Advanced Computer Network

LECTURE 5

OSPF(Open Shortest Path First)

What is OSPF?

- OSPF is a link-state routing protocol that was developed as an alternative for the distance vector Routing Information Protocol (RIP).
- OSPF has significant advantages over RIP in that it offers faster convergence and scales to much larger network implementations.
- OSPF uses the concept of areas. A network administrator can divide the routing domain into distinct areas that help control routing update traffic.
  - A link is an interface on a router, a network segment that connects two routers, or a stub network such as an Ethernet LAN that is connected to a single router.
  - Information about the state of a link is known as a link-state. All link-state information includes the network prefix, prefix length, and cost.

OSPF Basics

- All routing protocols share similar components. They all use routing protocol messages to exchange route information. The messages help build data structures, which are then processed using a routing algorithm.
- Routers running OSPF exchange messages to convey routing information using five types of packets:

1. Hello packet
2. Database description packet
3. Link-state request packet
4. Link-state update packet
5. Link-state acknowledgment packet

- These packets are used to discover neighboring routers and also to exchange routing information to maintain accurate information about the network.

OSPF Basics

OSPF messages are used to create and maintain three OSPF databases, as follows:

Database	Table	Description
Adjacency Database	Neighbor Table	•List of all neighbor routers to which a router has established bi-directional communication. •This table is unique for each router. •Can be viewed using the **`show ip ospf neighbor`** command.
Link-state Database (LSDB)	Topology Table	•Lists information about all other routers in the network. •The database represents the network LSDB. •All routers within an area have identical LSDB. •Can be viewed using the **`show ip ospf database`** command.
Forwarding Database	Routing Table	•List of routes generated when an algorithm is run on the link-state database. •Each router's routing table is unique and contains information on how and where to send packets to other routers. •Can be viewed using the **`show ip route`** command.

OSPF Basics

- The router builds the topology table using results of calculations based on the Dijkstra shortest-path first (SPF) algorithm. The SPF algorithm is based on the cumulative cost to reach a destination.
- The SPF algorithm creates an SPF tree by placing each router at the root of the tree and calculating the shortest path to each node. The SPF tree is then used to calculate the best routes. OSPF places the best routes into the forwarding database, which is used to make the routing table.

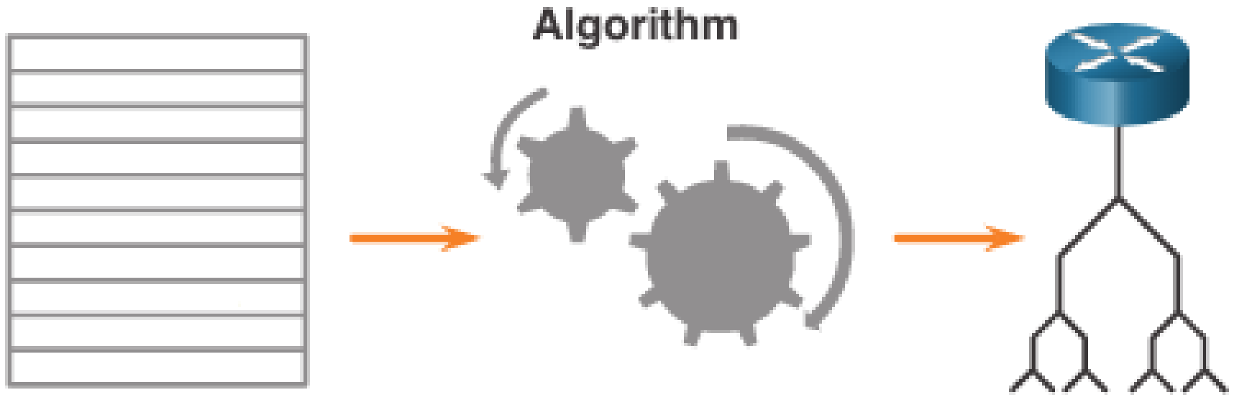

Link State Operation

To maintain routing information, OSPF routers complete a generic link-state routing process to reach a state of convergence. The following are the link-state routing steps that are completed by a router:

1. Establish Neighbor Adjacencies
2. Exchange Link-State Advertisements
3. Build the Link State Database
4. Execute the SPF Algorithm
5. Choose the Best Route

Single Area vs Multiarea OSPF

To make OSPF more efficient and scalable, OSPF supports hierarchical routing using areas. An OSPF area is a group of routers that share the same link-state information in their LSDBs. OSPF can be implemented in one of two ways, as follows:

- **Single-Area OSPF** - All routers are in one area. Best practice is to use area 0.
- **Multiarea OSPF** - OSPF is implemented using multiple areas, in a hierarchical fashion. All areas must connect to the backbone area (area 0). Routers interconnecting the areas are referred to as Area Border Routers (ABRs).

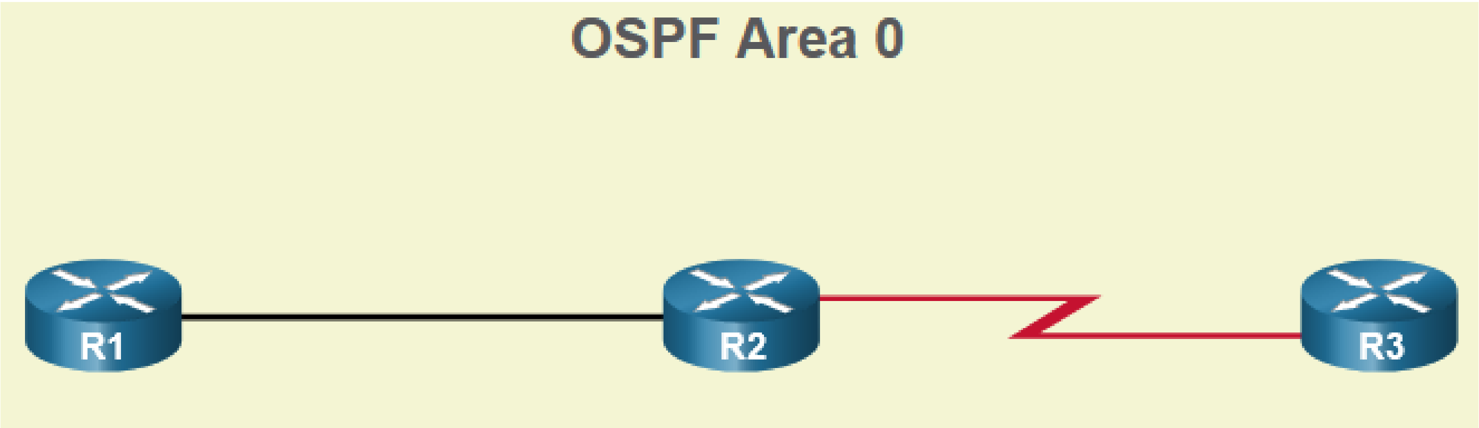

Multiarea OSPF

The hierarchical-topology design options with multiarea OSPF can offer the following advantages.

- **Smaller routing tables** - Tables are smaller because there are fewer routing table entries. This is because network addresses can be summarized between areas. Route summarization is not enabled by default.
- **Reduced link-state update overhead** - Designing multiarea OSPF with smaller areas minimizes processing and memory requirements.
- **Reduced frequency of SPF calculations** -– Multiarea OSPF localize the impact of a topology change within an area. For instance, it minimizes routing update impact because LSA flooding stops at the area boundary.

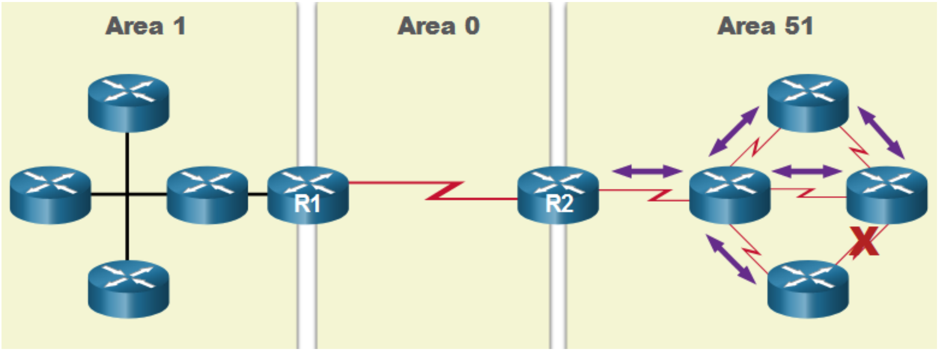

OSPFV3

- OSPFv3 is the OSPFv2 equivalent for exchanging IPv6 prefixes. OSPFv3 exchanges routing information to populate the IPv6 routing table with remote prefixes.
  - With the OSPFv3 Address Families feature, OSPFv3 includes support for both IPv4 and IPv6. OSPF Address Families is beyond the scope of this course.
- OSPFv3 has the same functionality as OSPFv2, but uses IPv6 as the network layer transport, communicating with OSPFv3 peers and advertising IPv6 routes. OSPFv3 also uses the SPF algorithm as the computation engine to determine the best paths throughout the routing domain.
- OSPFv3 has separate processes from its IPv4 counterpart. The processes and operations are basically the same as in the IPv4 routing protocol, but run independently.

Neighbor Adjacencies

- To determine if there is an OSPF neighbor on the link, the router sends a Hello packet that contains its router ID out all OSPF-enabled interfaces. The Hello packet is sent to the reserved All OSPF Routers IPv4 multicast address 224.0.0.5. Only OSPFv2 routers will process these packets.
- The OSPF router ID is used by the OSPF process to uniquely identify each router in the OSPF area. A router ID is a 32-bit number formatted like an IPv4 address and assigned to uniquely identify a router among OSPF peers.
- When a neighboring OSPF-enabled router receives a Hello packet with a router ID that is not within its neighbor list, the receiving router attempts to establish an adjacency with the initiating router.

OSPF Databases Synchronization

The database synchronization process consists of three steps as follows:

- Decide first router: The router with the highest router ID sends its DBD (Database Descriptor Packet) first.
- Exchange DBDs: As many as needed to convey the database. The other router must acknowledge each DBD with an LSAck packet.
- Send an LSR (Link State Request): Each router compares the DBD information with the local LSDB. If the DBD has more current link information, the router transitions to the loading state.

After all LSRs have been exchanged and satisfied, the routers are considered synchronized and in a full state. Updates LSUs (Link State Updates) are sent:

- When a change is perceived (incremental updates)
- Every 30 minutes

Multiaccess Network Issues

•Multiaccess networks can create two challenges for OSPF regarding the flooding of LSAs (Link State Advertisements), as follows:

- **Creation of multiple adjacencies** - Ethernet networks could potentially interconnect many OSPF routers over a common link. Creating adjacencies with every router would lead to an excessive number of LSAs exchanged between routers on the same network.
- **Extensive flooding of LSAs** - Link-state routers flood their LSAs any time OSPF is initialized, or when there is a change in the topology. This flooding can become excessive.

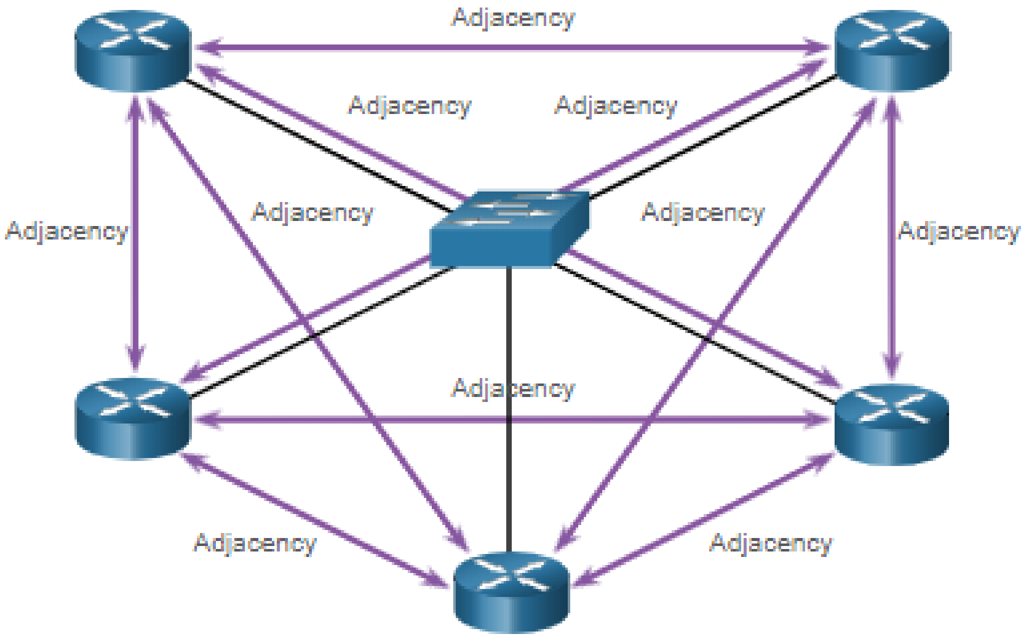

- Number of Adjacencies = n (n - 1) / 2
- n = number of routers
- Example: 5 (5 - 1) / 2 = 10 adjacencies

Router IDs

- An OSPF router ID is a 32-bit value, represented as an IPv4 address. It is used to uniquely identify an OSPF router, and all OSPF packets include the router ID of the originating router.
- Every router requires a router ID to participate in an OSPF domain. It can be defined by an administrator or automatically assigned by the router. The router ID is used by an OSPF-enabled router to do the following:
  - **Participate in the synchronization of OSPF databases** – During the Exchange State, the router with the highest router ID will send their database descriptor (DBD) packets first.
  - **Participate in the election of the designated router (DR)** - In a multiaccess LAN environment, the router with the highest router ID is elected the DR. The routing device with the second highest router ID is elected the backup designated router (BDR).

Router ID Assignment

- Instead of relying on physical interface, the router ID can be assigned to a loopback interface. Typically, the IPv4 address for this type of loopback interface should be configured using a 32-bit subnet mask (255.255.255.255).

```
R1(config-if)# interface Loopback 1
R1(config-if)# ip address 1.1.1.1 255.255.255.255
R1(config-if)# end
R1# show ip protocols | include Router ID
  Router ID 1.1.1.1
R1#
```

- Also we can use the **router-id** *rid* router configuration mode command to manually assign a router ID that is not related to an interface. In the example, the router ID 1.1.1.1 is assigned to R1. Use the **show ip protocols** command to verify the router ID.

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# end
*May 23 19:33:42.689: %SYS-5-CONFIG_I: Configured from console by console
R1# show ip protocols | include Router ID
  Router ID 1.1.1.1
R1#
```

OSPF Priority Settings

- The DR election process mentioned earlier depends on the value of the router interface priority and on the router ID.
- If the interface priorities are equal on all routers, the router with the highest router ID is elected the DR.
- Instead of relying on the router ID, it is better to control the election by setting interface priorities. This also allows a router to be the DR in one network and a DROTHER in another.
- To set the priority of an interface, use the command **ip ospf priority** *value*, where value is 0 to 255.
  - A value of 0 does not become a DR or a BDR.
  - A value of 1 to 255 on the interface makes it more likely that the router becomes the DR or the BDR.

The example shows the commands being used to change the R1 G0/0/0 interface priority from 1 to 255 and then reset the OSPF process.

```
R1(config)# interface GigabitEthernet 0/0/0
R1(config-if)# ip ospf priority 255
R1(config-if)# end
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R1# *Jun 5 03:47:41.563: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on GigabitEthernet0/0/0
from FULL to DOWN, Neighbor Down: Interface down or detached
```

OSPF Configuration

OSPFv2 is enabled using the **router ospf** *process-id* global configuration mode command. The *process-id* value represents a number between 1 and 65,535 and is selected by the network administrator.

The *process-id* value is locally significant and represents the instant of the OSPF process in case of more than one process operating simultaneously. It is considered best practice to use the same *process-id* on all OSPF routers.

Observe possible commands after starting OSPF process in the figure below:

```
R1(config)# router ospf 10
R1(config-router)# ?
  area                                OSPF area parameters
  auto-cost                           Calculate OSPF interface cost according to bandwidth
  default-information         Control distribution of default information
  distance                            Define an administrative distance
  exit                                Exit from routing protocol configuration mode
  log-adjacency-changes               Log changes in adjacency state
  neighbor                            Specify a neighbor router
  network                             Enable routing on an IP network
  no                                  Negate a command or set its defaults
  passive-interface           Suppress routing updates on an interface
  redistribute                Redistribute information from another routing protocol
  router-id                           router-id for this OSPF process
R1(config-router)#
```

OSPF Configuration

After starting the OSPF process, the next step is the network declaration. The command *network* is used followed by the network address and the wildcard mask (to be explained next).

Any active interface that is configured with an IPv4 address belonging to the declared network will participate in the OSPFv2 routing process.

- **Note**: Some IOS versions allow the subnet mask to be entered instead of the wildcard mask. The IOS then converts the subnet mask to the wildcard mask format.

OSPF Configuration

```
R1(config)# router ospf 10
R1(config-router)# network 10.10.1.0 0.0.0.255 area 0
R1(config-router)# network 10.1.1.4 0.0.0.3 area 0
R1(config-router)# network 10.1.1.12 0.0.0.3 area 0
R1(config-router)#
```

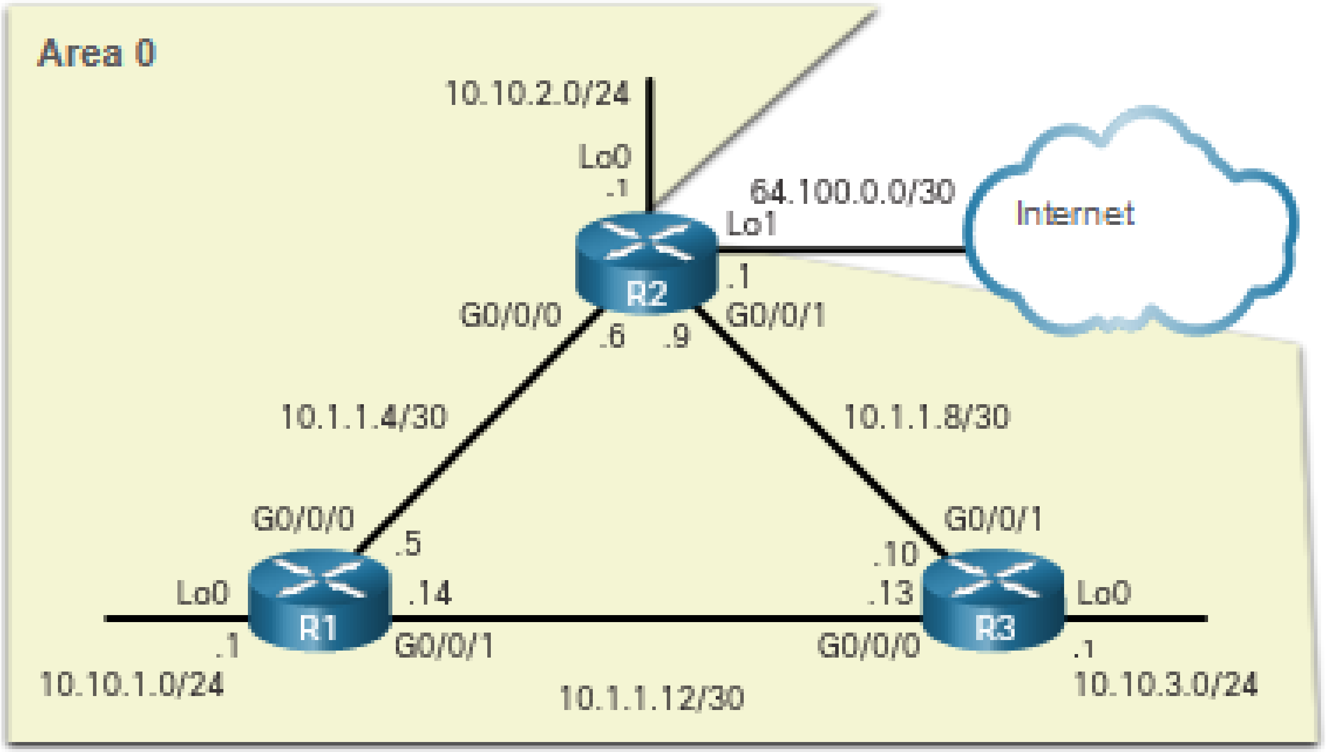

Wildcard Mask

- The wildcard mask is typically the inverse of the subnet mask configured on that interface.
- The easiest method for calculating a wildcard mask is to subtract the network subnet mask from 255.255.255.255, as shown for /24 and /26 subnet masks in the figure.

Calculating a Wildcard Mask for /24

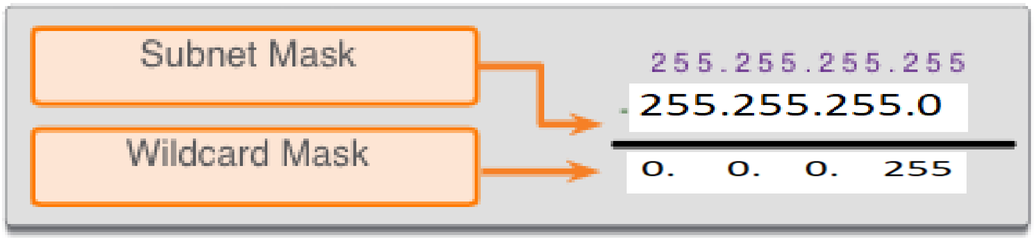

Calculating a Wildcard Mask for /26

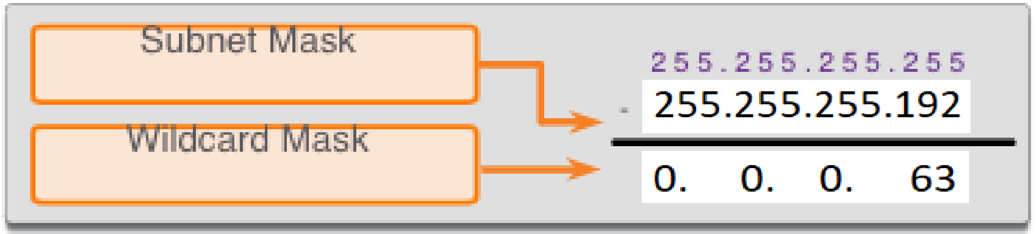

Passive Interfaces

By default, OSPF messages are forwarded out all OSPF-enabled interfaces. However, these messages only need to be sent out interfaces that are connecting to other OSPF-enabled routers.

- Sending out unneeded messages on a LAN affects the network in three ways:

- **Inefficient Use of Bandwidth** - Available bandwidth is consumed transporting unnecessary messages.
- **Inefficient Use of Resources** - All devices on the LAN must process and eventually discard the message.
- **Increased Security Risk** - Without additional OSPF security configurations, OSPF messages can be intercepted with packet sniffing software. Routing updates can be modified and sent back to the router, corrupting the routing table with false metrics that misdirect traffic.

Passive Interfaces

- To Solve this Problem we use the **passive-interface** router configuration mode command to prevent the transmission of routing messages through a router interface, but still allow that network to be advertised to other routers.
- The **show ip protocols** command may be used to verify that the interface is listed as passive and the problem is solved.

```
R1(config)# router ospf 10
R1(config-router)# passive-interface loopback 0
R1(config-router)# end
R1#
*May 23 20:24:39.309: %SYS-5-CONFIG_I: Configured from console by console
R1# show ip protocols
*** IP Routing is NSF aware ***
(output omitted)
Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
  Routing on Interfaces Configured Explicitly (Area 0):
    Loopback0
    GigabitEthernet0/0/1
    GigabitEthernet0/0/0
  Passive Interface(s):
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    3.3.3.3              110      01:01:48
    2.2.2.2              110      01:01:38
  Distance: (default is 110)
R1#
```

OSPF Metric (Cost)

- Routing protocols use a metric to determine the best path of a packet across a network. OSPF uses cost as a metric. A lower cost indicates a better path.
- The Cisco cost of an interface is inversely proportional to the bandwidth of the interface. Therefore, a higher bandwidth indicates a lower cost. The formula used to calculate the OSPF cost is:

Cost = reference bandwidth / interface bandwidth

- The default reference bandwidth is 10^8 (100,000,000); therefore, the formula is:

Cost = 100,000,000 bps / interface bandwidth in bps

- Because the OSPF cost value must be an integer, FastEthernet, Gigabit Ethernet, and 10 GigE interfaces share the same cost. To correct this situation, you can:
  - Adjust the reference bandwidth with the **auto-cost reference-bandwidth** command on each OSPF router.
  - Manually set the OSPF cost value with the **ip ospf cost** command on necessary interfaces.

OSPF Metric (Cost)

- Refer to the table for a breakdown of the cost calculation

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	100,000,000	÷	10,000,000,000	0.01 = 1
Gigabit Ethernet 1 Gbps	100,000,000	÷	1,000,000,000	0.1 = 1
Fast Ethernet 100 Mbps	100,000,000	÷	100,000,000	1
Ethernet 10 Mbps	100,000,000	÷	10,000,000	10

Same Costs due to reference bandwidth

Adjusting the Reference Bandwidth

- The cost value must be an integer. If something less than an integer is calculated, OSPF rounds up to the next integer. Therefore, the OSPF cost assigned to a Gigabit Ethernet interface with the default reference bandwidth of 100,000,000 bps would equal 1, because the next integer for 0.1 is 1.
- **Cost = 100,000,000 bps / 1,000,000,000 = 1**
- For this reason, all interfaces faster than Fast Ethernet will have the same cost value of 1 as a Fast Ethernet interface.
- To assist OSPF in making the correct path determination, the reference bandwidth must be changed to a higher value to accommodate networks with links faster than 100 Mbps.

Adjusting the Reference Bandwidth

- Changing the reference bandwidth does not actually affect the bandwidth capacity on the link; rather, it simply affects the calculation used to determine the metric.
- To adjust the reference bandwidth, use the **auto-cost reference-bandwidth** *Mbps* router configuration command.
  - This command must be configured on every router in the OSPF domain.
  - Notice in the command that the value is expressed in Mbps; therefore, to adjust the costs for Gigabit Ethernet, use the command **auto-cost reference-bandwidth 1000.** For 10 Gigabit Ethernet, use the command **auto-cost reference-bandwidth 10000.**
  - To return to the default reference bandwidth, use the **auto-cost reference-bandwidth 100** command.
- Another option is to change the cost on one specific interface using the **ip ospf cost** *cost* command.

Adjusting the Reference Bandwidth

- Whichever method is used, it is important to apply the configuration to all routers in the OSPF routing domain.
- The table shows the OSPF cost if the reference bandwidth is adjusted to accommodate 10 Gigabit Ethernet links. The reference bandwidth should be adjusted anytime there are links faster than FastEthernet (100 Mbps).
- Use the **show ip ospf interface** command to verify the current OSPFv2 cost assigned to the interface.

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	10,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	10,000,000,000	÷	1,000,000,000	10
Fast Ethernet 100 Mbps	10,000,000,000	÷	100,000,000	100
Ethernet 10 Mbps	10,000,000,000	÷	10,000,000	1000

OSPF Cost Computation

- The cost of an OSPF route is the accumulated value from one router to the destination network.
- Assuming the **auto-cost reference-bandwidth 10000** command has been configured on all three routers, the cost of the links between each router is now 10. The loopback interfaces have a default cost of 1.

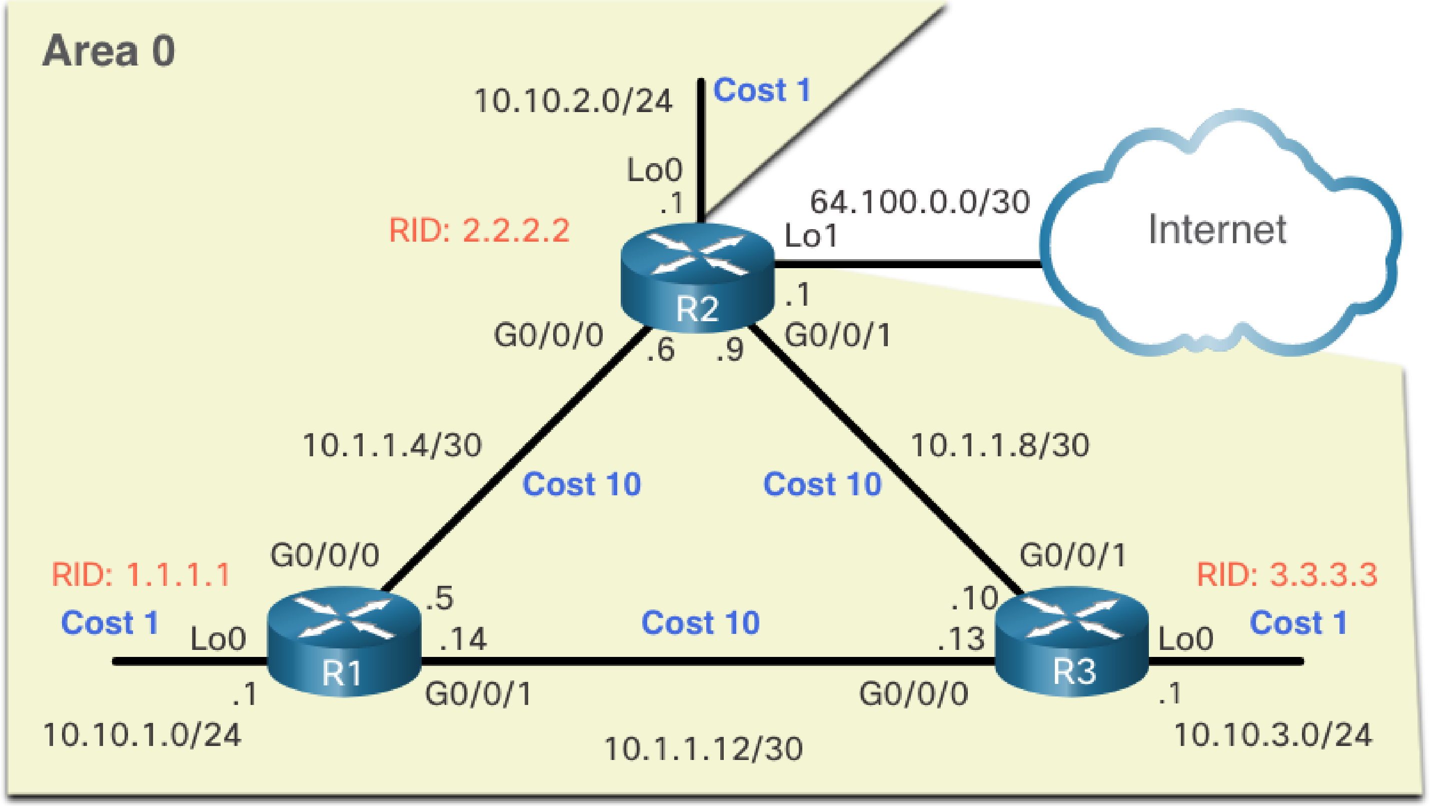

OSPF Cost Computation

- You can calculate the cost for each router to reach each network.
- For example, the total cost for R1 to reach the 10.10.2.0/24 network is 11. This is because the link to R2 cost = 10 and the loopback default cost = 1. 10 + 1 = 11.
- You can verify this with the **show ip route** command.

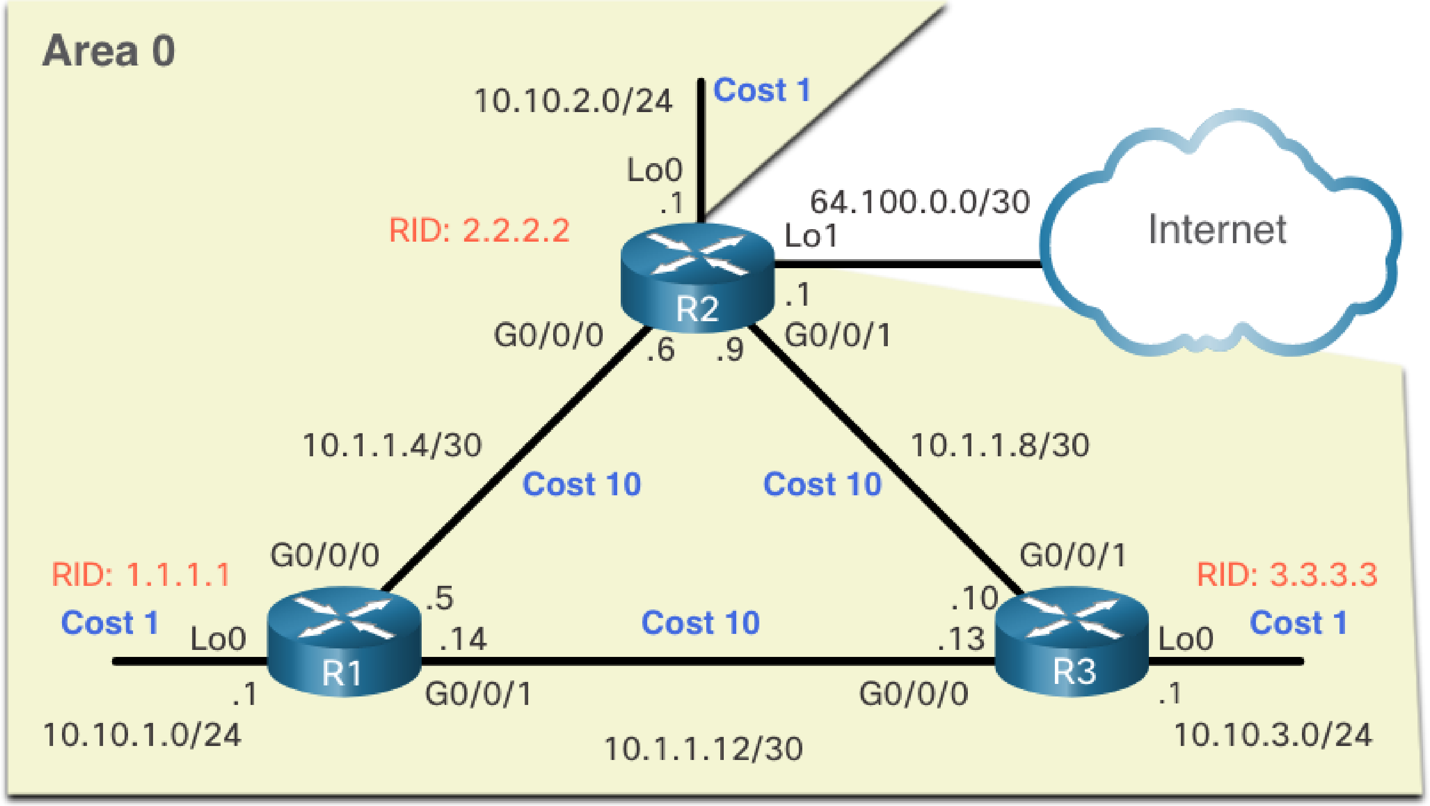

OSPF Cost Verification

Verifying the accumulated cost for the path to the 10.10.2.0/24 network:

```
R1# show ip route | include 10.10.2.0
O                10.10.2.0/24 [110/11] via 10.1.1.6, 01:05:02, GigabitEthernet0/0/0
R1# show ip route 10.10.2.0
Routing entry for 10.10.2.0/24
  Known via "ospf 10", distance 110, metric 11, type intra area
  Last update from 10.1.1.6 on GigabitEthernet0/0/0, 01:05:13 ago
  Routing Descriptor Blocks:
  * 10.1.1.6, from 2.2.2.2, 01:05:13 ago, via GigabitEthernet0/0/0
        Route metric is 11, traffic share count is 1
R1#
```

Manually Setting the Cost

Reasons to manually set the cost value include:

- The Administrator may want to influence path selection within OSPF, causing different paths to be selected than what normally would given default costs and cost accumulation.
- Connections to equipment from other vendors who use a different formula to calculate OSPF cost.

•To change the cost value reported by the local OSPF router to other OSPF routers, use the interface configuration command **ip ospf cost** *value*.

```
R1(config)# interface g0/0/1 R1(config-if)# ip
ospf cost 30 R1(config-if)# interface lo0
R1(config-if)# ip ospf cost 10 R1(config-if)# end
R1#
```

Propagating a Default Static Route

To propagate a default route, the edge router must be configured with the following:

- A default static route using the **ip route 0.0.0.0 0.0.0.0** [*next-hop-address* | *exit-intf*] command.
- The **default-information originate** router configuration command. This instructs R2 to be the source of the default route information and propagate the default static route in OSPF updates.

In the example, R2 is configured with a loopback to simulate a connection to the internet. A default route is configured and propagated to all other OSPF routers in the routing domain.

Note: When configuring static routes, best practice is to use the next-hop IP address. However, when simulating a connection to the internet, there is no next-hop IP address. Therefore, we use the *exit-intf* argument.

```
R2(config)# interface lo1
R2(config-if)# ip address 64.100.0.1 255.255.255.252
R2(config-if)# exit
R2(config)# ip route 0.0.0.0 0.0.0.0 loopback 1
%Default route without gateway, if not a point-to-point interface, may impact performance
R2(config)# router ospf 10
R2(config-router)# default-information originate
R2(config-router)# end
R2#
```

Verifying the Default Static Route

- You can verify the default route settings on R2 using the **show ip route** command. You can also verify that R1 and R3 received a default route.
- Notice that the route source on R1 is **O*E2**, signifying that it was learned using OSPFv2. The asterisk identifies this as a good candidate for the default route. The E2 designation identifies that it is an external route. The meaning of E1 and E2 is beyond the scope of this course.

```
R2# show ip route | begin Gateway
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
S*        0.0.0.0/0 is directly connected, Loopback1
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
(output omitted)
```

```
R1# show ip route | begin Gateway
Gateway of last resort is 10.1.1.6 to network 0.0.0.0
O*E2  0.0.0.0/0 [110/1] via 10.1.1.6, 00:11:08, GigabitEthernet0/0/0
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
(output omitted)
```

Thank you