

OSPF with Multi–Area Adjacency Configuration Example



Document ID: 118879

Contributed by Narendra Naukwal, Cisco TAC Engineer.
Mar 31, 2015

Contents

Introduction

Prerequisites

- Requirements

- Components Used

Background Information

Configure

- Network Diagram

- Initial Configurations

 - R1

 - R2

 - R3

 - R4

 - R5

- Default Behavior

- Multi–Area Adjacency Configuration

Verify

Troubleshoot

Introduction

This document describes how to configure the Open Shortest Path First (OSPF) link state routing protocol for Multi–Area Adjacency.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- OSPF
- Multi–Area Adjacency

Cisco also recommends that these requirements be met before you attempt the configuration that is described in this document:

- The OSPF link state routing protocol must be pre–configured in the network.
- Only two OSPF speakers use the interface between which OSPF Multi–Area functionality works. Multi–Area OSPF only works on Point–to–Point network types.

Components Used

The information in this document is based on Multi–Area OSPF.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.


Background Information

The OSPF link state routing protocol uses the concept of Areas, which are sub–domains within the OSPF domain. A router within an Area maintains the complete topology information of that Area. By default, an interface can only belong to one OSPF Area. This can not only cause sub–optimal routing in the network, but it can also lead to other issues if the network is not designed correctly.

When Multi–Area Adjacency is configured on an interface, the OSPF speakers form more than one Adjacency (ADJ) over that link. The Multi–Area interface is a logical, point–to–point interface over which the ADJ is formed. This document describes a scenario where Multi–Area OSPF ADJ can be used in order to work around a problem and meet the network requirements.

Configure

Network Diagram



R2 has a static route for 10.1.1.1/32 Prefix, which points to R1.
This static is redistributed in OSPF domain.

In this network diagram, a network/OSPF domain is used. The system requires that traffic from Router 5 (R5) to R1 (10.1.1.1) always passes through R3. Assume that R3 is a Firewall in the network through which all of the traffic should be routed, or that the link between R3 and R4 has more bandwidth than the link between R2 and R4. In either case, the system requires that the traffic must flow through R3 when it passes from R5 to R1 (10.1.1.1/32 prefix).

Initial Configurations

This section describes the initial configurations for R1 through R5.

R1

```
!  
interface Ethernet0/0  
ip address 192.168.12.1 255.255.255.0  
end  
  
!  
interface Loopback0  
ip address 10.1.1.1 255.255.255.255  
end  
  
!  
ip route 0.0.0.0 0.0.0.0 192.168.12.2  
!
```

R2

```
!  
interface Ethernet0/0  
ip address 192.168.12.2 255.255.255.0  
end  
  
!  
interface Ethernet0/1  
ip address 192.168.23.2 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 0  
end  
  
!  
interface Ethernet0/2  
ip address 192.168.24.2 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
end  
  
!  
interface Loopback0  
ip address 10.2.2.2 255.255.255.255  
end  
  
!  
ip route 10.1.1.1 255.255.255.255 192.168.12.1  
  
!  
router ospf 1  
router-id 0.0.0.2  
redistribute static metric-type 1 subnets  
!
```

R3

```
!  
interface Ethernet0/0  
ip address 192.168.34.3 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
end
```

```
!  
interface Ethernet0/1  
ip address 192.168.23.3 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 0  
end  
  
!  
interface Loopback0  
ip address 10.3.3.3 255.255.255.255  
end  
  
!  
router ospf 1  
router-id 0.0.0.3  
!
```

R4

```
!  
interface Ethernet0/0  
ip address 192.168.34.4 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
end  
  
!  
interface Ethernet0/1  
ip address 192.168.45.4 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
end  
  
!  
interface Ethernet0/2  
ip address 192.168.24.4 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
end  
  
!  
interface Loopback0  
ip address 10.4.4.4 255.255.255.255  
end  
  
!  
router ospf 1  
router-id 0.0.0.4  
!
```

R5

```
!  
interface Ethernet0/1  
ip address 192.168.45.5 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
end  
  
!  
interface Loopback0  
ip address 10.5.5.5 255.255.255.255  
end
```

```
!  
router ospf 1  
router-id 0.0.0.5  
!
```

Default Behavior

With the previous configurations in place, this section describes the default router behaviors.

Here is trace from R5 to 10.1.1.1. Notice that the traffic passes through R2, not R3:

```
R5#traceroute 10.1.1.1  
Type escape sequence to abort.  
Tracing the route to 10.1.1.1  
VRF info: (vrf in name/id, vrf out name/id)  
1 192.168.45.4 6 msec 6 msec 6 msec <<< R4  
2 192.168.24.2 6 msec 6 msec 8 msec <<< R2  
3 192.168.12.1 8 msec * 3 msec <<< R1
```

In this network, the router R4 must make the decision and should route the traffic to R3, not to R2 directly, as per the system requirement.

Here is an example of the routing table on R4:

```
R4#show ip route 10.1.1.1  
Routing entry for 10.1.1.1/32  
Known via "ospf 1", distance 110, metric 30, type extern 1  
Last update from 192.168.24.2 on Ethernet0/2, 00:14:33 ago  
Routing Descriptor Blocks:  
* 192.168.24.2, from 0.0.0.2, 00:14:33 ago, via Ethernet0/2 <<< Towards R2  
Route metric is 30, traffic share count is 1
```

A metric of **30** is associated with this route for prefix **10.1.1.1/32**. This is due to a default metric of 20 that is used by the Autonomous System Boundary Router (ASBR) (R2) and a cost of 10 on the interface Eth0/2 on R4.

The path from R4 to the 10.1.1.1/32 prefix via R3 is longer. Here, the cost for Interface Ethernet 0/2 on R4 (the path towards R2) is altered in order to verify whether it changes the behavior:

```
!  
interface Ethernet0/2  
ip address 192.168.24.4 255.255.255.0  
ip ospf network point-to-point  
ip ospf 1 area 99  
ip ospf cost 100  
end
```

Here is the trace from R5 and the **show ip route** command output from R4:

```
R5#traceroute 10.1.1.1  
Type escape sequence to abort.  
Tracing the route to 10.1.1.1  
VRF info: (vrf in name/id, vrf out name/id)  
1 192.168.45.4 4 msec 9 msec 8 msec <<< R4  
2 192.168.24.2 8 msec 9 msec 10 msec <<< R2  
3 192.168.12.1 10 msec * 5 msec <<< R1  
  
R4#show ip route 10.1.1.1  
Routing entry for 10.1.1.1/32  
Known via "ospf 1", distance 110, metric 120, type extern 1
```

```
Last update from 192.168.24.2 on Ethernet0/2, 00:01:50 ago
Routing Descriptor Blocks:
* 192.168.24.2, from 0.0.0.2, 00:01:50 ago, via Ethernet0/2
  Route metric is 120, traffic share count is 1
```

As the trace shows, the traffic from R5 takes the same path, and traffic does not flow via R3. Also, as shown in the output of the **show ip route 10.1.1.1** command on R4, the cost of 100 that was added on R4 (Interface Ethernet 0/2) takes effect and the cost of the route to the prefix is **120** (opposed to 30). However, the path still did not change and the requirement for traffic to flow across R3 is not yet met.

In order to determine the cause of this behavior, here is the R4 **show ip ospf border-routers** command output (the cost on the R4 Interface Ethernet 0/2 is still set to 100):

```
R4#show ip ospf border-routers
      OSPF Router with ID (0.0.0.4) (Process ID 1)
        Base Topology (MTID 0)
      Internal Router Routing Table
Codes: i - Intra-area route, I - Inter-area route

i 0.0.0.2 [100] via 192.168.24.2, Ethernet0/2, ABR/ASBR, Area 99, SPF 3
i 0.0.0.3 [10] via 192.168.34.3, Ethernet0/0, ABR, Area 99, SPF 3
```

On R4, you can see that there are two Area Border Routers (ABRs) (0.0.0.2, which is R2, and 0.0.0.3, which is R3) and that R2 is the ASBR. This output also shows the Intra-Area (*i*) information for the ASBR.

Now, the Interface Ethernet 0/2 is shut down on R4 in order to determine whether the traffic flows via R3 and in order to see how the **show ip ospf border-routers** command output appears:

```
interface Ethernet0/2
ip address 192.168.24.4 255.255.255.0
ip ospf network point-to-point
ip ospf 1 area 99
ip ospf cost 100
shutdown
end
```

Here is the trace from R5 and the **show ip route** command output from R4:

```
R5#traceroute 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.45.4 7 msec 7 msec 8 msec <<< R4
 2 192.168.34.3 9 msec 8 msec 8 msec <<< R3
 3 192.168.23.2 9 msec 9 msec 7 msec <<< R2
 4 192.168.12.1 8 msec * 4 msec <<< R1

R4#show ip route 10.1.1.1
Routing entry for 10.1.1.1/32
Known via "ospf 1", distance 110, metric 40, type extern 1 <<< Metric 40
Last update from 192.168.34.3 on Ethernet0/0, 00:01:46 ago <<< Traffic to R2
Routing Descriptor Blocks:
* 192.168.34.3, from 0.0.0.2, 00:01:46 ago, via Ethernet0/0
  Route metric is 40, traffic share count is 1
```

As shown, when Interface Ethernet 0/2 is shut down on R4, traffic passes through R3. Also, the cost that is associated with the route towards R3 is only **40**, while the cost towards 10.1.1.1/32 via R2 was 120. The OSPF protocol still prefers to route the traffic via R2 instead of R3, even though the cost to reach 10.1.1.1/32 is lower via R3.

Here is the output of *show ip ospf border-routers* once again on R4:

```
R4#show ip ospf border-routers
      OSPF Router with ID (0.0.0.4) (Process ID 1)
        Base Topology (MTID 0)
Internal Router Routing Table
Codes: i - Intra-area route, I - Inter-area route

I 0.0.0.2 [20] via 192.168.34.3, Ethernet0/0, ASBR, Area 99, SPF 4
i 0.0.0.3 [10] via 192.168.34.3, Ethernet0/0, ABR, Area 99, SPF 4
```

The information that is required in order to reach the ASBR is the *Inter–Area* information. However, the *Intra–Area* information that details how to reach ASBR is preferred over *Inter–Area* information irrespective of the OSPF cost that is associated with the two paths.

For this reason, the path via R3 was not preferred, even though the cost via R3 is lower.

Here, the Interface Ethernet 0/2 is brought back up on R4:

```
interface Ethernet0/2
no shutdown
end
!
```

The trace from R5 indicates that routing actions return to those previously observed (traffic does not flow via R3):

```
R5#traceroute 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.45.4 6 msec 7 msec 7 msec   <<< R4
 2 192.168.24.2 7 msec 8 msec 7 msec   <<< R2
 3 192.168.12.1 8 msec * 12 msec      <<< R1
```

There are various ways that you can resolve this issue (this list is not exhaustive):

- Change the Area between R2 and R3 to **90**, and then modify the cost.
- Add another link between R2 and R3 and configure it to be in **Area 99**.
- Use Multi–Area ADJ.

Refer to the next section in order to see the way in which Multi–Area OSPF ADJ works and how it can solve this issue at hand.

Multi–Area Adjacency Configuration

As previously mentioned, Multi–Area ADJ can be used in order to form multiple point–to–point logical adjacencies over a single link. The requirement is that there must be only two OSPF speakers on the link, and in a broadcast network, you must manually change the OSPF network type to **Point–to–Point** on the link.

This feature allows a single physical link to be shared by multiple Areas and creates an Intra–Area path in each of the Areas that share the link.

In order to meet this requirement, you must configure OSPF Multi–Area ADJ between R2 and R3 over link Ethernet 0/1, which is currently only in **Area 0**.

Here is the configuration for R2:

```
!  
interface Ethernet0/1  
ip address 192.168.23.2 255.255.255.0  
ip ospf network point-to-point  
  ip ospf multi-area 99  
ip ospf 1 area 0  
end
```

Here is the configuration for R3:

```
!  
interface Ethernet0/1  
ip address 192.168.23.3 255.255.255.0  
ip ospf network point-to-point  
  ip ospf multi-area 99  
ip ospf 1 area 0  
end
```

The OSPF ADJ comes up over the Virtual Link:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on OSPF_MA0 from LOADING to FULL, Loading Done
```

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on OSPF_MA0 from LOADING to FULL, Loading Done
```

Here is the newly formed ADJ:

```
R2#show ip ospf neighbor 0.0.0.3
```

```
<Snip>
```

```
Neighbor 0.0.0.3, interface address 192.168.23.3  
  In the area 99 via interface OSPF_MA0  
  Neighbor priority is 0, State is FULL, 6 state changes  
  DR is 0.0.0.0 BDR is 0.0.0.0  
  Options is 0x12 in Hello (E-bit, L-bit)  
  Options is 0x52 in DBD (E-bit, L-bit, O-bit)  
  LLS Options is 0x1 (LR)  
  Dead timer due in 00:00:39  
  Neighbor is up for 00:03:01  
  Index 2/3, retransmission queue length 0, number of retransmission 0  
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)  
  Last retransmission scan length is 0, maximum is 0  
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

```
R3#show ip ospf neighbor 0.0.0.2
```

```
<Snip>
```

```
Neighbor 0.0.0.2, interface address 192.168.23.2  
  In the area 99 via interface OSPF_MA0  
  Neighbor priority is 0, State is FULL, 6 state changes  
  DR is 0.0.0.0 BDR is 0.0.0.0  
  Options is 0x12 in Hello (E-bit, L-bit)  
  Options is 0x52 in DBD (E-bit, L-bit, O-bit)  
  LLS Options is 0x1 (LR)  
  Dead timer due in 00:00:39  
  Neighbor is up for 00:01:41  
  Index 2/3, retransmission queue length 0, number of retransmission 0  
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)  
  Last retransmission scan length is 0, maximum is 0  
  Last retransmission scan time is 0 msec, maximum is 0 msec
```


Verify

In order to verify whether your configuration works properly, enter the *show ip ospf border-routers* command on R4:

```
R4#show ip ospf border-routers
      OSPF Router with ID (0.0.0.4) (Process ID 1)
      Base Topology (MTID 0)
Internal Router Routing Table
Codes: i - Intra-area route, I - Inter-area route

i 0.0.0.3 [10] via 192.168.34.3, Ethernet0/0, ABR, Area 99, SPF 10
i 0.0.0.2 [20] via 192.168.34.3, Ethernet0/0, ABR/ASBR, Area 99, SPF 10
```

As shown, the Intra–Area information that is used in order to route traffic to R2 (0.0.0.2)/ASBR is via R3. This should resolve the previously mentioned problem.

Here is the trace from R5:

```
R5#traceroute 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.45.4 8 msec 9 msec 8 msec   <<< R4
 2 192.168.34.3 8 msec 8 msec 8 msec   <<< R3
 3 192.168.23.2 7 msec 8 msec 8 msec   <<< R2
 4 192.168.12.1 8 msec * 4 msec       <<< R1
```

As shown, the traffic from R5 that is destined to 10.1.1.1 properly flows via R3, and the system requirement is met.

Enter the *show ip ospf neighbor* command on R2, R3, and R4 in order to verify whether the ADJs are established:

```
R2#show ip ospf neighbor
Neighbor ID Pri  State    Dead Time  Address      Interface
0.0.0.3      0  FULL/-  00:00:39  192.168.23.3  Ethernet0/1
0.0.0.4      0  FULL/-  00:00:37  192.168.24.4  Ethernet0/2
0.0.0.3      0  FULL/-  00:00:33  192.168.23.3  OSPF_MAO
```

```
R3#show ip ospf neighbor
Neighbor ID Pri  State    Dead Time  Address      Interface
0.0.0.2      0  FULL/-  00:00:34  192.168.23.2  Ethernet0/1
0.0.0.2      0  FULL/-  00:00:35  192.168.23.2  OSPF_MAO
0.0.0.4      0  FULL/-  00:00:39  192.168.34.4  Ethernet0/0
```

```
R4#show ip ospf neighbor
Neighbor ID Pri  State    Dead Time  Address      Interface
0.0.0.2      0  FULL/-  00:00:32  192.168.24.2  Ethernet0/2
0.0.0.5      0  FULL/-  00:00:32  192.168.45.5  Ethernet0/1
0.0.0.3      0  FULL/-  00:00:35  192.168.34.3  Ethernet0/0
```

Note: In these outputs, the *Ethernet0/1* Interface entries indicate the ADJ over Area 0, and the *OSPF_MAO* Interface entries indicate the Multi–Area ADJ over Area 99.

The R4 Interface Ethernet 0/2 still has a cost of 100, and the path via R3 is preferred on R4. If this cost is removed, then R4 routes traffic directly to R2 as before.

Here is the configuration and the *show ip route* command output on R4 with the IP OSPF cost of 100 still configured on R4 Interface Ethernet 0/2:

```
!  
interface Ethernet0/2  
 ip address 192.168.24.4 255.255.255.0  
 ip ospf network point-to-point  
 ip ospf 1 area 99  
 ip ospf cost 100  
!
```

```
R4#show ip route 10.1.1.1  
Routing entry for 10.1.1.1/32  
Known via "ospf 1", distance 110, metric 40, type extern 1  
Last update from 192.168.34.3 on Ethernet0/0, 00:28:45 ago  
Routing Descriptor Blocks:  
* 192.168.34.3, from 0.0.0.2, 00:28:45 ago, via Ethernet0/0  
  Route metric is 40, traffic share count is 1
```

Here is the configuration and the *show ip route* command output on R4 when you remove the cost:

```
interface Ethernet0/2  
 ip address 192.168.24.4 255.255.255.0  
 ip ospf network point-to-point  
 ip ospf 1 area 99  
end
```

```
R4#show ip route 10.1.1.1  
Routing entry for 10.1.1.1/32  
Known via "ospf 1", distance 110, metric 30, type extern 1  
Last update from 192.168.24.2 on Ethernet0/2, 00:00:13 ago  
Routing Descriptor Blocks:  
* 192.168.24.2, from 0.0.0.2, 00:00:13 ago, via Ethernet0/2 <<< Route changed back to R2  
  Route metric is 30, traffic share count is 1
```

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.