

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.

Router types used in the lab

Cisco 7206VXR under Dynamips

Address Space Allocation

See the [IP Address Plan](#) for details of your addresses.

You will need to replace 'X' with the number of your campus group!

Exercises

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

<code>show arp</code>	: Shows IPv4 ARP cache
<code>show interface <int></code>	: Shows interface state and configuration
<code>show ip interface</code>	: Shows interface IPv4 state and config
<code>show ipv6 interface</code>	: Shows interface IPv6 state and config

Dynamic Routing with OSPFv2 & OSPFv3

Configure a new OSPF routing process for both IPv4 and IPv6 and 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**.

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 neighbouring router. However, it is strongly recommended that you use the same number throughout your network. Most

people use their Autonomous System (AS) number (although the OSPF process number 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 FastEthernet0/0
!
ipv6 router ospf 41
  router-id 100.68.X.242
  log-adjacency-changes
  passive-interface default
  no passive-interface FastEthernet0/0
!
interface Loopback0
  ip address 100.68.X.242 255.255.255.255
  ipv6 address 2001:DB8:X:2::242/128
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet0/0
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet0/1.41
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet0/1.51
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet0/1.61
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet1/0.42
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet1/0.52
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
interface FastEthernet1/0.62
  ip ospf 41 area 0
  ipv6 ospf 41 area 0
!
```

```
interface FastEthernet1/1
 ip ospf 41 area 0
 ipv6 ospf 41 area 0
!
```

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 0/0 has been marked as “no passive”.

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:

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

```
r1-bdr-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
```

From:

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

Permanent link:

<https://workshops.nsrc.org/dokuwiki/doku.php?id=2016:nsrc-tein-lernet:ospf>

Last update: **2016/03/14 04:57**

