OSPF Open Shortest Path First Protocol

OSPF Limerick

There once was a routing protocol
OSPF was its name, it was droll
It spread updates with care
To routers everywhere
So that packets would never grow stale.

- Most widely used interior gateway routing protocol
- Open standard (non-proprietary)
- Fast convergence
- It's been around since 1989.

How does it differ from other routing protocols?

Distance Vector Protocols

Distance Vector routing protocols base their decisions on the best path to a given destination based on the distance.

Distance is usually measured in hops, though the distance metric could be delay, packets lost, or something similar. If the distance metric is hop, then each time a packet goes through a router, a hop is considered to have traversed. The route with the least number of hops to a given network is concluded to be the best route towards that network.

The vector shows the direction to that specific network. Distance vector protocols send their entire routing table to directly connected neighbors.

* source - Pluralsight - Dynamic Routing Protocols: Distance Vector and Link State

Link State Protocols

Link state protocols are also called shortest-path-first protocols. Link state routing protocols have a complete picture of the network topology. Hence they know more about the whole network than any distance vector protocol.

Three separate tables are created on each link state routing enabled router. One table is used to hold details about directly connected neighbors, one is used to hold the topology of the entire internetwork and the last one is used to hold the actual routing table.

Link state protocols send information about directly connected links to all the routers in the network.

* source - Pluralsight - Dynamic Routing Protocols: Distance Vector and Link State

RIP and IGRP	**OSPF** and IS-IS (EIGRP is considered a hybrid protocol because it uses both distance vector and link state.)
Only communicates with neighbor routers	Communicates with **all** other routers in an area
Calculates path based on distance and vector	Calculates shortest path based on link-state parameters
Passes entire routing table to directly connected neighbor routers	Passes link-state routing updates to other routers

Link-State Routing

- Routers share their entire routing table via link-state advertisements (LSAs)
- LSAs include router and network information
- LSAs are stored in the routers' Link-State Database (LSDB)
- LSDB can be thought of like a network topology map.
- LSAs are the networking information NOT the packet. The packet that carries the LSA information is called the Link-State Update (LSU).

1. LSR Router A send a link-state request (LSR) to gather route information from the neighbor router -- Router B	2. LSU Router B replies with a Link -State Update (LSU) containing the requested Link-State Advertisement (LSA) information.	3. LSAck Router A receives the LSU with LSA attached and sends a Link-State Acknowledgement (LSAck) back to Router B

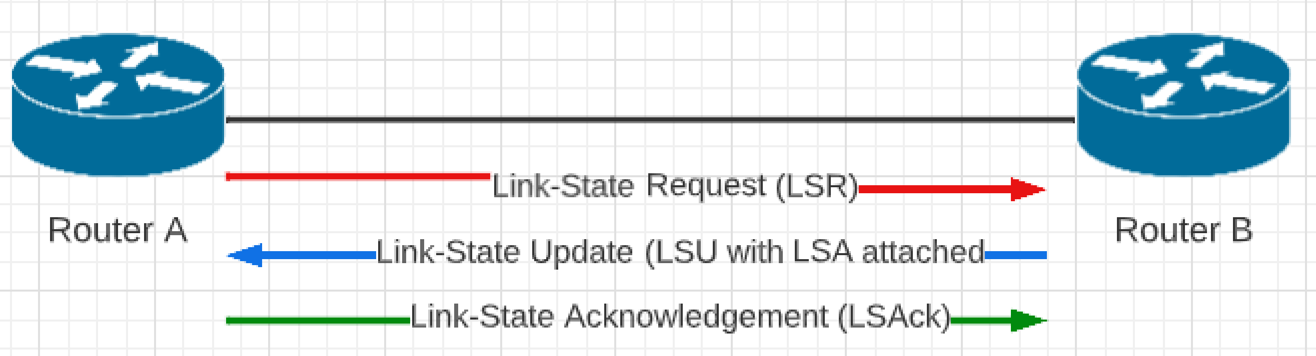

Dijkstra's Algorithm

The speed of OSPF can be attributed to the Dijkstra Algorithm (Shortest Path First)

Dijkstra's algorithm is a method for finding the shortest path between two points in a graph. The graph could represent, for example, a map with cities as nodes and roads as edges, where the length of the edges represents the distance between the cities. The algorithm starts at one node, called the "source" node, and explores all the paths to the other nodes, called "destination" nodes, in the graph. It keeps track of the shortest path to each destination node that it has found so far, and updates this information as it explores new paths.

At each step, the algorithm selects the destination node that can be reached with the shortest known path from the source node, and explores the paths to all the other nodes that can be reached from this node. This process is repeated until the algorithm has found the shortest path to the destination node that you are interested in.

The algorithm is called Dijkstra's algorithm because it was invented by a Dutch computer scientist named Edsger W. Dijkstra in the 1950s. Dijkstra's algorithm is also famous because it's the first algorithm that uses a priority queue(heap) rather than a queue to proceed from one node to the other.

It is often used for finding the shortest path in a navigation or transportation network, or for routing packets of data in a computer network

Image not found or type unknown	Illustration of Dijkstra's algorithm finding a path from a start node (lower left, red) to a goal node (upper right, green) in a robot motion planning problem. Open nodes represent the "tentative" set (aka set of "unvisited" nodes). Filled nodes are the visited ones, with color representing the distance: the greener, the closer. Nodes in all the different directions are explored uniformly, appearing more-or-less as a circular wavefront as Dijkstra's algorithm uses a heuristic identically equal to 0. * source: https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm

OSPF Path Selection

- Link-state information is flooded throughout the network, from router to router within an area, enabling all the routers within the area to have a synchronized and identical map of the network area.
- The Dijkstra's Algorithm (Shortest Path First) is then applied to this map of information to calculate the best (shortest) path.
- Once the Dijkstra Algorithm is applied and discovers the best routes the routes are applied to the routing table in the routers.

OSFP History

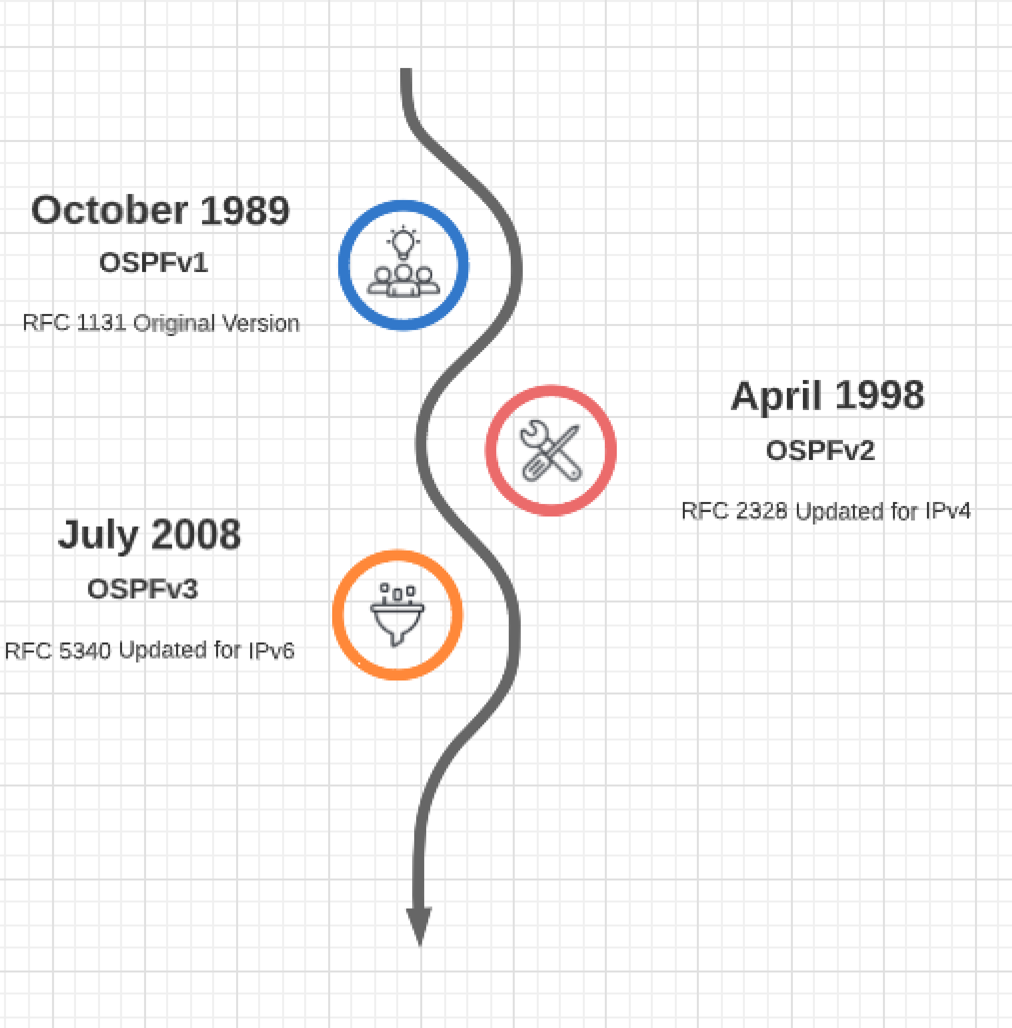

OSPF Areas

An OSPF Area is a network where routers all have the same routing information (LSAs) Network updates are localized per area and routers share topology information within the second table mentioned above (topology table). To reduce the size of the topology table in very large OSPF implementations the areas can be broken up into different OSPF area limits. Thus, the network updates will only be applied to the specifically identified area.

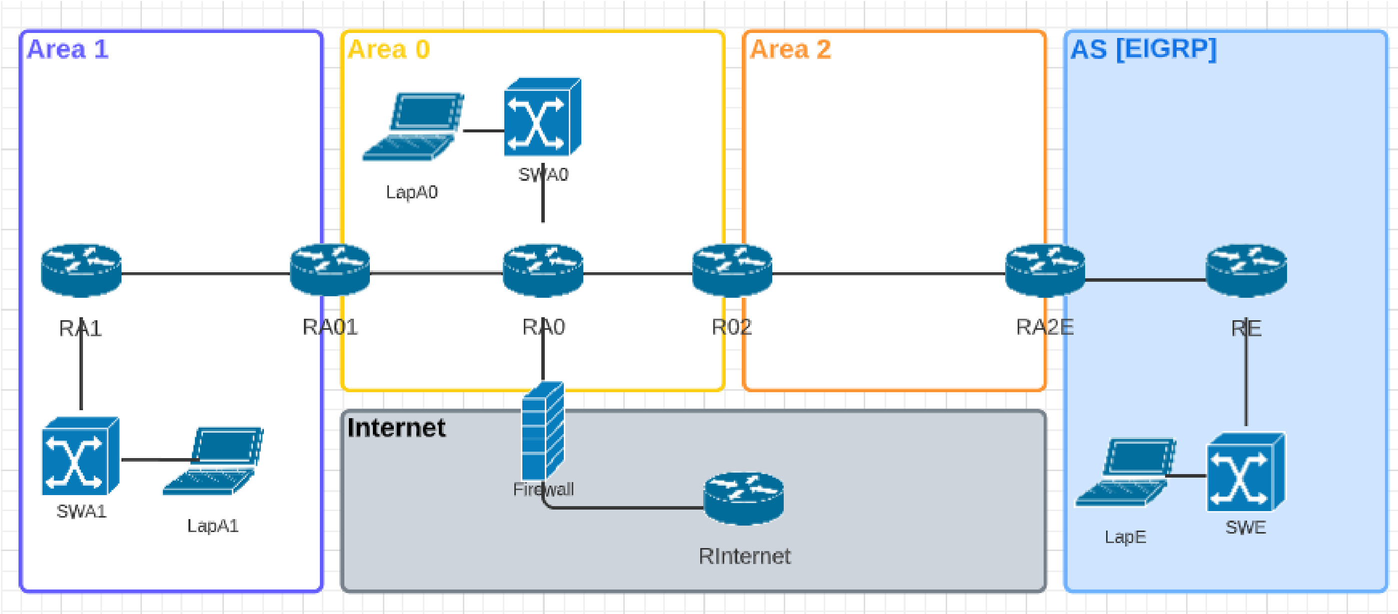

Area 0 is the Backbone and all other areas must connect back to Area 0

Adding an Area3 through Area 1 (as in the diagram below) would not be a correct or allowed configuration as Area 3 is NOT connected to Area 0, the Backbone.

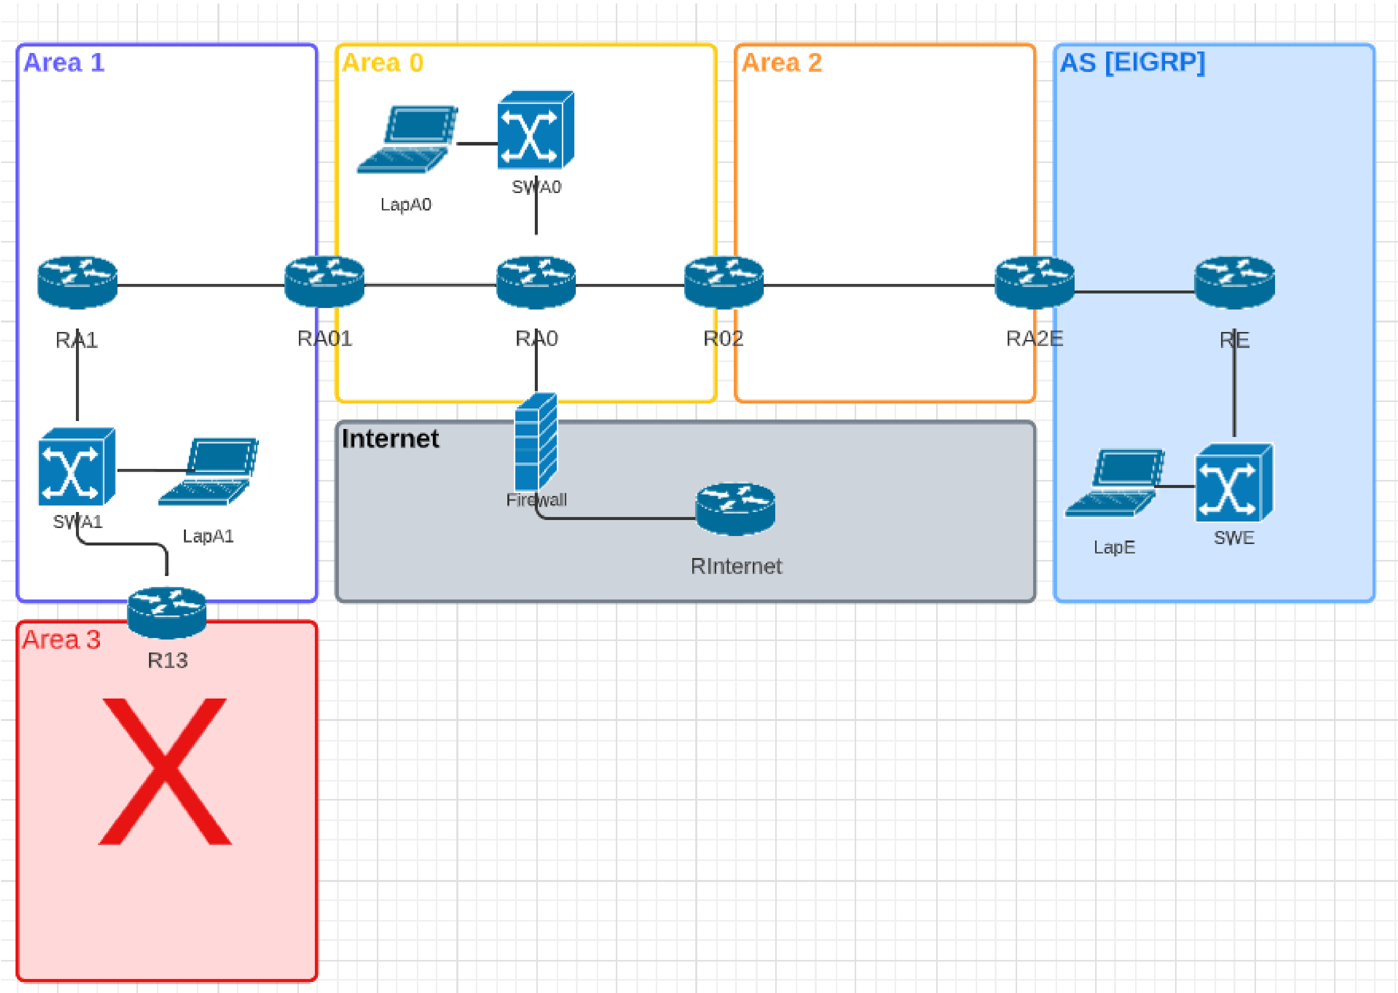

Adding Area 3 by connecting it to the backbone is the correct method for adding another OSPF Area as in the updated diagram below.

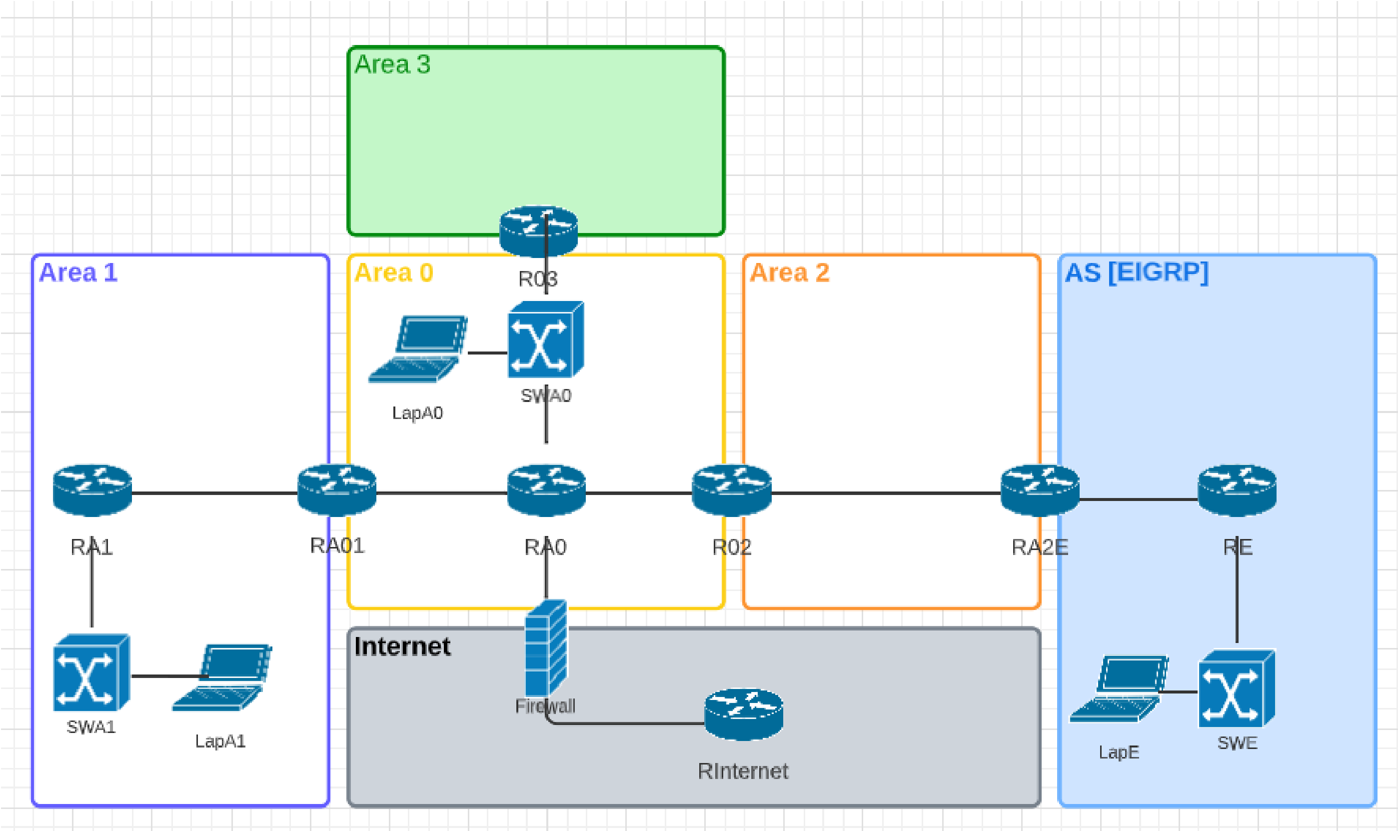

OSPF Neighbor Requirements

- Must be on the same Area
- Must match the following fields:
  - **Hello timer** - the interval in seconds that a router sends hello messages out of an OSPF-enabled interface
  - **Dead timer** - the time in seconds that an OSPF enabled interface will wait to receive a hello message from a neighbor before considering that neighbor to be down.
  - **Authentication**
  - **Maximum transmission unit (MTU)**
  - **Stub flags**

Hello timer Defaults: Sent every 10 seconds on broadcast or P2P networks. Sent every 30 seconds on non-broadcast multiple assess networks (NBMA) (i.e.; Frame Relay). Hello packets are sent to all other routers to the following broadcast addresses depending on IP version

- IPv4 - 224.0.0.5
- IPv6 - FF02::5

When sending Hello packets to designated routers the following addresses are used. What is a designated router? See Below.

- IPv4 - 224.0.0.6
- IPv6 - FF02::6

Dead timer Defaults: 4 times the hello timer.

```
RA0#show ip ospf interface g0/1
GigabitEthernet0/1 is up, line protocol is up
  Internet Address 22.22.22.2/30, Area 0, Attached via Network Statement
  Process ID 1, Router ID 33.33.33.1, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown     Topology Name
        0           1         no          no           Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 33.33.33.1, Interface address 22.22.22.2
  Backup Designated router (ID) 22.22.22.1, Interface address 22.22.22.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:05
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 2, maximum is 4
  Last flood scan time is 1 msec, maximum is 1 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 22.22.22.1  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

Wait timer is the number of seconds a router waits for the designated router or backup designated router to be advertised before beginning an election.

Retransmit is the number of seconds a router waits before retransmitting an OSPF packet that has not been acknowledged.

How to Change the Various Intervals

```
RA0(config-if)#ip ospf ?
  <1-65535>             Process ID
  adjacency             Adjacency staggering
  authentication        Enable authentication
  authentication-key    Authentication password (key)
  bfd                   Enable BFD on this interface
  cost                  Interface cost
  database-filter       Filter OSPF LSA during synchronization and flooding
  dead-interval         Interval after which a neighbor is declared dead
  demand-circuit        OSPF Demand Circuit
  flood-reduction       OSPF Flood Reduction
  hello-interval        Time between HELLO packets
  lls                   Link-local Signaling (LLS) support
  message-digest-key    Message digest authentication password (key)
  mtu-ignore            Ignores the MTU in DBD packets
  multi-area            Set the OSPF multi-area ID
  network               Network type
  prefix-suppression    OSPF prefix suppression
  priority              Router priority
  resync-timeout        Interval after which adjacency is reset if oob-resync is
                        not started
  retransmit-interval   Time between retransmitting lost link state
                        advertisements
  shutdown              Set OSPF protocol's state to disable under current
                        interface
  transmit-delay        Link state transmit delay
  ttl-security          TTL security check
```

Router0(config)# interface gigabitEthernet 0/1
Router0(Config-if)#ip ospf hello-interval 20

Designated Router (DR) and Backup Designated Router (BDR)

Router adjacencies are neighbor routers that share LSUs and database description packets.

A large network could have a huge number of adjacencies as show by this formula:

```
[n*(n-1]/2
```

So for 4 routers the number of adjacencies would be [4*(4-1)]/2 = 6. And for 10 the number of adjacencies would be [10*(10-1)]/2 = 45. As you can see the number of adjacencies becomes large very quickly. And a router will not be 'close' to all these adjacencies. Instead, an OSPF router is close to a specific router (neighbor) and that router is deemed the designated router (DR). And they also form an adjacency with a backup designated router (BDR). This is to reduce the adjacency volume.

```
Router1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.1.1        1   FULL/DR         00:00:31    172.16.1.1      GigabitEthernet0/0
3.3.3.3           1   FULL/BDR        00:00:32    172.16.2.2      GigabitEthernet0/1
```

How does a router select its DR and BDR?

- Priority level (default is 1) The higher the priority the better.
  - To adjust the priority level go to the interface global config and enter

    Router1(config-if)# ip ospf priority [#]

  - If you do not want a router to be considered as a DR or BDR set the priority level to zero (0) and it will no longer participate in the election.
- Router ID
  - Router ID can be configured

    Router1(config)#router ospf 1
    Router1(config-router)#router-id [id]

- Loopback interface IP address
- Interface IP address

```
Router1#show ip ospf
 Routing Process "ospf 1" with ID 2.2.2.2
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x000000
 Number of opaque AS LSA 0. Checksum Sum 0x000000
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 2 normal 0 stub 0 nssa
 External flood list length 0
    Area BACKBONE(0)
        Number of interfaces in this area is 3
        Area has no authentication
        SPF algorithm executed 8 times
        Area ranges are
        Number of LSA 7. Checksum Sum 0x02b477
        Number of opaque link LSA 0. Checksum Sum 0x000000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Not all network types conduct a DR/BDR election

Network Type	DR/BDR Election?
Point to Point	No
Broadcast	Yes
Non broadcast multiple access (NMBA)	Yes
Point to Multipoint	No

Area Border Router (ABR)

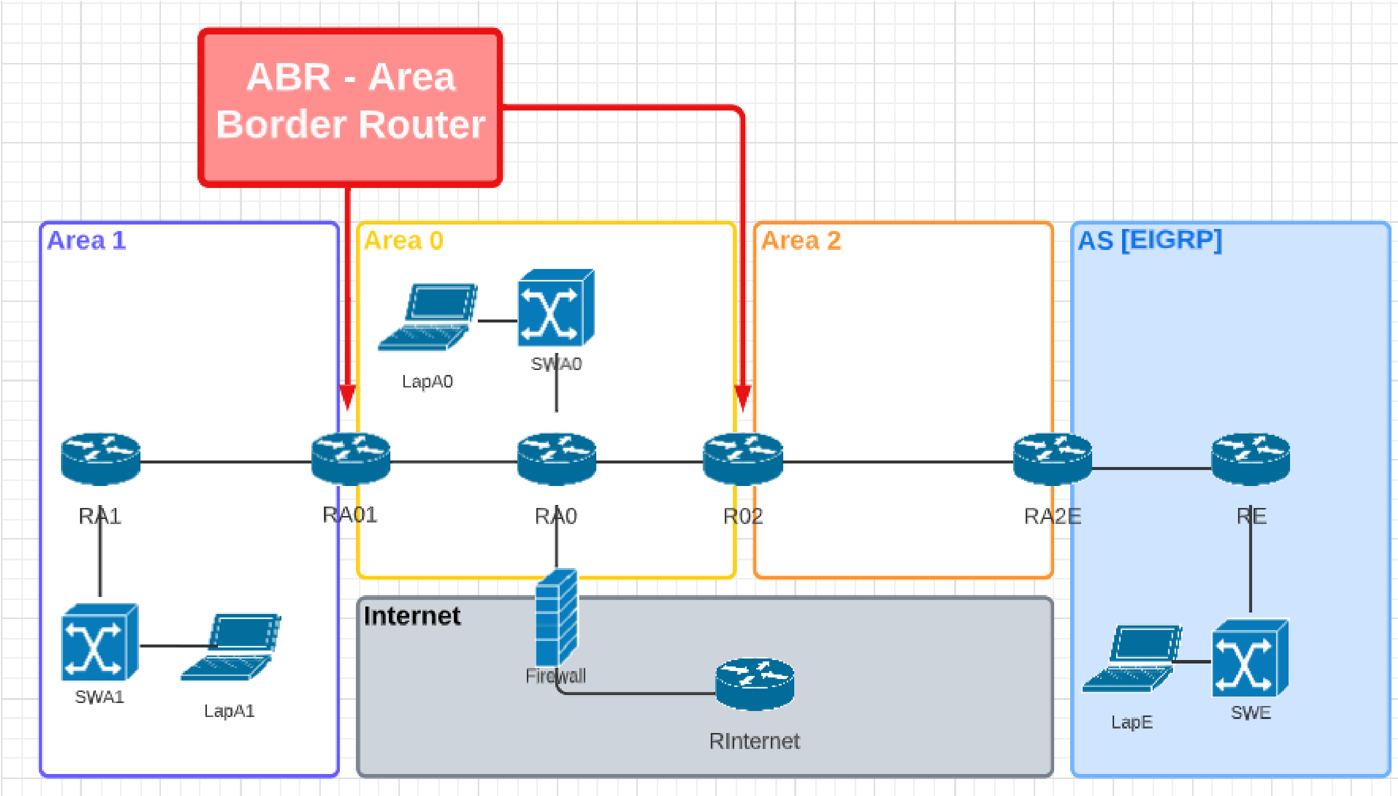

Typically an ABR has more processing power that an non ABR since they will have to store the routing tables for multiple OSPF Areas.

Autonomous System Boundary Router (ASBR)

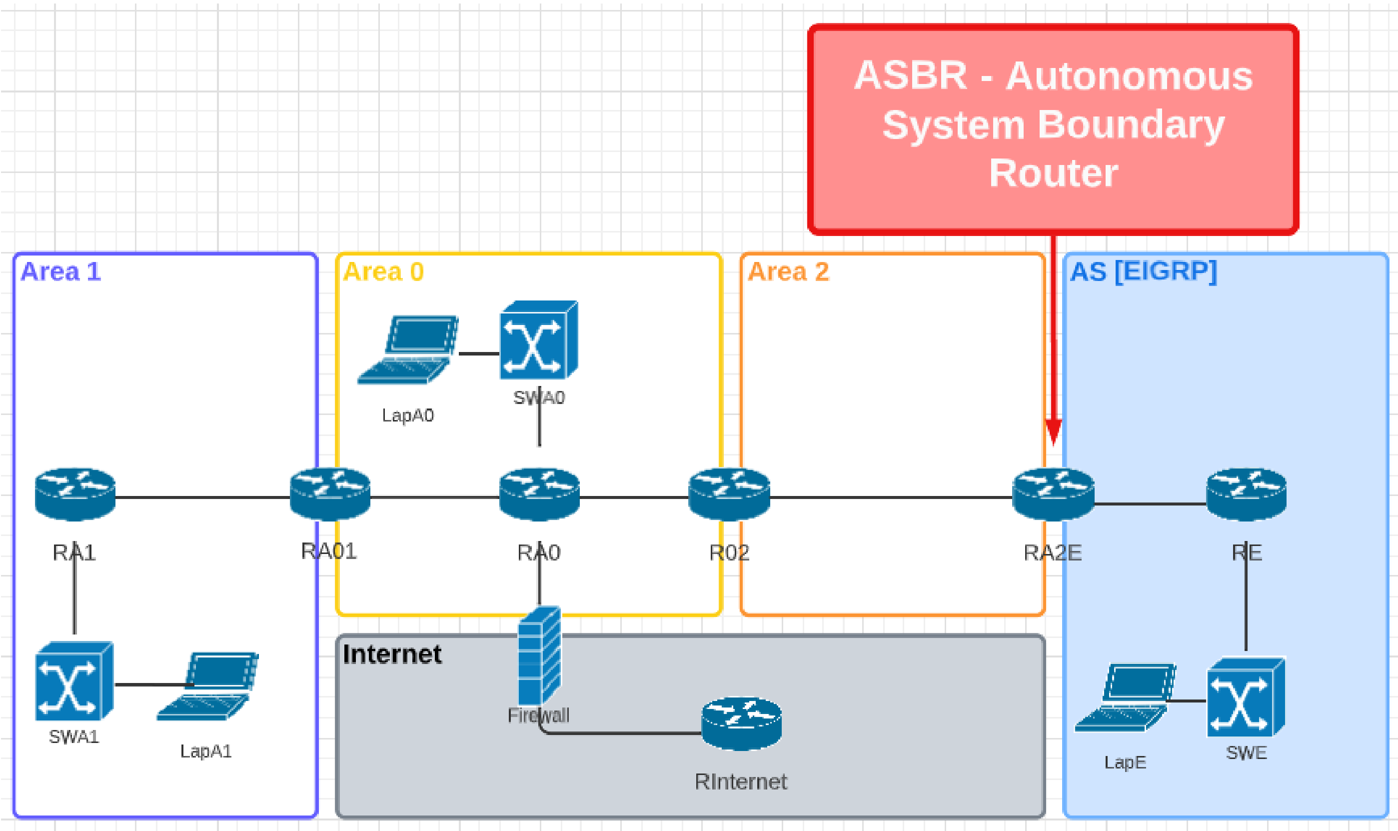

This router has one interface connected to OSPF and another interface with an entirely different routing protocol (i.e.; EIGRP)

LSA Types

- Type 1: Router LSA - this LAS type is what the routers use to advertise directly attached networks.
- Type 1: Network LSA - used in transit networks for DR/BDR elections
    - No need for Type 2 Network LSA in Point to Point (except Frame Relay) or Point to Multipoint
    - No need for Type 2 Network LSA of links going to end devices
- Type 3: Summary LSA - these LSAs come into play for Area Border Routers (ABR)
- Type 4 and Type 5 LSAs are created when an OSPF network is connected to another autonomous system (i.e.; EIGRP)
    - Type 4: Summary ASBR LSA
    - Created by an ABR to tell an Area how to reach and ASBR

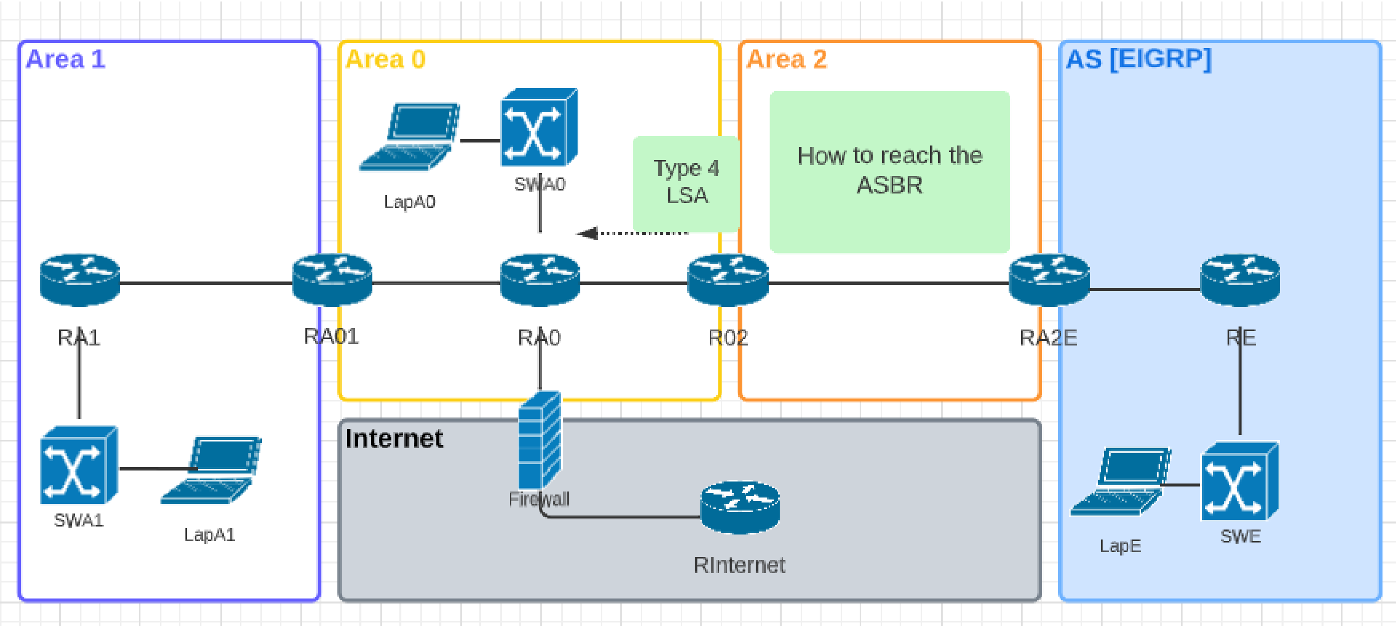

    - Type 5: AS External LSA
    - Type 5 is created by the ASBR to advertise networks in different autonomous systems

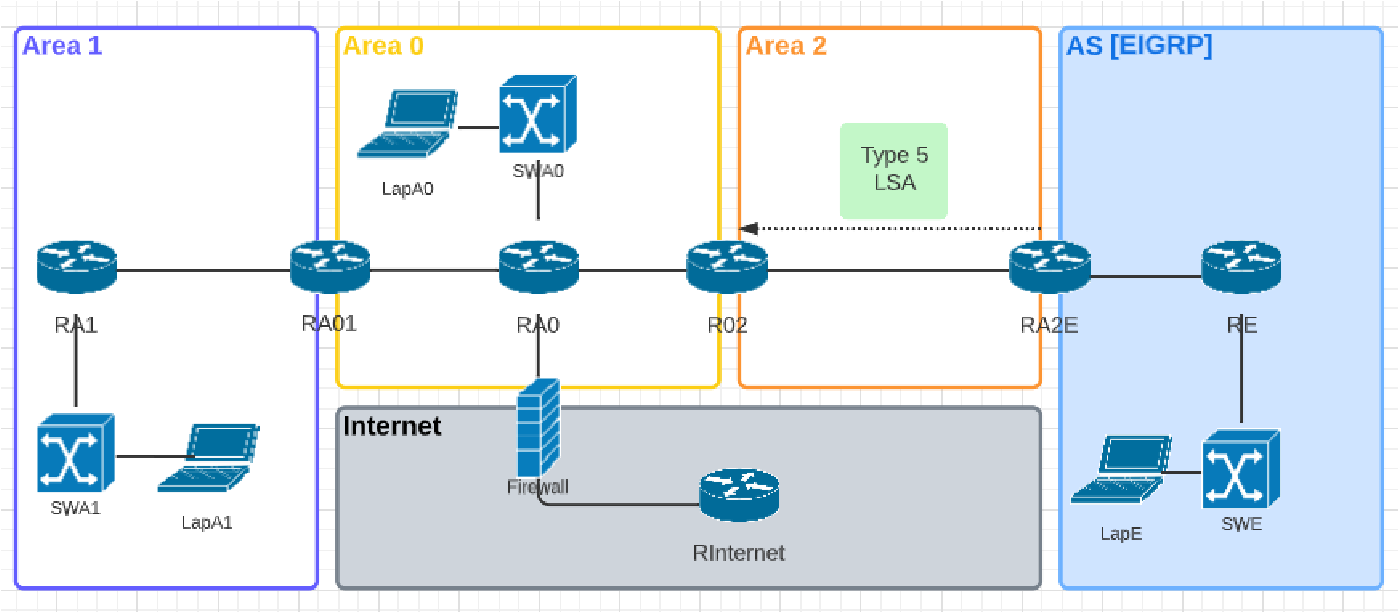

Router0#show ip ospf database

```
RA0#show ip ospf database

            OSPF Router with ID (33.33.33.1) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
22.22.22.1      22.22.22.1      177         0x80000006 0x00F628 1
33.33.33.1      33.33.33.1      164         0x80000009 0x00A24C 3
44.44.44.1      44.44.44.1      228         0x80000007 0x006CE9 1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
22.22.22.2      33.33.33.1      164         0x80000004 0x008564
33.33.33.2      44.44.44.1      228         0x80000004 0x0085DF

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.0.0.0        22.22.22.1      177         0x80000004 0x00FCED
11.11.11.0      22.22.22.1      177         0x80000004 0x00D502
44.44.44.0      44.44.44.1      228         0x80000004 0x001B17

                Summary ASB Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
55.55.55.1      44.44.44.1      228         0x80000004 0x008784

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
55.55.55.0      55.55.55.1      63          0x80000004 0x001A4D 0
192.168.55.0    55.55.55.1      63          0x80000004 0x00DE8A 0
```

Stub Areas

An area that is only connected to the backbone and is not connected to any other autonomous systems

Route Summarization

In IPv4 networking, route summarization, also known as route aggregation, is a technique used to reduce the number of routing table entries on a router. This is done by grouping together several smaller networks into a single, larger network, and then advertising the summary route to other routers instead of the individual smaller networks.

For example, imagine a network with several subnets, such as 10.1.1.0/24, 10.1.2.0/24, 10.1.3.0/24, and so on. Instead of having a separate routing table entry for each of these subnets, route summarization allows a router to group them together into a single network, such as 10.1.0.0/22. This single entry will take the place of all the individual entries, reducing the size of the routing table and making it more efficient.

Route summarization is done by using a mask (also called a subnet mask) that is longer than the original subnet mask. This allows a group of subnets to be represented by a single IP address and mask.

The benefit of this technique is that it can help reduce the number of routing table entries and improve the performance of the routing process. By reducing the size of the routing table, routers can make faster and more efficient routing decisions. It also can help to reduce the amount of routing information that needs to be exchanged between routers.

It's important to note that while route summarization can help to improve the scalability and performance of a network, it can also create potential problem like black hole routing, where packets might be dropped if they are not delivered to the correct destination. Therefore, it's important to be mindful of the address space being used and make sure that the summary route being advertised not overlap with other routes that exist in the network.

Example

Decimal	Octet 1	Octet 2	Octet 3	Octet 4
192.168.8.0	11000000	10101000	00001000	00000000
192.168.9.0	11000000	10101000	00001001	00000000
192.168.10.0	11000000	10101000	00001010	00000000
192.168.11.0	11000000	10101000	00001011	00000000
SUMMARY IP	11000000	10101000	00001000	00000000
	192	168	8	0
Subnet Mask	11111111	11111111	11111100	00000000
/22	255	255	252	0

Summarization Result = 192.168.8.0 with subnet 255.255.252.0 (/22)

Revision #4
Created 11 January 2023 15:39:42 by Glen Taylor
Updated 24 January 2023 22:45:52 by Glen Taylor