Chapter 08

Dynamic Routing - OSPF

Introduction to OSPF

OSPF Message Encapsulation

- OSPF packet type

  There exist 5 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

Introduction to OSPF

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

Introduction to OSPF

Hello Protocol

- OSPF Hello Packet

  - Purpose of 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 **D**esignated **R**outer and a **B**ackup **D**esignated **R**outer

Introduction to OSPF

- Hello Packets continued
  - 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 and every 10 on multi-access 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

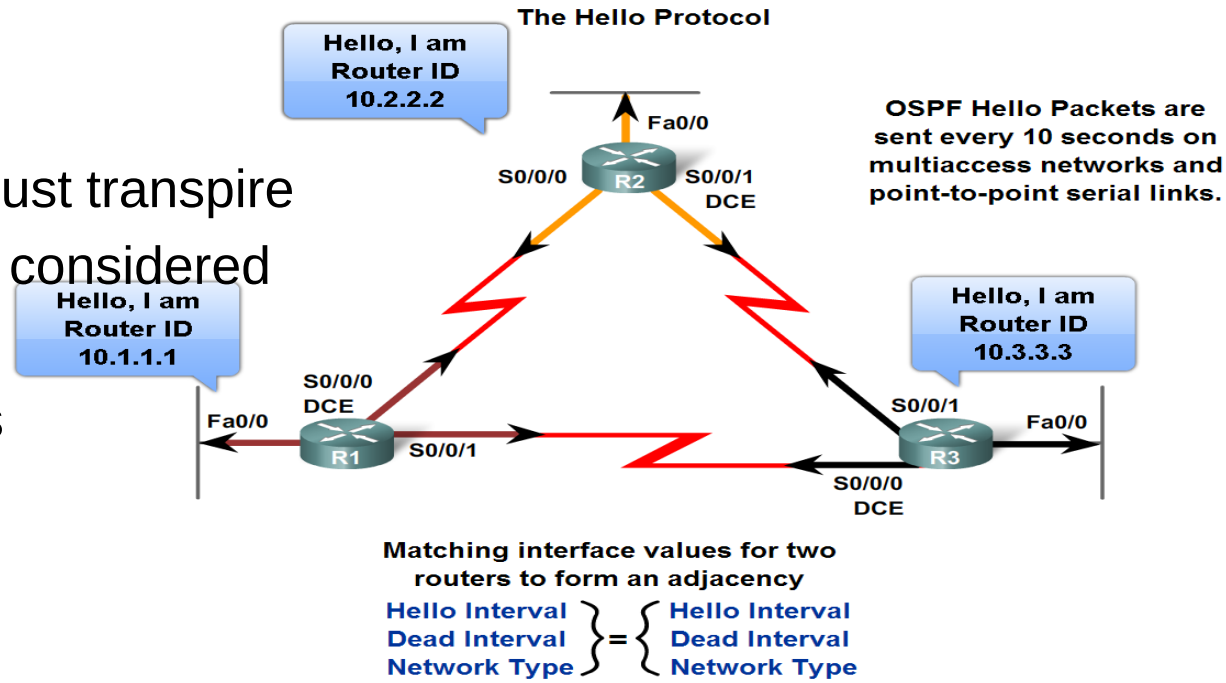

Introduction to OSPF

- Hello protocol packets contain information that is used in electing

  - Designated Router (DR)

    - DR is responsible for updating all other OSPF routers

  - Backup Designated Router (BDR)

    - This router takes over DR’s responsibilities if DR fails

Introduction to OSPF

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

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

Introduction to OSPF

OSPF 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

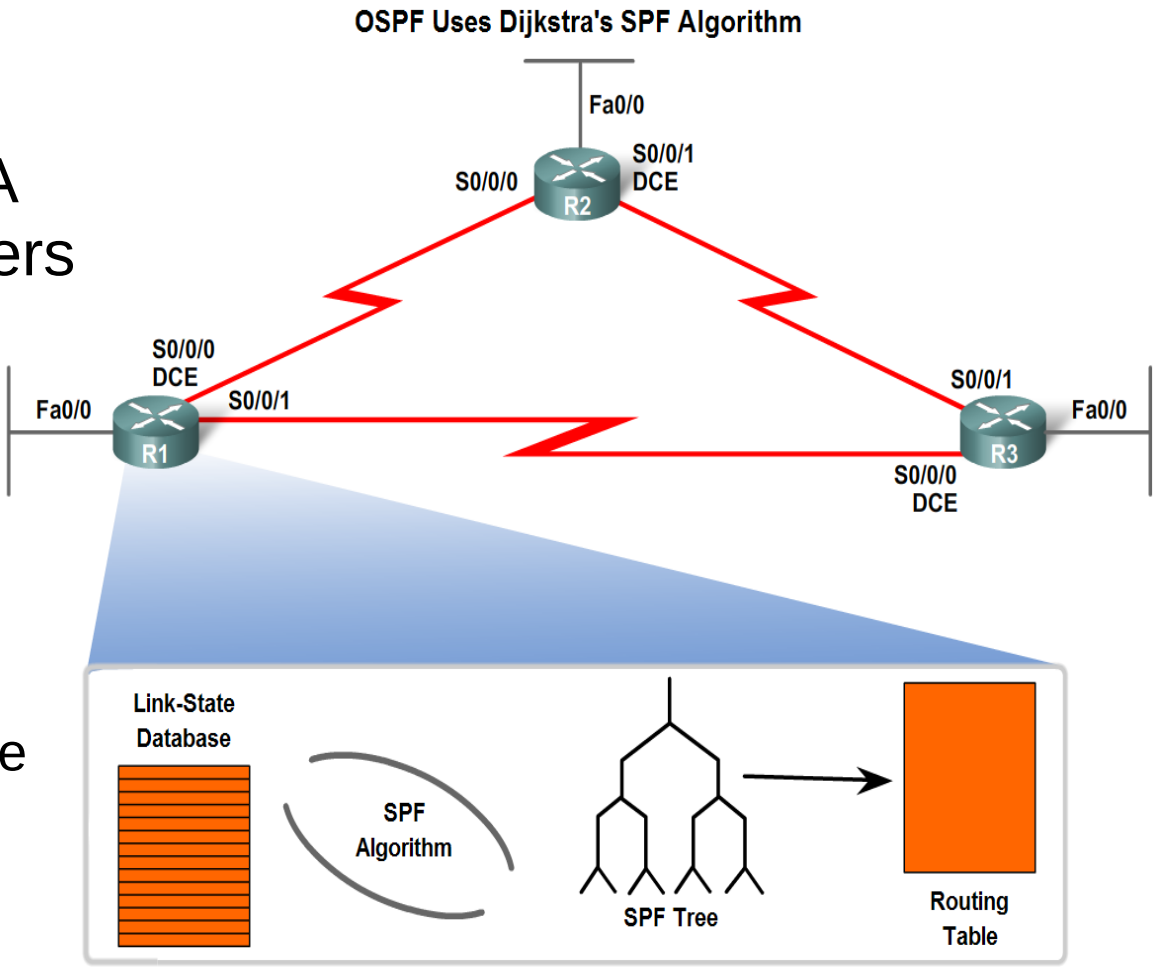

Introduction to OSPF

Administrative Distance

- Default Administrative Distance for OSPF is 110

Default Administrative Distances

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

Introduction to OSPF

- 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

Basic OSPF Configuration

The router ospf command

- To enable OSPF on a router use the following command

  R1(config)#**router ospf *process-id***

  Process id

  - A locally significant number between **1** and **65535**

    -this means it does not have to match other OSPF routers

```
R1(config)#router ospf 1
R1(config-router)#
```

```
R2(config)#router ospf 1
R2(config-router)#
```

```
R3(config)#router ospf 1
R3(config-router)#
```

Basic OSPF Configuration

- OSPF network command

-Requires entering: **network address**

wildcard mask - the inverse of the subnet mask

area-id - area-id refers to the OSPF area. OSPF area is a group of routers that share link state information

-Example: Router(config-router)#**network** network-address wildcard-ask **area** area-id

```
R1(config)#router ospf 1
R1(config-router)#network 172.16.1.16 0.0.0.15 area 0
R1(config-router)#network 192.168.10.0 0.0.0.3 area 0
R1(config-router)#network 192.168.10.4 0.0.0.3 area 0
```

```
R2(config)#router ospf 1
R2(config-router)#network 10.10.10.0 0.0.0.255 area 0
R2(config-router)#network 192.168.10.0 0.0.0.3 area 0
R2(config-router)#network 192.168.10.8 0.0.0.3 area 0
```

Basic OSPF Configuration

- Router ID
  - This is an IP address used to identify a router
  - 3 criteria for deriving the router ID
    - Use IP address configured with OSPF *router-id* command

      - Takes precedence over loopback and physical interface addresses
    - If router-id command not used then router chooses highest IP address of any loopback interfaces
    - If no loopback interfaces are configured then the highest IP address on any active interface is used

Basic OSPF Configuration

OSPF Router ID

- Commands used to verify current router ID
  - Show ip protocols
  - Show ip ospf
  - Show ip ospf interface

```
R1#show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 192.168.10.5
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
```

```
R2#show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 192.168.10.9
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
```

```
R3#show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 192.168.10.10
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
```

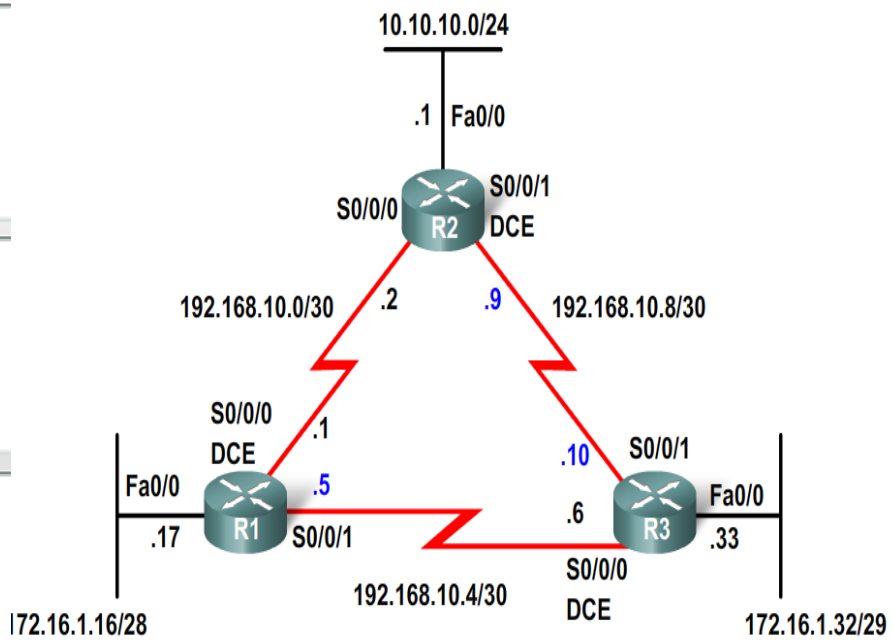

Basic OSPF Configuration

OSPF Router ID

```
R1(config)#interface loopback 0
R1(config-if)#ip add 10.1.1.1 255.255.255.255
```

- 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
- The OSPF router-id command
  - Introduced in IOS 12.0
  - Command syntax
    - Router(config)#router ospf *process-id*
    - Router(config-router)#router-id *ip-address*
- Modifying the Router ID
  - After that you should use the command Router#clear ip ospf process

Basic OSPF Configuration

Verifying OSPF

- Use the show ip ospf command to verify & troubleshoot OSPF networks

  Command will display the following:

  - Neighbor adjacency

    -No adjacency indicated by -

    - Neighboring router's Router ID is not displayed
    - A state of **full** is not displayed

    -Consequence of no adjacency-

    - No link state information exchanged
    - Inaccurate SPF trees & routing tables

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.3.3.3          1   FULL/  -        00:00:30    192.168.10.6    Serial0/0/1
10.2.2.2          1   FULL/  -        00:00:33    192.168.10.2    Serial0/0/0
```

Basic OSPF Configuration

Verifying OSPF - Additional Commands

Command	Description
Show ip protocols	Displays OSPF process ID, router ID, networks router is advertising & administrative distance
Show ip ospf	Displays OSPF process ID, router ID, OSPF area information & the last time SPF algorithm calculated
Show ip ospf interface	Displays hello interval and dead interval

Basic OSPF Configuration

Examining the routing table

- Use the show ip route command to display the routing table

  - An “O’ at the beginning of a route indicates that the router source is OSPF

  - Note OSPF does not automatically summarize at major network boundaries

```
R1#show ip route

Codes: <some code output omitted>
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

Gateway of last resort is not set

     192.168.10.0/30 is subnetted, 3 subnets
C       192.168.10.0 is directly connected, Serial0/0/0
C       192.168.10.4 is directly connected, Serial0/0/1
O       192.168.10.8 [110/128] via 192.168.10.2, 14:27:57, Serial0/0/0
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
O       172.16.1.32/29 [110/65] via 192.168.10.6, 14:27:57, Serial0/0/1
C       172.16.1.16/28 is directly connected, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O       10.10.10.0/24 [110/65] via 192.168.10.2, 14:27:57, Serial0/0/0
C       10.1.1.1/32 is directly connected, Loopback0
```

OSPF Metric

- 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
    - can be modified using
      - **Router(config-router)#auto-cost reference-bandwidth**

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

OSPF Metric

- Cost of an OSPF route

  Is the accumulated value from one router to the next

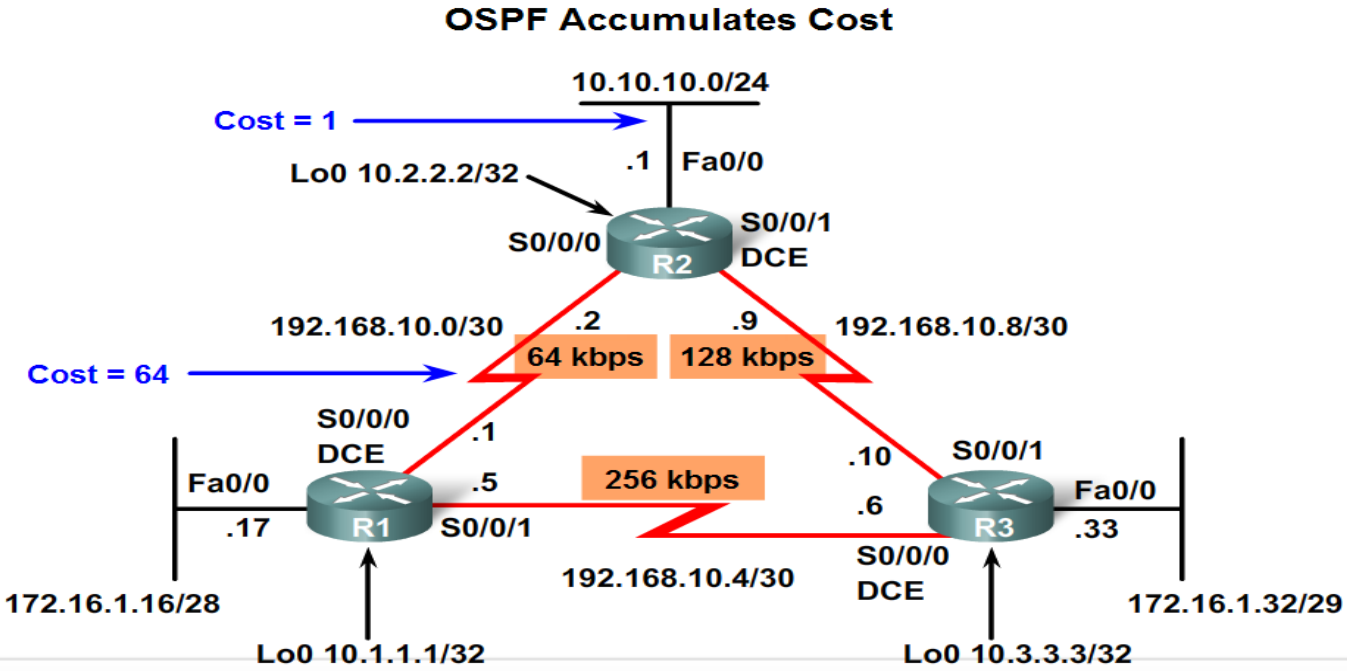


```
R1#show ip route
Codes: <some code output omitted>
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

<route ouput omitted.
O       10.10.10.0/24 [110/65] via 192.168.10.2, 14:27:57, Serial0/0/0
```

Accumulated Cost = 65

OSPF in Multiaccess Networks

- Solution to LSA flooding issue is the use of
  - Designated router (DR)
  - Backup designated router (BDR)
  - DRothers
- 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

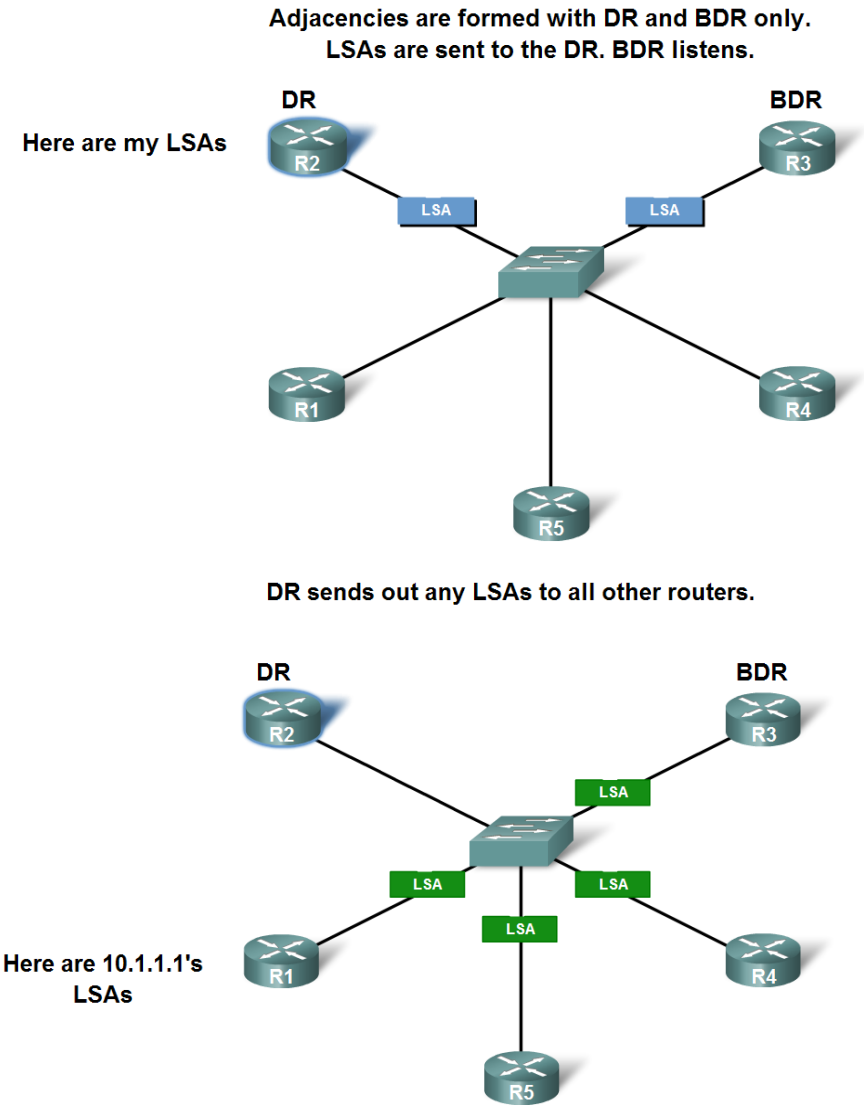

OSPF in Multiaccess Networks

DR/BDR Election Process

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

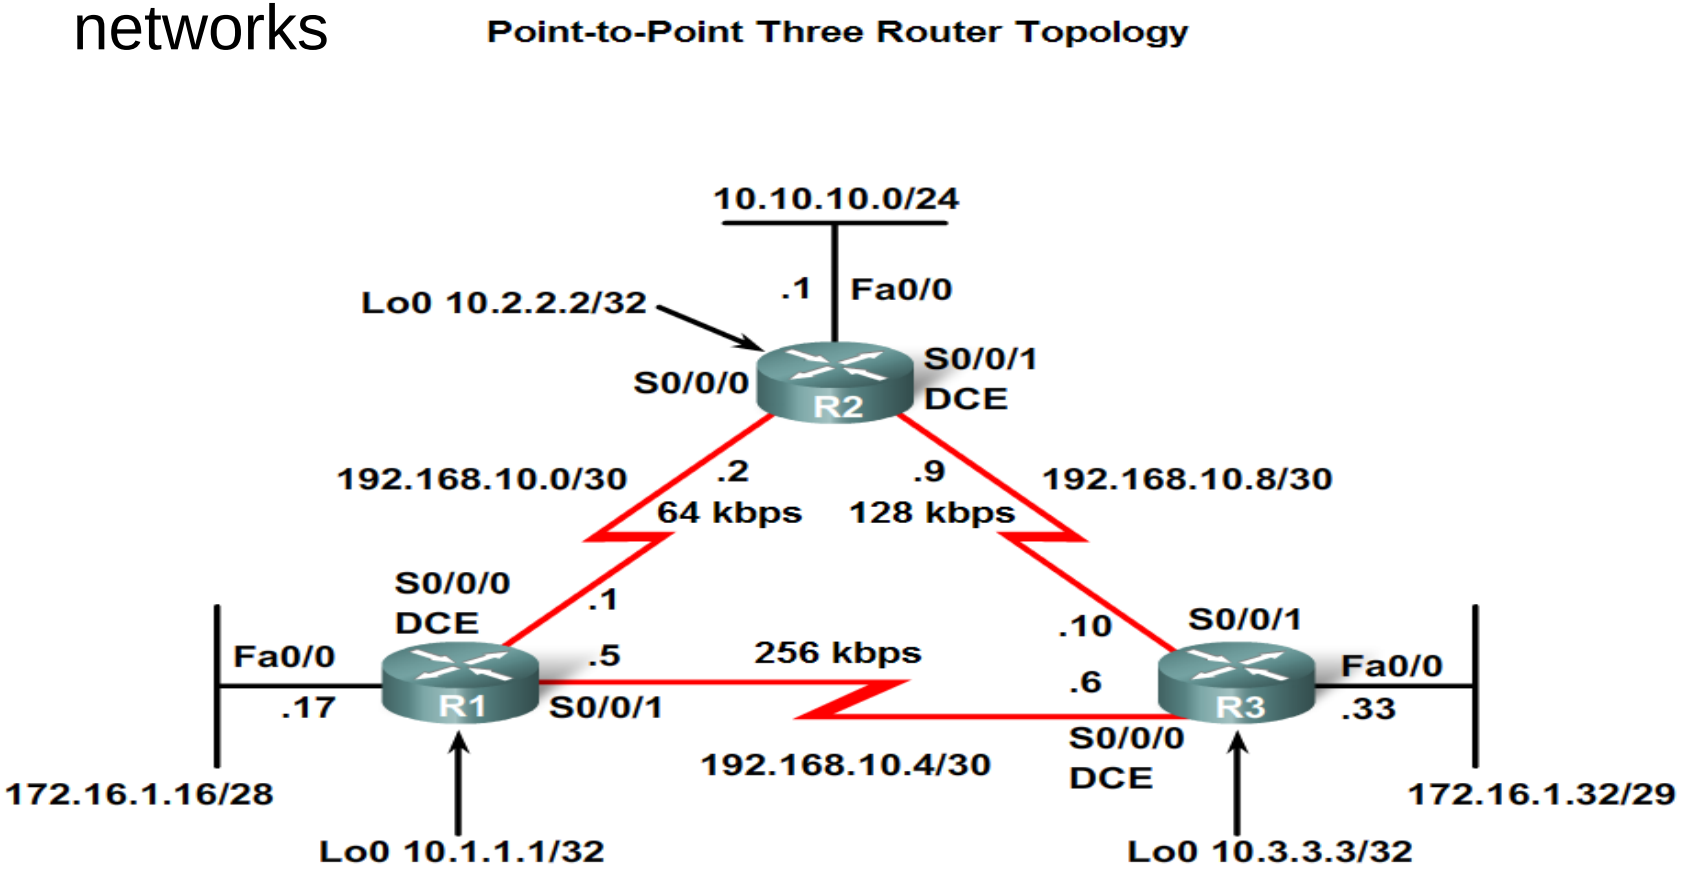

OSPF in Multiaccess Networks

- DR/BDR elections **will take place on multiaccess networks only** as shown below

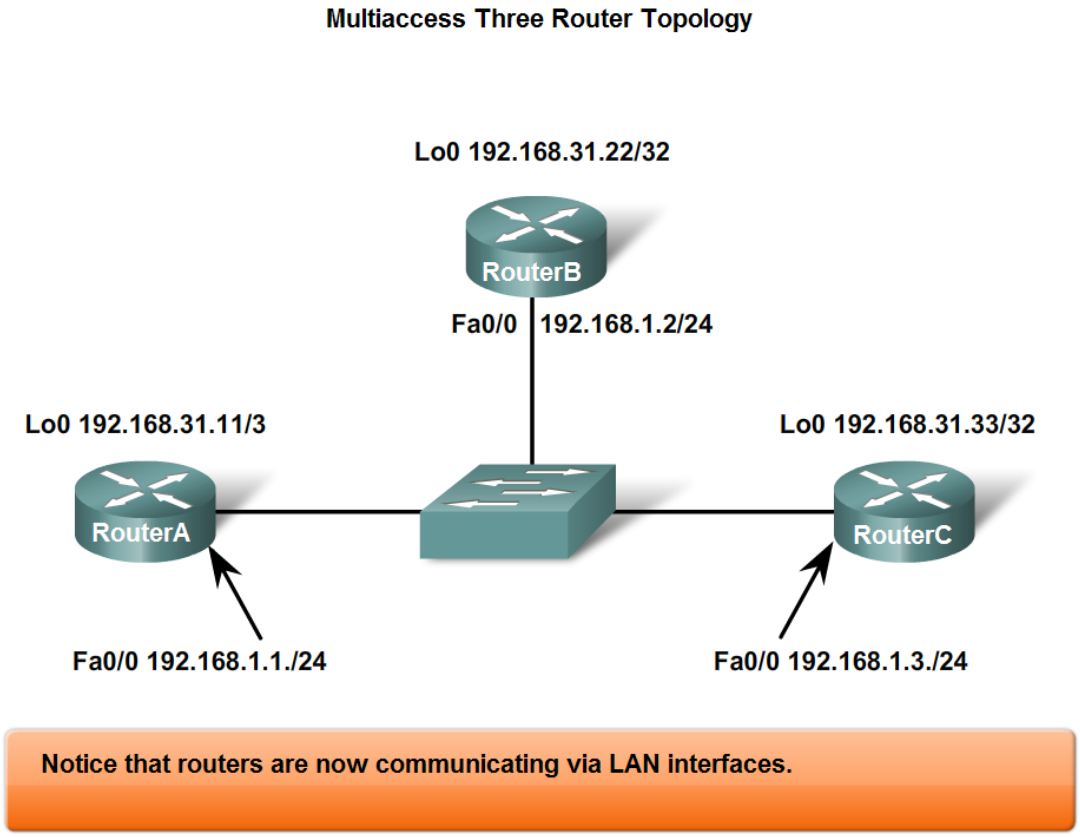

OSPF in Multiaccess Networks

- 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.

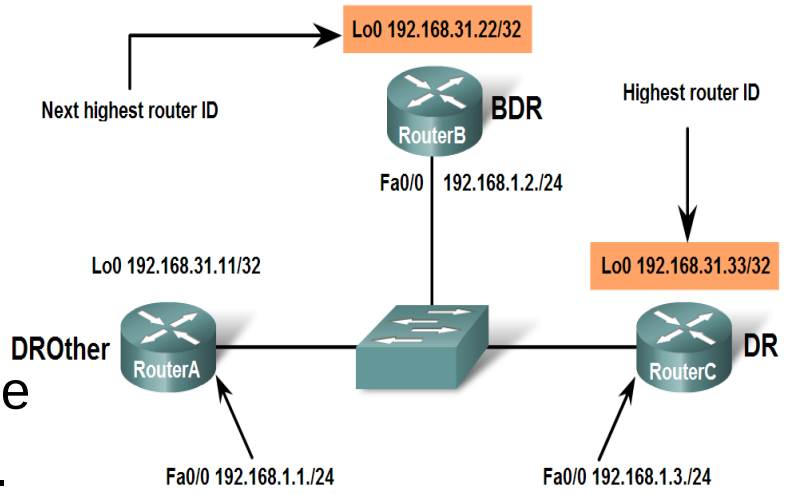


```
RouterA#show ip ospf neighbor

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.

```
RouterA#show ip ospf interface fastethernet 0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 192.168.1.1/24, Area 0
  Process ID 1, Router ID 192.168.31.11, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 192.168.31.33, Interface address 192.168.1.3
  Backup Designated router (ID) 192.168.31.22, Interface address 192.168.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 192.168.31.22  (Backup Designated Router)
    Adjacent with neighbor 192.168.31.33  (Designated Router)
  Suppress hello for 0 neighbor(s)
```

More OSPF Configuration

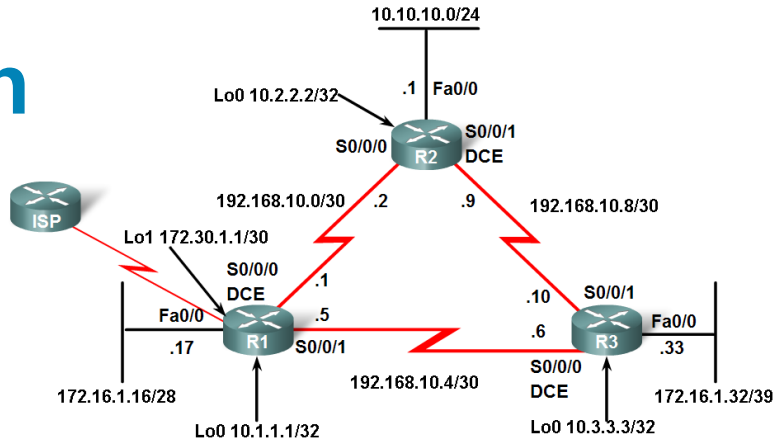

Redistributing an OSPF Default Route

- Topology includes a link to ISP
  - Router connected to ISP
    - Called an autonomous system border router
    - Used to propagate a default route
      - Example of static default route

        R1(config)#**ip route 0.0.0.0 0.0.0.0 loopback 1**
      - Requires the use of the **default-information originate** command
      - Example of default-information originate command

        R1(config-router)#**default-information originate**