

OSPF

October 4, 2018 5:21 PM

OSPF traffic is multicast (unicast in NBMA)

SPF algorithm

The AD for OSPF is 110.

OSPFv2/OSPFv3

Packet Types

Type 1: Hellos

Type 2: Database description packet

Type 3: Link-State Request LSR

Type 4: Link-State Update (holds LSAs of other routers) LSU

Type 5: Link-State Acknowledgment

LSA Types (sent 30mins)

- Type 1 - Router Link State (Internal Routers (All routers)) [router ID | Router ID]
- Type 2 - Net Link State (DR) [IP | RouterID]
- Type 3 - Summary Net Link-state (ABR) O IA [IP | RouterID]
- Type 4 - Summary ASB Link State (ASBR) [router ID | Router ID]
- Type 5 - Type-5 AS External Link States (ASBR)

Type 1 is sent by everyone containing info about the router, networks, configurations (timers, cost, metric, prefix, mask, length, etc.) (Internal and ABR will send type 1 to their areas)

Type 2 is only advertise in multi-access areas because it needs a DR. Any area thats a multi-access can have a DR however DR only exist if it neighbors with another router in the multi-access area. If there is no OSPF neighbor then the network/interface is considered to be a stub network.

Type 3: ABR only sends type 3 towards the backbone and other areas beyond the backbone.

Type 4 are generated by ABR only when theres an ASBR in the same area (ASBR doesn't have to be in area 0 - backbone).

The ASBR will send a Type 1 with an E-bit and when the ABR receives that it will generate a type 4 which will send it out both ways (its area and the backbone). This is used to advertise that there is an ASBR (gateway to another protocol). And that you should go through me to get to the ASBR

Type 5 originates from ASBR and is flooded throughout the entire OSPF domain.

Show IP Route

O	Route that are learnt within the same area.
O IA	OSPF routes from another area. ABR advertises this and are LSA type 2 packets.
O E1	Routes outside of the OSPF domain and originating from the ASBR with metric calculated by OSPF. The metric is an additional cumulative cost that gets added by each router starting from ASBR.
O E2	(default) Outside of OSPF domain. Cost to destination is the same across source to destination. The metric never changes. Default metric is 20 and does not change when the packet goes through the network. The metric is the cost it takes to get from one end to the other.

How is the topology made?

The topology is made based on the SPF algorithm. Routers will send LSA Type 1-5 and the routers will build a database from that (show ip ospf database). The topology is built in a hierarchy topological way where the DR from area 0 is the king (only in multi access networks). And from there path connect to other areas and other areas will have there own paths where the ABR will transmit inter area ip routes.

These are the steps I think we should take to figure out the topology:

1. Distinguish all the areas
2. Distinguish all the routers and their Router ID if possible!!!
3. Find the ABR and associated them with the areas
4. Focus on routers in the same area
5. Connect them by identifying the neighbors (use router ID AND IP)
 - a. Careful not to mix IP/router-ID

Only in multi-access network will you find:

- DR
- BDR
- DROTHER

Election process for DR

This is only available in a multi-access network. The DR is the one with the highest priority.

0	Will never get elected
1	Default
1-255	Highest priority wins

The order for the DR elections is:

1. Highest priority (default is 1)
2. Highest Router-ID

OSPF Terminology

Database

The database contains OSPF neighbors in the entire OSPF domain. You can build the topology just by viewing some show commands and analyzing.

- Show ip ospf neighbors
- Show ip ospf database

SPF

This is the algorithm OSPF uses to find the shortest path. It uses the k-values to calculate the metric and then the FC.

LSU

This is a packet containing all the LSA.

LSA

Sent by every router advertising their configs and what networks they are able to reach.

Area

OSPF is divided into small areas for the network. This is to keep convergence time low and reduce the amount of traffic generated and routing entries. OSPF divides the topology into smaller areas which makes maintenance easier and waste less network resource

Backbone router/area

This is the area that interconnects all the areas. All routing info is sent to this area and then it is sent out to everyone.

ABR

This is the border router that connects areas to the backbone (multi-area OSPF). "O IA" will be displayed in the routing table for routes coming from the ABR.

ASBR

This is a border router in OSPF that connects different routing protocol domains. This is usually the router that connects an enterprise to the ISP.

DR (224.1.1.6)

The DR is a designated router that only exists inside a multi-access network (broadcast domain - ethernet connection). Routers in OSPF flood the network with Type 1 packets to advertise itself. When it's in a multi-access network, it sends broadcast traffic which floods the entire network which is bad. To prevent this, a Designated router is used to be the only one to generate traffic. A DR is elected before any LSA is sent out. Once a DR is elected all internal routers inside a multi-access network will send their LSA Type 1 to the DR and then the DR will send out LSA Type 2 to every router in the area. This reduces the amount of traffic being sent during convergence.

BDR

This is the backup DR which is elected in the process.

Internal (224.1.1.5)

All the routers that are participating in OSPF. Internal routers are independent from other areas meaning other areas may not know they exist.

Backbone (Transit)

This is the most important area of multi-area OSPF because it is in charge of connecting every area together. (virtual link can bypass this but it's not best practice).

Multi-access network

Advantages: large internetworks are divided into smaller internetworks called areas which reduces the amount of traffic (LSA) generated for convergence. Convergence only occurs inside areas and only has to calculate networks within an area. This means the database only has to contain network info on the area only so it doesn't care about other areas except for the ABR. SPF, flooding and database is restricted to an area.

OSPF Configurations

```
R1(config)# ospf version 2
```

1. Basic Config

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# network 192.168.0.0 0.0.0.255 area 0
R1(config-router)# network 192.168.1.0 0.0.0.3 area 1
R1(config-router)# end
```

2. Enable the interface for OSPF

```
R1(config)# int s0/0/0
R1(config-if)# ip ospf 10 area 0
```

ABR summary route

```
ABR(config)#router ospf 10
ABR(config-router)# area 0 range 10.0.0.0 255.0.0.0
```

ASBR Summary Route

```
ASBR (config-router)# summary-address 10.0.1.0 255.255.255.0
```

Passive Interface

```
R1(config)# router ospf 1
R1(config-if)# passive interface g0/0
```

To change priority:

```
R1(config)# clear ip ospf process
R1(config)# int s0/0/0
R1(config-if)# ip ospf priority <0-255>
```

OSPF Neighbors

Enabling OSPF on interfaces

1. Basic Config

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# network 192.168.0.0 0.0.0.255 area 0
R1(config-router)# network 192.168.1.0 0.0.0.3 area 1
R1(config-router)# end
```

2. Enable the interface for OSPF

```
R1(config)# int s0/0/0
R1(config-if)# ip ospf 10 area 0
```

Passive Interface

A router may have an interface where we do not want to have any LSA to go through. The router is participating in OSPF however the LAN connected to it (Stub) doesn't need to advertise any routes through the ABR. OSPF is good for redundancy stuff. So in order to prevent any LSA from being sent out to a stub (to preserve network resource we can set an interface to become passive. So the interface will still receive, however it will not send it out the passive interface. OSPF is prevented from sending any payload through that interface.

*NOTE: if the configured IP/network of a passive interface is included in a network statement, it will be advertised through OSPF as a stub network.

Command

```
R1(config)# router ospf 10
R1(config-router)# passive-interface g0/0
```

What must match to neighbor

For OSPF to function, the routers on both ends have to agree on the configurations to become neighbors. (Same Configs)

- Hello/Dead Intervals (10 seconds default)
- Network IP & Type
- Area ID
- Options (stub type & capability)
- Authentication (if any)
- *IP MTU - important to match
 - o ip ospf mtu-ignore -> this is used to ignore the MTU (Don't do this)

Optimizing convergence

Hello/dead timers should be identical on both ends to be safe. If an interval surpasses the dead timer then the link will have difficulties and will probably result in a down link. Hello/dead should be 3x about (i.e. 5/15). If one side sends packet at larger interval it may pass the dead timer limit.

Unique Router ID

Each router has to have a unique 32 bit router ID. The router ID is used to identify the router. This has to exist for OSPF to work. The process to determine the ID is (in this order):

1. Manually configure Router ID (**router-id x.x.x.x | clear ip ospf process**)
2. Highest Loopback address; else
3. Highest IP sent on the router

Same MTU

The MTU must be identical on both ends of the link.

Virtual Links

Why? (Concept)

A virtual link is an OSPF feature that creates a logical extension of the backbone area to a remote area, without actually adding any physical interfaces between the two areas. The area doesn't have to be connected to area 0. Technically speaking, this means that the routes from the remote area will not be learnt throughout the topology because routes have to be connected to the backbone to propagate routes between area. It may happen that a topology may have a network that cannot reach the backbone and in that case, virtual-links can be used to connect them.

Configurations

To achieve this, you must identify the router ID of the transit area between the remote area to the backbone (the routers it need to go through from one end to the other). The routers can only be ABR and 1 must be connected to the backbone. The other must be connected to the remote area. Intermediate routers can be there too as long as they are ABR.

```
(area 10)-(x1)-(area 0)-(x2)-(area 20)-(x3)-(area 100)
```

In this example 100 is not connected so you need to use a virtual link to connect 100 to the backbone.

Step 1: Identify the router ID for the transit; x2 and x3.

Step 2:

```
X3(config)#router ospf 10
X3(config-router)# area 100 virtual-link 2.2.2.2

X2(config)#router ospf 10
X2(config-router)# area 100 virtual-link 3.3.3.3
```

What do you expect to see

- From "show ip route" we should be able to see routes to the remote area as "O IA"
- From "show ip ospf neighbor" we should be able to see the neighbor connected as a OSPF_VL0 (Note: if the router appears twice, one will be a virtual link)
- From "show ip ospf interface" it will identify its virtual interface
- From "show ip ospf database" virtual links will be "DNA" (DoesNotAge)

Route Summarization

ABR summary route command

(Example 192.168.100.0/24 - 192.168.103.0/24)

```
ABR(config)#router ospf 10
ABR(config-router)# area 0 range 192.168.100.0 255.255.252.0
```

Consolidating Type 3 LSA

We will only see 192.168.100.0 being advertised by the connected router.

What do we see in the database summary?

When looking in "show ip ospf database" we will only see 192.168.100.0 because it's a summary. This actually represents 192.168.100.0/22. (check routing table for proof)

What do we see in the routing table?

However, if we look in the routing table we can see 192.168.100.0/22. This is expected because it's the summary and it will be a Null (to protect against looping). You will also see 192.168.100.0/24 which indicates it knows routes there.

Manual Summarization at ABRs

ASBR summary route command

```
ASBR(config)#router ospf 10
ASBR (config-router)# summary-address 10.0.1.0 255.255.255.0
```

Type 5 LSA

The ASBR will send a single LSA type 5 throughout the entire OSPF domain advertising that it knows how to reach this network (multiple network because it's a summary).

Default route

Done from the ASBR router to advertise it has a default route. You do not have to write the default route. You only have to do the following command and it will create it automatically.

Default route command

```
ASBR (config) #router ospf 10
ASBR (config-router) # default-information originate always
```

From routing table

```
O*E2 0.0.0.0/0 [110/1].....
```

Stub Areas

Only one exit and can't be a virtual-link transit area.

*All routers in the stub area must be configured as a stub.

**ONLY TSA will the ABR need to be configured differently

Stubby Area (stub)

- Only sees routes in its own AS (will see routes to other networks/areas)
- Accept Type 1, 2, 3 LSA
- Does not accept Type 4, 5
- A default route to the ABR is propagated throughout the stub area

Stub router command

```
R1 (config) #router ospf 10
R1 (config-router) # area 1 stub
```

```
ABR (config) #router ospf 10
ABR (config-router) # area 1 stub
```

Totally Stubby Area

- Only sees routes in its own area
- Does not get anything from other areas
- Does not accept 3, 4, 5 (except default from ABR)
- Default route is sent automatically by the ABR

Stub router command

```
R1 (config) #router ospf 10
R1 (config-router) # area 1 stub
```

```
ABR (config) #router ospf 10
ABR (config-router) # area 1 stub no-summary
```

Not-so-Stubby Area (NSSA)

- It's a stubby area with an ASBR
- ASBR can send Type 4 and 5 through the stub into the backbone
- The ABR will not allow Type 5 into the NSSA
- No default route going inside the NSSA

NSSA router command

```
R1 (config) #router ospf 10
R1 (config-router) # area 1 nssa
```

Best Practices

- A router should not be in more than 3 area
- Not more than 50 routers per area
- A router should not have more than 60 neighbors