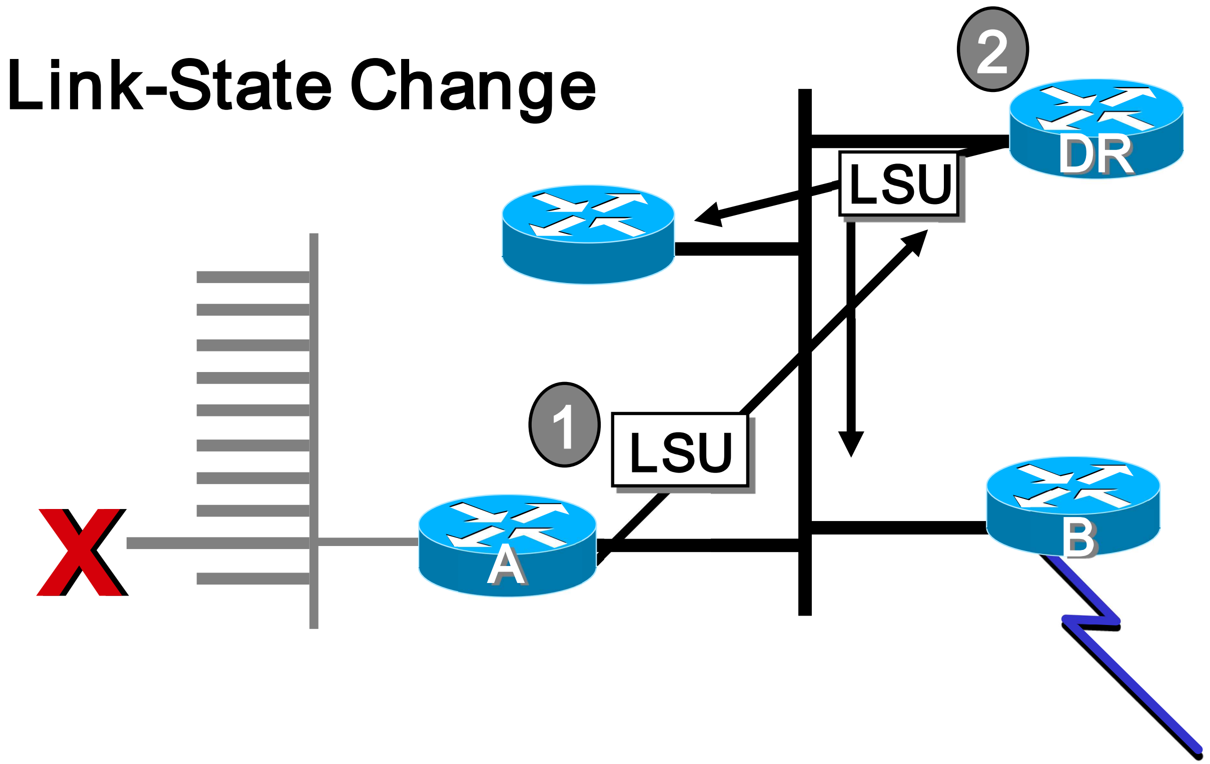

- Router A tells all OSPF DRs on 224.0.0.6
- DR tells all others on 224.0.0.5

Maintaining Routing Information

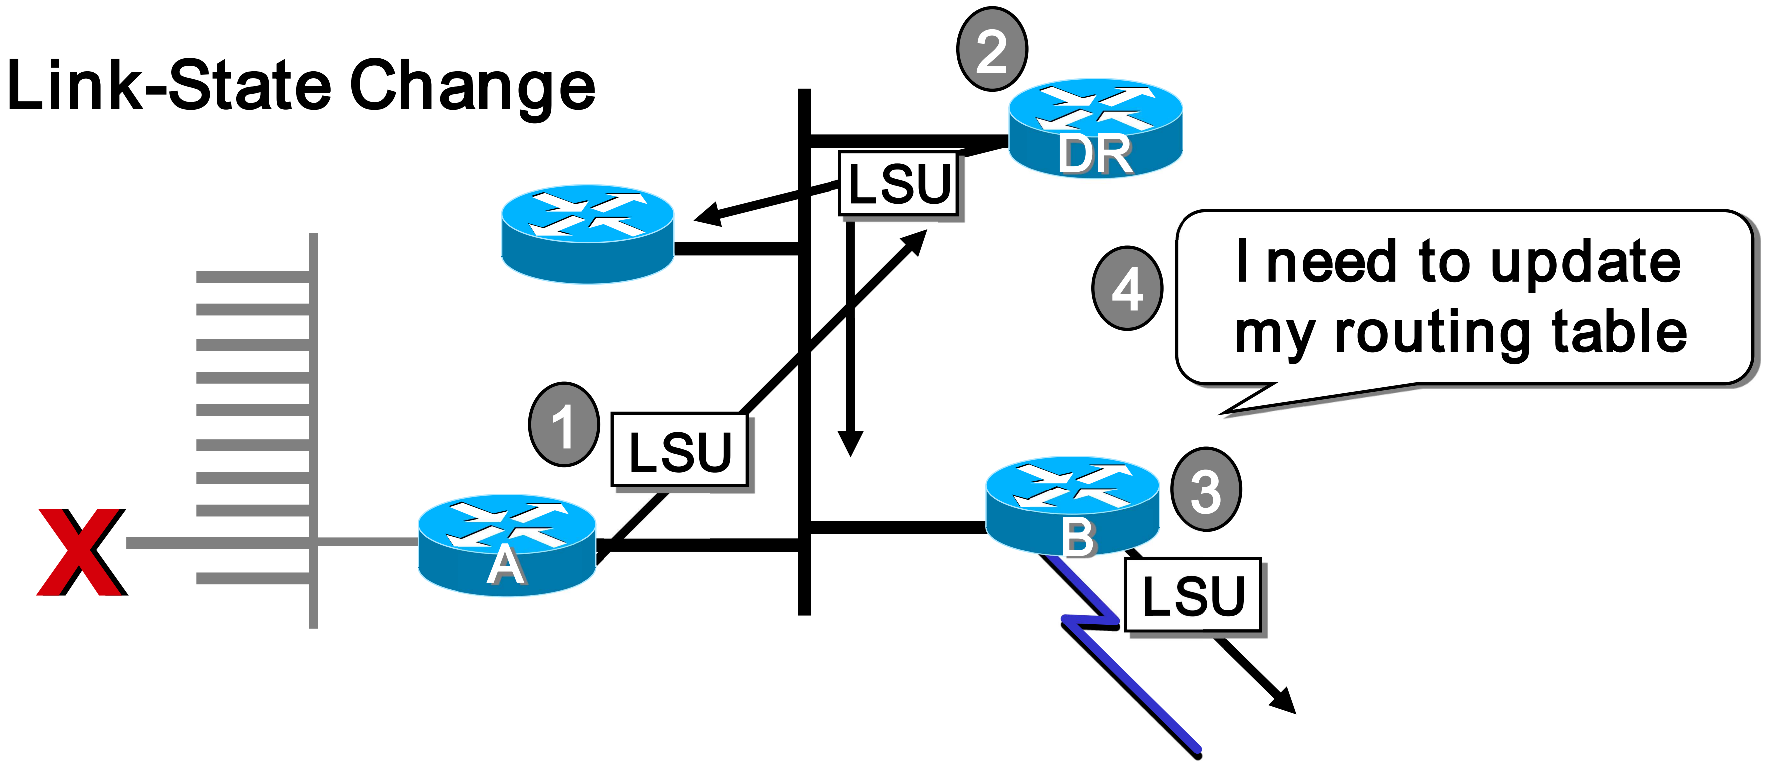

- Router A tells all OSPF DRs on 224.0.0.6
- DR tells all others on 224.0.0.5

Maintaining Routing Information

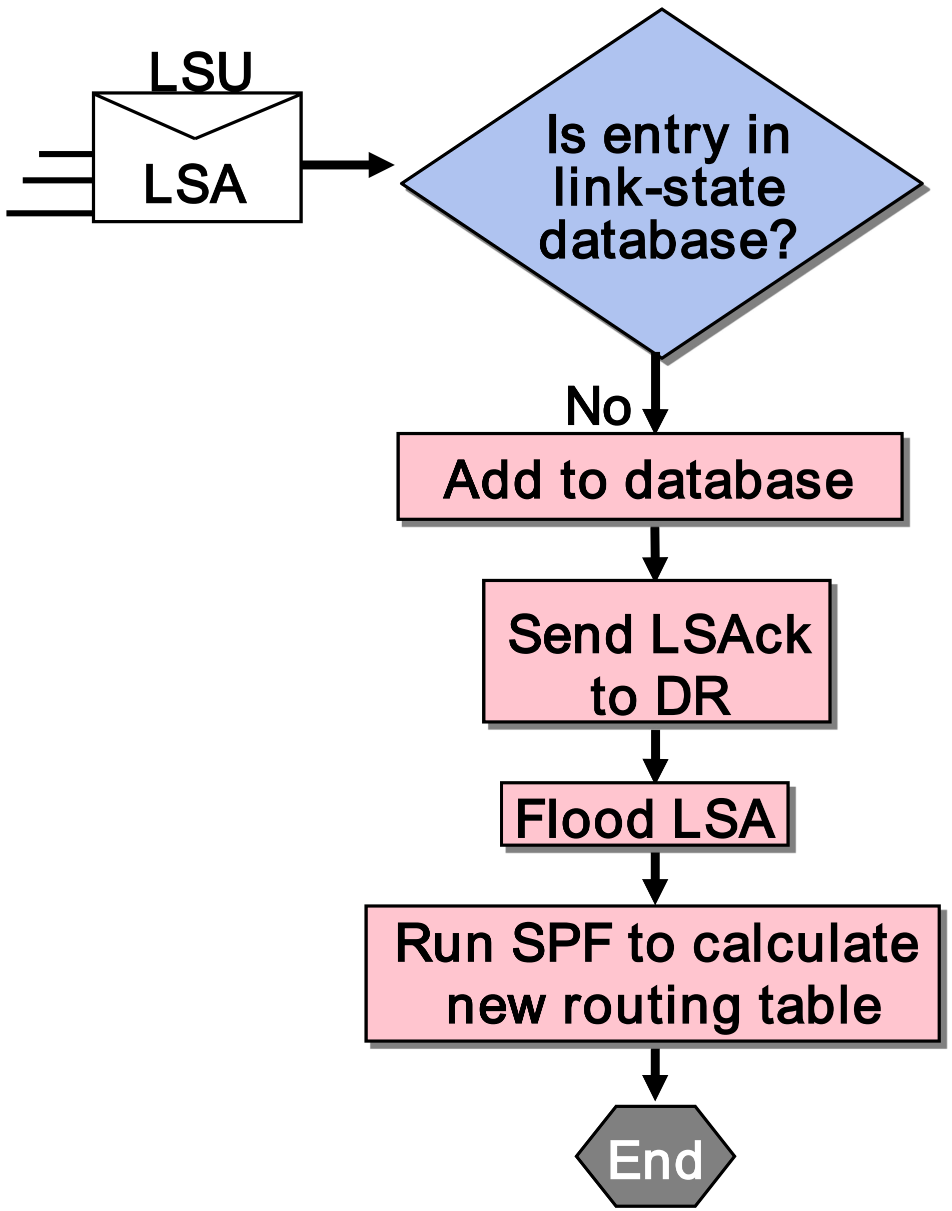

Maintaining Routing Information

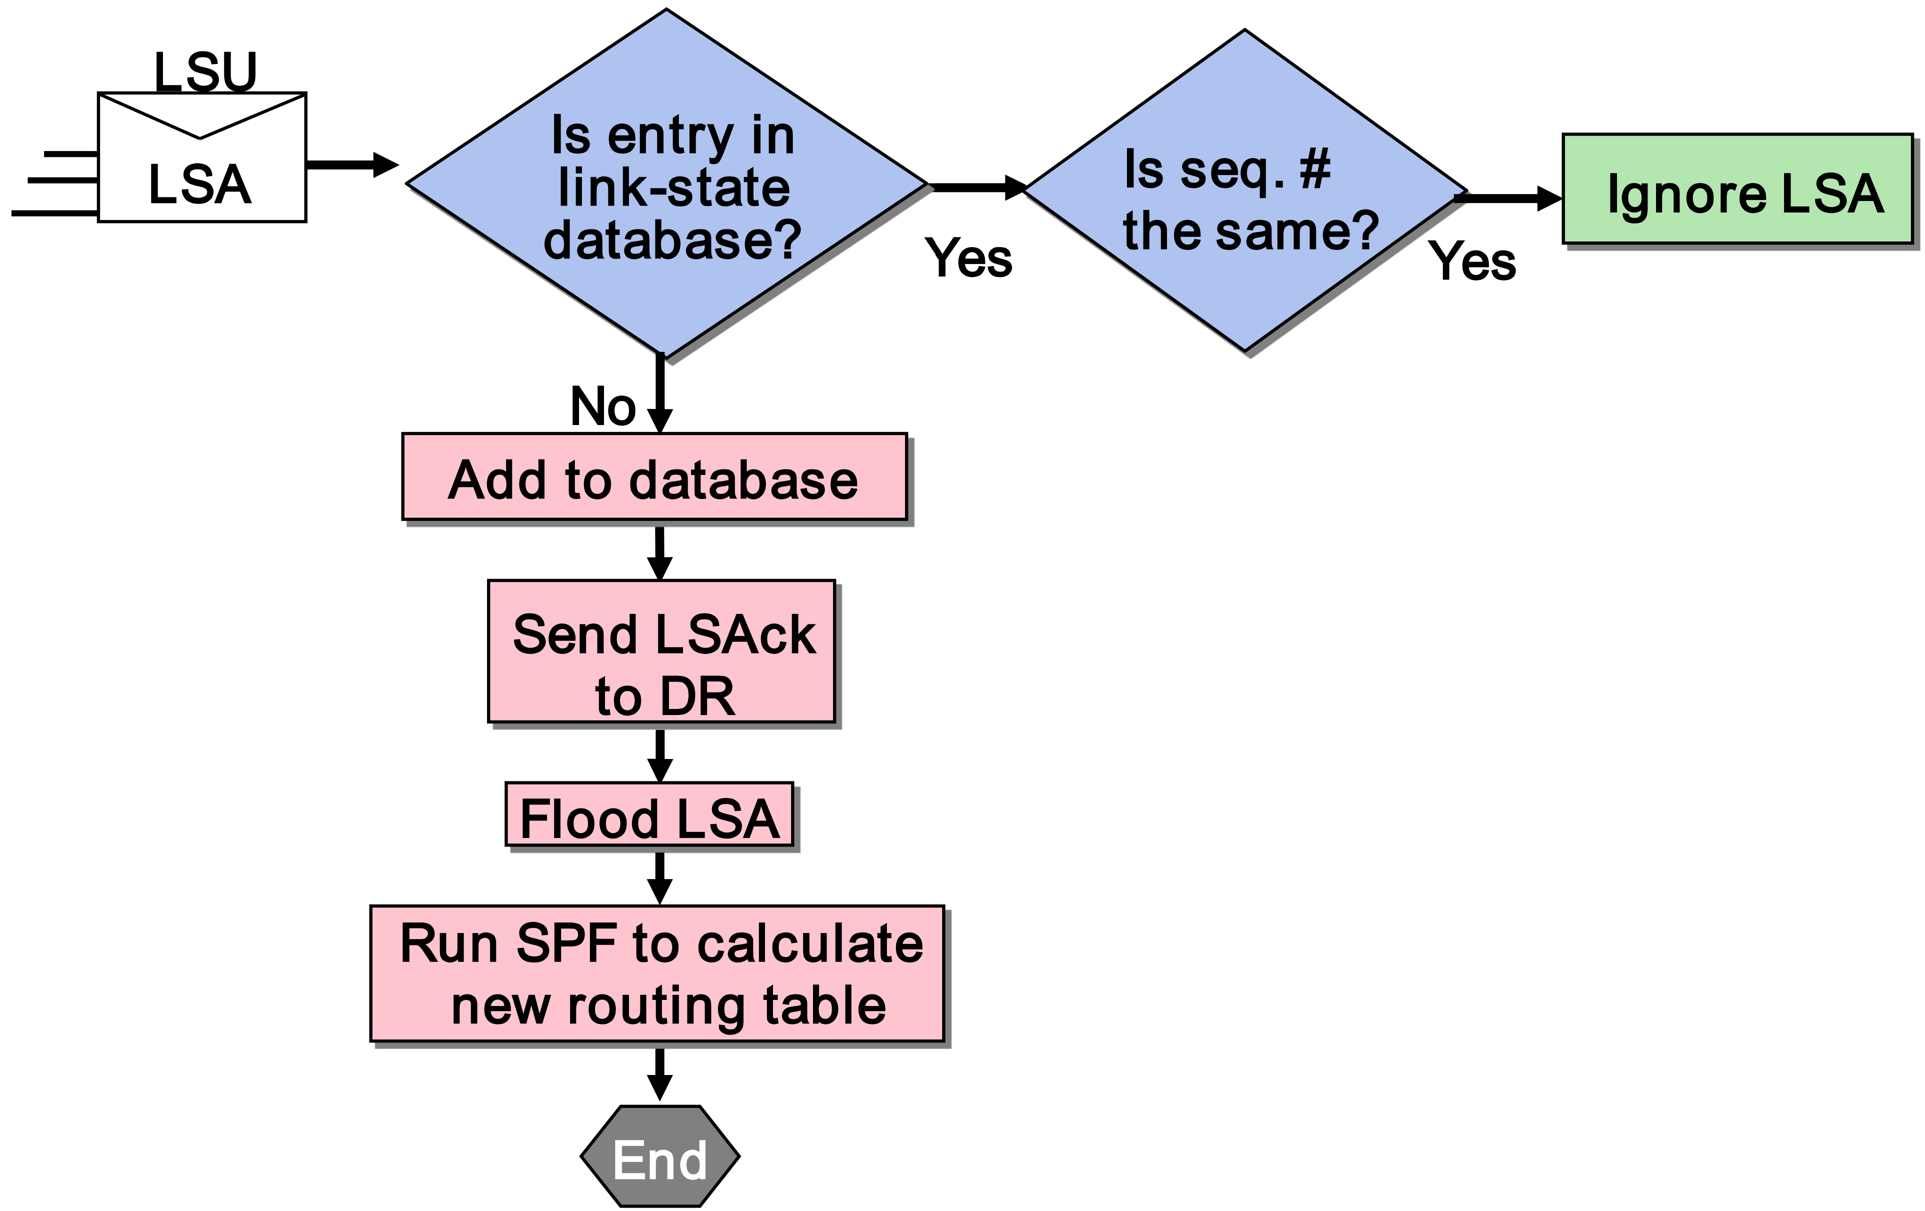

Maintaining Routing Information

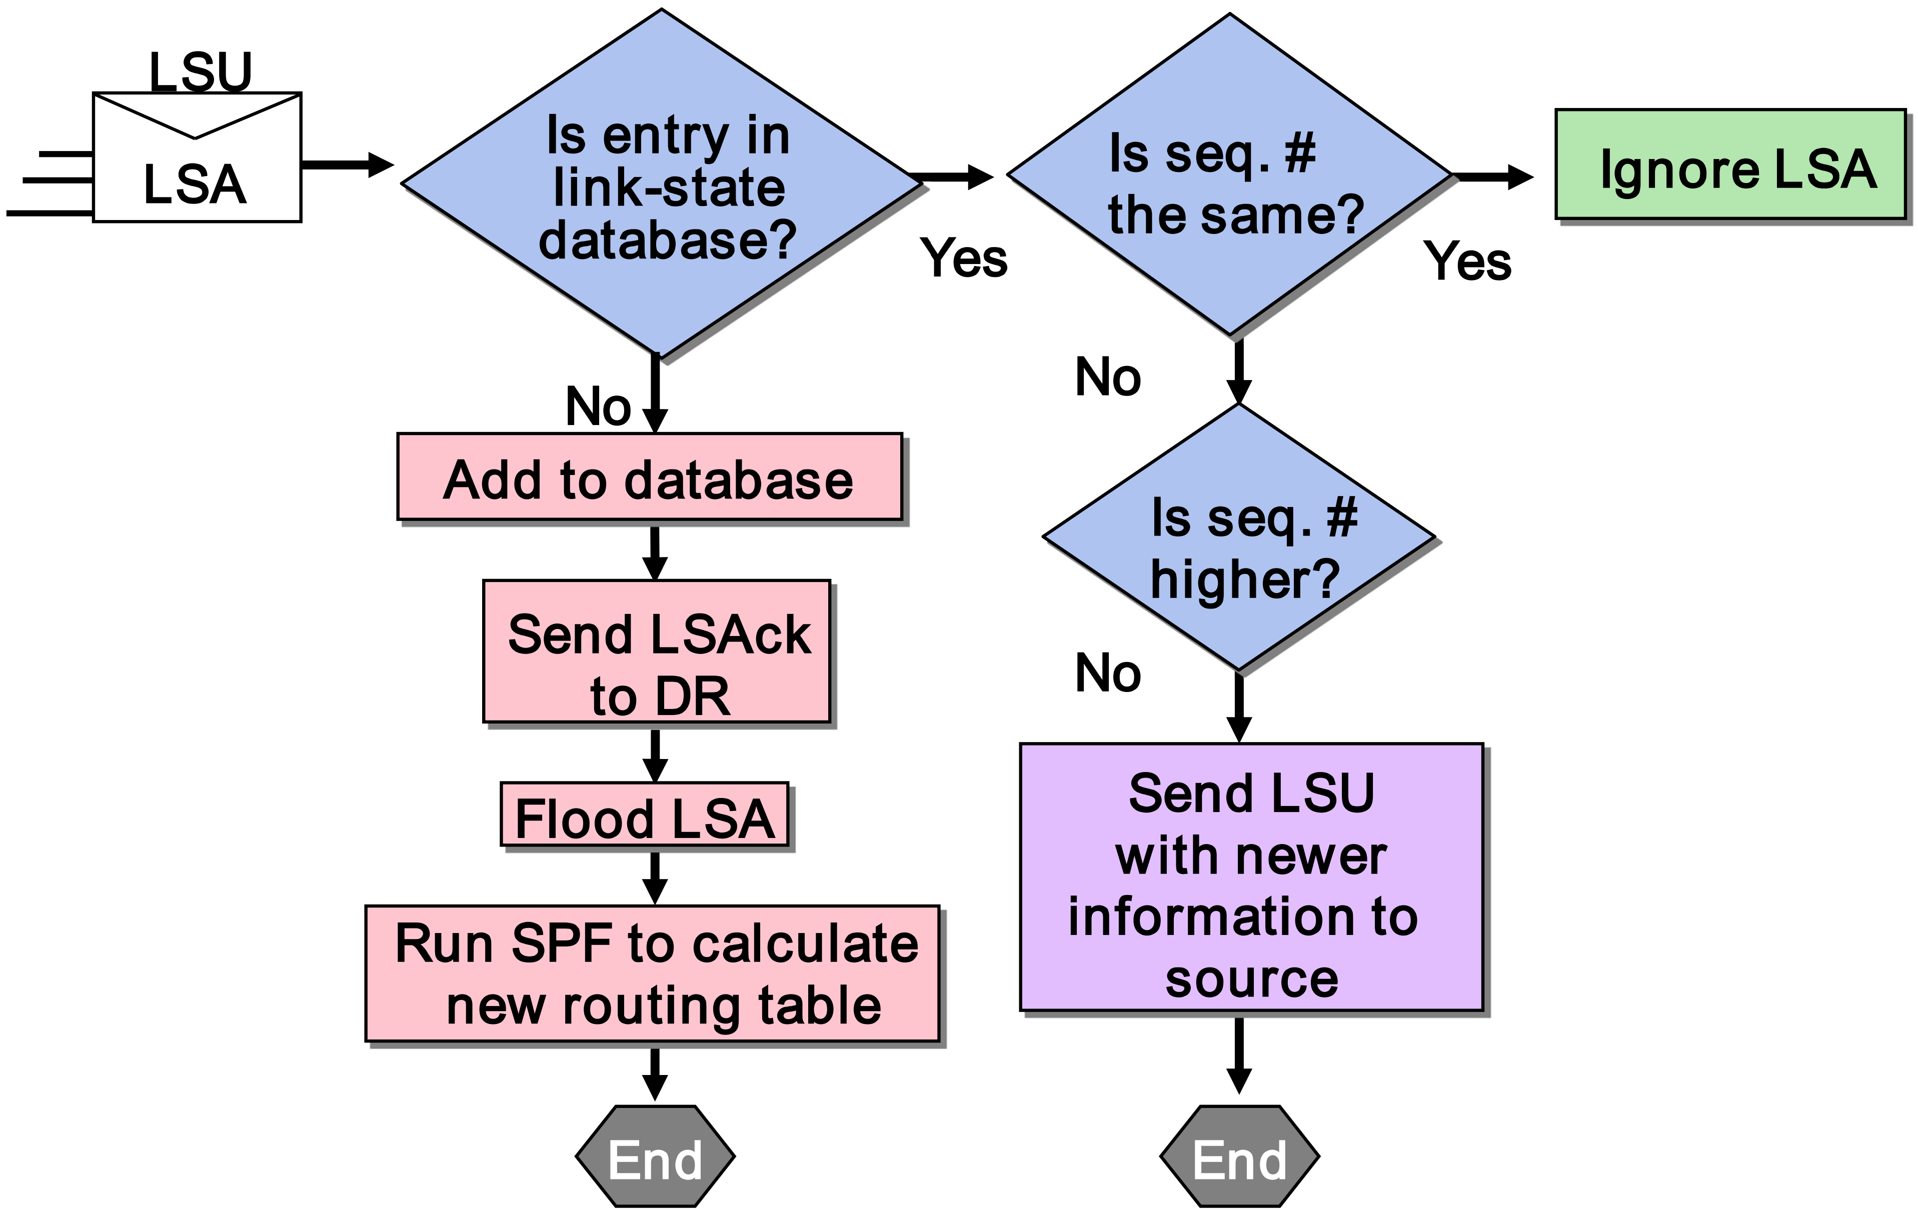

Maintaining Routing Information

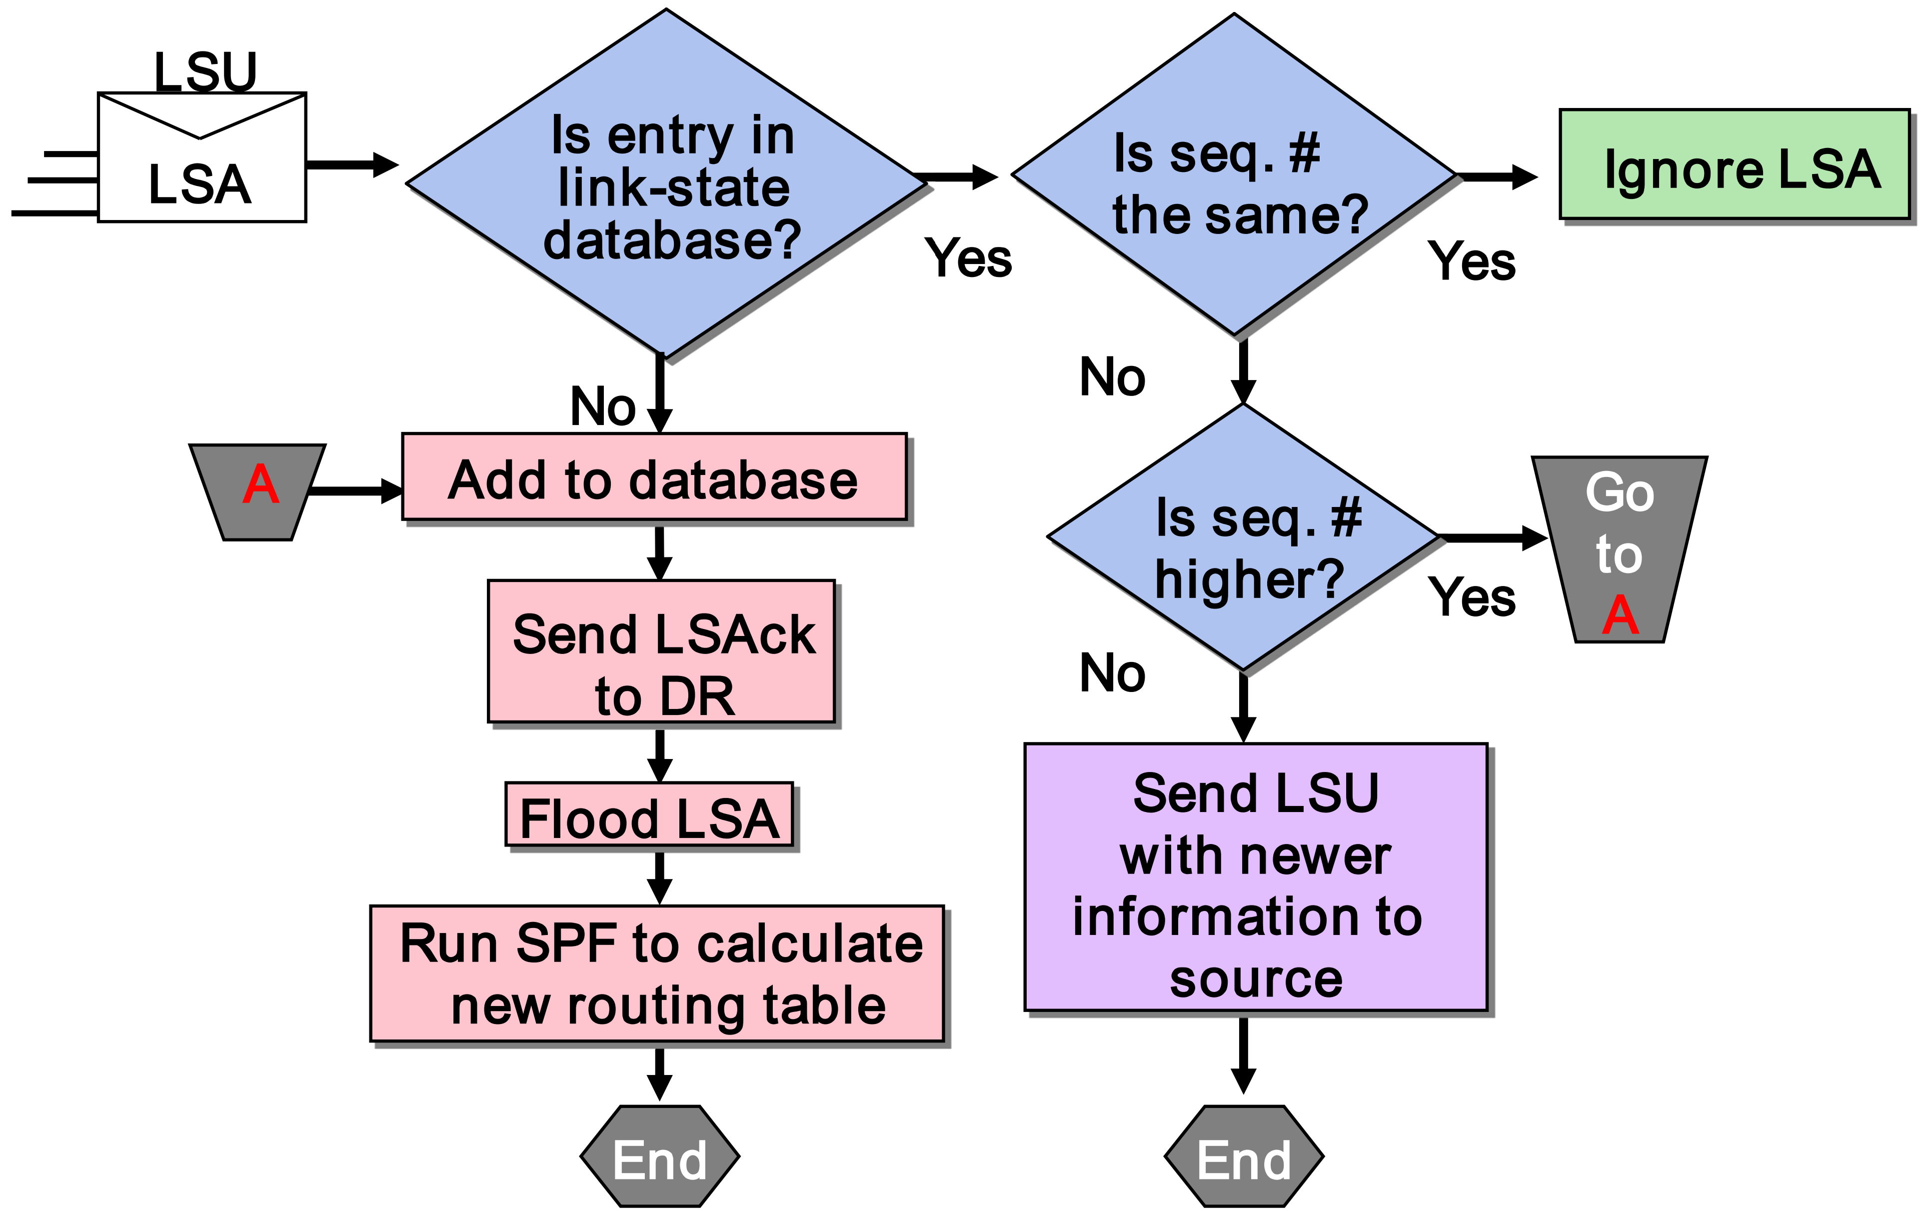

Steps to OSPF Operation

1. *Establishing router adjacencies*
2. *Electing DR and BDR*
3. *Discovering Routes*
4. *Choosing Routes*
5. *Maintaining Routing Information*

OSPF States (Finite State Machine)

- Every OSPF router represents its communications with other OSPF routers in the form of neighbor data structures.
- Every neighbor can be in one of the following states:
  - ***Down State***
  - ***Attempt State***
  - ***Init State***
  - ***Two-way State***
  - ***ExStart State***
  - ***Exchange State***
  - ***Loading State***
  - ***Full Adjacency State***

Steps to OSPF Operation (OSPF states)

1. Establishing router adjacencies
   - Down State
   - Init State
   - Two-way State
   - (ExStart State unless DR/BDR election needed)
2. Electing DR and BDR
   - ExStart State with DR and BDR
   - Two-way State with all other routers
3. Discovering Routes
   - ExStart State
   - Exchange State
   - Loading State
   - Full State
4. Choosing Routes
5. Maintaining Routing Information

OSPF packet types

Type	Description
1	Hello (establishes and maintains adjacency relationships with neighbors) **OSPF Type-1 (Hello)**
2	Database description packet (describes the contents of an OSPF router's link-state database) **OSPF Type-2 (DBD)**
3	Link-state request (requests specific pieces of a neighbor router's link-state database) **OSPF Type-3 (LSR)**
4	Link-state update (transports link-state advertisements (LSAs) to neighbor routers) **OSPF Type-4 (LSU)**
5	Link-state acknowledgement (Neighbor routers acknowledge receipt of the LSAs) **OSPF Type-5 (LSAck)**

Discovering Routes and reaching Full State

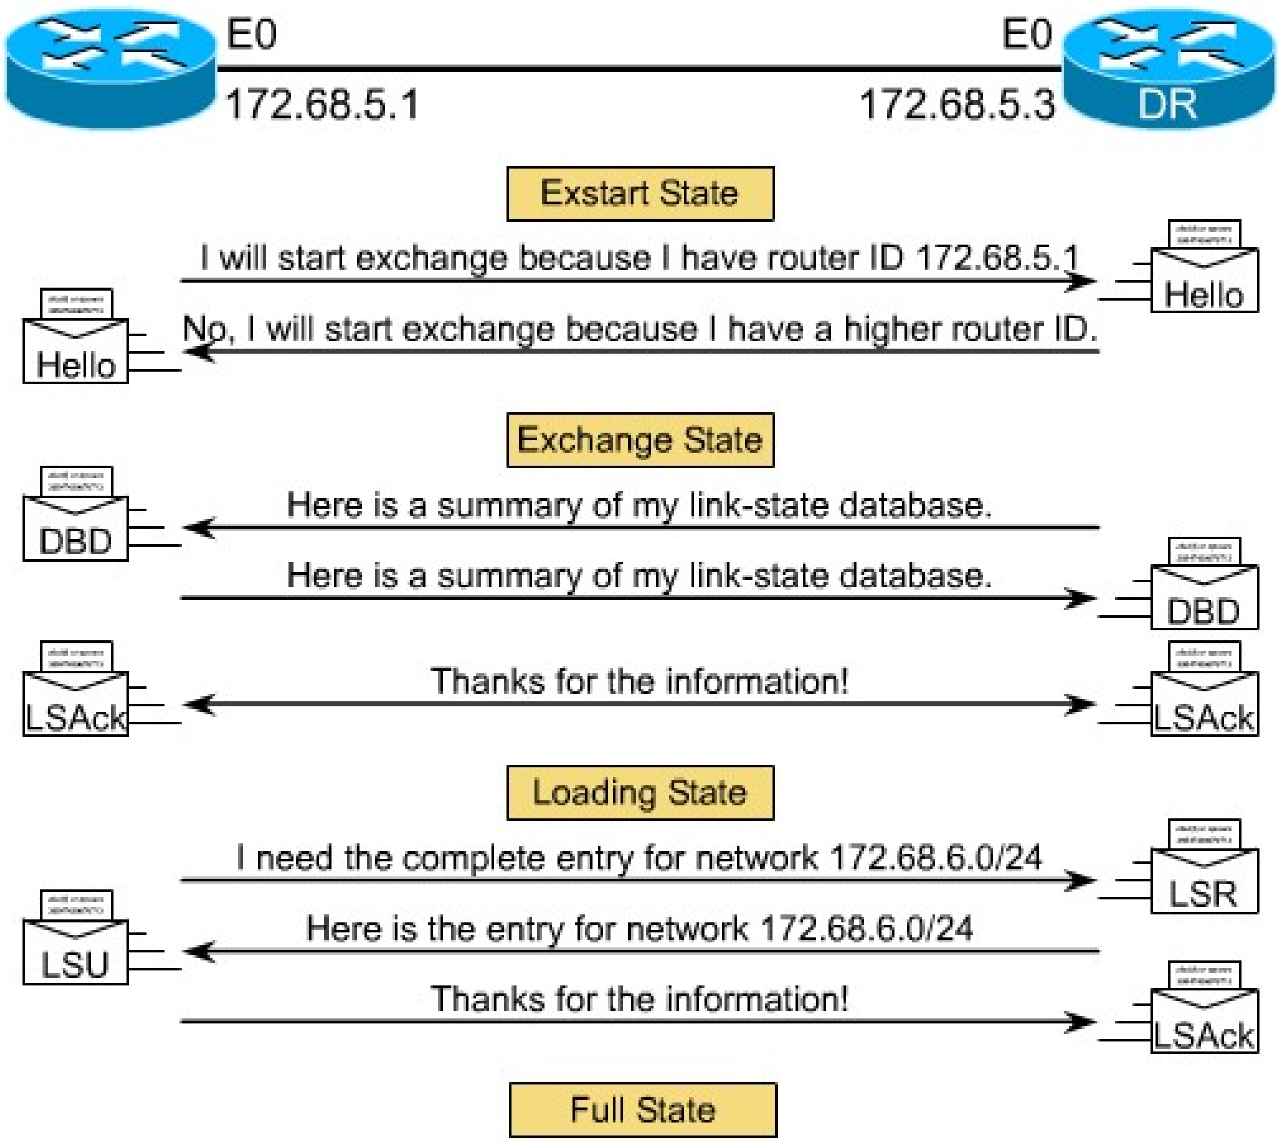

"adjacent"

OSPF Type-2 (DBD)

OSPF Type-2 (DBD)

OSPF Type-2 (DBD)

OSPF Type-2 (DBD)

OSPF Type-5 (LSAck)

OSPF Type-3 (LSR)

OSPF Type-4 (LSU)

OSPF Type-5 (LSAck)

"full adjacency"

Calculating the Routing Table

The LSAs that build the database contain three important pieces of generic information: **RouterID** of the sender of the LSA, the **NeighborID**, and **cost** of the link between the Router and the neighbor (i.e the state of the link or link-state).

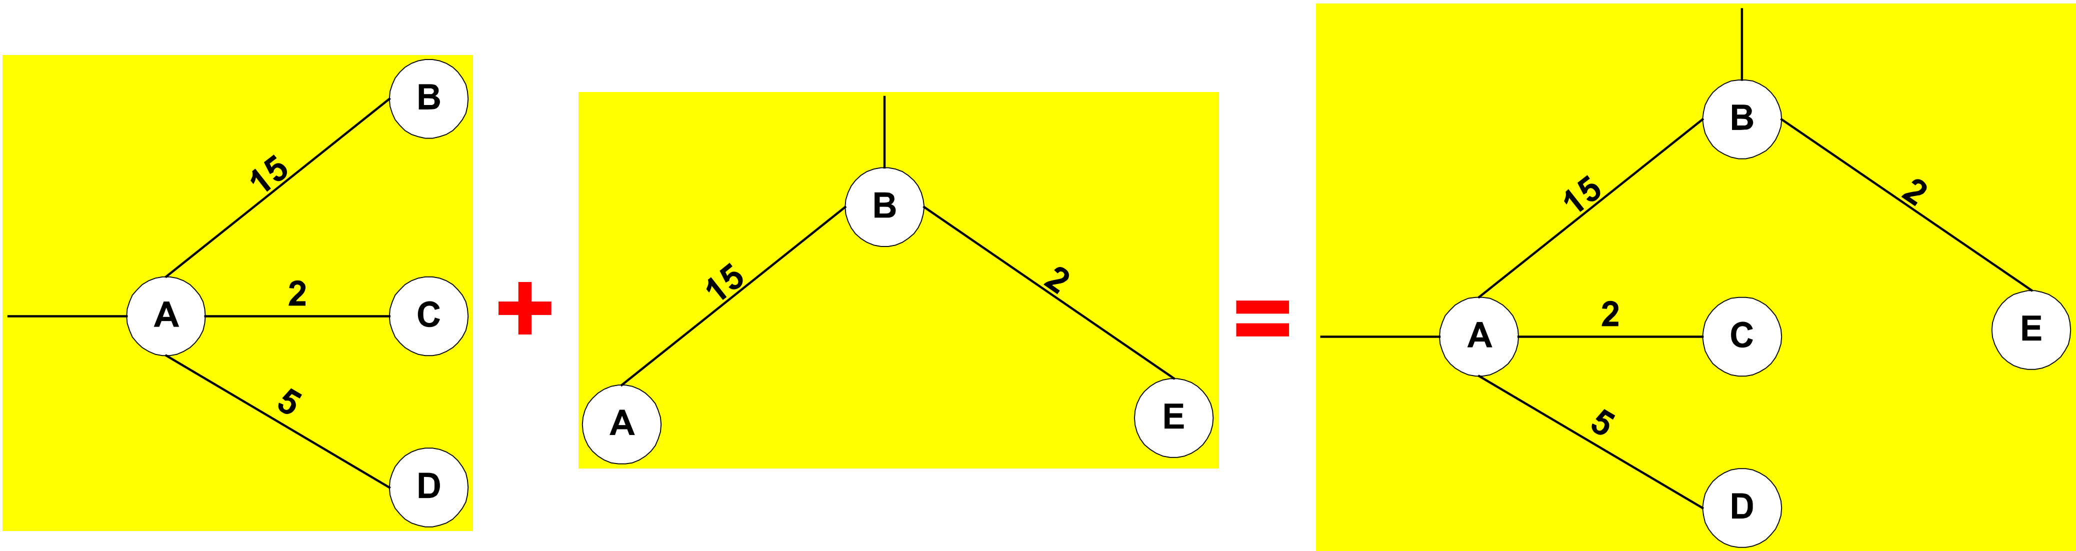

Dijkstra's Algorithm

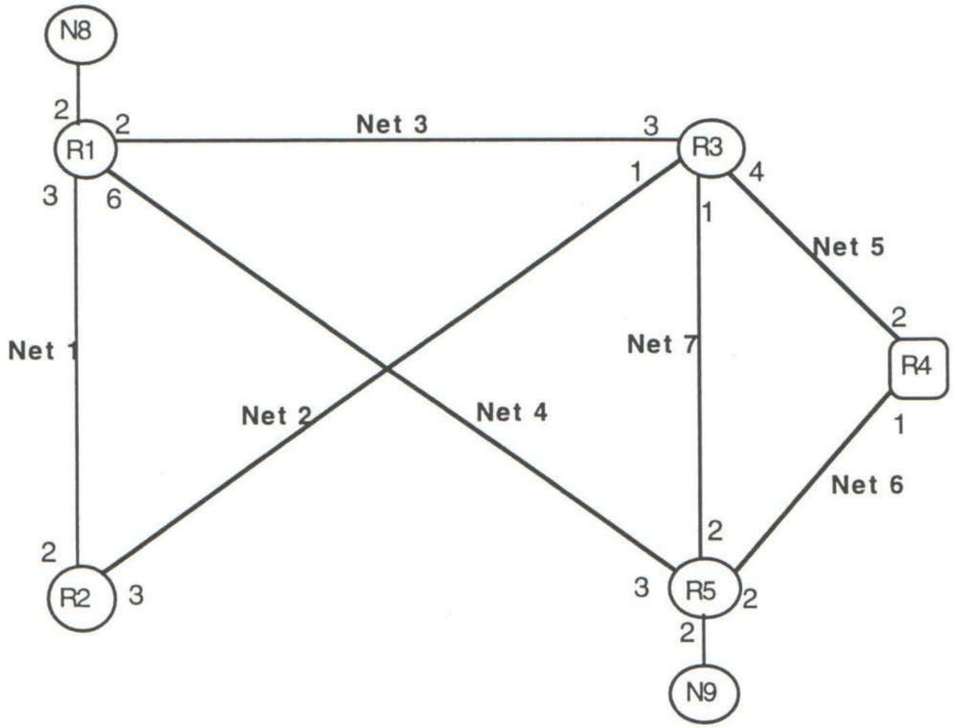

Link State Database (all routers)	
R2	Net 1/2, Net 2/3
R1	Net 1/3, Net 3/2, Net 4/6, Net 8/2
R3	Net 3/3, Net 2/1, Net 7/1, Net 5/4
R4	Net 5/2, Net 6/1
R5	Net 4/3, Net 7/2, Net 6/2, Net 9/2

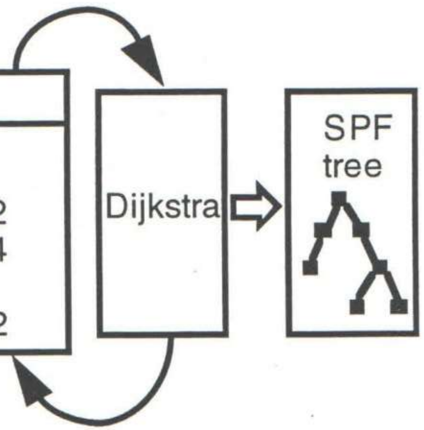

Dijkstra's Algorithm—Step 1

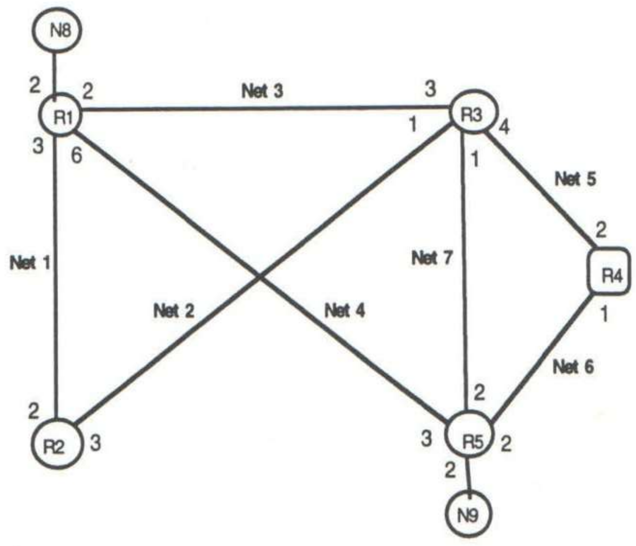

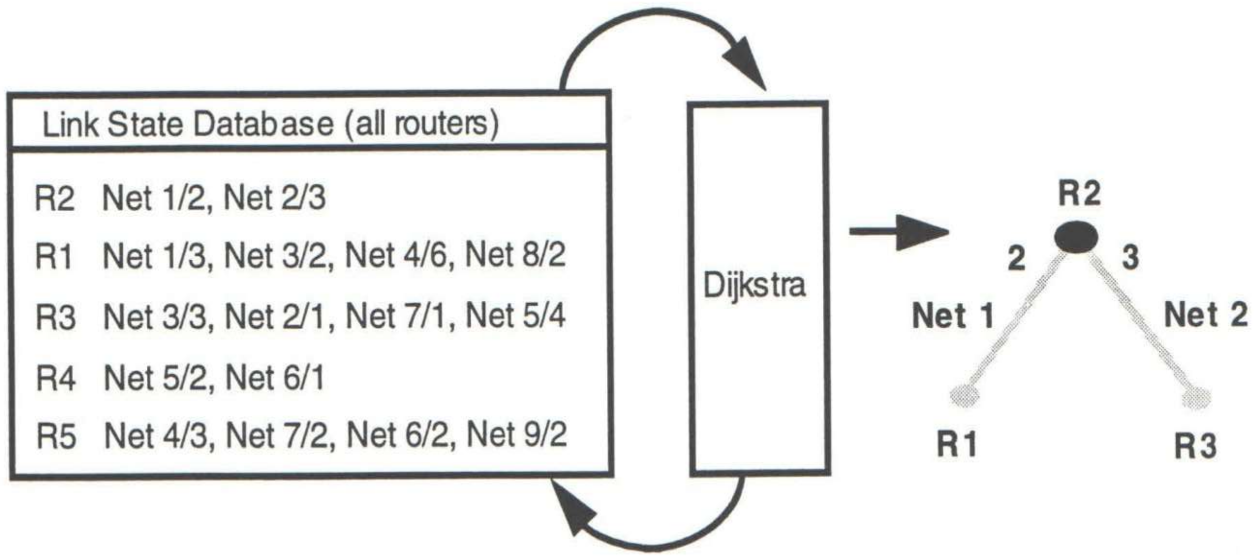

Dijkstra's Algorithm—Step 2

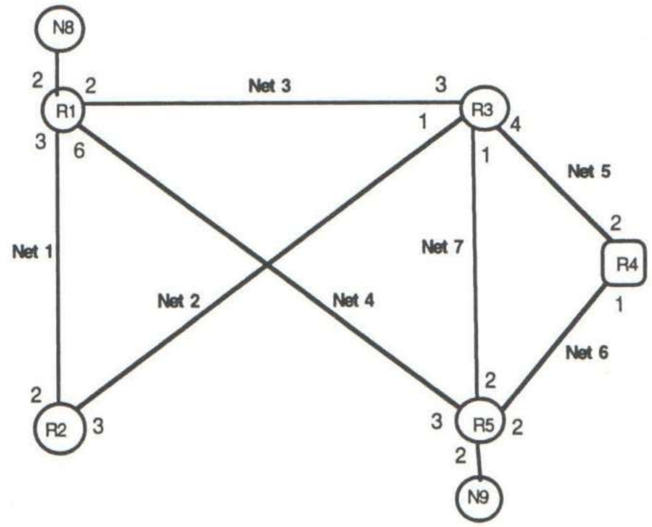

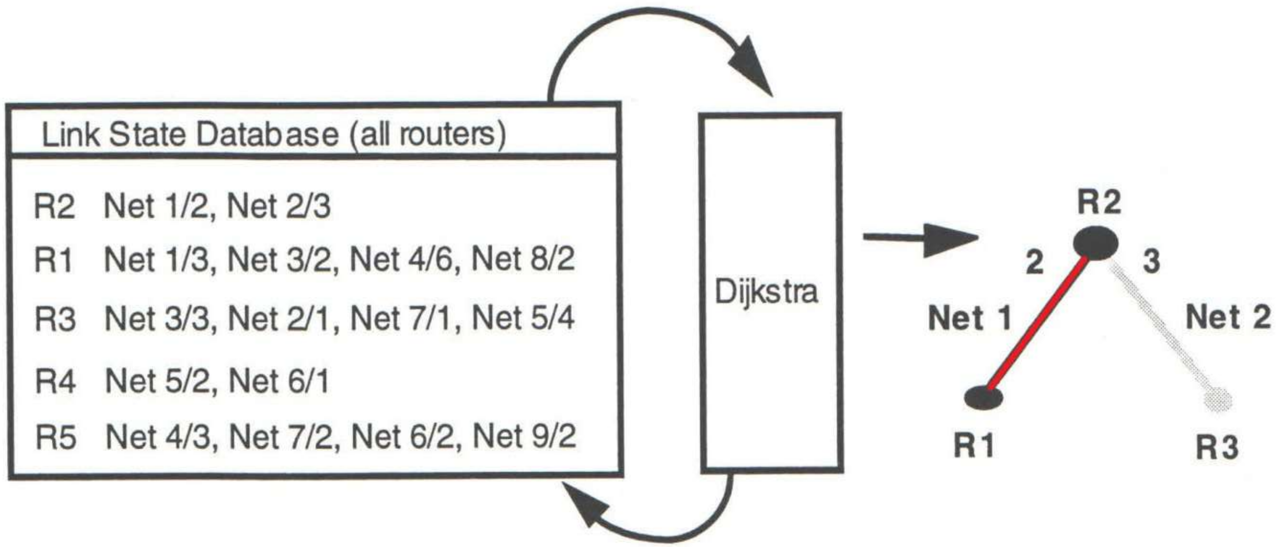

Dijkstra's Algorithm—Step 3

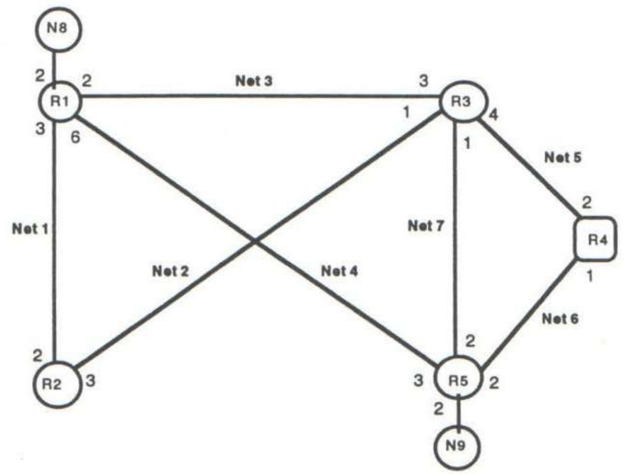

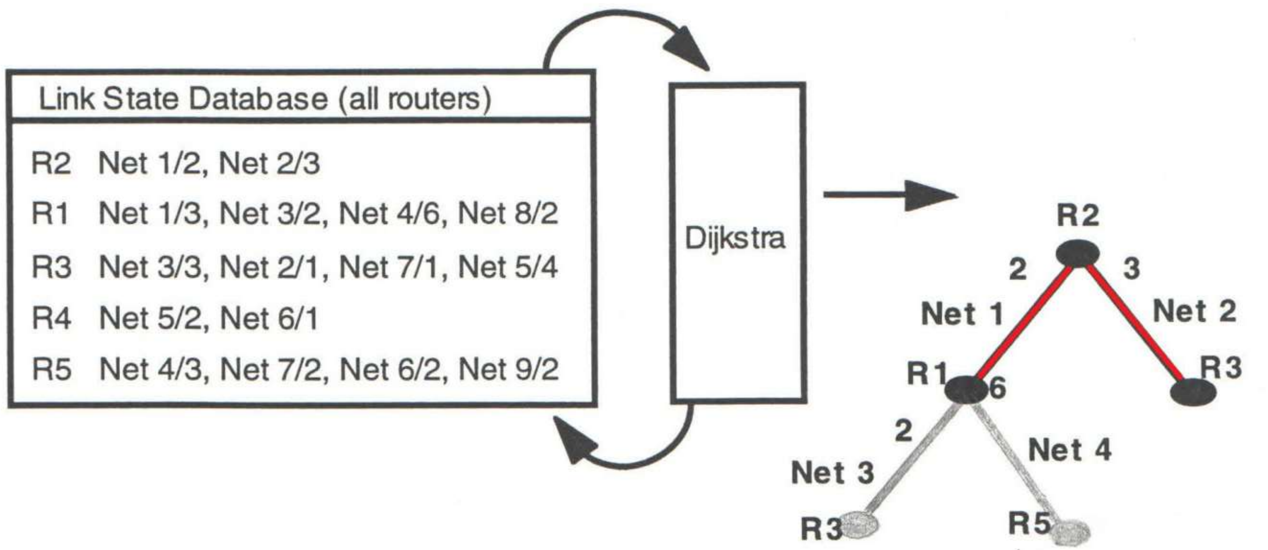

Dijkstra's Algorithm—Step 4

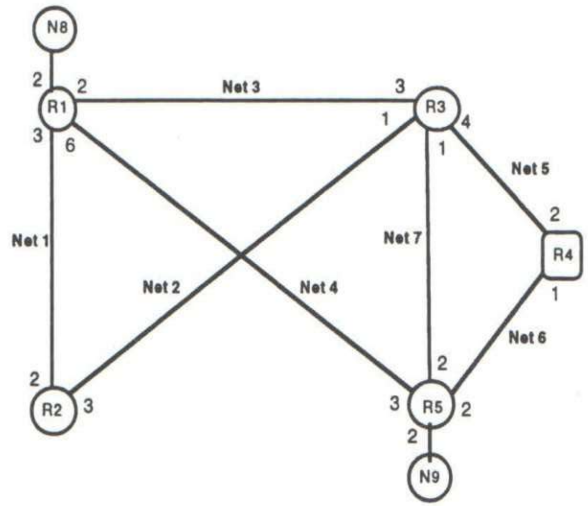

Link State Database (all routers)

R2 Net 1/2, Net 2/3

R1 Net 1/3, Net 3/2, Net 4/6, Net 8/2

R3 Net 3/3, Net 2/1, Net 7/1, Net 5/4

R4 Net 5/2, Net 6/1

R5 Net 4/3, Net 7/2, Net 6/2, Net 9/2

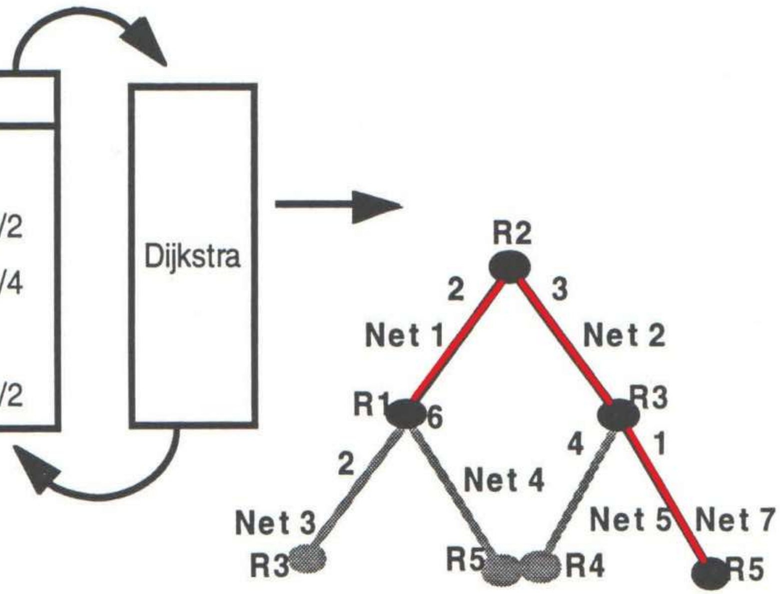

Dijkstra's Algorithm—Step 5

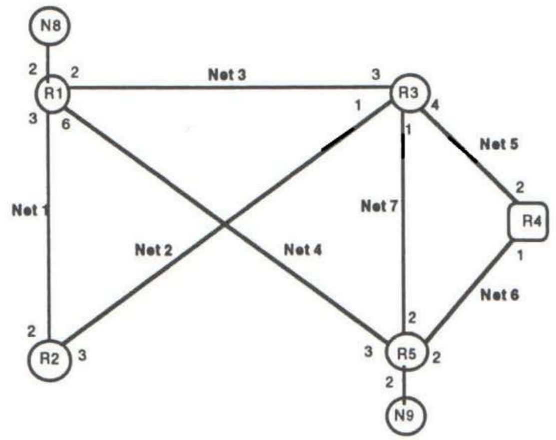

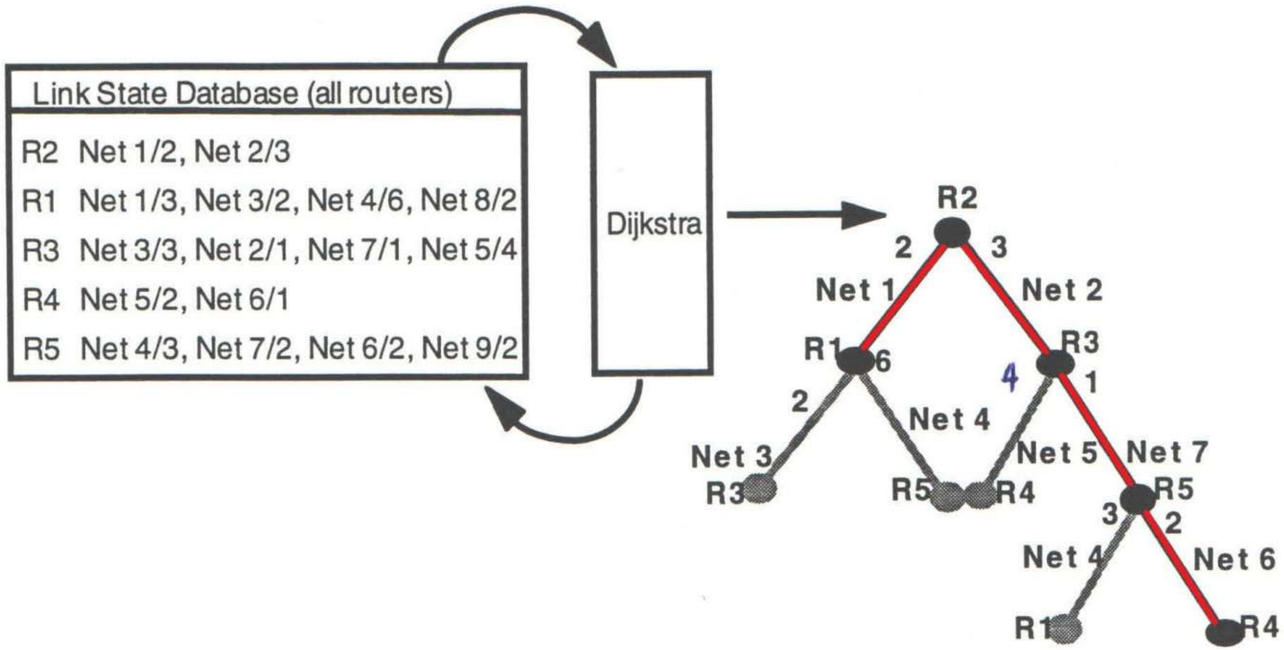

Dijkstra's Algorithm—Step 6

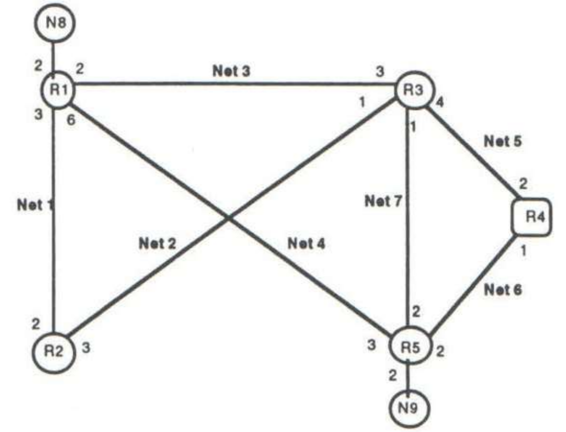

Link State Database (all routers)	
R2	Net 1/2, Net 2/3
R1	Net 1/3, Net 3/2, Net 4/6, Net 8/2
R3	Net 3/3, Net 2/1, Net 7/1, Net 5/4
R4	Net 5/2, Net 6/1
R5	Net 4/3, Net 7/2, Net 6/2, Net 9/2

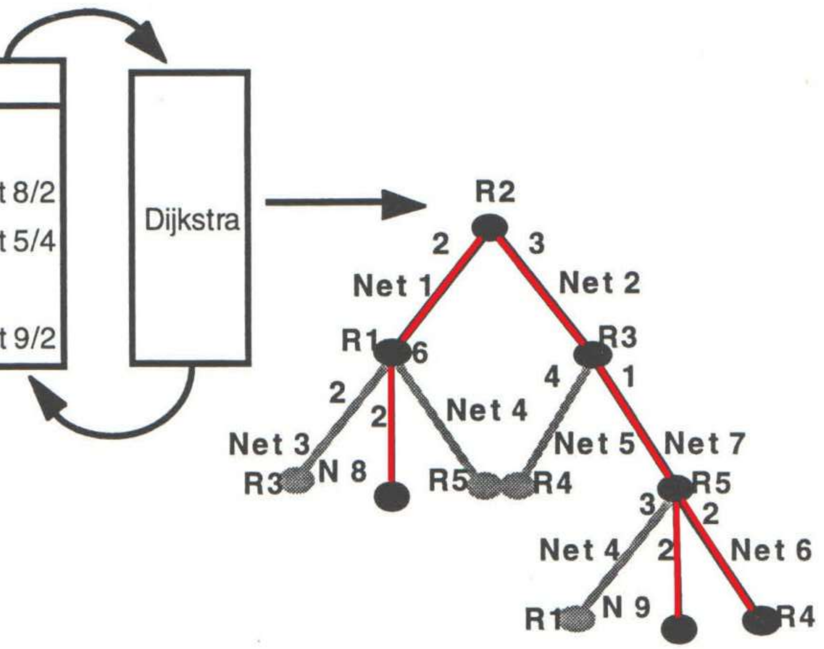

Routing Protocols

Questions ?