Routing Information Protocol (RIP v1, RIP v2) Open Shortest Path First (OSPF)

Routing Information Protocol - RIPv1

RIP Characteristics:

- a classful, Distance Vector (DV) routing protocol;
- Metric = hop count;
- routes with a hop count > 15 are unreachable;
- updates are broadcast every 30 seconds.

Routing Information Protocol - RIPv1

RIP uses 2 message types:

- Request message - sent out on startup by each RIP enabled interface. Requests all RIP enabled neighbors to send routing table.
- Response message - sent to requesting router containing routing table

Routing Information Protocol - RIPv1

RIP message encapsulated into a UDP segment

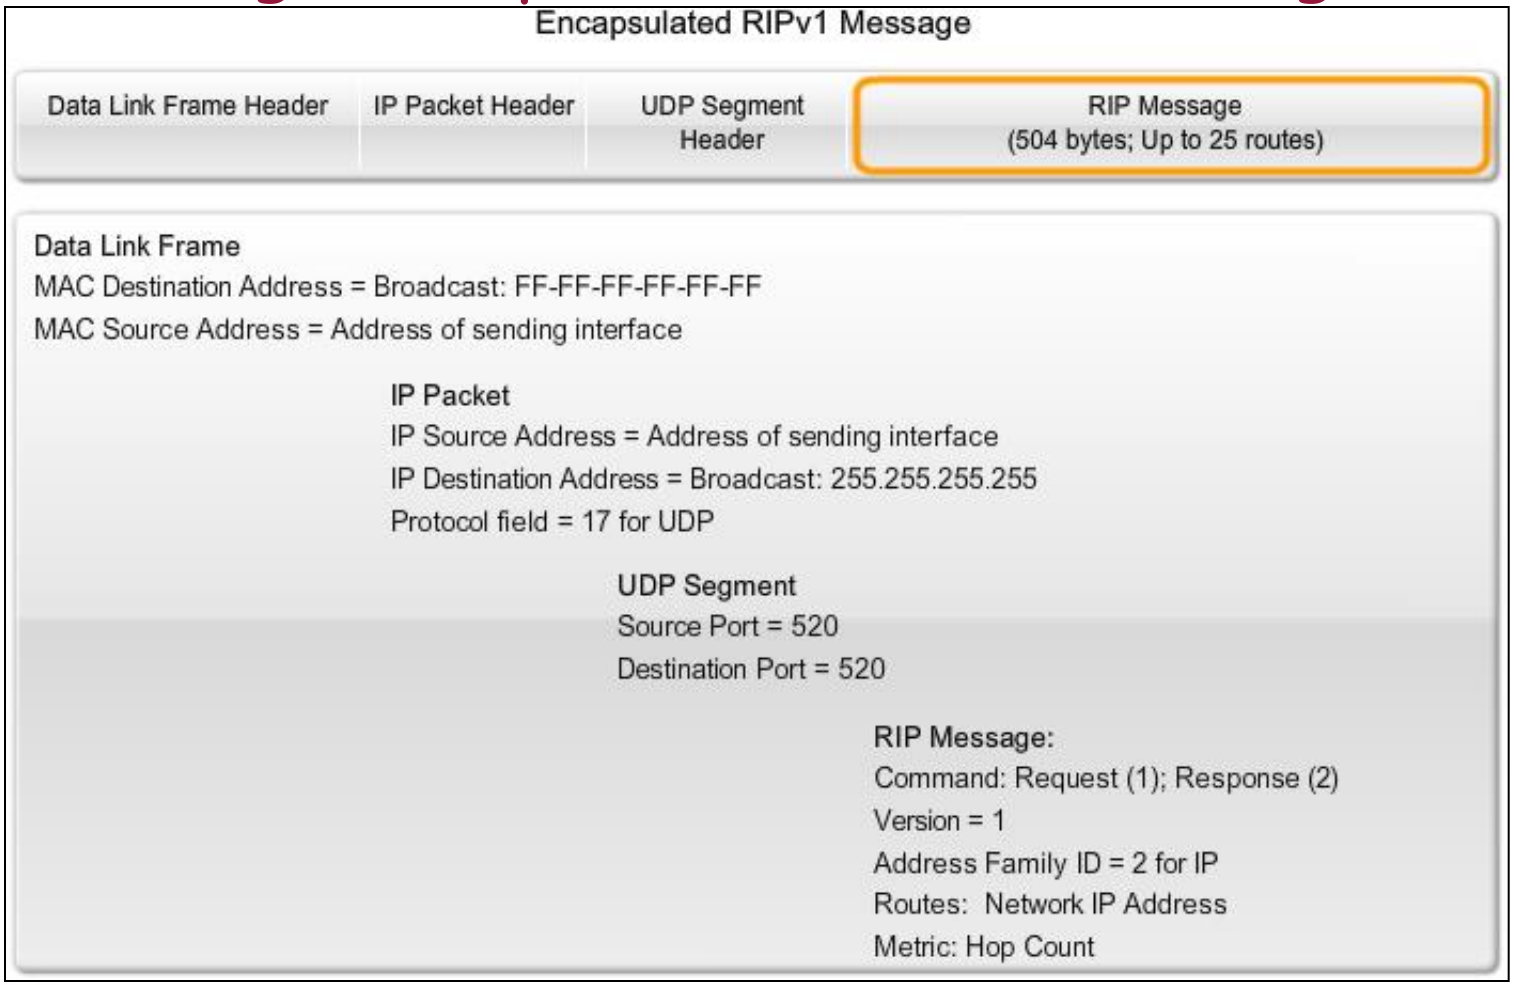

Routing Information Protocol - RIPv1

RIP Message Format

- RIP header - divided into 3 fields
  - Command field
  - Version field
  - Must be zero

RIP Message Format

- Route Entry - composed of 3 fields
  - Address family identifier
  - IP address
  - Metric

RIPv1 Message Format

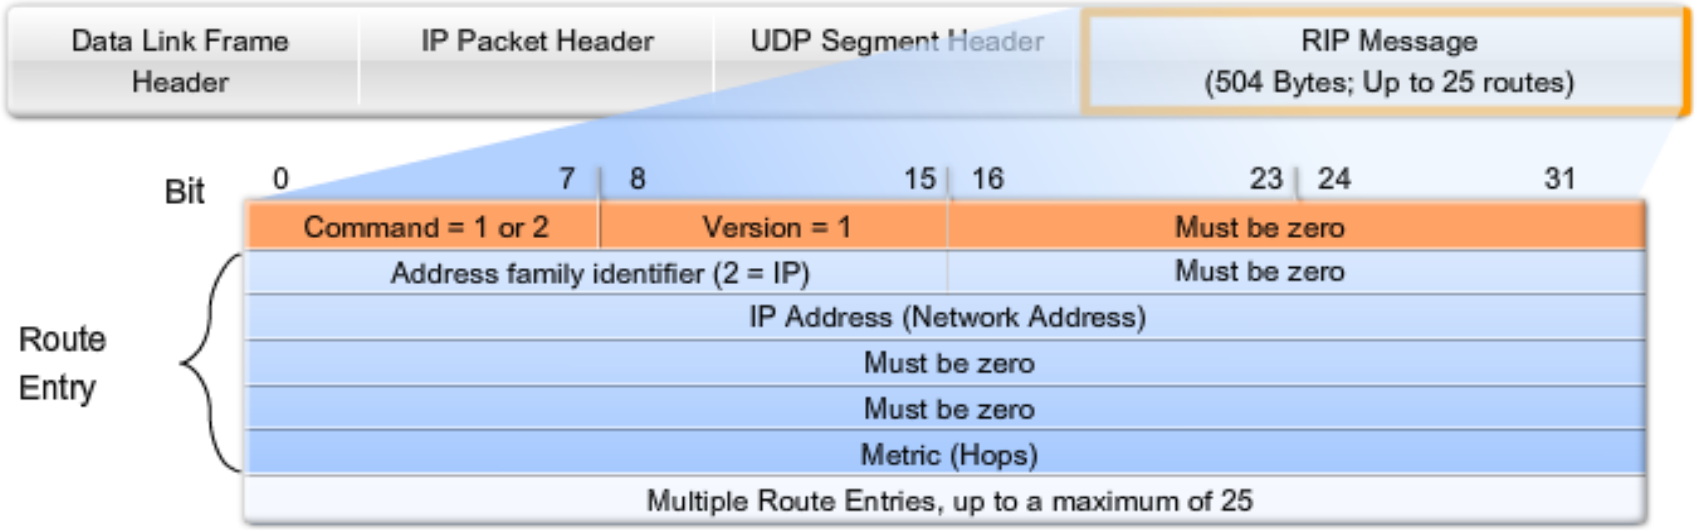

Routing Information Protocol - RIPv1

RIP Message Format

Command	1 for a Request or 2 for a Reply.
Version	1 for RIP v 1 or 2 for RIP v 2.
Address Family Identifier	2 for IP unless a Request is for the full routing table in which case, set to 0.
IP Address	The address of the destination route, which may be a network, subnet, or host address.
Metric	Hop count between 1 and 16. Sending router increases the metric before sending out message.

Administrative Distance & Route Preference of RIP v1

AD in the Cisco IOS

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

Route Preference in the Juniper JUNOS

Route Preference Values

Source	Default Preference
Direct	0
Local	0
Static	5
OSPF internal	10
RIP	100
Aggregate	130
OSPF AS external	150
BGP (both EGBP and IBGP)	170

RIPv1 Route in Routing Table

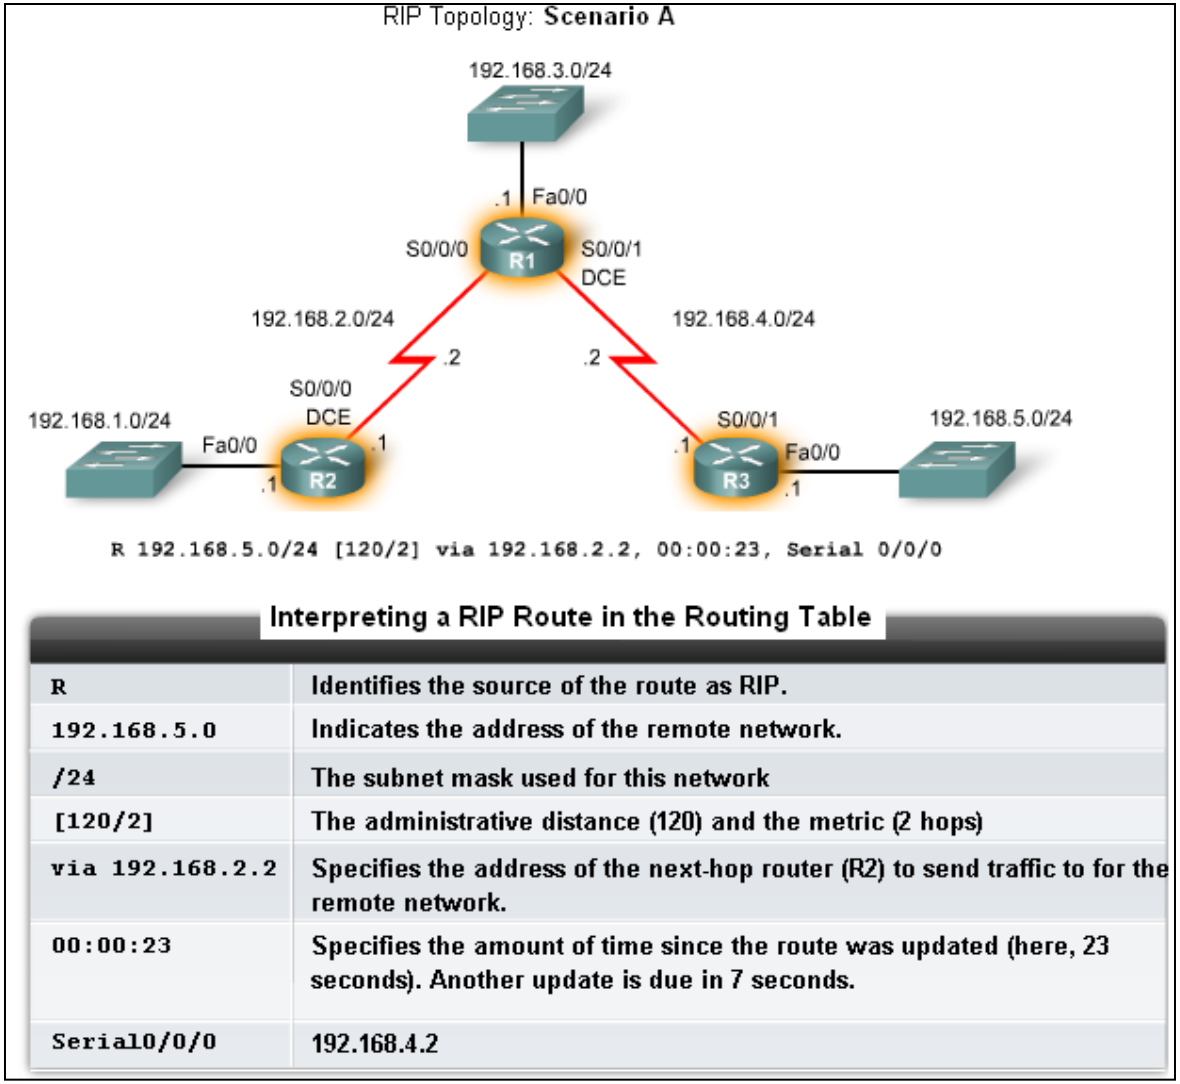

Interpreting a RIP Route in the Routing Table

`R`	Identifies the source of the route as RIP.
`192.168.5.0`	Indicates the address of the remote network.
`/24`	The subnet mask used for this network
`[120/2]`	The administrative distance (120) and the metric (2 hops)
`via 192.168.2.2`	Specifies the address of the next-hop router (R2) to send traffic to for the remote network.
`00:00:23`	Specifies the amount of time since the route was updated (here, 23 seconds). Another update is due in 7 seconds.
`Serial0/0/0`	192.168.4.2

Automatic Summarization

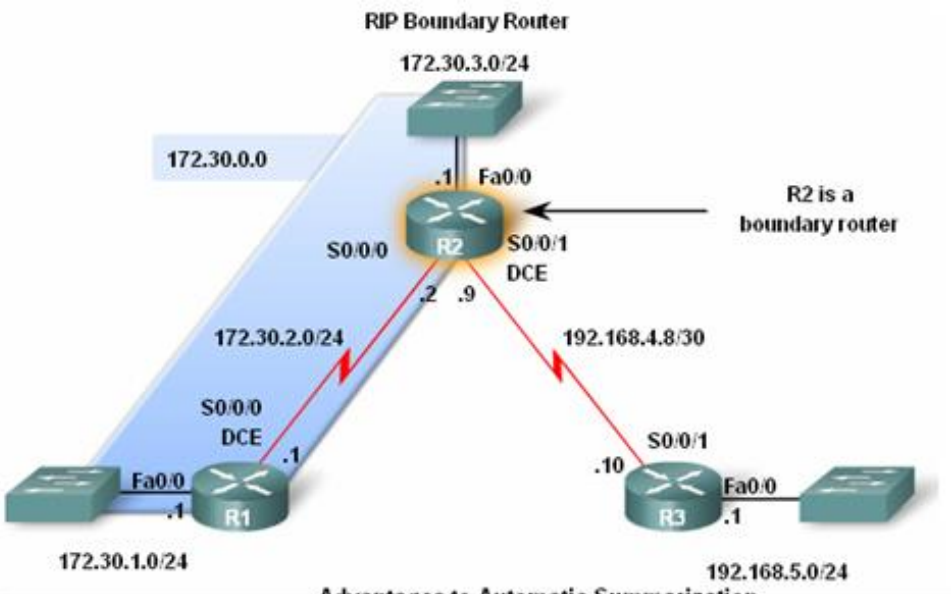

Advantages to Automatic Summarization

```
R3#show ip route
Codes: C - connected, S - static, I -IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSFP inter area
       N1 - OSPF NSSA externaal type 1, N2 - OSPF NSSA external type2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, II - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

R    172.30.0.0/16 [120/1] via 192.168.4.9, 00:00:15, Serial0/0/1
     192.168.4.0/30 is subnetted, 1 subnets
C       192.168.4.8 is directly connected, Serial0/0/1
C    192.168.5.0/24 is directly connected, FastEthernet0/0
```

R3 receives a single summarized route.

Boundary Routers

- RIP automatically summarizes classful networks;
- Boundary routers summarize RIP subnets from one major network to another.

Automatic Summarization

Disadvantage of Automatic Summarization

- does not support discontiguous network;

Discontiguous network

- comprises a major net that separates another major net.

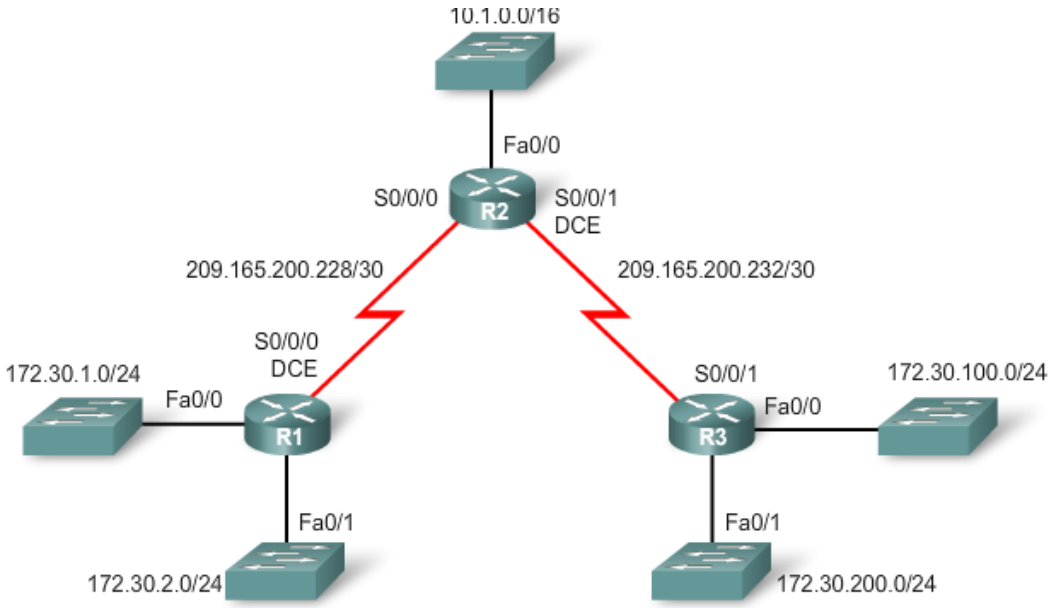

172.30.0.0/16 is a discontiguous network

- R1 and R3 both have subnets from the 172.30.0.0/16 major network, whereas R2 does not;
- R1 and R3 are boundary routers for 172.30.0.0/16 because they are separated by another major network, 209.165.200.0/24;
- This separation creates a discontiguous network, as two groups of 172.30.0.0/24 subnets are separated by at least one other major network.

RIPv1 Limitations

RIPv1

- a classful routing protocol;
- subnet mask are not sent in updates;
- summarizes networks at major network boundaries;
- if network is discontiguous and RIPv1 configured convergence will not be reached;
- does not support VLSM;

  (Reason: RIPv1 does not send subnet mask in routing updates)
- to determine which subnets to advertise:

  - does summarize routes to the Classful boundary;

  - or uses the Subnet mask of the outgoing interface.

Comparing RIPv1 & RIPv2

Difference between RIPv1 & RIPv2

RIPv1

- A classful distance vector routing protocol;
- Does not support discontiguous subnets;
- Does not support VLSM;
- Does not send subnet mask in routing update;
- Routing updates are broadcast.

RIPv2

- A classless distance vector routing protocol that is an enhancement of RIPv1's features;
- Next hop address is included in updates;
- Routing updates are multicast;
- The use of authentication is an option.

Comparing RIPv1 & RIPv2

Similarities between RIPv1 & RIPv2

- Use of timers to prevent routing loops;
- Use of split horizon or split horizon with poison reverse;
- Use of triggered updates;
- Maximum hop count of 15.

Comparing RIPv1 & RIPv2

Comparing RIPv1 & RIPv2 Message Formats

RIPv2 Message format:

- is similar to RIPv1;
- has 2 extensions:

  1st extension - subnet mask field;

  2nd extension - addition of next hop address.

Comparing RIPv1 & RIPv2

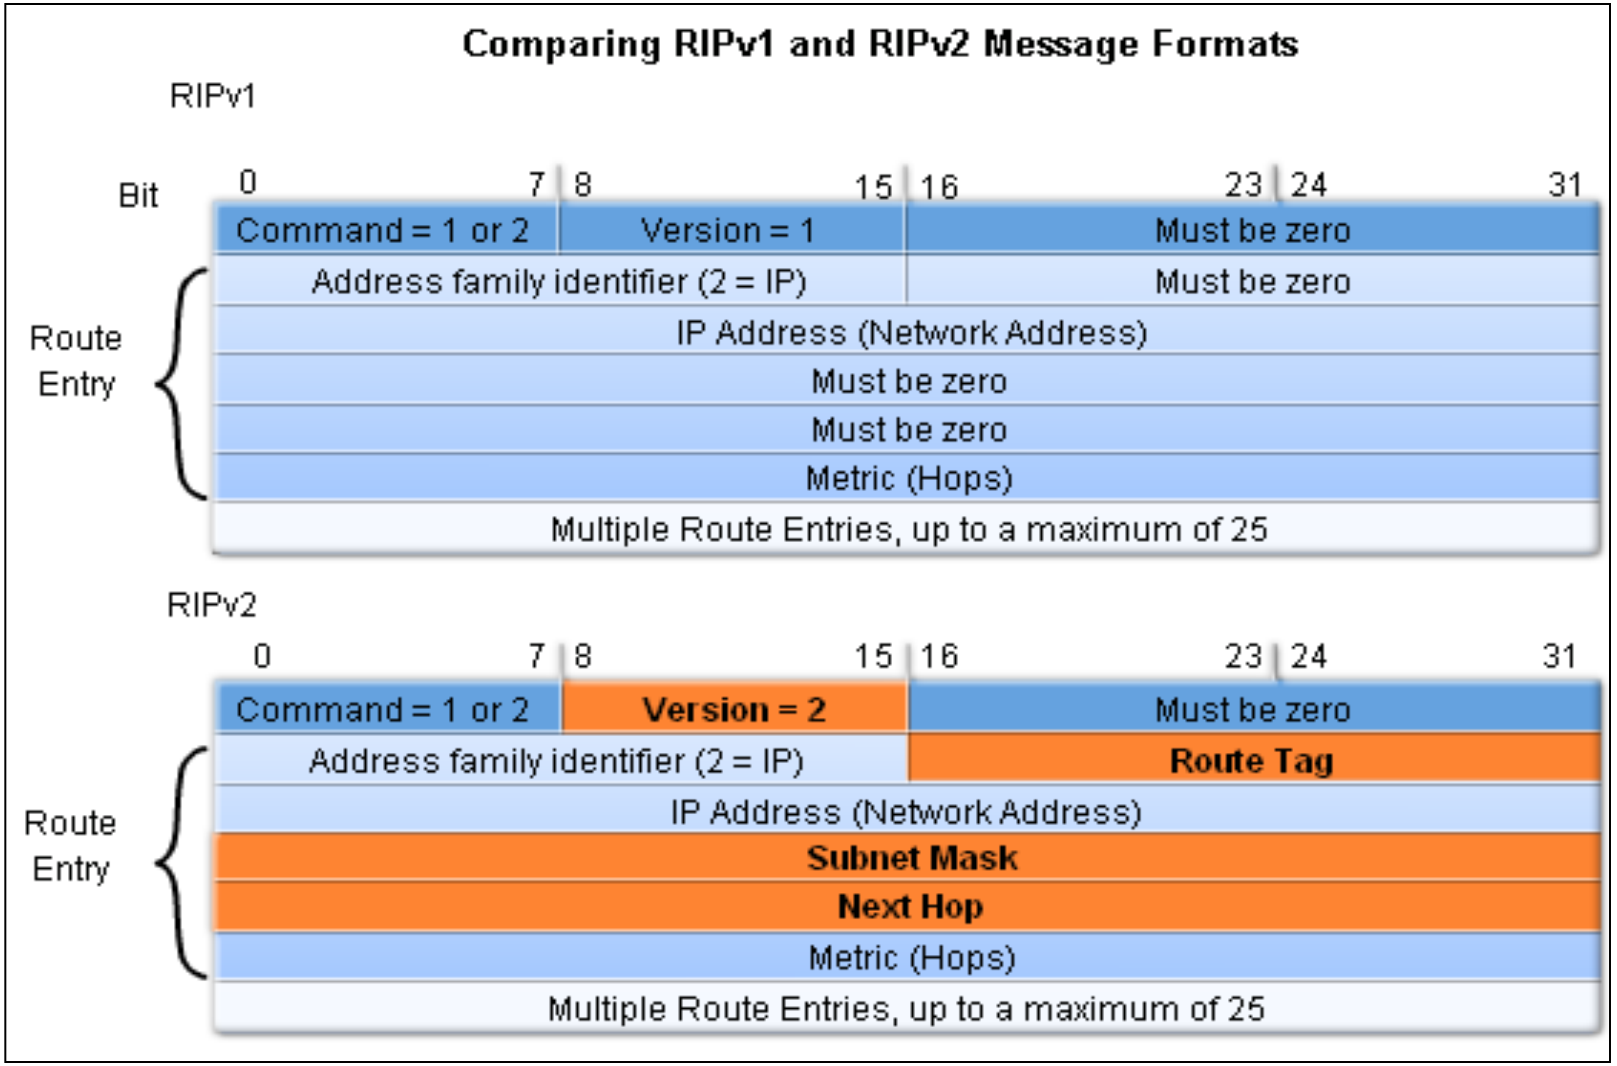

RIPv2

Auto-Summary & RIPv2

- RIPv2 will automatically summarize routes at major network boundaries
- can summarize routes with a subnet mask that is smaller than the classful subnet mask

VLSM & CIDR

RIPv2 and VLSM

- Networks using a VLSM IP addressing scheme use classless routing protocols (i.e. RIPv2) to disseminate network addresses and their subnet masks

CIDR & Supernetting

CIDR – Classless Inter-domain Routing

Supernetting

- is a bunch of contiguous classful networks that is addressed as a single network.

RIP Loop Protection Mechanisms

Split horizon

- router should not advertise a network through the interface through which the update came from.

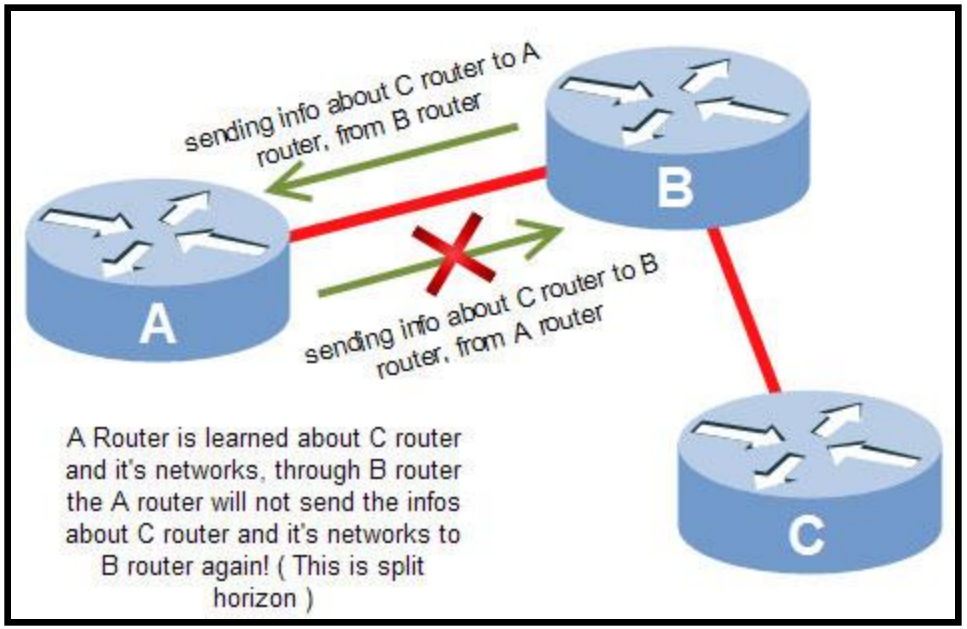

RIP Loop Protection Mechanisms

Split horizon with Poison Reverse

- Route poisoning is used to mark the route as unreachable in a routing update that is sent to other routers.
- Unreachable is interpreted as a metric that is set to the maximum (or maximum + 1).
  - For RIP, a poisoned route has a metric of 16.
- The concept of split horizon with poison reverse is that explicitly telling a router to ignore a route (sending poison route back to the sending router).

RIP Loop Protection Mechanisms

Split horizon with Poison Reverse

- The following process occurs:
- Network 10.4.0.0 becomes unavailable due to a link failure.
- R3 poisons the metric with a value of 16 and then sends out a triggered update stating that 10.4.0.0 is unavailable.
- R2 processes that update, invalidates the routing entry in its routing table, and immediately sends a poison reverse back to R3.

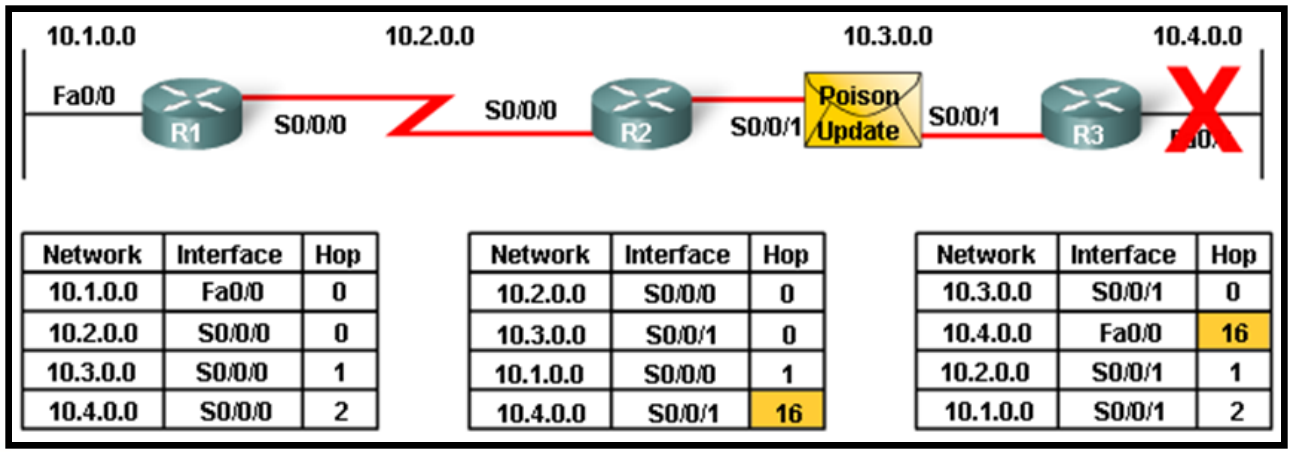

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1
10.4.0.0	S0/0/0	2

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	16

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	16
10.2.0.0	S0/0/1	1
10.1.0.0	S0/0/1	2

RIP Timers

RIP timers must be identical on all routers on the RIP network, otherwise massive instability will occur.

Update Timer

- default 30 seconds – indicates how often the router will send out a routing table update.

Invalid Timer

- default 180 seconds - indicates how long a route will remain in a routing table before being marked as invalid (but not removed from routing table), if no new updates are heard about this route.
- The invalid timer will be reset if an update is received for that particular route before the timer expires.
- Route is marked (and advertised) with a metric of 16 (unreachable) and placed in hold-down-state.

RIP Timers

Hold-down Timer

- default 180 seconds – indicates how long RIP will "suppress" a route that it has placed in a hold-down state.
- RIP will not accept any new updates for routes in a hold-down state, until the hold-down timer expires.
- A route will enter a hold-down state for one of three reasons:
  - The invalid timer has expired.
  - An update has been received from another router, marking that route with a metric of 16 (or unreachable).
  - An update has been received from another router, marking that route with a higher metric than the current metric in the routing table. This is to prevent loops.

RIP Timers

Flush Timer

- default 240 seconds – indicates how long a route can remain in a routing table before being flushed, if no new updates are heard about this route.
- The flush timer runs concurrently with the invalid timer, and thus will flush out a route 60 seconds after it has been marked invalid.

Timers	Default Value	Uses
Hold down timer	180 seconds	Used to hold the routing information for the specified time.
Invalid route timer	180 seconds	Used to keep track of discovered routes
Route update timer	30 seconds	Used to update routing information
Route flush timer	240 seconds	Used to set time interval for any route that becomes invalid and its deletion from the routing table.

Summary

Routing Protocol	Distance Vector	Classless Routing Protocol	Uses Hold-Down Timers	Use of Split Horizon or Split Horizon with Poison Reverse	Max Hop count = 15	Auto Summary	Support CIDR	Supports VLSM	Uses Authen-tication
RIPv1	Yes	No	Yes	Yes	Yes	Yes	No	No	No
RIPv2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Open Shortest Path First (OSPF)

Background of OSPF

- Began in 1987
- 1989 OSPFv1 released in RFC 1131 - experimental version & never deployed
- 1991 OSPFv2 released in RFC 1247
- 1998 OSPFv2 *updated* in RFC 2328
- 1999 OSPFv3 published in RFC 2740

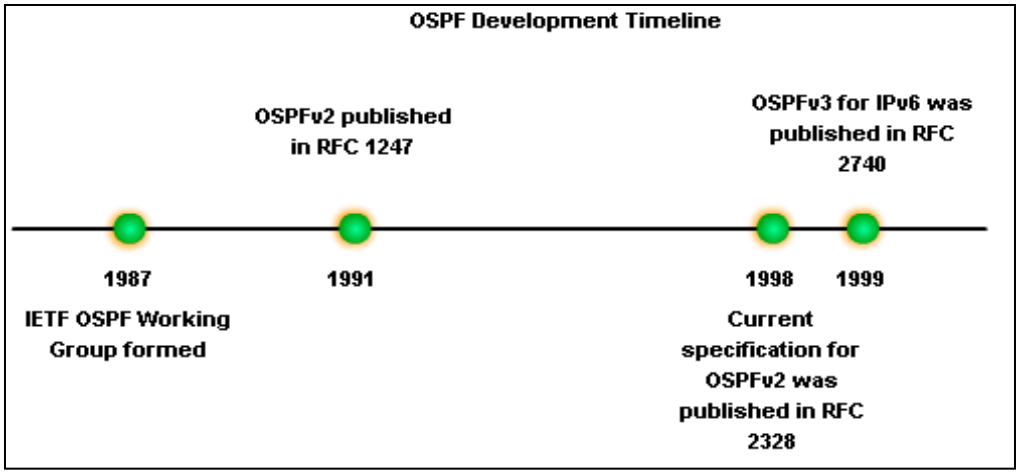

Encapsulation of OSPF Message

OSPF packet type

- 5 packet types

OSPF packet header

- Contains - Router ID and area ID and Type code for OSPF packet type

- IP packet header

  Contains - Source IP address, Destination IP address, & Protocol field set to 89

Encapsulation of OSPF Message

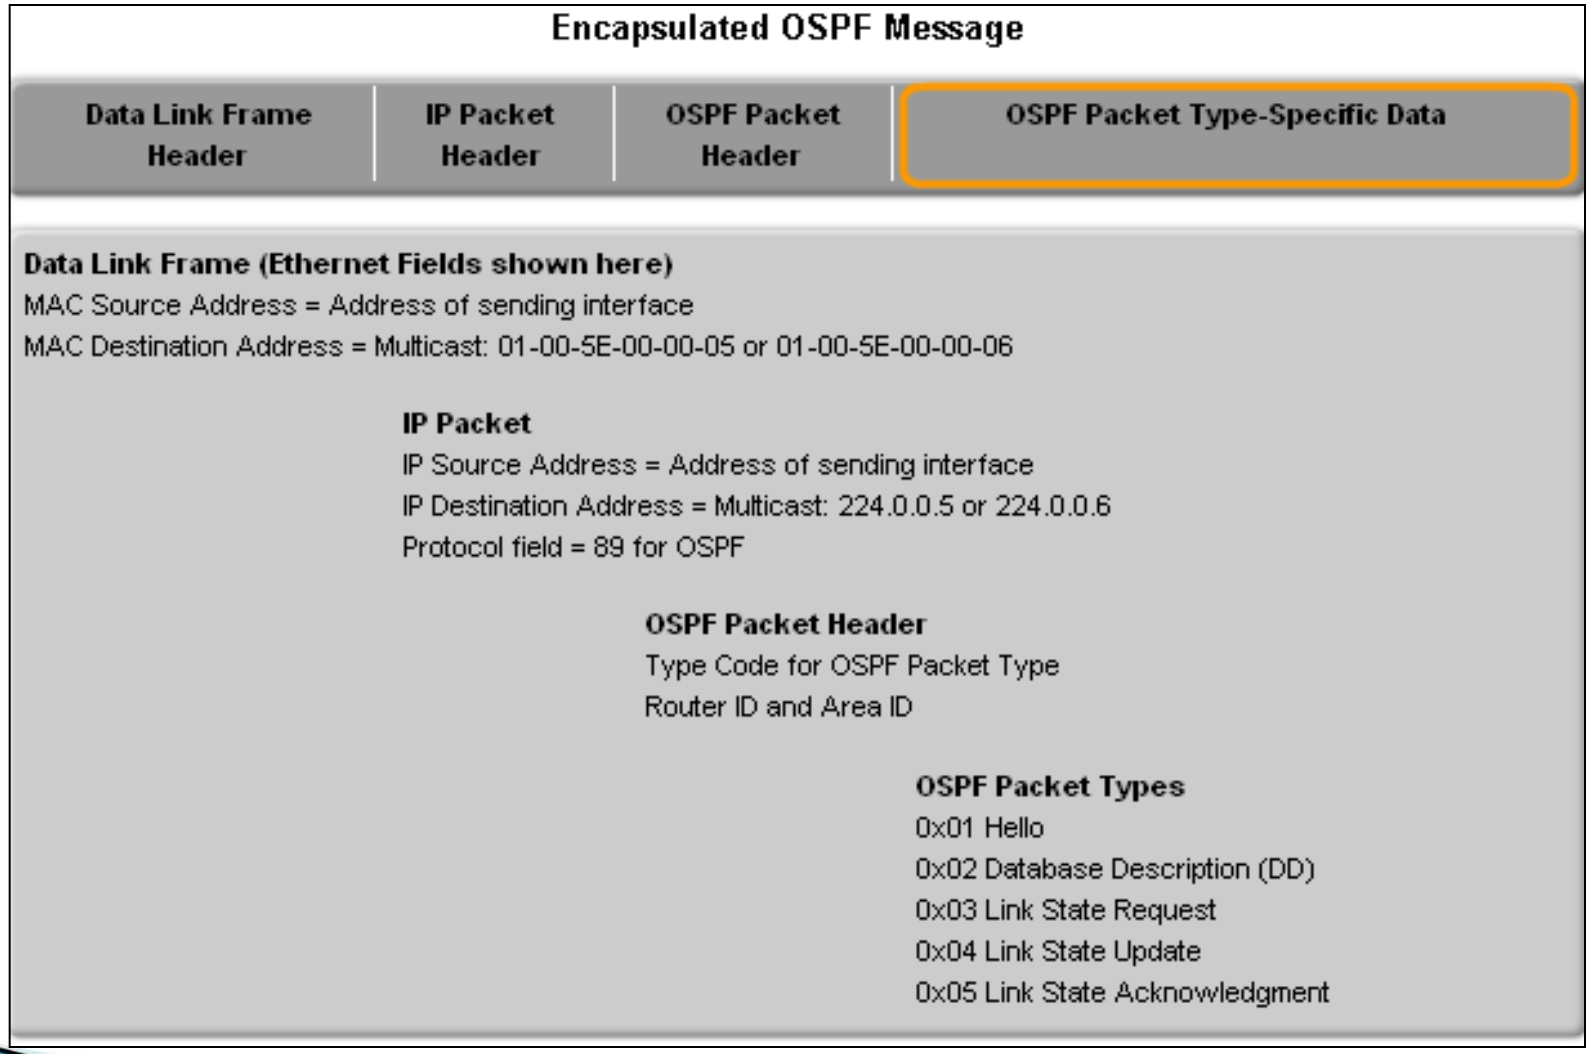

OSPF Packet Types

OSPF Packet Types

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	Database Description (DBD)	Checks for database synchronization between routers
3	Link-State Request (LSR)	Requests specific link-state records from router to router
4	Link-State Update (LSU)	Sends specifically requested link-state records
5	Link-State Acknowledgement (LSAck)	Acknowledges the other packet types

OSPF Packet Types

OSPF Hello Packet

- Discover OSPF neighbors & establish adjacencies;
- Advertise guidelines on which routers must agree to become neighbors;
- Used by multi-access networks to elect a Designated Router (DR) and a Backup Designated Router (BDR).

OSPF Packet Types

Contents of a Hello Packet

- Router ID of transmitting router.

OSPF Hello Intervals

- Usually multicast (224.0.0.5);
- Sent every 30 seconds for NBMA segments.

OSPF Dead Intervals

- This is the time that must transpire before the neighbor is considered down;
- Default time is 4 times the hello interval

OSPF Packet Types

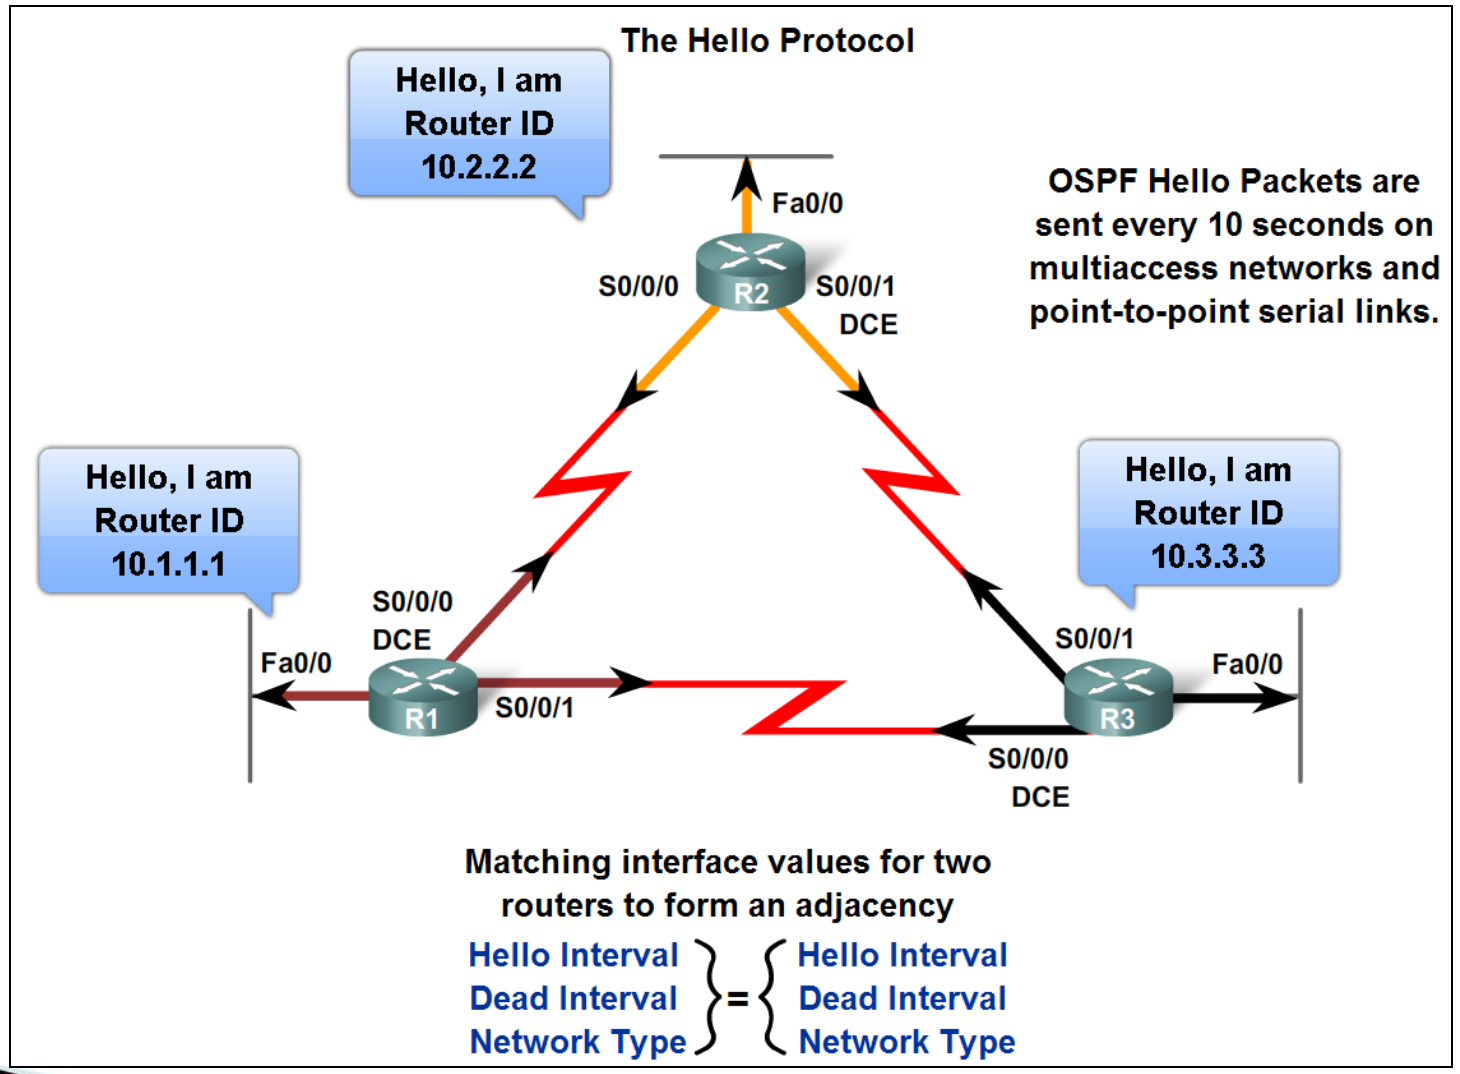

OSPF Packet Types

Hello protocol packets

- Contain information that is used in electing DR and BDR
- Designated Router (DR) is responsible for updating all other OSPF routers
- Backup Designated Router (BDR) takes over DR's responsibilities if DR fails

OSPF Packet Types

OSPF Link-state Updates

- Purpose of a Link State Update (LSU)

  Used to deliver link state advertisements

- Purpose of a Link State Advertisement (LSA)

  Contains information about neighbors & path costs

OSPF Packet Types

LSUs Contain Link-State Advertisements (LSAs)

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	DBD	Checks for database synchronization between router
3	LSR	Requests specific link-state records from router to router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types

The acronyms LSA and LSU are often used interchangeably.

An LSU contains one or more LSAs.

LSAs contain route information for destination networks.

LSA specifics are discussed in CCNP.

LSA Type	Description
1	Router LSAs
2	Network LSAs
3 or 4	Summary LSAs
5	Autonomous System Extrenal LSAs
6	Multicast OSPF LSAs
7	Defined for Not-So-Stubby Areas
8	External Attributes LSA for Border Gatway Protocol(BGP)
9,10,11	Opaque LSAs

OSPF - SPF Algorithm

OSPF routers

- Build & maintain link-state database containing LSA received from other routers;
- Information found in database is utilized upon execution of Dijkstra SPF algorithm;
- SPF algorithm used to create SPF tree ;
- SPF tree used to populate routing table.

OSPF - SPF Algorithm

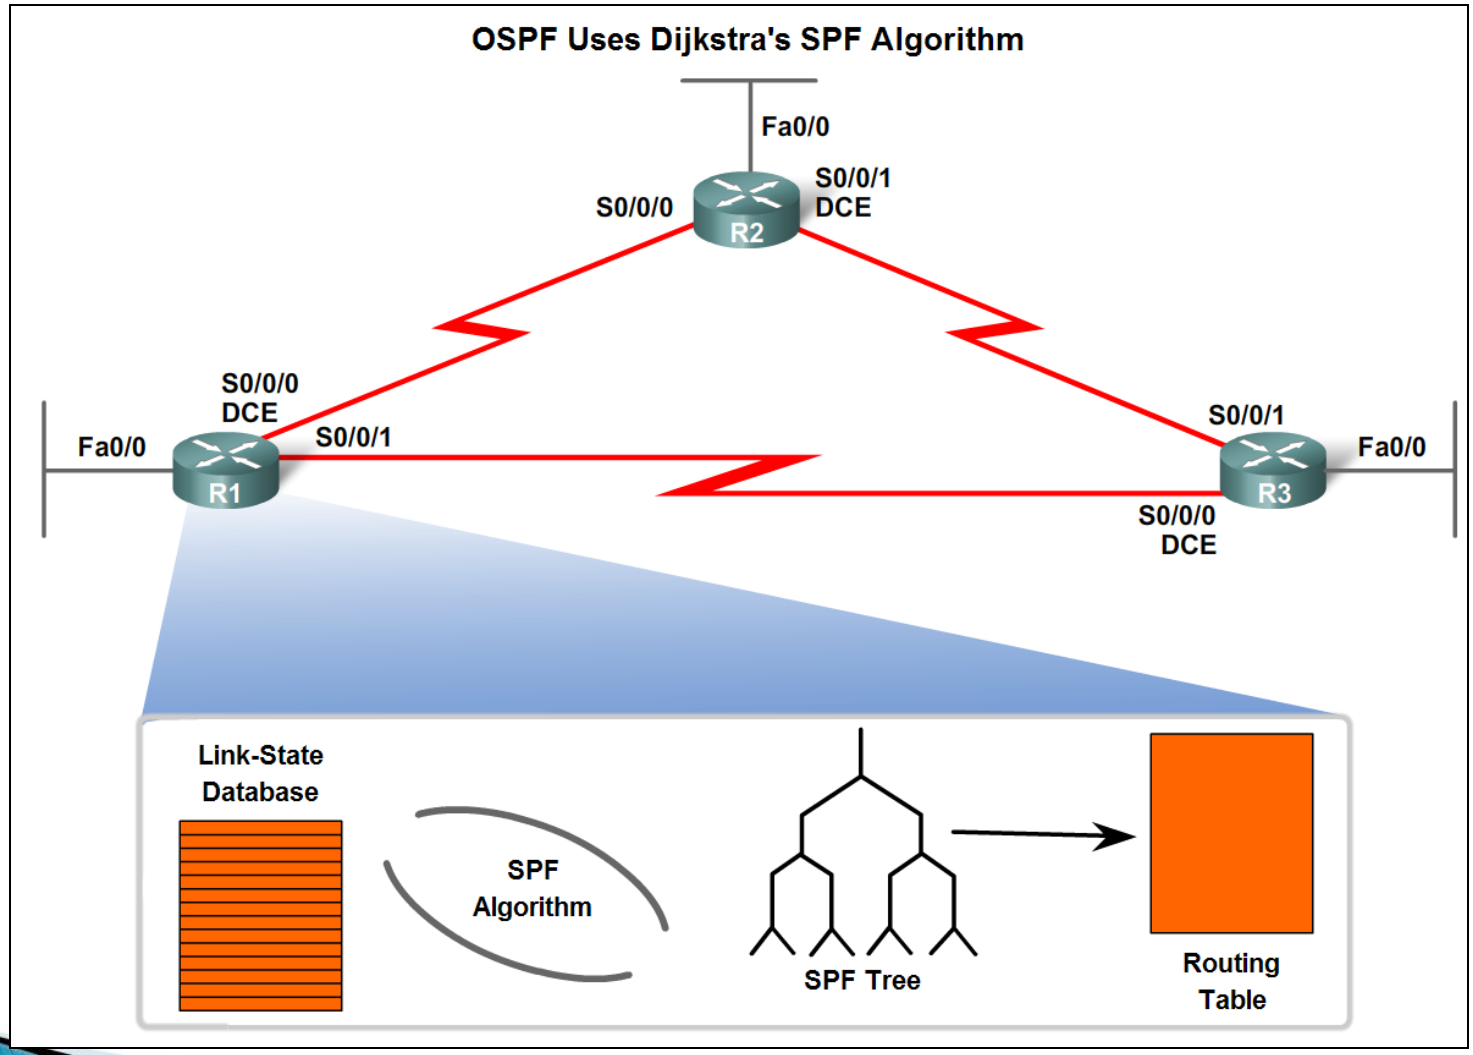

Administrative Distance & Route Preference of OSPF

AD in the Cisco IOS

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

Route Preference in the Juniper JUNOS

Route Preference Values

Source	Default Preference
Direct	0
Local	0
Static	5
OSPF internal	10
RIP	100
Aggregate	130
OSPF AS external	150
BGP (both EGBP and IBGP)	170

OSPF Authentication

OSPF Authentication

- Purpose is to encrypt & authenticate routing information;
- This is an interface specific configuration;
- Routers will only accept routing information from other routers that have been configured with the same password or authentication information.

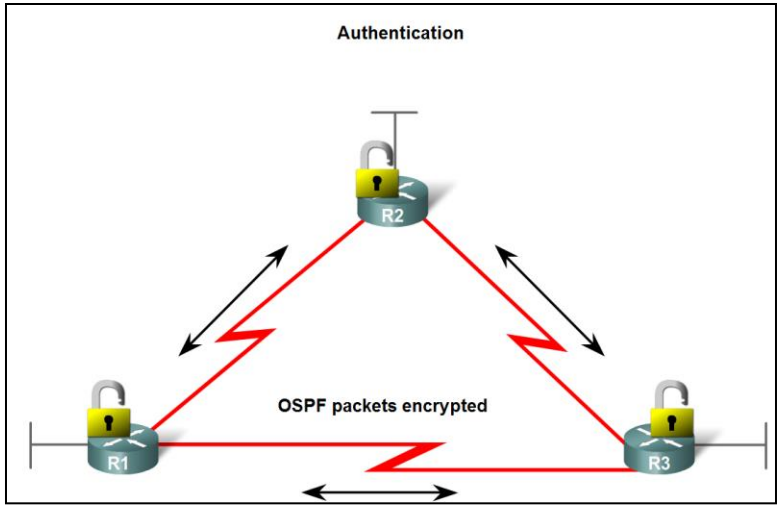

OSPF Router ID

Router ID

- is an IP address used to identify a router

Criteria for deriving the Router ID

- Router-id is set
  - uses IP address configured with OSPF router-id command;
  - takes precedence over loopback and physical interface addresses
- Router-id is not set
  - router chooses highest IP address of any loopback interfaces
- No loopback interfaces are configured
  - the highest IP address on any active interface is used

OSPF Router ID

Router ID & Loopback addresses

- Highest loopback address will be used as router ID if router-id command isn't used
- Advantage of using loopback address the loopback interface cannot fail → OSPF stability

OSPF area

- is a group of routers that share link state information

OSPF Timers

OSPF default Timers in different environments

- if a router stops receiving hello messages at regular intervals - the hello interval, from a neighbor, after a set period - the dead interval, the router will assume the neighbor has gone down.
- By default the dead interval is four times (4x) the hello interval.

OSPF Network Type	Default HelloInterval	Default RouterDeadInterval
Broadcast	10 seconds	40 seconds
Non-broadcast	30 seconds	120 seconds
Point-to-Point	10 seconds	40 seconds
Point-to-Multipoint	30 seconds	120 seconds
Point-to-Multipoint Non-broadcast	30 seconds	120 seconds
Loopback	N/A	N/A

OSPF Metric

OSPF cost

- OSPF uses cost as the metric for determining the best route:
  - The best route will have the lowest cost;
  - Cost is based on bandwidth of an interface;
  - Cost is calculated using the formula

$$10^8 / \text{bandwidth}$$

Reference bandwidth

- defaults to 100Mbps

OSPF Cost

Costs of different interfaces

Interface Type	10^8/bps = Cost
Fast Ethernet and faster	10^8/100,000,000 bps = 1
Ethernet	10^8/10,000,000 bps = 10
E1	10^8/2,048,000 bps = 48
T1	10^8/1,544,000 bps = 64
128 kbps	10^8/128,000 bps = 781
64 kbps	10^8/64,000 bps = 1562
56 kbps	10^8/56,000 bps = 1785

COST of an OSPF route

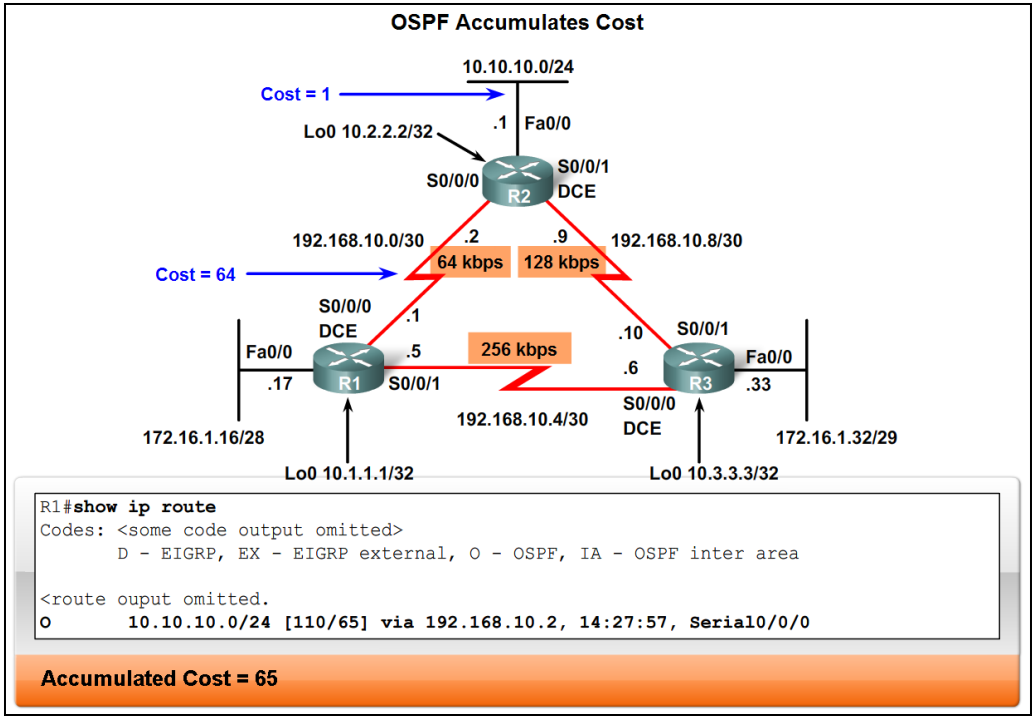

OSPF and Multiaccess Networks

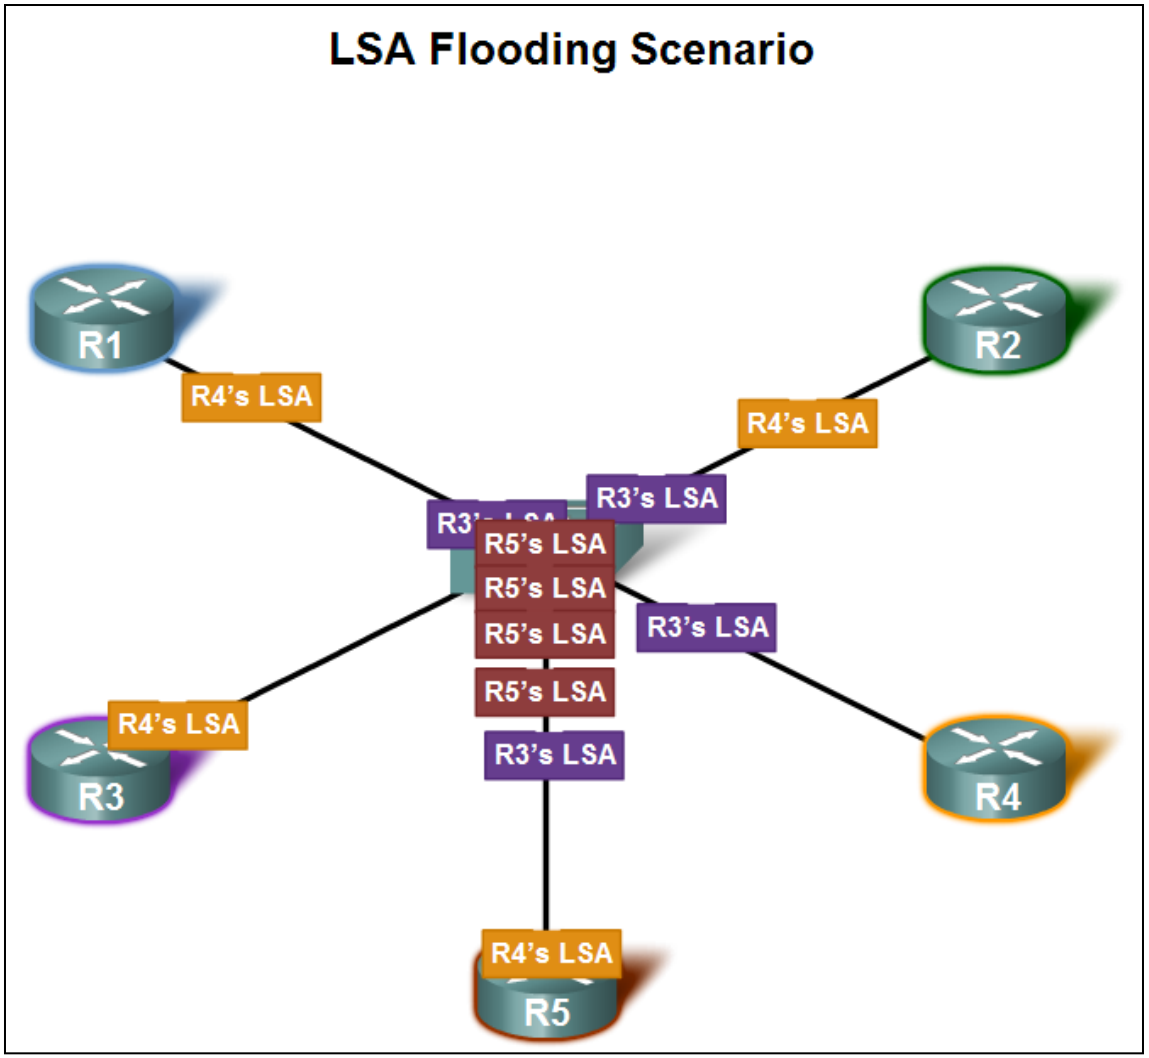

OSPF in Multiaccess Networks

Solution to LSA flooding:

Designated router (DR)

Backup designated router (BDR)

- DR & BDR selection
  - Routers are elected to send & receive LSA
- Sending & Receiving LSA
  - DRothers send LSAs via multicast 224.0.0.6 to DR & BDR
  - DR forward LSA via multicast address 224.0.0.5 to all other routers

OSPF and Multiaccess Networks

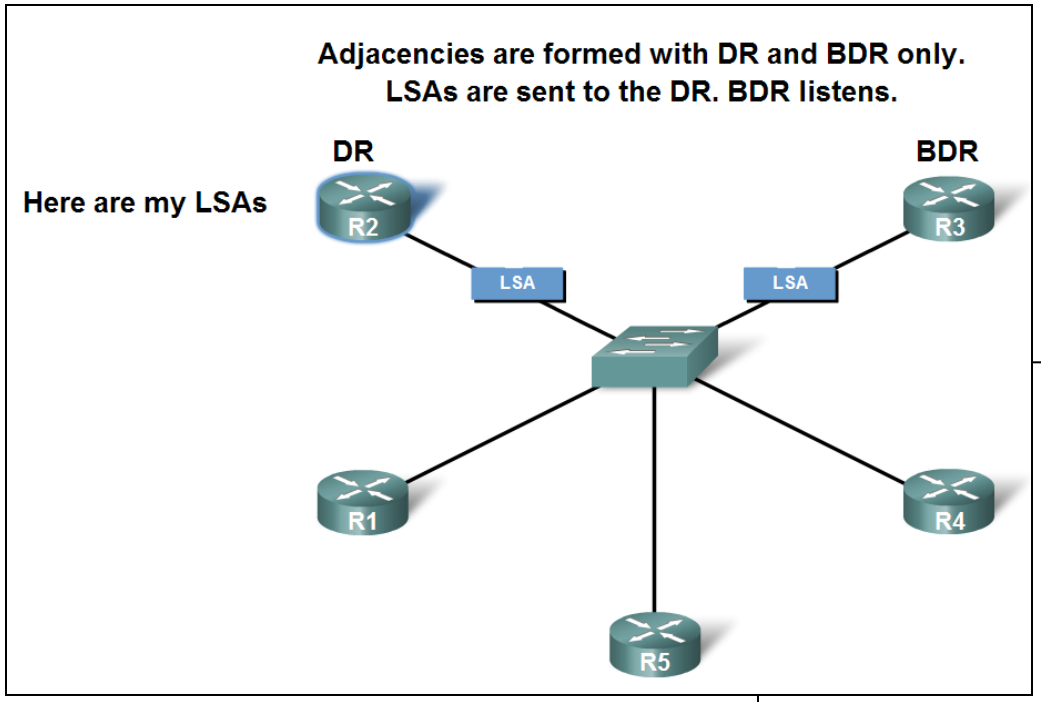

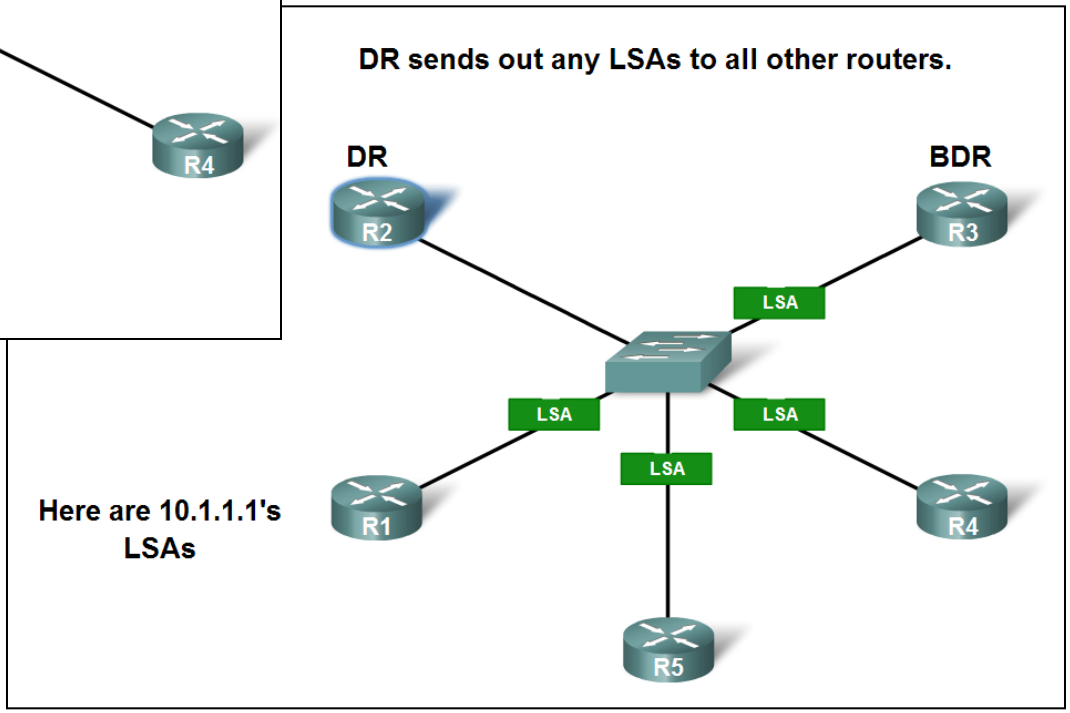

OSPF in Point to Point Networks

DR/BDR elections DO NOT occur in point to point networks

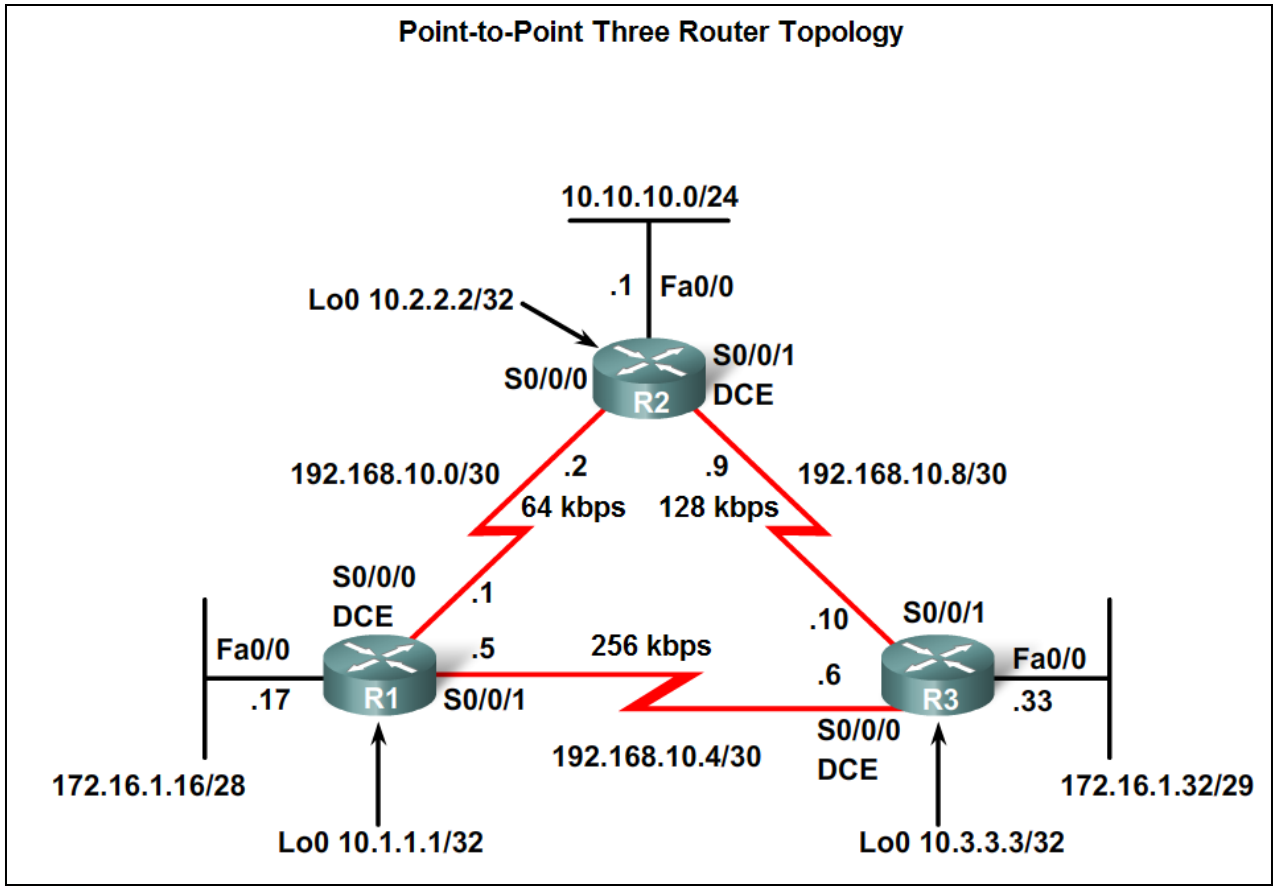

OSPF in Multiaccess Networks

DR/BDR elections in multiaccess networks

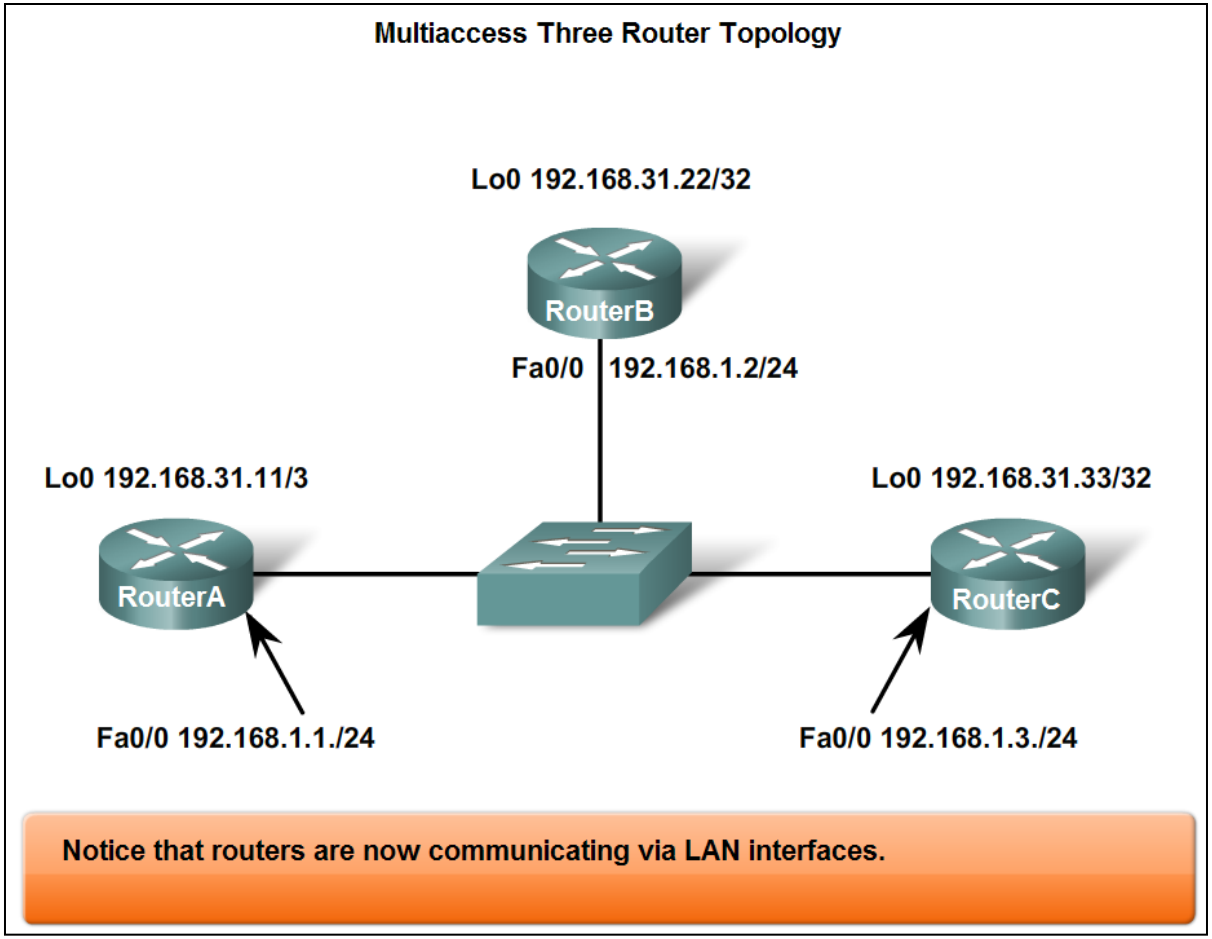

OSPF in Multiaccess Networks

Criteria for getting elected DR/BDR

- 1. DR: Router with the highest OSPF interface priority.
- 2. BDR: Router with the second highest OSPF interface priority.
- 3. If OSPF interface priorities are equal, the highest router ID is used to break the tie.

OSPF in Multiaccess Networks

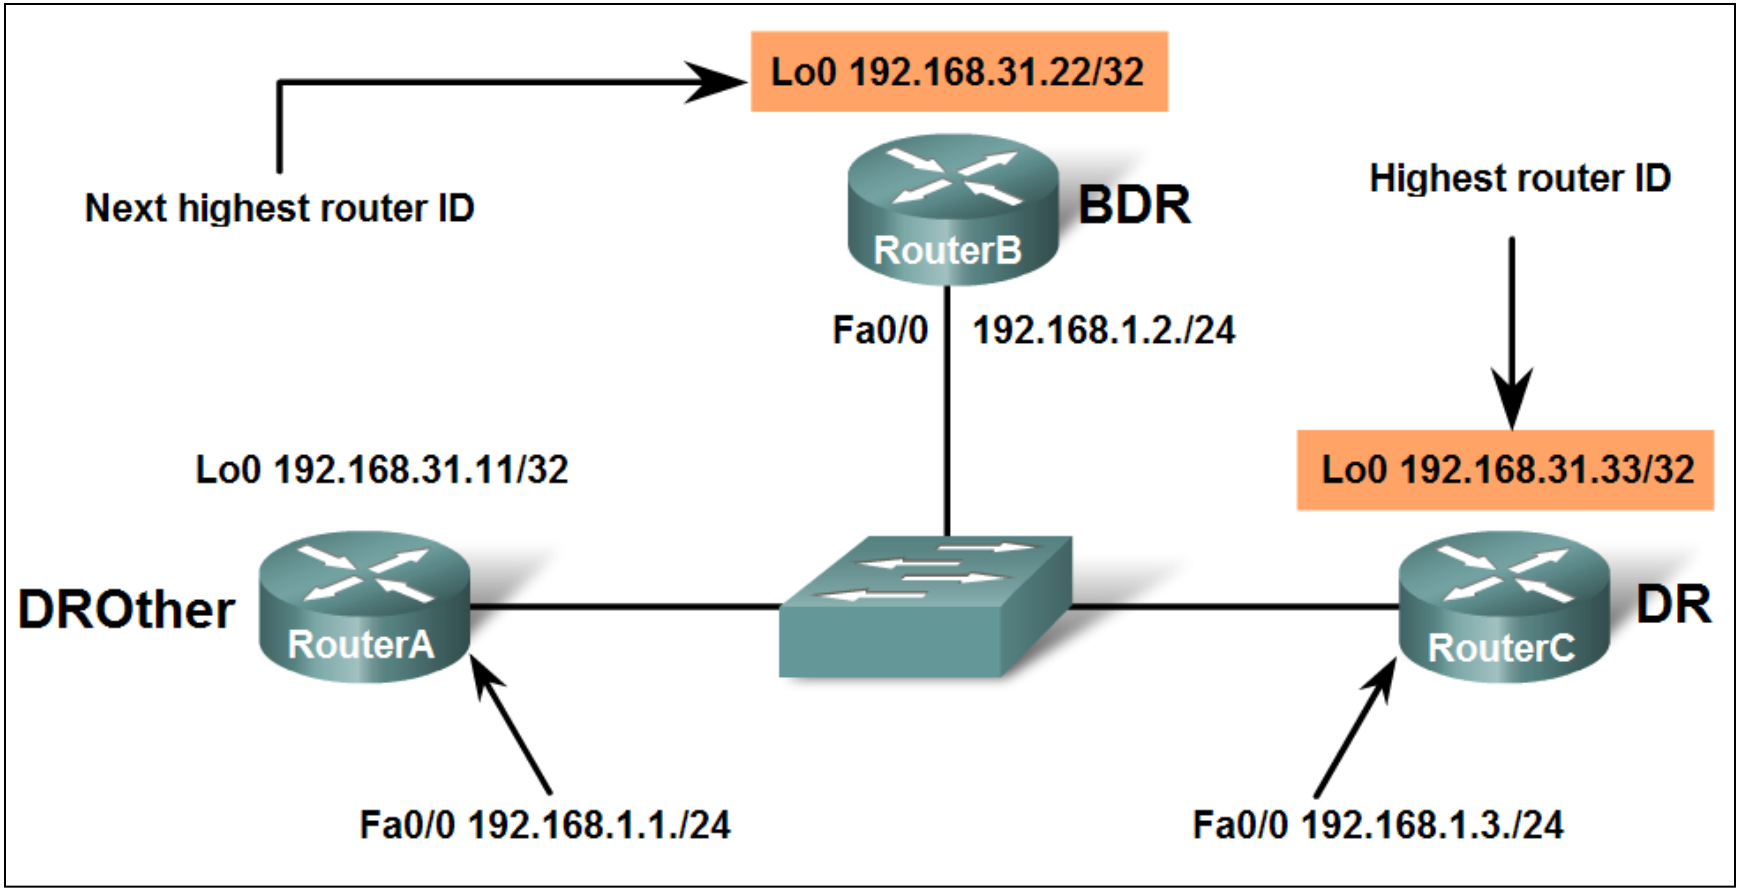

OSPF in Multiaccess Networks

```
RouterA#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.31.33     1   FULL/DR         00:00:39    192.168.1.3     FastEthernet0/0
192.168.31.22     1   FULL/BDR        00:00:36    192.168.1.2     FastEthernet0/0
```

```
RouterB#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.31.33     1   FULL/DR         00:00:34    192.168.1.3     FastEthernet0/0
192.168.31.11     1   FULL/DROTHER    00:00:38    192.168.1.1     FastEthernet0/0
```

```
RouterC#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.31.22     1   FULL/BDR        00:00:35    192.168.1.2     FastEthernet0
192.168.31.11     1   FULL/DROTHER    00:00:32    192.168.1.1     FastEthernet0
```

Priority is equal at the default value of 1.

OSPF in Multiaccess Networks

Timing of DR/BDR Election

Occurs as soon as 1st router has its interface enabled on multiaccess network

- When a DR is elected it remains as the DR until one of the following occurs:
  - The DR fails;
  - The OSPF process on the DR fails;
  - The multiaccess interface on the DR fails.

OSPF in Multiaccess Networks

Influence the election of DR & BDR

Do one of the following:

- Boot up the DR first, followed by the BDR, and then boot all other routers,

  OR

- Shut down the interface on all routers, followed by a no shutdown on the DR, then the BDR, and then all other routers.

OSPF in Multiaccess Networks

OSPF Interface Priority

- Priority number range 0 to 255
  - "0" means the router cannot become the DR or BDR
  - "1" is the default priority value

Single-Area OSPF Limitations

Single-area OSPF is useful in smaller networks. If an area becomes too big, the following issues must be addressed:

- Large routing table (no summarization by default)
- Large link-state database (LSDB)
- Frequent SPF algorithm calculations

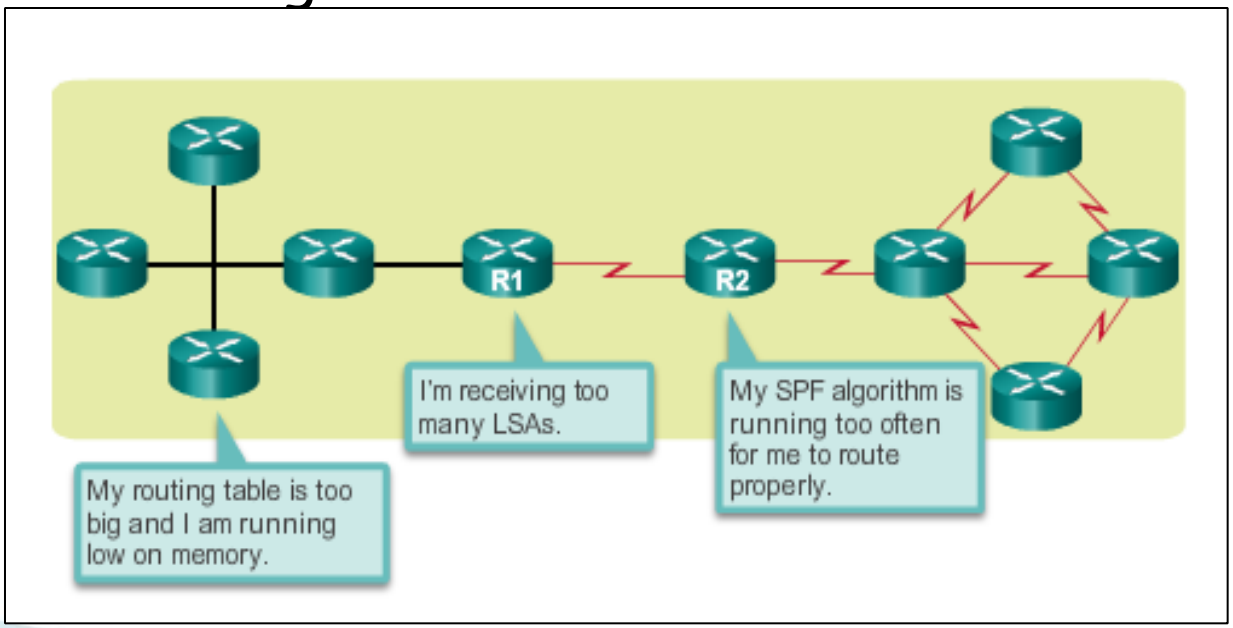

Multi-Area OSPF Router Types

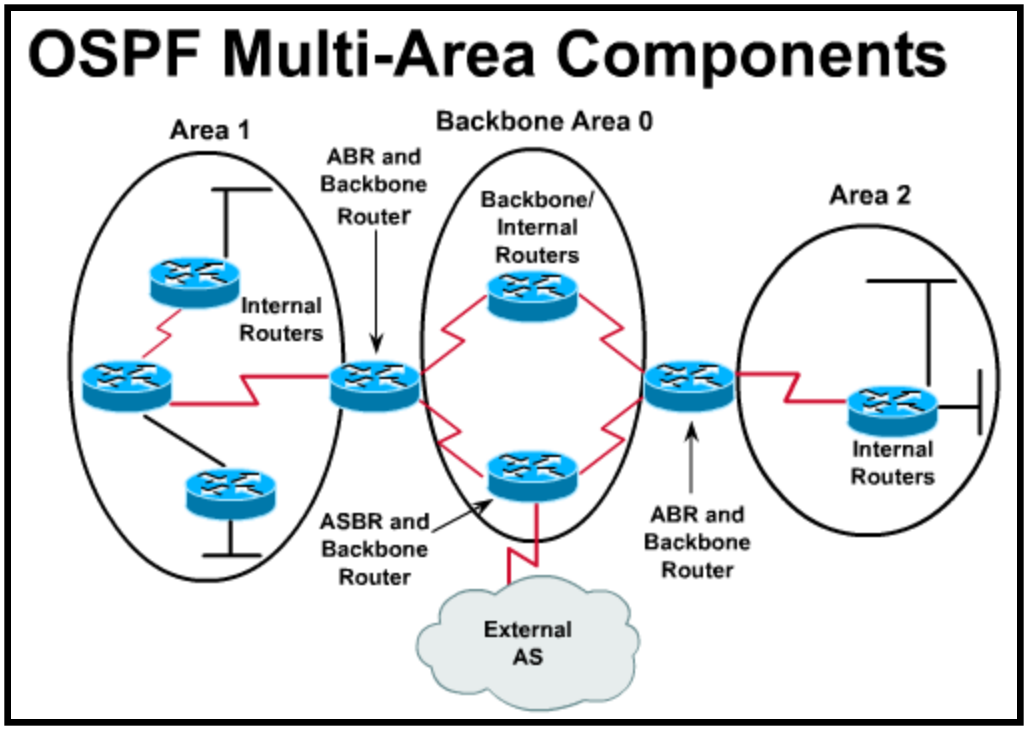

Backbone Router: Router with at least one interface in Area 0

Internal Router: Router with all interfaces in the same area

Area Border Router (ABR): Router with interfaces in two or more different areas

Autonomous System Boundary Router (ASBR): Router with at least one interface connected to a non-OSPF domain

Multi-Area OSPF Router Types

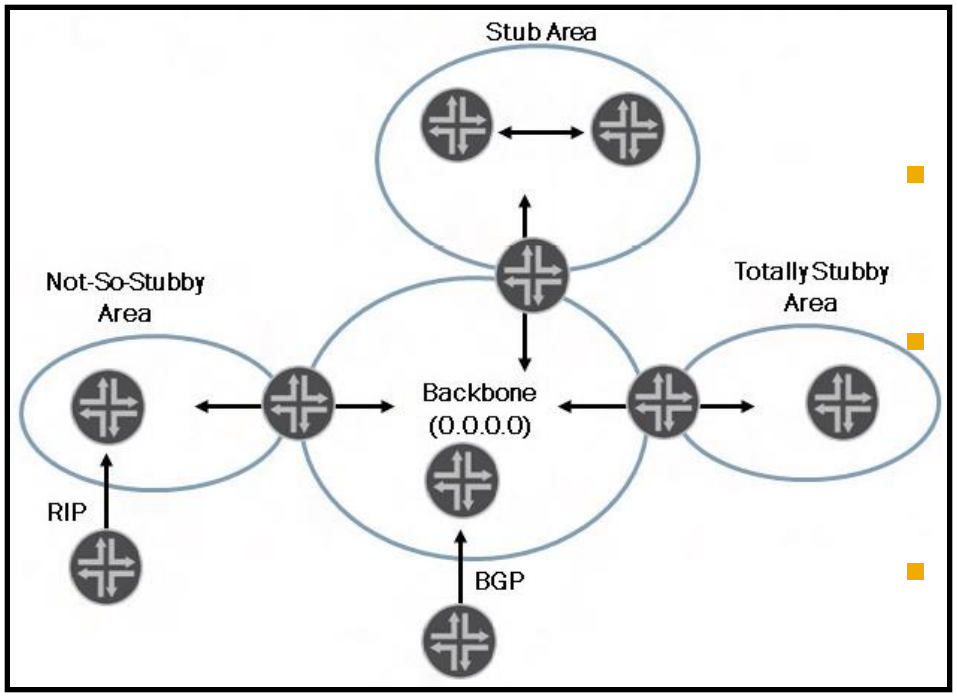

- **Backbone Area** – distributes routing information amongst all other areas and has to be connected to all areas.
- **Stub Area** - AS external advertisements are not flooded
- **Totally Stub Area** - receives only the default route from the backbone.
- **Not-So-Stubby-Area** - allows external routes to be flooded within the area and to be leaked into other areas, external routes from other areas do not enter the NSSA, receives only default route from the backbone.

OSPF LSA Packet Types

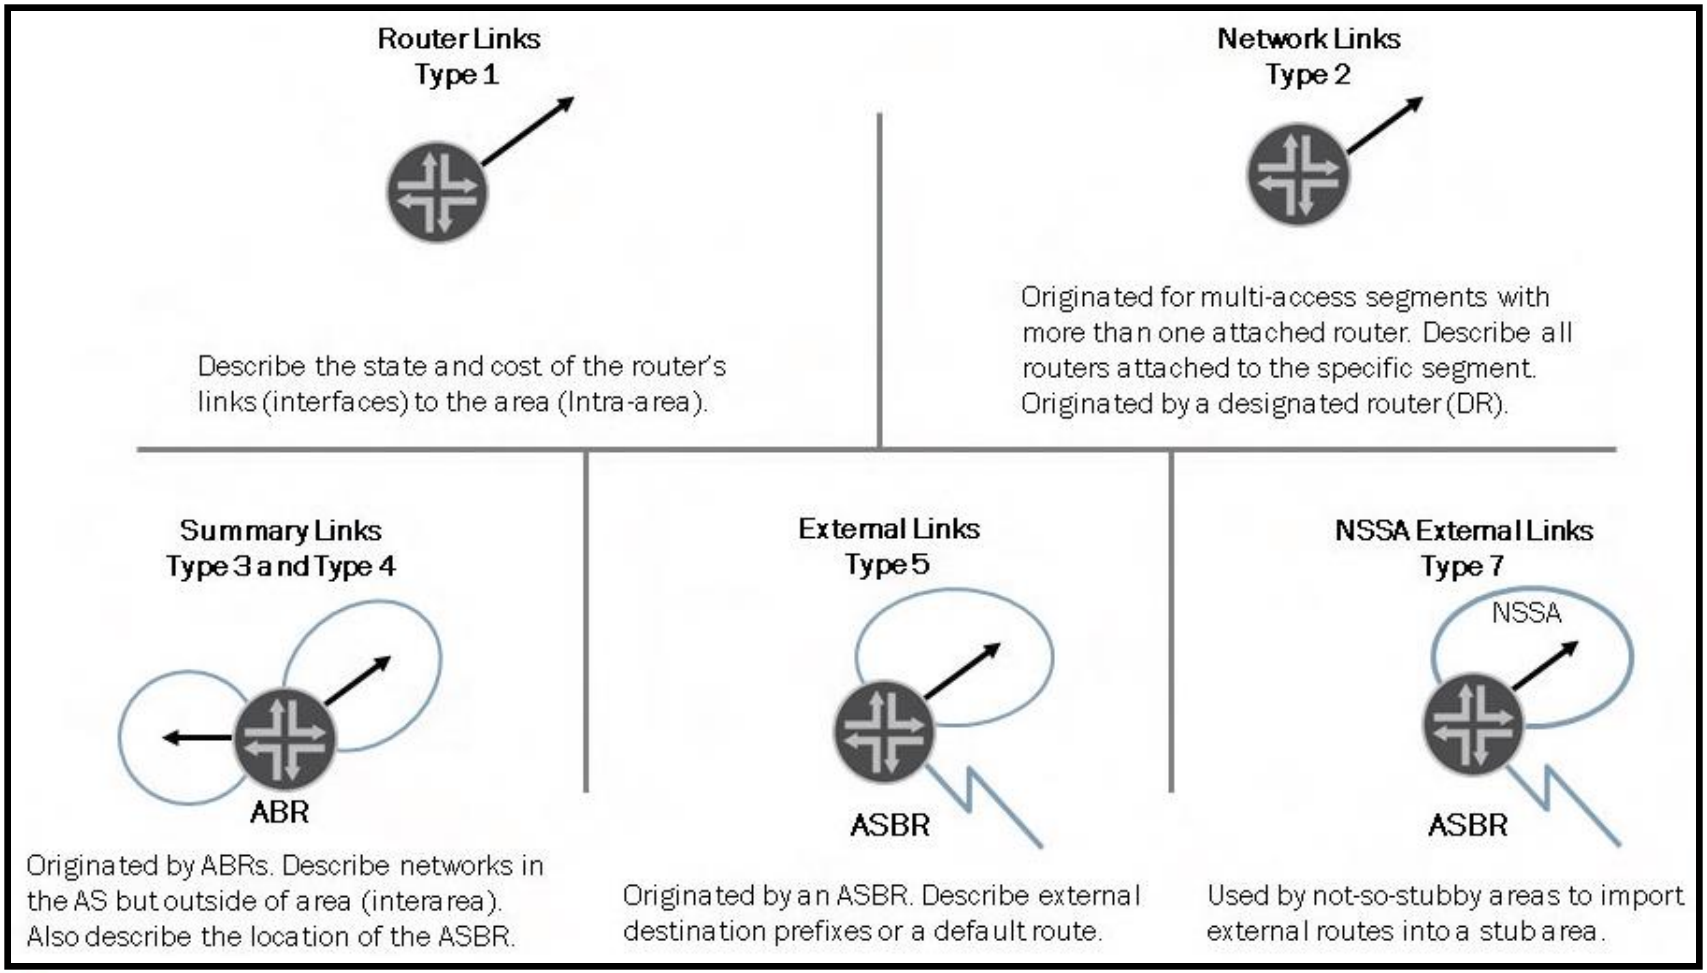

OSPF LSA Packet Types

- Type 1 - Router LSAs - describe the interfaces and neighbors of each OSPF router to all other OSPF routers in the same area (intra-area).
- Type 2 - Network LSAs - describe an Ethernet segment. These LSAs are sent by the designated router to other OSPF routers in the same area (intra-area).
- Type 3 - Summary LSAs - describe IP prefixes learned from Router and Network LSAs. These LSAs are sent by the ABR attached to the area from where the prefix information was learned and sent to other OSPF areas (interarea).

OSPF LSA Packet Types

- Type 4 - ASBR Summary LSAs - sent by the ABR attached to the area in which the ASBR is located to other OSPF areas (interarea).
- Type 5 - External LSAs - describe IP prefixes redistributed from other routing protocols, such as RIP, BGP, or even static routes and are sent by ASBRs.
- Type 7 - NSSA External LSAs - describe IP prefixes redistributed from other routing protocols, such as RIP, BGP, or even static routes and are sent by ASBRs in NSSA areas.

OSPF Routing Table Entries

```
R1# show ip route
Codes:L - local, C-connected, S-static, 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
      i - IS-IS, su-IS-IS summary, L1-IS-IS level-1, L2-IS-IS level-2
      ia - IS-IS inter area,*-candidate default,U-per-user static route
      o - ODR, P-periodic downloaded static route, H-NHRP, l-LISP
      + - replicated route, % - next hop override

Gateway of last resort is 192.168.10.2 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 192.168.10.2, 00:00:19, Serial0/0/0
      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C        10.1.1.0/24 is directly connected, GigabitEthernet0/0
L        10.1.1.1/32 is directly connected, GigabitEthernet0/0
C        10.1.2.0/24 is directly connected, GigabitEthernet0/1
L        10.1.2.1/32 is directly connected, GigabitEthernet0/1
O        10.2.1.0/24 [110/648] via 192.168.10.2, 00:04:34, Serial0/0/0
O IA  192.168.1.0/24 [110/1295] via 192.168.10.2, 00:01:48,Serial0/0/0
O IA  192.168.2.0/24 [110/1295] via 192.168.10.2, 00:01:48,Serial0/0/0
      192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
C        192.168.10.0/30 is directly connected, Serial0/0/0
L        192.168.10.1/32 is directly connect
O        192.168.10.4/30 [110/1294] via 192.
```

Cisco IOS

Juniper JUNOS

```
user@R2> show route 172.18.1.0/24

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.18.1.0/24       *[OSPF/150] 02:37:46, metric 0, tag 0
                     > to 172.26.1.1 via ge-0/0/3.0

user@R2> show route 172.26.4.0/30

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.26.4.0/30       *[OSPF/10] 02:24:29, metric 2
                     > to 172.26.2.2 via ge-0/0/1.0
```

External prefix injected by R1

Remote subnet connecting R3 and R4 is reachable through desired path.

Thank You!