Chapter 11

Open Shortest Path First

OSPF

Note for Instructors

- These presentations are the result of a collaboration among the instructors at St. Clair College in Windsor, Ontario.
- Thanks must go out to Rick Graziani of Cabrillo College. His material and additional information was used as a reference in their creation.
- If anyone finds any errors or omissions, please let me know at:
  - tdame@stclaircollege.ca.

Introduction to OSPF

	Distance Vector Routing Protocols (Interior Gateway Protocols)		Link State Routing Protocols (Interior Gateway Protocols)		Path Vector (Exterior Gateway Protocols)
Classful	RIP	IGRP			EGP
Classless	RIPv2	EIGRP	OSPFv2	IS-IS	BGPv4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGPv4 for IPv6

Introduction to OSPF

	Distance Vector Routing Protocols (Interior Gateway Protocols)		Link State Routing Protocols (Interior Gateway Protocols)		Path Vector (Exterior Gateway Protocols)
Classful	RIP	IGRP			EGP
Classless	RIPv2	EIGRP	OSPFv2	IS-IS	BGPv4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGPv4 for IPv6

- OSPF is a classless, link-state routing protocol that uses the concept of areas for scalability.
- RFC 2328 defines the OSPF metric as an arbitrary value called cost.
  - Cisco IOS software uses bandwidth to calculate the OSPF cost metric.

Development Timeline

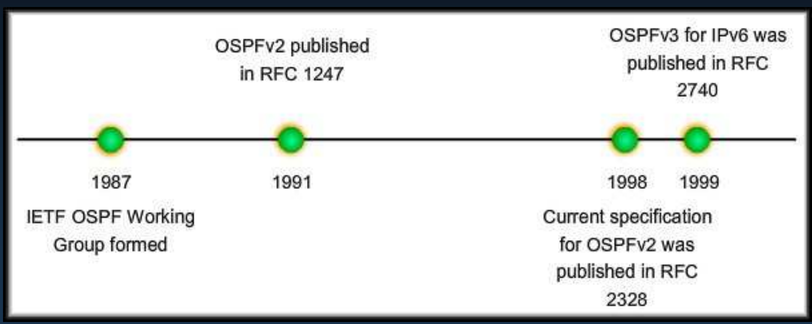

OSPF Message Encapsulation (Ethernet)

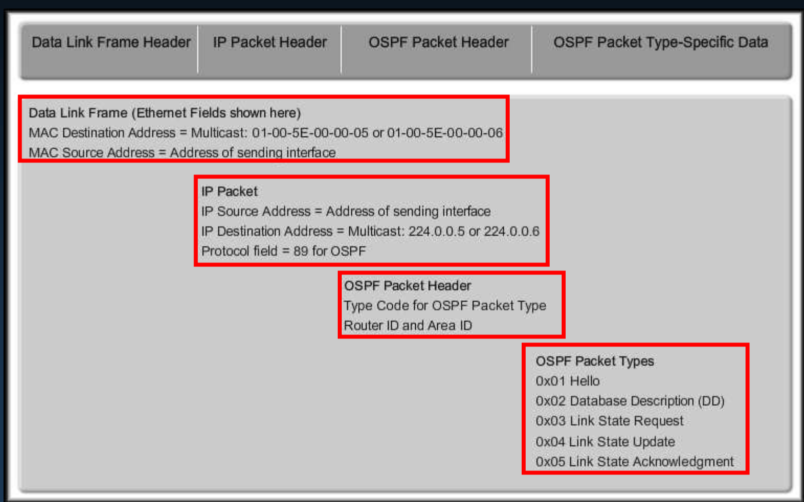

OSPF Packet Types

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

Hello Protocol

Hello Packet – Type 1

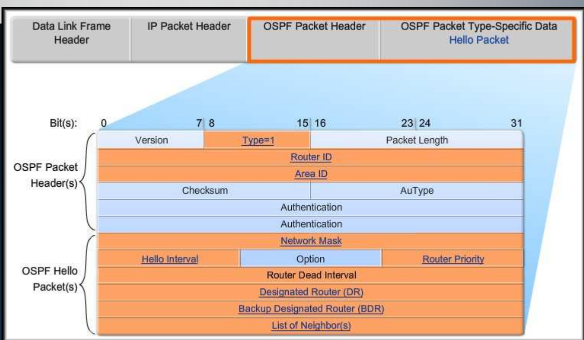

- Discover OSPF **neighbors** and establish neighbor adjacencies.
- **Advertise parameters** on which two routers must agree to become neighbors.
- Elect the **Designated Router and Backup Designated Router** on **multi-access networks** such as Ethernet and Frame Relay.

Hello Protocol: Important Fields

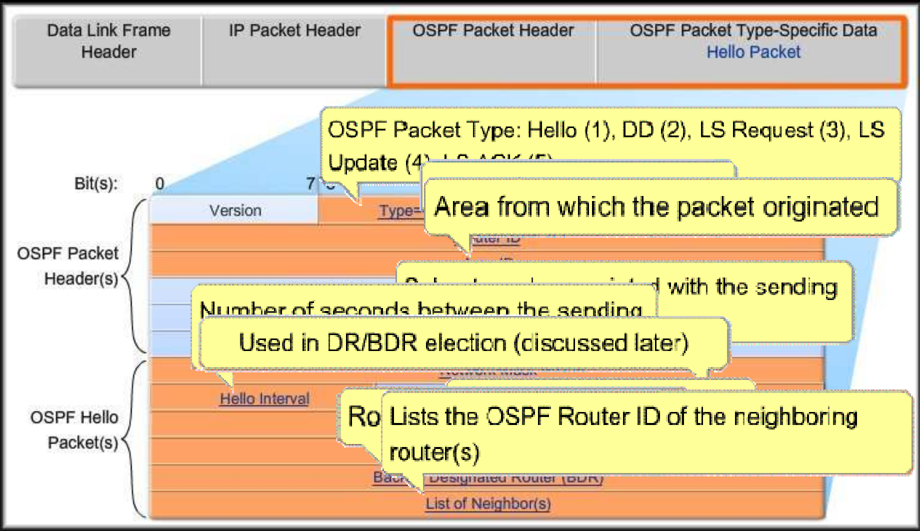

Hello Protocol: Neighbour Establishment

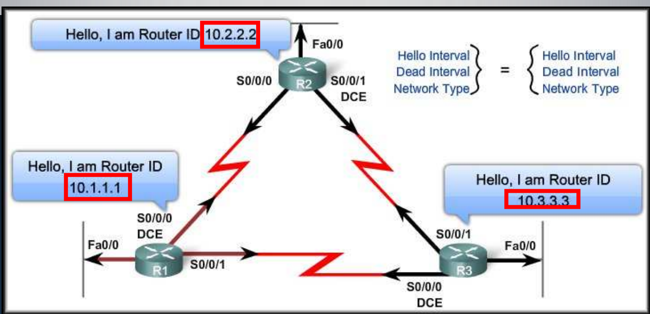

- Before an OSPF router can flood its link states to other routers, it determines whether there are any other OSPF neighbors on any of its links by sending Hello packets out all configured interfaces.
- The information in the OSPF Hello includes the OSPF Router ID of the router sending the Hello packet.

Hello Protocol: Neighbour Establishment

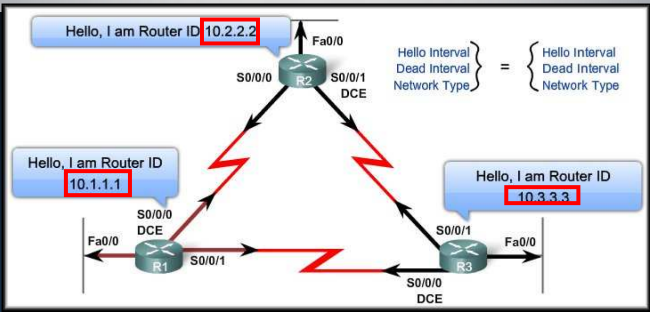

- Receiving an OSPF Hello packet on an interface confirms for a router that there is another OSPF router on this link.
- OSPF then establishes adjacency with the neighbor.
- Full adjacency happens after both routers have exchanged any necessary LSUs and have identical link-state databases.

Hello Protocol: Neighbour Establishment

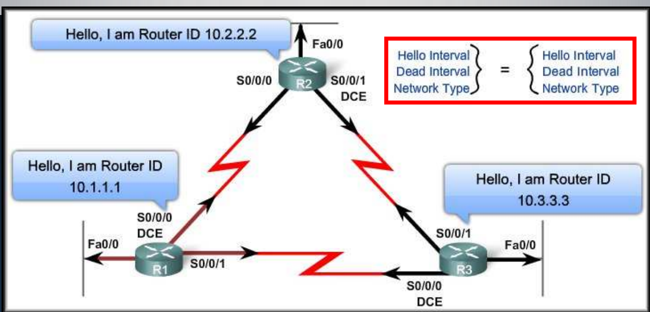

- Before two routers can form an OSPF neighbor adjacency, they must also agree on three values:
  - Hello interval
  - Dead interval
  - Network type

Hello Protocol: Hello Interval

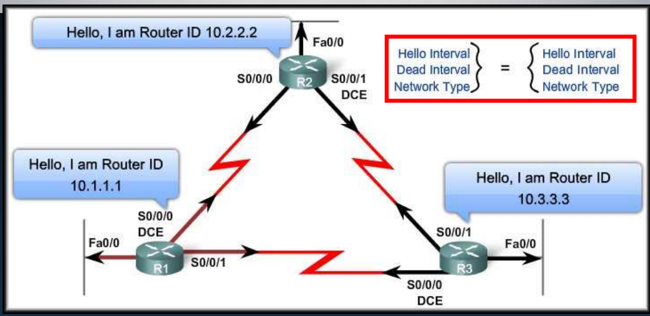

- The OSPF **Hello interval** indicates **how often** an OSPF router transmits its Hello packets.
  - Every 10 seconds on multi-access and point-to-point segments.
  - Every 30 seconds on non-broadcast multi-access **(NBMA)** segments (Frame Relay, X.25, ATM).

Hello Protocol: Dead Interval

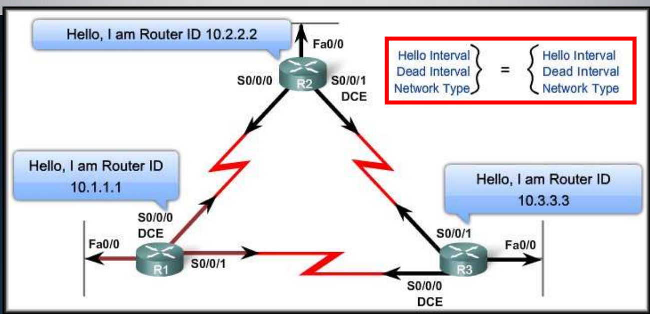

- The OSPF **Dead interval** is the period, expressed in seconds, that the router will wait to receive a Hello packet before declaring the neighbor "down."
- Cisco uses a default of four times the Hello interval.
  - 40 seconds for multi-access and point-to-point links.
  - 120 seconds for NBMA networks.

Hello Protocol: Dead Interval

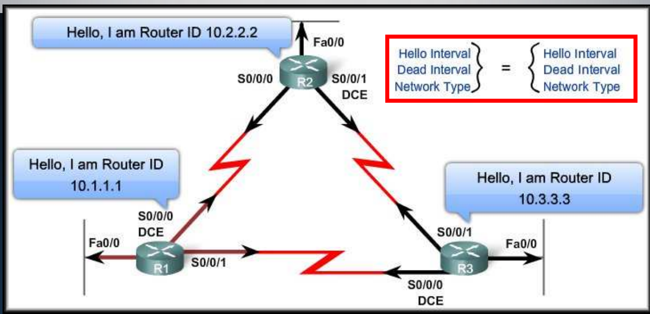

- If the Dead Interval expires before the routers receive a Hello packet, OSPF removes that neighbor from its link-state database.
- The router floods the link-state information about the “down” neighbor out all OSPF-enabled interfaces.
- Stay tuned for Network Type.

Hello Protocol: Electing a DR and BDR

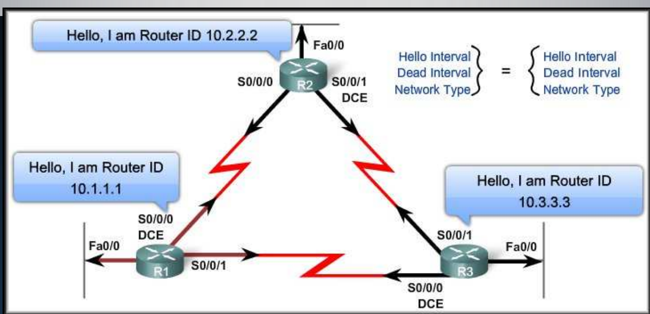

- To reduce the amount of OSPF traffic on multi-access networks, OSPF elects a Designated Router (DR) and Backup Designated Router (BDR).
- The DR updates all other OSPF routers when a change occurs in the multi-access network.
- The BDR takes over as DR if the current DR fails.

Hello Protocol: Electing a DR and BDR

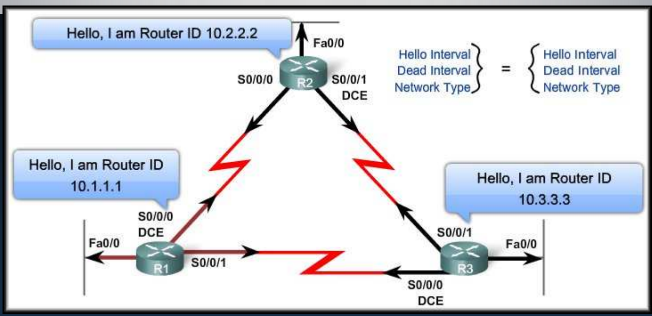

- R1, R2, and R3 are connected through point-to-point links so there is no DR/BDR election.
- The DR/BDR election and processes is discussed later.
- The topology will be changed to a multi-access network.

OSPF LSUs

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

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

- The acronyms LSA and LSU are often used interchangeably.
- An LSU contains one or more LSAs.
- LSAs contain route information for destination networks.

OSPF Algorithm

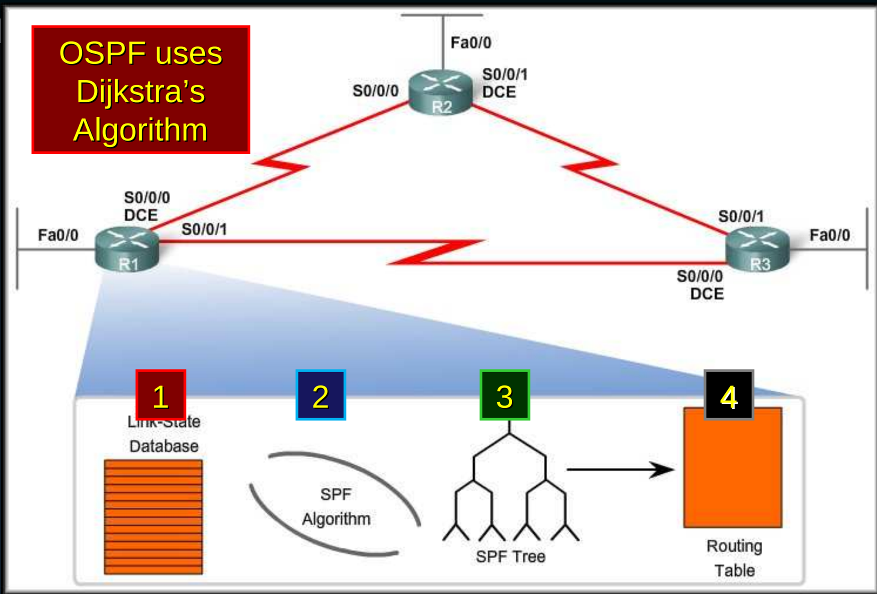

Administrative Distance

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

- Administrative Distance (AD) is the trustworthiness (or preference) of the route source.

Authentication

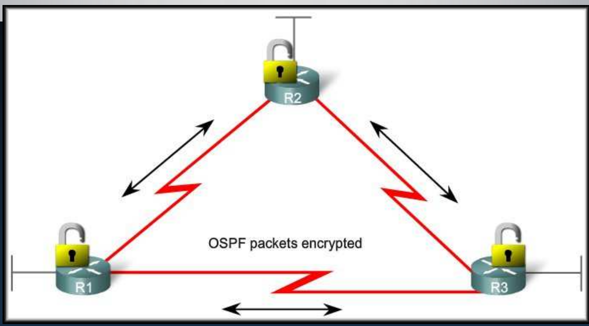

- Like other routing protocols, OSPF can be configured for authentication.
- Authentication ensures that routers will only accept routing information from other routers that have been configured with the same password or authentication information.

OSPF

Basic OSPF Configuration

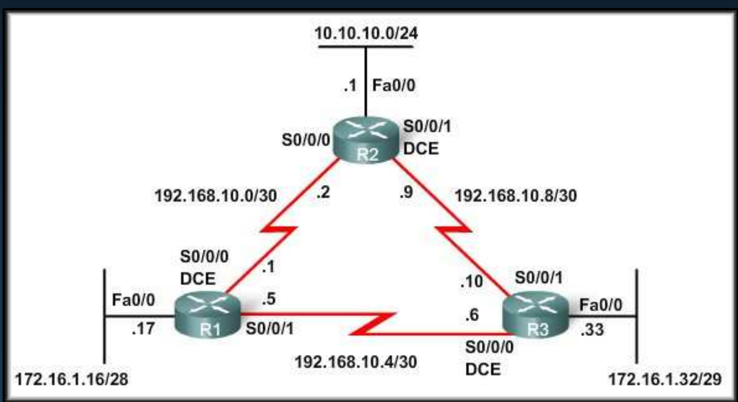

Topology

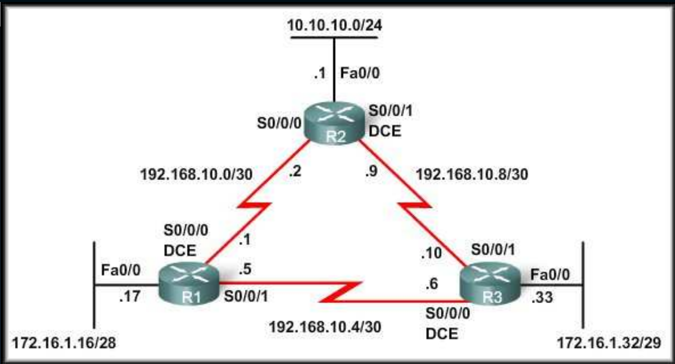

- Notice that the addressing scheme is discontiguous.
- OSPF is a classless routing protocol so we will include the subnet mask as part of our configuration.
- There are three serial links of various bandwidths and each router has multiple paths to each remote network.

The **router ospf** Command

```
R1(config)# router ospf 1

R1(config-router)#
```

- Like EIGRP, OSPF requires a Process ID.
  - Between 1 and 65,535 and chosen by the network administrator.
- Unlike EIGRP, the process ID is only locally significant.
  - Does not have to match on all OSPF routers.
- For consistency, we will enable OSPF on all three routers using the same Process ID of 1.

The **network** Command

```
Router(config-router)# network [network-address]
                               [wildcard-mask]
                          area [area-id]
```

- The network command used with OSPF has the same function as when used with other IGP routing protocols:
  - Any interfaces on a router that match the network address in the network command will be enabled to send and receive OSPF packets.
  - This network (or subnet) will be included in OSPF routing updates.

The **network** Command

```
Router(config-router)# network [network-address]
                               [wildcard-mask]
                          area [area-id]
```

```
  255.255.255.255
- 255.255.255.240 Subtract the subnet mask
---------------
    0.  0.  0. 15 Wildcard mask
```

- OSPF requires the wildcard mask.
  - The network address, along with the wildcard mask, is used to specify the interface or range of interfaces that will be enabled for OSPF using this network command.

The `network` Command

```
Router(config-router)# network [network-address]
                               [wildcard-mask]
                          area [area-id]
```

- The **area area-id** refers to the OSPF area.
  - An OSPF area is a group of routers that share link-state information.
  - All OSPF routers in the same area must have the same link-state information in their link-state databases.

Hierarchical Design

- Link-state routing protocols such as OSPF and IS-IS use the concept of areas.

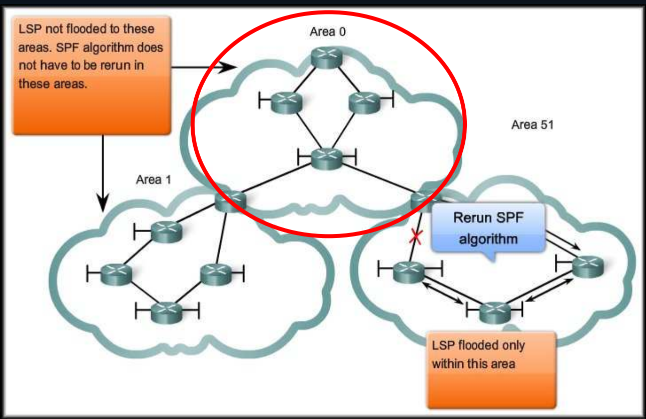

- Multiple areas create a hierarchical design to networks, allowing better route aggregation (summarization) and the isolation of routing issues within an area.

The **network** Command

```
Router(config-router)# network [network-address]
                               [wildcard-mask]
                          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

R3(config)# router ospf 1
R3(config-router)# network 172.16.1.32 0.0.0.7 area 0
R3(config-router)# network 192.168.10.4 0.0.0.3 area 0
R3(config-router)# network 192.168.10.8 0.0.0.3 area 0
```

- **network** commands for all three routers, enabling OSPF on all interfaces.
- At this point, all routers should be able to ping all networks.

OSPF Router ID

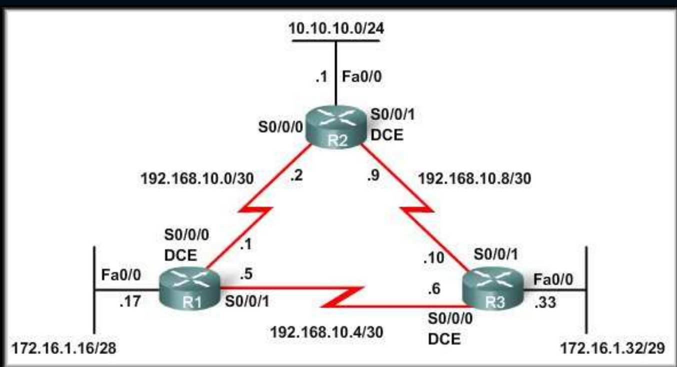

- The OSPF Router ID plays an important role in OSPF.
- Later ,you will see how the Router ID is used in the DR and BDR process.
  - It is used to uniquely identify each router in the OSPF routing domain.

OSPF Router ID

- Highest Active IP Address:
  - A Router ID is simply an IP address.
  - Cisco routers derive the router ID based on three criteria and with the following precedence:

1. Use the IP address configured with the OSPF **`router-id`** command.
2. If the Router ID is not configured, the router chooses the highest IP address of any of its loopback interfaces.
3. If no loopback interfaces are configured, the router chooses the highest active IP address of any of its physical interfaces.

OSPF Router ID: Highest Active IP Address

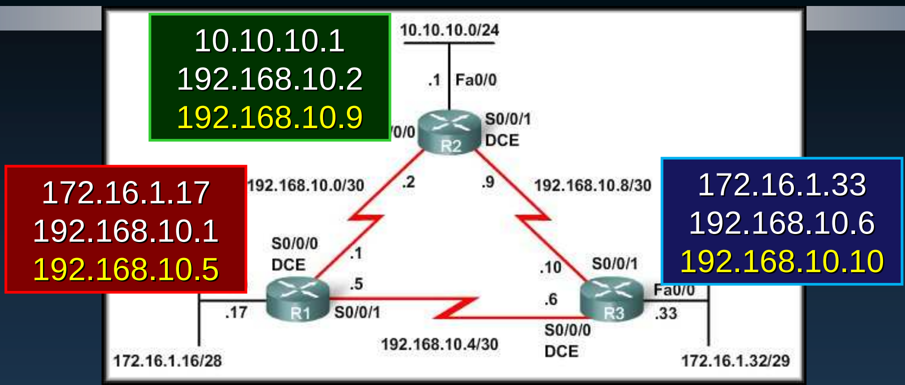

- Because we have not configured Router IDs or loopback interfaces on our three routers, the Router ID for each router is determined by the third criterion in the preceding list.
  - The highest active IP address on any of the router's physical interfaces.

OSPF Router ID: Highest Active IP Address

- Verifying the Router ID:

```
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
<output omitted>

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
<output omitted>

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
<output omitted>
```


OSPF Router ID: Loopback Address

- The advantage of using a loopback interface is that, unlike physical interfaces, it cannot fail.

```
Router(config)# interface loopback [number]

Router(config-if)# ip address [ip-address] [subnet-mask]
```

  - "Virtual" interface with a /32 subnet mask.
  - Automatically "up" and "up".
  - Very useful in setting Router IDs as they never go down.
  - Also useful to configure virtual networks that you can ping and route as if they were attached networks.

OSPF Router ID: Loopback Address

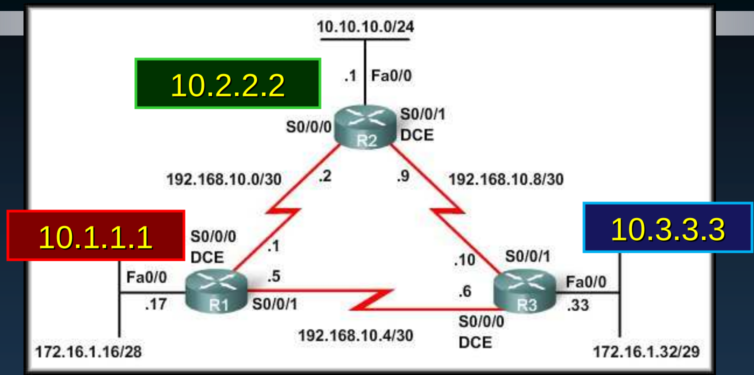


```
R1(config)# interface loopback 0
R1(config-if)# ip address 10.1.1.1 255.255.255.255

R2(config)# interface loopback 0
R2(config-if)# ip address 10.2.2.2 255.255.255.255

R3(config)# interface loopback 0
R3(config-if)# ip address 10.3.3.3 255.255.255.255
```


OSPF Router ID: **router-id** Command

```
Router(config)# router ospf [process-id]

Router(config-router)# router-id [ip-address]
```

- The OSPF **router-id** command is a fairly recent addition to Cisco IOS software.
  - The OSPF **router-id** command was introduced in Cisco IOS Software Release 12.0(T) and takes precedence over loopback and physical interface IP addresses for determining the router ID.
- Because this command is a recent addition, it is more common to find loopback addresses used for configuring OSPF Router IDs.

OSPF Router ID

- Modifying the Router ID:
  - The router ID can be modified with the IP address from a subsequent OSPF **router-id** command by reloading the router or by using the following command:

```
Router# clear ip ospf process
```

  - Modifying a router ID with a new loopback or physical interface IP address may require reloading the router.

OSPF Router ID

- Duplicate Router IDs:
  - When two routers have the same router ID in an OSPF domain, routing might not function properly.
  - If the router ID is the same on two neighboring routers, the neighbor establishment might not occur.
  - When duplicate OSPF router IDs occur, Cisco IOS software displays the following message.

```
%OSPF-4-DUP_RTRID1: Detected router with duplicate router ID
```

OSPF Router ID

- Verifying Router IDs – Loopback Interfaces:

```
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 10.1.1.1
<output omitted>

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 10.2.2.2
<output omitted>

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 10.3.3.3
<output omitted>
```


Verifying OSPF

- The **show ip ospf neighbor** command:

The Router ID of the neighbouring router.

The amount of time remaining that the router will wait to receive an OSPF Hello packet.

```
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
```

FULL state means that the routers are fully adjacent and have identical link-state databases.

The neighbor's IP Address and Interface used to form the adjacency.

Verifying OSPF

- Two routers may not form an OSPF adjacency if any one of the following conditions is present:
  - The subnet masks do not match, causing the routers to be on separate networks.
  - OSPF Hello or Dead timers do not match.
  - OSPF network types do not match.
  - There is a missing or incorrect OSPF network command.
- Other powerful troubleshooting tools:

```
show ip protocols

show ip ospf

show ip ospf interface [interface]
```


Examining The Routing Table

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

    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
```

- The quickest way to verify OSPF convergence is to look at the routing table for each router.
- Loopback interfaces are included.
- Unlike RIPv2 and EIGRP, OSPF does not automatically summarize at major network boundaries.

OSPF

The OSPF Metric

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

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

- The OSPF metric is called **cost**.

  From RFC 2328: *"A cost is associated with the output side of each router interface. This cost is configurable by the system administrator. The lower the cost, the more likely the interface is to be used to forward data traffic."*

OSPF Metric

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

- The Cisco IOS uses the cumulative bandwidths of the outgoing interfaces from the router to the destination network as the cost value.

Cisco IOS Cost for OSPF = 10^8/bandwidth in bps

OSPF Accumulates Costs

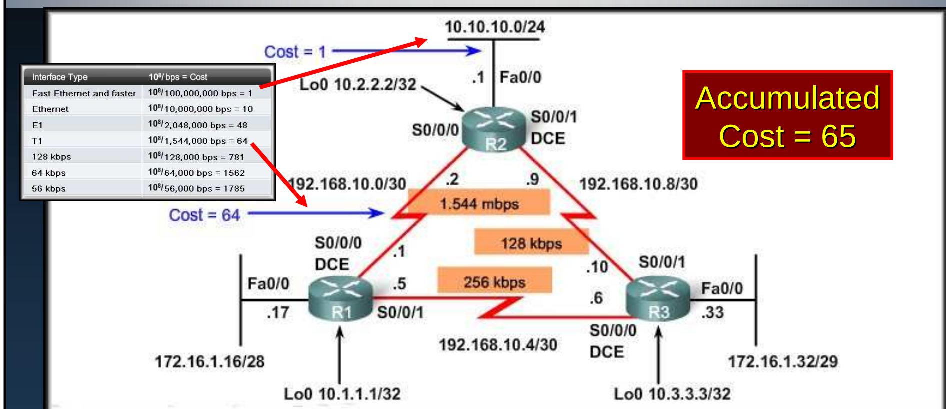


```
R1#show ip route
Codes: <some code output omitted>
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

***output omitted***
O        10.10.10.0/24 [110/65] via 192.168.10.2, 14:27:57, Serial0/0/0
```

Default Bandwidth on Serial Interfaces

```
R1# show interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Description: Link to R2
  Internet address is 192.168.10.1/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```

- On Cisco routers, the bandwidth value on many serial interfaces defaults to T1 (1.544 Mbps) and some may default to 128 Kbps.
- Never assume that OSPF is using any particular bandwidth value to calculate cost.
- Always check the default value with the **`show interface`** command.

Default Bandwidth on Serial Interfaces

```
R1# show interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Description: Link to R2
  Internet address is 192.168.10.1/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```


- Changing the bandwidth value does not affect the speed of the link.
  - It is this value that is used in calculating the metric.
  - It is important that the bandwidth value reflect the actual speed of the link so that the routing table has accurate best path information.

Default Bandwidth on Serial Interfaces

```
R1# show interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Description: Link to R2
  Internet address is 192.168.10.1/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```

Both WAN links have defaulted to 1544 Kbit and have a cost of 128.

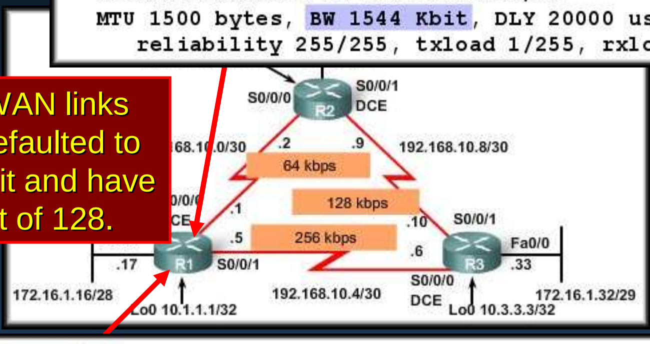


```
R1# show ip route
<route ouput omitted>
O     192.168.10.8 [110/128] via 192.168.10.6, 14:27:57, Serial0/0/1
                   [110/128] via 192.168.10.2, 14:27:57, Serial0/0/0
```


Default Bandwidth on Serial Interfaces

```
R1# show ip route
<route ouput omitted>
O     192.168.10.8 [110/128] via 192.168.10.6, 14:27:57, Serial0/0/1
                   [110/128] via 192.168.10.2, 14:27:57, Serial0/0/0
```

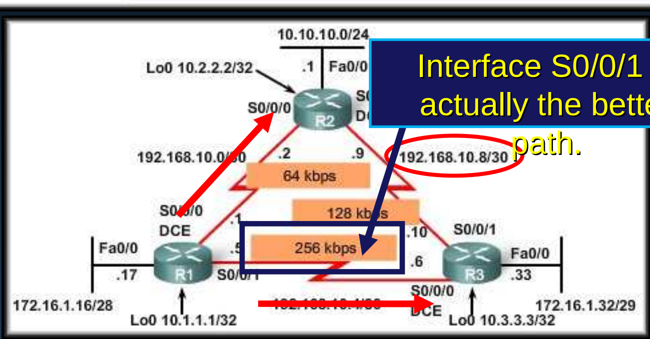

Interface S0/0/1 is actually the better path.

R1 believes it has two equal cost paths to network 192.168.10.8/30.

Default Bandwidth on Serial Interfaces

```
R1# show ip ospf interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 192.168.10.1/30, Area 0
  Process ID 1, Router ID 10.1.1.1, Network Type POINT_TO_POINT, Cost: 64
<output omitted>
```

- The calculated OSPF cost of an interface can be verified with the **show ip ospf interface** command.
  - As we have seen, this is NOT the cost of a 64 Kbps link.

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
64 kbps	10^8/64,000 bps = 1562
56 kbps	10^8/56,000 bps = 1785

```
Router(config-if)# bandwidth [bandwidth-kbps]
```


Modifying the Cost of a Link

```
R2(config)# inter serial 0/0/0
R2(config-if)# bandwidth 64

R2(config-if)# inter serial 0/0/1
R2(config-if)# bandwidth 128
```

```
R1(config)# inter serial 0/0/0
R1(config-if)# bandwidth 64

R1(config-if)# inter serial 0/0/1
R1(config-if)# bandwidth 256
```

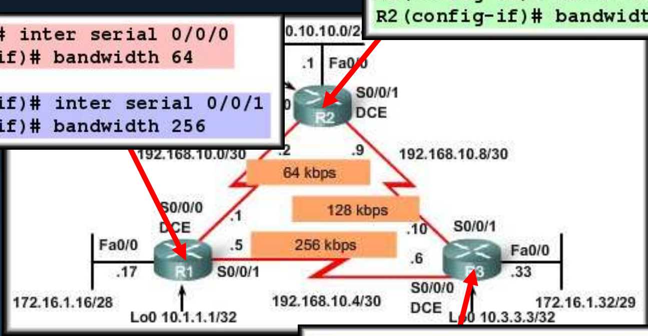


```
R3(config)# inter serial 0/0/0
R3(config-if)# bandwidth 256

R3(config-if)# inter serial 0/0/1
R3(config-if)# bandwidth 128
```

Modifying the Cost of a Link

```
R1# show ip route
<route ouput omitted>
O     192.168.10.8 [110/390] via 192.168.10.6, 14:27:57, Serial0/0/1
                   [110/1562] via 192.168.10.2, 14:27:57, Serial0/0/0
```

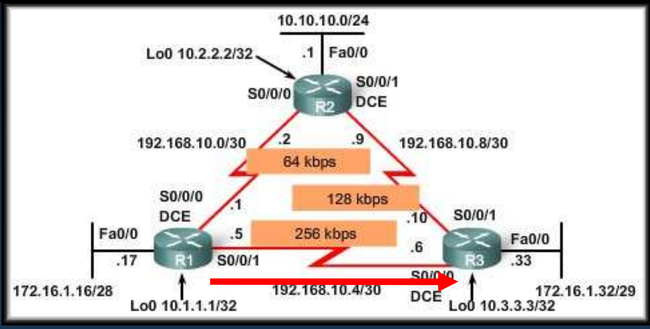

The faster 256 Kbps link becomes the preferred route.

Modifying the Cost of a Link

- An alternative method to using the **bandwidth** command is to use the **ip ospf cost** command, which allows you to directly specify the cost of an interface.

```
R1(config)# inter serial 0/0/0
R1(config-if)# bandwidth 64

R1(config-if)# inter serial 0/0/1
R1(config-if)# bandwidth 256
```

Cost = 1562

```
R1(config)# interface serial 0/0/0
R1(config-if)# ip ospf cost 1562
```

Modifying the Cost of a Link

Useful In multivendor environments.

bandwidth Commands		ip ospf cost Commands
Router R1 R1(config)#**interface serial 0/0/0** R1(config-if)#**bandwidth 64**	=	**Router R1** R1(config)#**interface serial 0/0/0** R1(config-if)#**ip ospf cost 1562**
R1(config)#**interface serial 0/0/1** R1(config-if)#**bandwidth 256**	=	R1(config)#**interface serial 0/0/1** R1(config-if)#**ip ospf cost 390**
Router R2 R2(config)#**interface serial 0/0/0** R2(config-if)#**bandwidth 64**	=	**Router R2** R2(config)#**interface serial 0/0/0** R2(config-if)#**ip ospf cost 1562**
R2(config)#**interface serial 0/0/1** R2(config-if)#**bandwidth 128**	=	R2(config)#**interface serial 0/0/1** R2(config-if)#**ip ospf cost 781**
Router R3 R3(config)#**interface serial 0/0/0** R3(config-if)#**bandwidth 256**	=	**Router R3** R3(config)#**interface serial 0/0/0** R3(config-if)#**ip ospf cost 390**
R3(config)#**interface serial 0/0/1** R3(config-if)#**bandwidth 128**	=	R3(config)#**interface serial 0/0/0** R3(config-if)#**ip ospf cost 781**

Both methods produce the same result.

Modifying the Cost of a Link

Interface Type	OSPF Cost
Fast Ethernet and Faster (100Mbps or Greater)	1
Ethernet (10 Mbps)	10
E1 (Europe) 2048 Kbps	46
T1 (North America/Japan) 1544 Kbps	64
768 Kbps	130
512 Kbps	195
256 Kbps	390
128 Kbps	781
64 Kbps	1562
56 Kbps	1785

OSPF

OSPF and Multi-access Networks

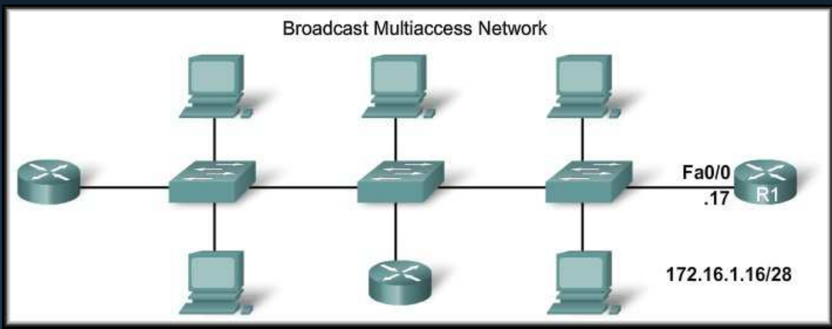

Challenges in Multi-access Networks

- Multi-access Broadcast network because a single device is capable of sending a single frame that has all devices on the network as its destination.

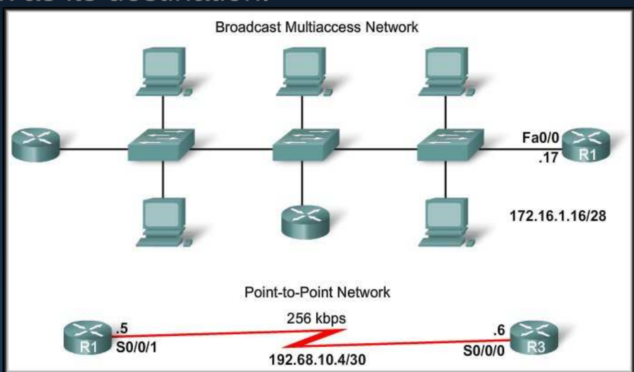

- As opposed to the point-to-point links we have been using.

Challenges in Multi-access Networks

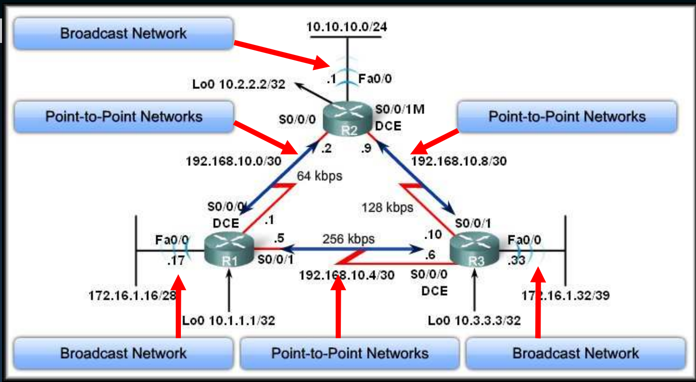

- OSPF defines five network types:
  - Point-to-Point
  - Broadcast Multi-access
  - Non-broadcast Multi-access (NBMA)
    - (Frame Relay - MIT446)
  - Point to Multipoint
  - Virtual links

Challenges in Multi-access Networks

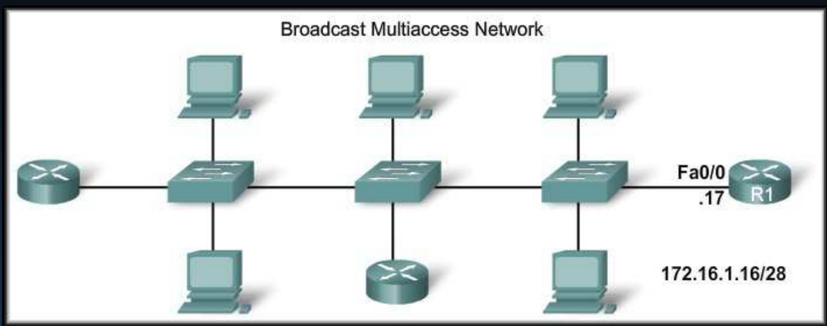

- Multi-access networks can create two challenges for OSPF regarding the flooding of LSAs:
  - Creation of multiple adjacencies, one adjacency for every pair of routers.
  - Extensive flooding of LSAs.

Multiple Adjacencies

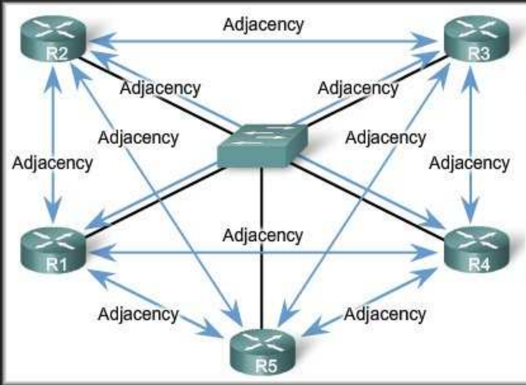

Routers	Adjacencies
n	n(n-1)/2
5	10
10	45
20	190
100	4,950

- The creation of an adjacency between every pair of routers in a network would cause the number of adjacencies to grow exponentially.

Multiple Adjacencies

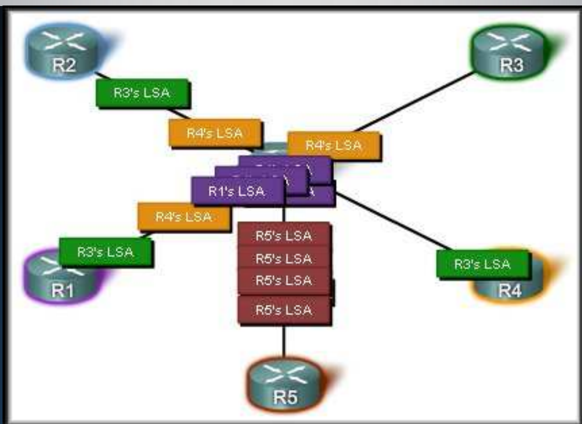

- Link-state routers flood their link-state packets when OSPF is initialized or when there is a change in the topology.
- This would lead to an excessive number of LSAs.

Solution: Designated Router

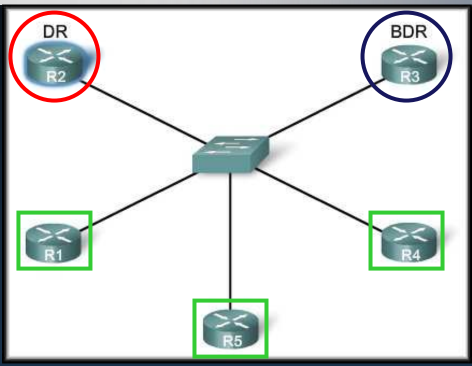

- The solution to managing the number of adjacencies and the flooding of LSAs on a multi-access network is the **Designated Router** (DR).
- On multi-access networks, OSPF **elects a DR** to be the **collection and distribution point** for LSAs sent and received.
- A **Backup Designated Router** (BDR) is also elected in case the DR fails.
- All other routers become **DROthers**.

Solution: Designated Router

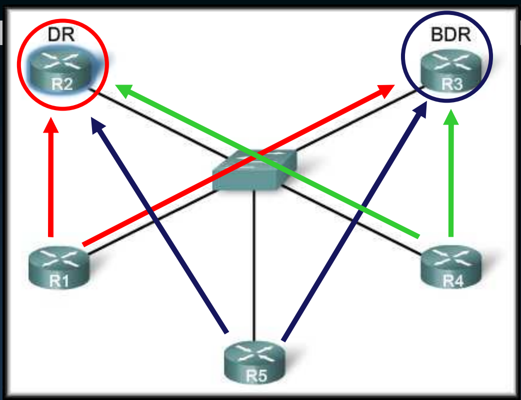

- **DROthers** only form full adjacencies with the DR and BDR in the network.
- Instead of flooding LSAs to all routers in the network, **DROthers** only send their LSAs to the DR and BDR using the multicast address **224.0.0.6**.

Solution: Designated Router

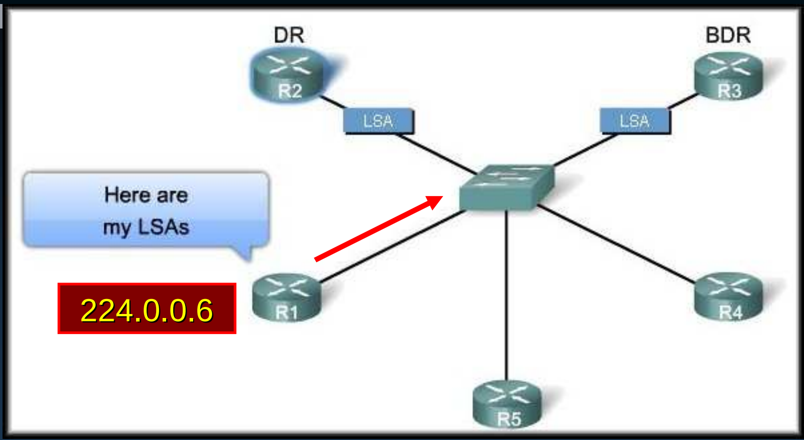

- Only the DR and BDR are listening for LSA's on multicast address 224.0.0.6.
- When R1 sends it's LSAs, both receive them.

Solution: Designated Router

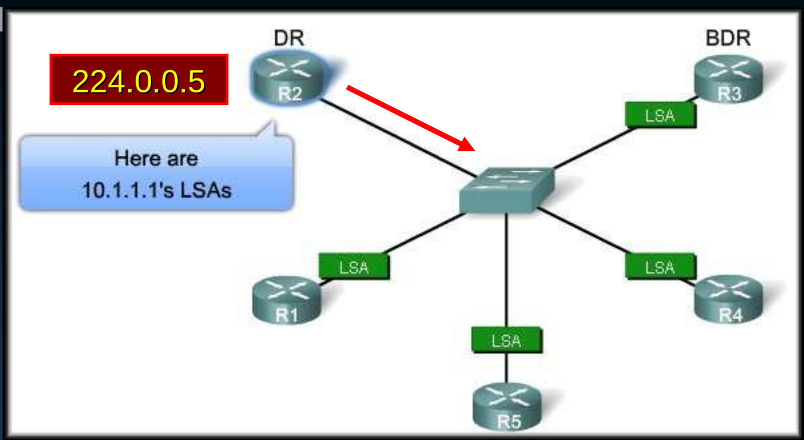

- The DR then forwards the LSAs from R1 to all other routers using the multicast address **224.0.0.5**.
- **Result:** *Only one router flooding all LSAs*.

DR/BDR Election Process

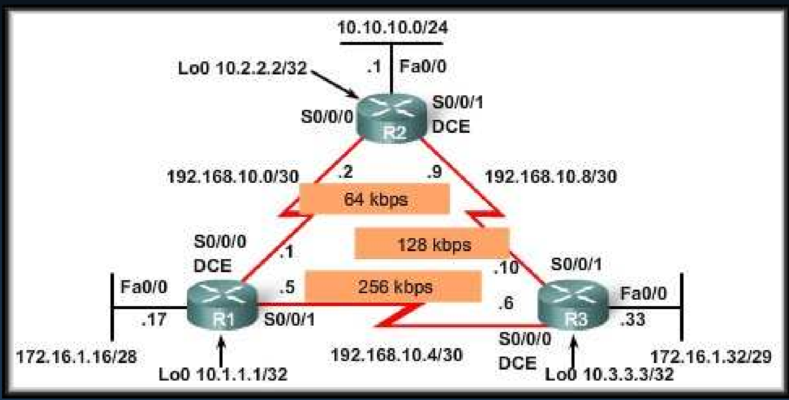

DR/BDR elections do not occur in point-to-point networks.

DR/BDR Election Process

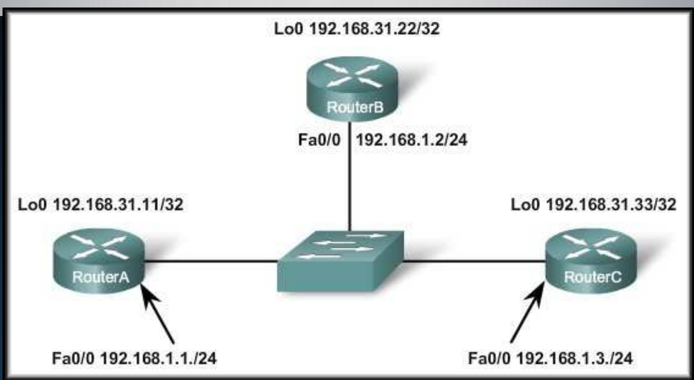

- In this new topology, we have three routers sharing a common Ethernet multiaccess network, 192.168.1.0/24.
- Each router is configured with an IP address on the Fast Ethernet interface and a loopback address for the router ID.

DR/BDR Election Process

- How do the DR and BDR get elected?

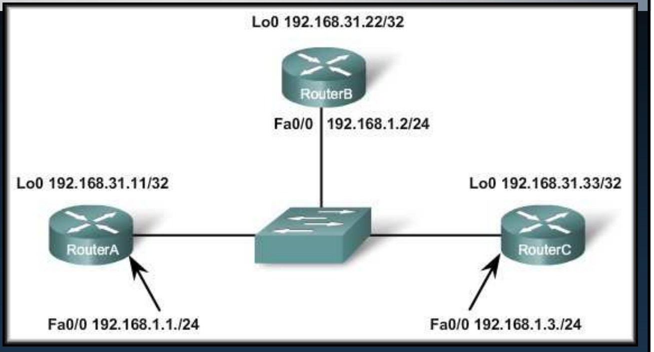

- The following criteria are applied:

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.

DR/BDR Election Process

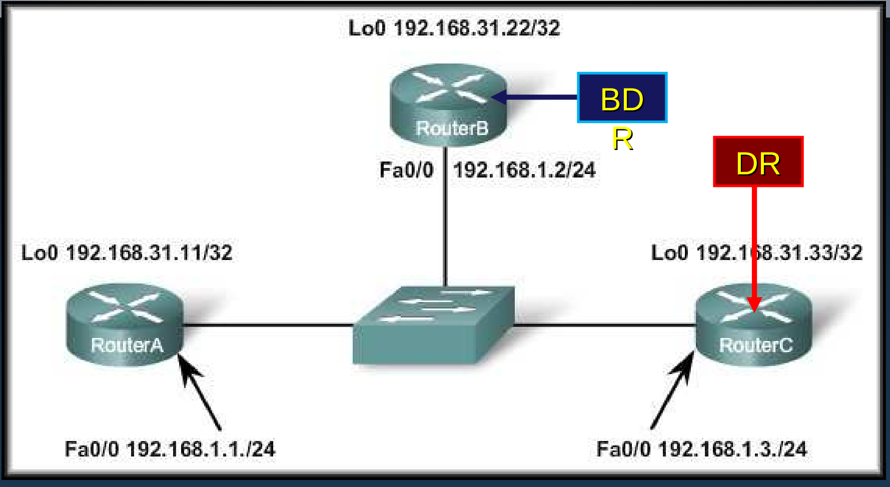

- Assuming that the router priority is the default, what is the result of the election?

DR/BDR Election Process

```
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
```

- The **show ip ospf neighbor** command displays the neighbour adjacencies of each router.
- Notice that the interface priority has not been changed so the Router ID (loopback address) will be used.

DR/BDR Election Process

```
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
```

- DROthers only form full adjacencies with the DR and BDR but will still form a neighbor adjacency with any DROthers that join the network.
- This means that they will still receive Hello packets so that they are aware of all routers in the network.

DR/BDR Election Process

- Interface Priority:
- *Like other elections, this one can be rigged.*

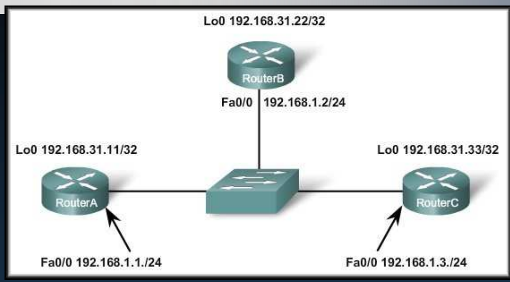

- The router's priority field (0-255) can be set to either ensure that it becomes the DR or prevent it from being the DR.

```
Rtr(config-if)# ip ospf priority <0-255>
```

  - Higher priority becomes DR/BDR
  - Default = 1
  - 0 = Ineligible to become DR/BDR

DR/BDR Election Process

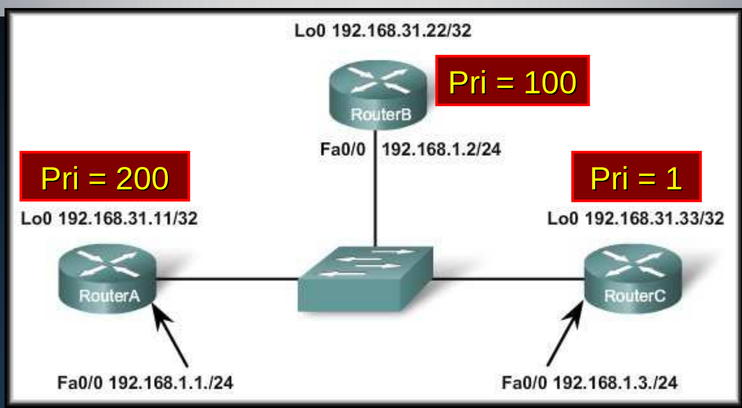


```
RouterA(config)# interface fastethernet 0/0
RouterA(config-if)# ip ospf priority 200

RouterB(config)# interface fastethernet 0/0
RouterB(config-if)# ip ospf priority 100
```

DR/BDR Election Process

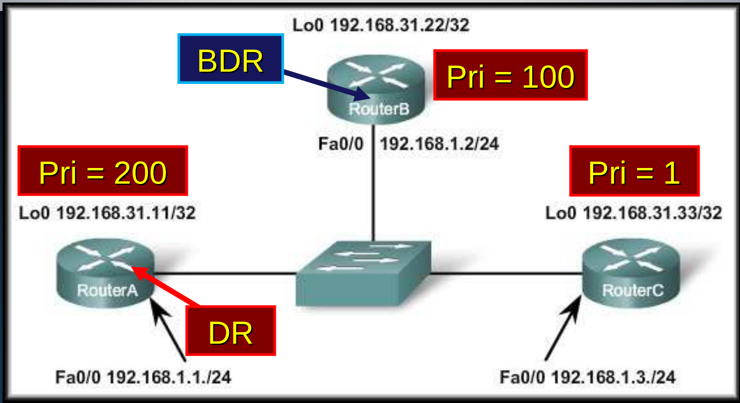


```
RouterA(config)# interface fastethernet 0/0
RouterA(config-if)# ip ospf priority 200

RouterB(config)# interface fastethernet 0/0
RouterB(config-if)# ip ospf priority 100
```


OSPF

More OSPF Configuration

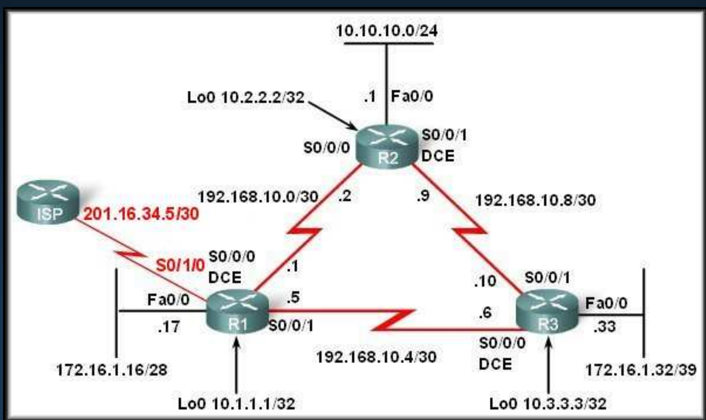

Redistributing an OSPF Default Route

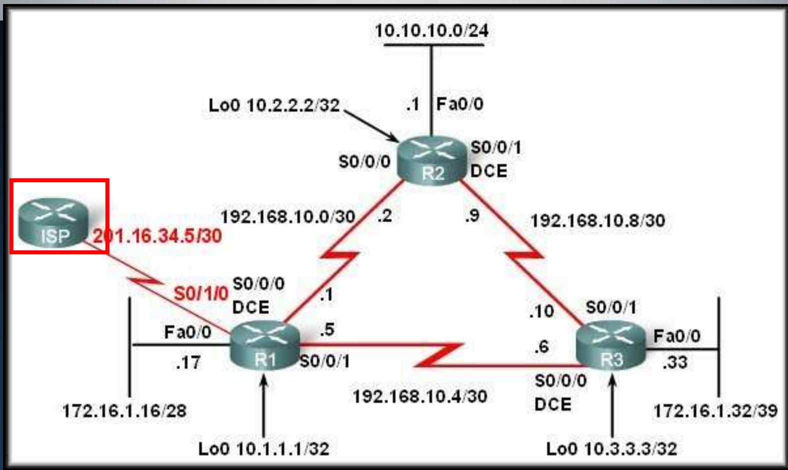

- Let's return to the earlier topology, which now includes a new link to an ISP.

Redistributing an OSPF Default Route

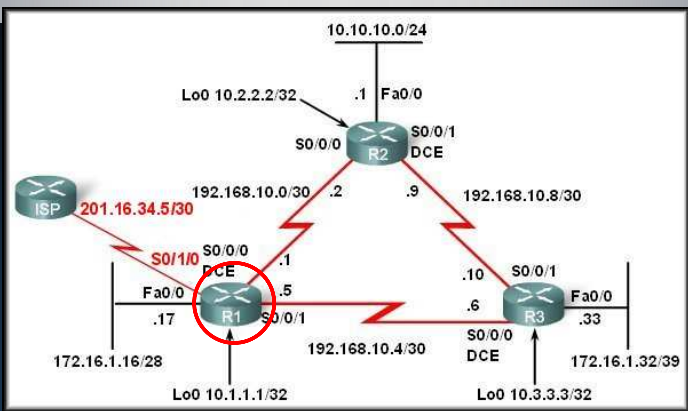

- The router connected to the Internet is used to propagate a default route to other routers in the OSPF routing domain.
- This router is sometimes called the edge, entrance, or gateway router.

Redistributing an OSPF Default Route

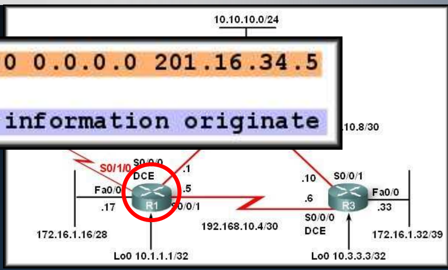


```
R1(config)# ip route 0.0.0.0 0.0.0.0 201.16.34.5
R1(config)# router ospf 1
R1(config-router)# default-information originate
```

- Like RIP, OSPF requires the use of the **`default-information originate`** command to advertise the 0.0.0.0/0 static default route to the other routers in the area.
- If the **`default-information originate`** command is not used, the default "quad zero" route will not be propagated to other routers in the OSPF area.

Redistributing an OSPF Default Route

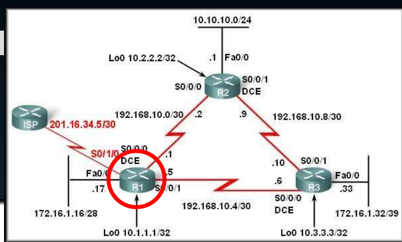


```
R1# show ip route

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

   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 [110/1171] via 192.168.10.6, 00:00:58, Serial0/0/1
   172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
O     172.16.1.32/29 [110/391] via 192.168.10.6, 00:00:58, Serial0/0/1
C     172.16.1.16/28 is directly connected, FastEthernet0/0
   201.16.34.0 is subnetted, 1 subnets
C     201.16.34.4 is directly connected, Serial0/1/0
   10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O     10.10.10.0/24 [110/1172] via 192.168.10.6, 00:00:58, Serial0/0/1
C     10.1.1.1/32 is directly connected, Loopback0
S*  0.0.0.0/0 is directly connected, 201.16.34.5
```

Redistributing an OSPF Default Route

E2 denotes an OSPF External type 2 route.

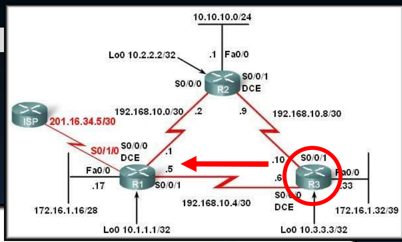


```
R3# show ip route

Gateway of last resort is 192.168.10.5 to network 0.0.0.0

     192.168.10.0/30 is subnetted, 3 subnets
O       192.168.10.0 [110/1952] via 192.168.10.5, 00:00:38, S0/0/0
C       192.168.10.4 is directly connected, Serial0/0/0
C       192.168.10.8 is directly connected, Serial0/0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.1.32/29 is directly connected, FastEthernet0/0
O       172.16.1.16/28 [110/391] via 192.168.10.5, 00:00:38, S0/0/0
     10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.3.3.3/32 is directly connected, Loopback0
O       10.10.10.0/24 [110/782] via 192.168.10.9, 00:00:38, S0/0/1
O*E2 0.0.0.0/0 [110/1] via 192.168.10.5, 00:00:27, Serial0/0/0
```


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

- 100,000,000 (10^8) is the default bandwidth referenced when the actual bandwidth is converted into a cost metric.
- As you know from previous studies, we now have link speeds that are much faster than Fast Ethernet speeds, including Gigabit Ethernet and 10GigE.
- Using a reference bandwidth of 100,000,000 results in interfaces with bandwidth values of 100 Mbps and higher having the same OSPF cost of 1.

Reference Bandwidth

Interface Type	10^8/ bps = Cost
128 kbps	10^8/128,000 bps = 781
64 kbps	10^8/64,000 bps = 1562
56 kbps	10^8/56,000 bps = 1785

```
R1(config-router)# auto-cost reference-bandwidth ?
1-4294967 The reference bandwidth in terms of Mbits per second.

R1(config-router)# auto-cost reference-bandwidth 10000
```

Bandwidth of 10GigE

- To obtain more accurate cost calculations, it might be necessary to adjust the reference bandwidth value.
- The reference bandwidth can be modified to accommodate these faster links by using the OSPF command **`auto-cost reference-bandwidth`**.
- When it is used, use it on all routers. The cost calculations will produce a proper metric.

Modifying OSPF Intervals

```
Router(config-if)# ip ospf hello-interval seconds

Router(config-if)# ip ospf dead-interval seconds
```

- It might be desirable to change the OSPF timers so that routers will detect network failures in less time.
- *Before changing any timer default values, be sure to give it careful consideration and understand the effects of making those changes.*
- Also keep in mind that OSPF requires that the Hello and Dead intervals match between two routers for them to become adjacent.
- This differs from EIGRP, where the hello and hold-down timers do not need to match for two routers to form an EIGRP adjacency.