

OSPF Lab

Introduction


The purpose of this exercise is to learn how to configure OSPF between the Border and Core routers so that they exchange network reachability information and maintain their own routing tables dynamically.

In the last lab we configured some static routes to forward packets between the Border and Core routers. While it's possible to run the network like this it can be difficult to keep the static routes up to date as your network changes.

Make sure to take the examples and adapt them to your own router, network topology and addressing scheme.

Lab Layout

As a reminder, the following diagram shows the layout of the devices and all the links for each campus:



Our campus network consists of two routers, **bdr1-campusX** and **core1-campusX** as well as six switches.

The following table shows the connections between each device in the campus:

Device	Interface	Remote Device	Remote Interface
dist1-bY-campusX	FastEthernet1/12	edge1-bY-campusX	FastEthernet1/14
	FastEthernet1/13	edge1-bY-campusX	FastEthernet1/15
	FastEthernet1/14	edge2-bY-campusX	FastEthernet1/15
core1-campusX	FastEthernet1/0	bdr1-campusX	FastEthernet0/1
	FastEthernet1/1	dist1-b1-campusX	FastEthernet1/15
	FastEthernet1/2	dist1-b2-campusX	FastEthernet1/15
	FastEthernet1/15	pc1-campusX	

Replace **Y** with your building number and **X** with your campus number.

Accessing the Lab

The Workshop Instructors will let you know what the lab environment is. It will either be run on a Virtual Platform, or on real physical switches provided in the Training Room.

Refer to the **correct** document below for information about logging into the devices that have been assigned to you:

VIRTUAL ENVIRONMENT: [Lab Access Instructions - Virtualised Platform](#)

PHYSICAL HARDWARE: [Lab Access Instructions - Physical Hardware](#)

Summary Commands

During all exercises, verify the output of the following commands:

```
show arp           : Shows ARP cache
show interface <int> : Shows interface state and configuration
show ip interface  : Shows interface IPv4 state and configuration
show ipv6 interface : Shows interface IPv6 state and configuration
show run interface <int> : Shows the configuration of the interface
```

Dynamic Routing with OSPFv2 & OSPFv3

We will now prepare the border and core router to introduce OSPF to them. We will be replacing the static routes we configured earlier with the OSPF, where the routes on each router are announced to their neighbour. In future this will mean that we do not need to introduce a static route when we introduce new subnets in our campus.

Loopback Interface

First of all we will set up a Loopback interface on both routers. It is industry convention and best practice to create a Loopback interface on all L3 devices. The loopback is used to generate the router-id for routing processes on Cisco routers. It is used for many other things too, as we will see in the remainder of this workshop.

On the Border router:

```
interface Loopback0
 ip address 100.68.X.241 255.255.255.255
 ipv6 address 2001:DB8:X:2::241/128
!
```

On the Core router:

```
interface Loopback0
 ip address 100.68.X.242 255.255.255.255
 ipv6 address 2001:DB8:X:2::242/128
!
```

Setting up the OSPF routing process

Configure a new OSPF routing process for both IPv4 and IPv6

Notice that we will use the number “41” as the OSPF process number for the routers. This number is local to the router, so it doesn't need to match the process number of a neighboring router. However, it is strongly recommended that you use the same number throughout your network. Most people use their Autonomous System number (although OSPF has nothing to do with the BGP AS).

On the Core router:

```
router ospf 41
 router-id 100.68.X.242
 log-adjacency-changes
 passive-interface default
 no passive-interface FastEthernet1/0
!
ipv6 router ospf 41
 router-id 100.68.X.242
 log-adjacency-changes
 passive-interface default
 no passive-interface FastEthernet1/0
!
```

On the Border router:

```
router ospf 41
  router-id 100.68.X.241
  log-adjacency-changes
  passive-interface default
  no passive-interface FastEthernet0/1
!
ipv6 router ospf 41
  router-id 100.68.X.241
  log-adjacency-changes
  passive-interface default
  no passive-interface FastEthernet0/1
!
```

We mark all interfaces as “passive” by default (which means that we do not look for adjacencies on the interfaces), and then only mark the interfaces where we do expect to find an adjacency as “no passive”. The Core router will only expect to see the Border router, hence why FastEthernet 1/0 has been marked as “no passive”.

Notice that we have explicitly configured the router-id for both the IPv4 and IPv6 OSPF processes - the router-id is a 32-bit integer, and Cisco’s default is for OSPF to automatically use the Loopback interface address. However, we have simply adopted the Loopback interface address for the router-id without making use of the builtin default. (This is more future proof for when networks no longer need IPv4 and there would be no more IPv4 addresses used on campus networks.)

Adding the OSPF process to interfaces

We will now configure OSPF on the interfaces **where adjacencies need to be established**, and also on **any interface that needs to have its subnets advertised by OSPF**.

On the Core router:

```
interface Loopback0
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet1/0
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface Vlan 10
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface Vlan 11
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface Vlan 12
```

```
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
interface Vlan 20
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
interface Vlan 21
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
interface Vlan 22
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
interface FastEthernet1/15
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
```

Note that we must put OSPF configuration commands on each interface which has a subnet that we expect to put into OSPF. While Cisco IOS has a shortcut for IPv4, it does not have this shortcut for IPv6, hence us using the format in the example above.

On the Border router:

```
interface Loopback0
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
interface FastEthernet0/1
ip ospf 41 area 0
ipv6 ospf 41 area 0
!
```

STOP - Checkpoint.

```
show ip ospf neighbor      : show adjacencies
show ip route              : show routes in routing table
show ip ospf               : shows general OSPF information
show ip ospf interface     : shows the status of OSPF in an interface
show ip ospf rib           : shows the OSPF Routing Information Base (RIB)
```

The status commands for IPv6 are very similar: simply replace “ip” in the above show commands with “ipv6”.

Question: How many routes do you see on the Border Router for each access network? Do you see the campus subnets appearing in OSPF now?

Removing Static Routes

The OSPF Lab is normally conducted after completing the static routing exercise. There still will be lots of static routes within the network so that each group can reach the others.

We will now remove these static routes, carefully.

Ensure that Border router is announcing a default route by OSPF into the core network. To do this, we use the following command.

On the border router:

```
router ospf 41
 default-information originate
!
ipv6 router ospf 41
 default-information originate
!
```

This will originate a default route into OSPF (which means that it will be distributed to OSPF neighbours) **as long as a default route exists in the router's Global RIB**. The default is already in the Border Router Global RIB because of the static default route we set up in the Static Routing Lab exercise.

Once the Border router is announcing the default route by OSPF, we can remove this on the Core router using:

```
no ip route 0.0.0.0 0.0.0.0 100.68.X.1
no ipv6 route ::/0 2001:DB8:X::1
```

Check that routing to the other groups is still working using:

```
show ip route
show ipv6 route
```

You should see the default route in the table as an OSPF announcement. There should now be no static routes remaining in the core router.

The Border router can now remove the static routes pointing to the Core Router for each of the Staff, Student and Management VLANs. These routes are now learned by OSPF from those routers.

To finish off, all teams should now check the Routing Table. Document the output of:

```
show ip route
show ipv6 route
```

and be prepared to show this to the workshop instructors. The group network is now using a dynamic routing protocol to share routing information within the group - a much more scalable solution than the effort that was required to set up the static routes in the previous lab exercise.

Dropping packets for networks that are not in use

We are only using some of the subnets in our network address allocation. If we receive packets from outside our network at the Border router to these ranges:

```
100.68.X.0/24
172.2X.0.0/16
2001:DB8:X::/48
```

we should never forward packets to the Core router unless we have learnt a route for a subnet. For example, what should we do with a packet addressed to **100.68.X.67**?

On the Border router try running the command:

```
bdr1-campus1#show ip route 100.68.1.67
% Subnet not in table
```

A: The only entry in the routing table that matches this address is the default route.

It's good practice to have routes in place that drop traffic like this. We do this on the Border router only using:

```
ip route 100.68.X.0 255.255.255.0 Null0
ip route 172.2X.0.0 255.255.0.0 Null0
ipv6 route 2001:DB8:X::/48 Null0
```

The more specific routes we learn on the Border router via OSPF make sure that only traffic for networks that are in use are sent to the Core router.

And because we are now sending traffic destined for unannounced routes towards the Null interface, we need to set up the Null interface to not respond with ICMP unreachable messages (this is so that a stream of traffic to unannounced routes does not result in the router sending “unreachable” responses).

```
interface Null0
 no ip unreachable
```

Final Checks

Check the routing table you see on both the border and core routers now. Do you have full Internet connectivity for IPv4 - check by trying to ping or traceroute to 8.8.8.8. If you have followed the steps above you'd have migrated your campus from static routes to dynamic routing using OSPF without any break in connectivity.

Here is an example of the routing table from the border router in campus1 from a previous Campus Network Design Workshop:

```
bdr1-campus1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 100.68.0.1 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 100.68.0.1
      100.0.0.0/8 is variably subnetted, 8 subnets, 4 masks
C      100.68.0.0/30 is directly connected, FastEthernet0/0
L      100.68.0.2/32 is directly connected, FastEthernet0/0
S      100.68.1.0/24 is directly connected, Null0
C      100.68.1.0/28 is directly connected, FastEthernet0/1
L      100.68.1.1/32 is directly connected, FastEthernet0/1
O      100.68.1.128/28 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
C      100.68.1.241/32 is directly connected, Loopback0
O      100.68.1.242/32 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
      172.21.0.0/16 is variably subnetted, 7 subnets, 2 masks
S      172.21.0.0/16 is directly connected, Null0
O      172.21.10.0/24 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
O      172.21.11.0/24 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
O      172.21.12.0/24 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
O      172.21.20.0/24 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
O      172.21.21.0/24 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
O      172.21.22.0/24 [110/2] via 100.68.1.2, 00:02:00, FastEthernet0/1
```

How well does this compare with what you see?

What about IPv6?

From:

<https://workshops.nsrc.org/dokuwiki/> - Workshops

Permanent link:

<https://workshops.nsrc.org/dokuwiki/2016/nsrc-tein-bdren/ospf>

Last update: **2016/10/19 09:22**

