Route Filtering With Open Shortest Path First (OSPF) in IPv4

At first glance it may seem more of an art than a science to filter routes within OSPF. Novice network engineers learn the hard way that an outbound distribute-list does not suppress routes advertised to OSPF neighbors. Do routes have to be removed from the link state database (LSD) to be dropped from the routing table? What is the impact of filtering a route in OSPF? This document will demonstrate how to filter routes anywhere within the OSPF routing domain. It will also describe the impact of different filtering methods on routing tables and link state databases (LSDs).

The attached diagram illustrates the most commonly used OSPF areas and types of routes. In the diagram, Router-1 and Router-2 are area 0 (backbone) routers. Router-4 and Router-5 are Area Border Routers (ABRs). Router-3 and Router-9 are Autonomous System Boundary Routers (ASBRs). All routers in the diagram are connected to Ethernet segments, and all except Router-3, Router-4 and Router-5 have three loopback interfaces*.

OSPF has built-in controls over route propagation. OSPF routes are permitted or denied into different OSPF areas based on area type, such as backbone area, normal (non-stub) area, stub area, Not-So-Stubby Area (NSSA), and totally stubby area. OSPF is a link state protocol that populates the LSD to give routers within the same area an identical perspective of the OSPF routing domain. That perspective is tempered by the type of area the routers are in.

OSPF ABRs limit the advertisement of different types of routes into different OSPF areas based on the type of Link State Advertisement (LSA) associated with each network. An Area Border Router (ABR) bordering any type of OSPF stub area prevents the advertisement of external routes into the stub area. Exceptions are those generated within any type of NSSA by an Autonomous System Boundary Router (ASBR) within the same area as the ABR. The ABR in a stub or totally-stubby area advertises a default route as an inter-area route (type 3 Network Summary LSA). An ABR to a totally stubby area prevents advertisement of any other inter-area routes as well as external routes into that area. The same ABR can border multiple types of areas. An ABR could also be an ASBR, a combination not represented in this scenario.

External routes could include redistributed static or connected routes, or routes redistributed from another routing protocol or OSPF process. The metric type associated with an OSPF external route can be the default of External Type 2 (E2) or External Type 1 (E1). When an external route is defined as an E1, ABRs generate a type 4 (ASBR Summary) LSA into non-backbone, non-stub areas. An ABR bordering the NSSA will only generate E1 routes with type 5 LSAs as they are converted from type 7 LSAs. The type 4 LSA reflects the cost from that area's ABR to the ASBR (itself) that redistributed that route into the OSPF domain. This cost is added to the area router's cost to reach its ABR plus the metric of the external LSA.

Inter-area routes are those originating in another OSPF area within the same routing domain as the local router, with the exception of the default route generated by the ABR into stub and totally stubby areas. This default route is not propagated outside of a stub or totally stubby area. In the descriptions that follow, this default route is the exception to the rules of how inter-area routes are handled by ABRs.

OSPF routers create type 1 (Router) and type 2 (Network) LSAs for networks within an area. These detailed LSAs cannot leave their respective areas. The flooding of LSAs within an area can be prevented with the `ip ospf database-filter all out` command, which can be applied to an interface within broadcast, non-broadcast, and point-to-point networks, or to a specific neighbor over point-to-multipoint networks.

A recent enhancement to IOS v12.4(15)T, called prefix suppression, provides a mechanism to exclude all IP prefixes on the local router from the LSD. When prefix suppression is applied globally to the OSPF process it does not filter prefixes that are associated with loopbacks, secondary IP addresses, and passive interfaces. There is also a command that enables prefix suppression on an interface on which it is otherwise not enabled by the global command.

Type 3 (Network Summary) LSAs are created by the ABRs to represent the networks to areas outside the originating area within the same OSPF routing domain. Type 3 LSAs can be filtered on ABRs to prevent the advertisement of routes between the areas of an OSPF routing domain.

Several methods filter routes on the local router, whether the router is in the same or a different area than the originator of the routes. Most filtering does not remove the networks from the LSD. The routes are removed from the routing table, which prevents the local router from using them to forward traffic. The filters have no impact on the presence of routes in the routing table of any other router in the OSPF routing domain.

Cost and administrative distance can be adjusted to remove routes from the local router's routing table. The final example in this document demonstrates the use of an infinite administrative distance.

In the absence of filtering, all intra-area, inter-area, and external routes are advertised into area 0, the backbone area within a multi-area OSPF routing domain.

Area 1 is a normal (non-stub) OSPF area. In the absence of filtering, all intra-area, inter-area, and external routes are advertised within this area. This is the only area type that could have type 4 (ASBR Summary) LSAs in the LSD.

Area 2 is a Not-So-Stubby Area, an NSSA. There is at least one ASBR in this area, Router-9, redistributing an Enhanced Interior Gateway Protocol (EIGRP) route into OSPF. The advertisement of an external network specified by a type 5 (AS External) LSA is not permitted into a stub or totally stubby area, or NSSA. It is allowed using a type 7 (NSSA External) LSA, generated by the ASBR. Only the ASBR in the NSSA can generate type 7 LSAs for the external routes that it redistributes into OSPF. Within the NSSA, when the type 7 LSA reaches the ABR, the LSA is changed to a type 5 and propagated into the backbone. The route now appears as an ordinary external route to the routers in non-stub areas outside of the NSSA. The Routing Information Protocol (RIP) route redistributed into OSPF by Router-3 is not in the LSD or routing tables of the routers in area 2.

Router-4 is an ABR to areas 0, 1, and 2. The advertisement of the external route, `192.168.2.0/24`, into Area 2 by the ASBR is an N2 route in the routing tables of all of the other NSSA routers, including the ABR**.

```
router-4#sh ip route | include 192.168
O E1 192.168.50.0/24 [110/22] via 10.20.20.18, 00:10:07, FastEthernet2/0.3
O N2 192.168.2.0/24 [110/20] via 10.2.2.3, 00:10:07, FastEthernet2/0.2
```

In contrast to what occurs within stubby or totally stubby areas, a default route is not automatically generated into an NSSA. A special statement on the NSSA ASBR, `area` ***area_num*** `default-information originate`, will advertise a default route (if there is one in the ASBR's routing table) into the NSSA with a type 7 LSA. This default route is propagated into non-stub areas in the rest of the OSPF routing domain, by the same rules that apply to any other external route.

Area 3 is a stub area. Native intra-area and inter-area routes are advertised into this area, but not external routes. In place of external routes, the ABR, Router-5, automatically advertises a default route into the stub area as an inter-area route (IA).

```
router-10#sh ip route
O IA    172.16.4.16 [110/12] via 10.3.3.5, 00:54:12, Ethernet0/0
O IA    172.16.1.16 [110/16] via 10.3.3.5, 00:54:12, Ethernet0/0
O       172.16.3.0 [110/12] via 10.3.3.5, 00:54:13, Ethernet0/0
O*IA 0.0.0.0/0 [110/11] via 10.3.3.5, 00:54:15, Ethernet0/0

router-11#sh ip route
O IA    10.30.30.8/30 [110/3] via 10.3.3.1, 00:54:50, FastEthernet0/1
O IA    10.20.20.8/30 [110/5] via 10.3.3.1, 00:54:50, FastEthernet0/1
O IA    10.4.4.8/29 [110/3] via 10.3.3.1, 00:54:50, FastEthernet0/1
O       10.3.3.8/29 [110/3] via 10.3.3.1, 00:54:50, FastEthernet0/1
O*IA 0.0.0.0/0 [110/2] via 10.3.3.1, 00:54:52, FastEthernet0/1
```

Area 4 is a totally stubby area. Only intra-area routes (OSPF routes originating within area 4) and a default route are advertised within this area. This default route is advertised into the totally stubby area in the same manner as it was advertised into the stub area. The ABR, Router-5, automatically advertises a default route as an inter-area route (IA).

```
router-12#sh ip route
O       172.16.4.16 [110/12] via 10.4.4.5, 00:01:00, Ethernet0
O       172.16.4.8 [110/12] via 10.4.4.5, 00:01:00, Ethernet0
O       172.16.4.0 [110/12] via 10.4.4.5, 00:01:00, Ethernet0
O*IA 0.0.0.0/0 [110/11] via 10.4.4.5, 00:01:01, Ethernet0

router-13#sh ip route
O       10.4.4.8/29 [110/3] via 10.4.4.1, 00:00:03, FastEthernet0/1
O       10.4.4.4/30 [110/2] via 10.4.4.1, 00:00:03, FastEthernet0/1
O*IA 0.0.0.0/0 [110/2] via 10.4.4.1, 00:00:04, FastEthernet0/1
```

Here are two methods to filter the type 3 LSAs on the ABR, which removes them from the OSPF LSD.

The first method is to apply a filter-list to networks from one area into or out of another area. The use of the command: `area x filter-list prefix` *name* `[in | out]` will filter type 3 LSAs on an ABR. This command can be applied to inbound or outbound routes.

In this example, Router-5 is the ABR supporting areas 0, 3 and 4. To filter the loopback networks of Router-10 so they are not advertised into area 0, add the following statements on Router-5:

```
router-5(config)#ip prefix-list FROM_router-10 seq 5 deny 10.3.3.0/27 le 30
router-5(config)#ip prefix-list FROM_router-10 seq 10 permit 0.0.0.0/0 le 32
router-5(config)#router ospf 1
router-5(config-router)#area 0 filter-list prefix FROM_router-10 in
```

Using this method, filtering routes into area 0 eliminates their advertisement within all OSPF areas in the routing domain except the originating area. Filtering routes into a non-backbone area eliminates their advertisement only within the target area.

The other method for filtering type 3 LSAs is to use an `area range` statement with the "`not-advertise`" parameter. This statement is comparable to the `filter-list prefix` statement with the `out` parameter. The `area range` statement is more commonly used to advertise a range of networks to other areas within an OSPF routing domain. The use of the `not-advertise` parameter suppresses the advertisement of a range of networks to the other areas***.

In this example, an area range statement is configured on Router-4, an ABR supporting areas 0, 1 and 2. Router-7's loopback networks will not be advertised outside of area 1.

```
router-4(config)#router ospf 1
router-4(config-router)#area 1 range 172.16.1.0 255.255.255.224 not-advertise
```

Other filtering methods do not remove routes from the LSD of area routers. Routes are only removed from the routing table of the local router. Other routers in the same area that do not have filters applied will continue to advertise the routes. A possible result is a black hole in the routing domain. OSPF neighbors could forward traffic to a router that is filtering the route to which it has the lowest cost path.

The most basic filter is the inbound distribute-list. The outbound distribute-list is exclusively used to filter OSPF routes in route redistribution between OSPF and a different routing protocol or OSPF process. The `distribute-list in` can be applied to any router within the area. A standard access-list or a prefix-list may be used. As demonstrated below, a prefix-list is a concise method to specify a range of contiguous networks in a single statement. Conversely, an access-list concisely specifies a group of non-contiguous networks, such as all the 10.x.x.x subnets with an odd number in the third octet.

```
access-list 1 deny 10.0.1.0 0.255.254.255
access-list 1 permit any
```

This example demonstrates how to filter routes between routers within the same area, in this case the totally stubby area, area 4. Router-13's loopback networks are filtered from the routing table of Router-12.

```
router-12#ping 172.16.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/8 ms
router-12(config)#ip prefix-list AREA4-172-ROUTES seq 5 deny 172.16.4.0/27 le 30
router-12(config)#ip prefix-list AREA4-172-ROUTES seq 10 permit 0.0.0.0/0 le 32
router-12(config)#router ospf 1
router-12(config-router)#distribute-list prefix AREA4-172-ROUTES in

router-12#ping 172.16.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/8 ms
```

If the filter worked, why did the ping work? Router-12 used the default route advertised by Router-5, and Router-5 still has the `172.16.4.0/27` networks in its routing table. Router-12 does not have the routes in its routing table, although they are in the LSD.

In area 1 a distribute-list is used to filter the `10.1.1.0/27` networks advertised to Router-7. There is no change to the LSD on Router-7 after the distribute-list is applied.

```
router-7(config)#ip prefix-list AREA1-10-ROUTES seq 5 deny 10.1.1.0/27 le 30
router-7(config)#ip prefix-list AREA1-10-ROUTES seq 10 permit 0.0.0.0/0 le 32
router-7(config)#router ospf 1
router-7(config-router)#distribute-list prefix AREA1-10-ROUTES in
```

Since Router-7 is not in a stub area, there is no default route to bail it out once the distribute-list is applied. As soon as the distribute-list is removed, Router-7 regains access to the `10.1.1.0/27` networks since they are back in Router-7's routing table.

```
router-7#ping 10.1.1.25

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.25, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
router-7(config)#router ospf 1
router-7(config-router)#no distribute-list prefix AREA1-10-ROUTES in

router-7#ping 10.1.1.25

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.25, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

There is vastly improved flexibility when combining a distribute-list with a route-map. There are more options to match, such as metric, route-type, and next hop.

What follows in the next few paragraphs is an example using three different options within one route-map to filter routes that would otherwise be visible to backbone Router-1.

The most common method of selecting routes is by subnet. Referring to the example below, the `172.16.3.8/29` network advertised by Router-11 is filtered. A prefix-list could be used in place of a standard access-list.

Another method of selecting routes uses the source of the route. In this example, all the networks advertised by Router-8 are filtered on Router-1. When using OSPF, the route source is the router ID of the advertising router, regardless of whether the router ID is advertised or on the same data-link segment. Refer to the output of `show ip route 192.168.50.0`, which indicates that the route originated from `10.15.3.254`, Router-3. The router ID of Router-8 is `10.15.8.254`, which is the address used in the standard access-list `FILTER_ROUTE_SOURCE`.

When redistributing routes from another routing process, the routes to be filtered can be tagged. In the illustration below the external route was tagged on Router-3. The tag is evident in the route description on Router-1. Route tags can be combined with other criteria to create complex conditions****.

```
router-1#sh ip route 192.168.50.0
Routing entry for 192.168.50.0/24
  Known via "ospf 1", distance 110, metric 22
  Tag 3, type extern 1
  Last update from 10.20.20.18 on FastEthernet2/0.3, 00:00:24 ago
  Routing Descriptor Blocks:
  * 10.20.20.18, from 10.15.3.254, 00:00:24 ago, via FastEthernet2/0.3
      Route metric is 22, traffic share count is 1
      Route tag 3
```

These statements are used to configure and apply the route-map to the OSPF process on Router-1.

```
ip access-list standard FILTER_SUBNET
 permit 172.16.3.8 0.0.0.7
!
ip access-list standard FILTER_ROUTE_SOURCE
 permit 10.15.8.254
!
route-map ROUTE_MAP_FILTERS deny 10
 match ip address FILTER_SUBNET
!
route-map ROUTE_MAP_FILTERS deny 20
 match ip route-source FILTER_ROUTE_SOURCE
!
route-map ROUTE_MAP_FILTERS deny 30
 match tag 3
!
route-map ROUTE_MAP_FILTERS permit 40
!
router ospf 1
distribute-list route-map ROUTE_MAP_FILTERS in
```

Another method of filtering OSPF inbound routes is to alter the administrative distance of the routes. Administrative distance is an arbitrary value reflecting the desirability of the route relative to how it is learned by the local router. The route could be learned via multiple routing protocols, connected networks, or static routes. If a route has multiple sources, the one with the lowest administrative distance is inserted into the routing table. An OSPF route is considered unreachable with an administrative distance of 255, the equivalent of infinity, and will not be inserted into the routing table.

As was the case with the route-map route-source example, the OSPF `distance` command uses the router ID as the route source. Router IDs that are not in the same address range may require a separate `distance` statement for each router that could advertise the route on a multi-access network. The route source could change after additional interfaces are configured and/or a router is rebooted. Explicit configuration of OSPF router IDs on all OSPF routers will ease the management of route filtering that relies on the router ID.

In the following example, the `10.30.30.0/27` networks advertised to Router-3 are filtered by setting an administrative distance of 255 to those networks. Since Router-2 and Router-5 are supplying routing information for this address range and the `10.30.30.20/30` subnet, there is a separate `distance` statement for each router. As illustrated immediately below, a single `distance` statement could have been used in this example. In the first `distance` statement the router IDs are in a range that is summarized. The second `distance` statement includes all possible sources.

```
distance 255 10.15.0.254 0.0.7.0 BB-ROUTES
```
or
```
distance 255 0.0.0.0 255.255.255.255 BB-ROUTES
```

The example shows the difference when applying the distance statement for the `10.30.30.20/30` route, currently advertised by Router-5.

```
router-3(config)#ip access-list standard BB-ROUTES
router-3(config-std-nacl#permit 10.30.30.0 0.0.0.31
router-3(config)#router ospf 1
router-3(config-router)#distance 255 10.15.2.254 0.0.0.0 BB-ROUTES
router-3(config-router)#distance 255 10.15.5.254 0.0.0.0 BB-ROUTES

router-3# ping 10.30.30.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.30.30.21, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

router-3(config)#router ospf 1
router-3(config-router)#no distance 255 10.15.5.254 0.0.0.0 BB-ROUTES

router-3#ping 10.30.30.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.30.30.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Here is an example of OSPF LSD output before and after the prefix-suppression command is applied globally on Router-6, condensed for brevity. This is the output of `show ip ospf database router` prior to the application of the global `prefix-suppression` command under the OSPF process.

```
Number of Links: 4

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 10.1.1.8
     (Link Data) Network Mask: 255.255.255.248

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 10.1.1.16
     (Link Data) Network Mask: 255.255.255.248

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 10.1.1.25
     (Link Data) Network Mask: 255.255.255.255

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.1.1.2
     (Link Data) Router Interface address: 10.1.1.2
```

Here is the output afterward:

```
Number of Links: 2

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 10.1.1.25
     (Link Data) Network Mask: 255.255.255.255

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.1.1.2
     (Link Data) Router Interface address: 10.1.1.2
```

We expect to see 10.1.1.25 in the OSPF database but where are the other loopbacks? A loopback interface in OSPF is a host address (32 bit mask) by default. The other loopback interfaces are configured with `ip ospf network point-to-point`. The prefix-suppression feature only excludes loopbacks with 32 bit masks. The remaining loopback disappears after prefix-suppression is applied directly to the interface.

```
router-6(config)#int l2
router-6(config-if)#ip ospf prefix-suppression

  Number of Links: 1

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.1.1.1
     (Link Data) Router Interface address: 10.1.1.2
```

In summary, filtering routes within an OSPF domain takes many forms. OSPF has built-in conventions for filtering different types of routes into different types of areas, such as stub, totally-stubby, and NSSA.

Several filtering methods remove LSAs from the LSD. These include prefix suppression and type 3 LSA filtering on an ABR with either the `area x filter-list` or the `area range not-advertise` statements.

Methods that only impact the routing table of the local router and not the LSD include cost/administrative distance adjustments and inbound distribute-lists with or without route-maps. These methods have no impact on the advertisement of routes to other routers in the OSPF domain. The routes are only removed from the local routing table, which prevents the use of the routes to forward traffic from the local router.

*Loopback subnets on routers are listed below. Router IDs on all routers are in the format 10.15.*router_number*.254.

Router-1	Router-2
10.20.20.4/30	10.30.30.4/30
10.20.20.8/30	10.30.30.8/30
10.20.20.12/30	10.30.30.12/30
Router-6	Router-7
10.1.1.8/29	172.16.1.0/29
10.1.1.16/29	172.16.1.8/29
10.1.1.24/29	172.16.1.16/29
Router-8	Router-9
172.16.2.0/29	10.3.3.8/29
172.16.2.8/29	10.3.3.16/29
172.16.2.16/29	10.3.3.24/29
Router-10	Router-11
172.16.2.0/29	10.3.3.8/29
172.16.2.8/29	10.3.3.16/29
172.16.2.16/29	10.3.3.24/29
Router-12	Router-13
10.4.4.8/29	172.16.4.0/29
10.4.4.16/29	172.16.4.8/29
10.4.4.24/29	172.16.4.16/29

**Routing table output is condensed for brevity.

*** A similar command is used on an ASBR to filter routes into an NSSA. The "`summary-address`" command with the "`not-advertise`" parameter will prevent a range of addresses from being advertised into an OSPF routing domain from another routing protocol or a different OSPF routing domain. This feature does not filter type 3 LSAs, but it is another method of filtering routes into OSPF.

****BGP Autonomous System (AS) paths can be preserved as they pass through an OSPF routing domain through a special use of route tags.

Route Filtering With Open Shortest Path First (OSPF)