

Understanding Redistribution of OSPF Routes into BGP

Document ID: 5242

Author: Syed Faraz Shamim

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Network Setup

Redistribution of Only OSPF Internal (Intra- and Inter-Area) Routes into BGP

Redistribution of Only OSPF External (Type 1 and 2) Routes into BGP

Redistribution of Only OSPF External Type 1 or Type 2 Routes into BGP

Redistribution of OSPF Internal and External Routes into BGP

Redistribution of OSPF NSSA-External Routes into BGP

Modify the Redistribution Option in OSPF

Unable to Redistribute iBGP Learnt Routes into an IGP such as EIGRP, OSPF

Related Information

Introduction

This Tech Note explains the behavior of Open Shortest Path First (OSPF) to Border Gateway Protocol (BGP) redistribution on Cisco routers. The behavior of OSPF to BGP redistribution is outlined in RFC 1403 .

There are several types of OSPF routes:

- Intra-Area In a multiarea OSPF network, routes, originated within an area, are known by the routers in the same area as Intra-Area routes. These routes are flagged as O in the **show ip route** command output..
- Inter-Area When a route crosses an OSPF Area Border Router (ABR), the route is known as an OSPF Inter-Area route. These routes are flagged as O IA in the **show ip route** command output.

Both Intra and Inter-Area routes are also called OSPF Internal routes, as they are generated by OSPF itself, when an interface is covered with the OSPF **network** command.

- External Type-2 or External Type-1 Routes which were redistributed into OSPF, such as Connected, Static, or other Routing Protocol, are known as External Type-2 or External Type-1. These routes are flagged as O E2 or O E1 in the **show ip route** command output.
- NSSA external type 2 or NSSA external type 1 When an area is configured as a Not-So-Stub Area (NSSA), and routes are redistributed into OSPF, the routes are known as NSSA external type 2 or NSSA external type 1. These routes are flagged as O N2 or O N1 in the **show ip route** command output.

Describing the differences between External and NSSA type 2 or 1 is beyond the scope of this document refer to the OSPF Design Guide for more information.

The default behavior is not to redistribute any routes from OSPF into BGP. Redistribution must be configured. You can use the **route-map** command to filter routes during OSPF to BGP redistribution. When redistributing, specific keywords like **internal**, **external** and **nssa-external** are required to redistribute

respective routes.

Prerequisites

Requirements

An understanding of OSPF route types is required before using this document.

Components Used

This document is not restricted to specific software and hardware versions.


The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Network Setup

There are four cases of redistributing OSPF routes into BGP discussed below. The network diagram applies to the first three cases. The diagram and setup for the fourth case can be found in the Redistribution of OSPF NSSA–External Routes into BGP section.



Redistribution of Only OSPF Internal (Intra– and Inter–Area) Routes into BGP

If you configure the redistribution of OSPF into BGP without keywords, by only OSPF intra–area and

inter-area routes are redistributed into BGP, by default. You can use the **internal** keyword along with the **redistribute** command under **router bgp** to redistribute OSPF intra- and inter-area routes.

This configuration is a new configuration of Router B that redistributes only the intra-area route (131.108.2.0/24) and inter-area route (131.108.1.0/24) into BGP and only OSPF Internal (intra-area and inter-area) routes are redistributed into BGP:

```

RTB
-----
hostname RTB
!
interface Ethernet0/0
 ip address 2.2.2.3 255.255.255.0
!
interface Serial1/0
 ip address 3.3.3.1 255.0.0.0
!
router ospf 1
 network 2.0.0.0 0.255.255.255 area 1
!
router bgp 100
 redistribute ospf 1

!-- This redistributes only OSPF intra- and inter-area routes into BGP.

neighbor 3.3.3.2 remote-as 200
!
end

RTB# show ip route
Codes: 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, 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

Gateway of last resort is not set

    2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Ethernet0/0
C       3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
    131.108.0.0/24 is subnetted, 2 subnets
O       131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet0/0
O IA   131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0
RTB#
```

Router B redistributes only OSPF Internal routes:

```

RTB# show ip bgp
BGP table version is 10, local router ID is 192.168.1.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 2.2.2.0/24       0.0.0.0           0             32768 ?
*> 131.108.1.0/24   2.2.2.2           84            32768 ?
*> 131.108.2.0/24   2.2.2.2           74            32768 ?
```

RTB#

Router C learning these routes from BGP:

```
RTC# show ip route
```

```
Codes: 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, 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
```

```
Gateway of last resort is not set
```

```
      2.0.0.0/24 is subnetted, 1 subnets
B       2.2.2.0 [20/0] via 3.3.3.1, 00:11:19
C       3.0.0.0/8 is directly connected, Serial0/0
      131.108.0.0/24 is subnetted, 2 subnets
B       131.108.2.0 [20/74] via 3.3.3.1, 00:03:56
B       131.108.1.0 [20/84] via 3.3.3.1, 00:03:28
RTC#
```

Redistribution of Only OSPF External (Type 1 and 2) Routes into BGP

Use the **external** keyword along with the **redistribute** command under **router bgp** to redistribute OSPF external routes into BGP. With the **external** keyword, you have three choices:

- redistribute both external type-1 and type-2 (Default)
- redistribute type-1
- redistribute type-2

Enter the commands in the configuration mode as described here:

```
RTB(config-router)# router bgp 100
RTB(config-router)# redistribute ospf 1 match external
```

In this configuration of Router B, we redistribute only OSPF External routes, but both type-1 and type-2:

```
RTB
-----
hostname RTB
!
interface Ethernet0/0
 ip address 2.2.2.1 255.0.0.0
!
interface Serial1/0
 ip address 3.3.3.1 255.0.0.0
!
router ospf 1
 network 2.0.0.0 0.255.255.255 area 1
!
router bgp 100
 redistribute ospf 1 match external 1 external 2

!--- This redistributes ONLY OSPF External routes,
!--- but both type-1 and type-2.
```

```
neighbor 3.3.3.3 remote-as 200
!
end
```

Note: The configuration shows **match external 1 external 2** and the command entered was **redistribute ospf 1 match external**. This is normal because OSPF automatically appends "external 1 external 2" in the configuration. It matches both OSPF external 1 and external 2 routes and it redistributes both routes into BGP.

Router B redistributes only OSPF External routes:

```
RTB# show ip bgp
BGP table version is 21, local router ID is 192.168.1.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 200.1.1.0        2.2.2.2             20           32768 ?
*> 200.2.2.0        2.2.2.2            104           32768 ?
RTB#
```

Router C learns about these two OSPF external routes from BGP:

```
RTC# show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

B    200.1.1.0/24 [20/20] via 3.3.3.1, 00:01:43
B    200.2.2.0/24 [20/0] via 3.3.3.1, 00:01:43
C    3.0.0.0/8 is directly connected, Serial0/0
```

Redistribution of Only OSPF External Type 1 or Type 2 Routes into BGP

Enter this command under **router bgp 100** on Router B to redistribute only OSPF external 1 routes:

```
router bgp 100
 redistribute ospf 1 match external 1
```

The Router B BGP table shows that it is only redistributing external 1 routes into BGP and all other OSPF routes are not redistributed into BGP:

```
RTB# show ip bgp
BGP table version is 24, local router ID is 192.168.1.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 200.2.2.0        2.2.2.2            104           32768 ?
RTB#
```

In the same way, enter this command under **router bgp 100** on Router B to redistribute only OSPF external 2 routes:

```
router bgp 100
  redistribute ospf 1 match external 2
```

Redistribution of OSPF Internal and External Routes into BGP

In this case, all OSPF routes are redistributed into BGP by using both the **internal** and **external** keywords, as shown in this Router B configuration:

RTB
<pre>hostname RTB ! interface Ethernet0/0 ip address 2.2.2.1 255.0.0.0 ! interface Serial1/0 ip address 3.3.3.1 255.0.0.0 ! router ospf 1 network 2.0.0.0 0.255.255.255 area 1 ! router bgp 100 redistribute ospf 1 match internal external 1 external 2 !--- This redistributes all OSPF routes into BGP. neighbor 3.3.3.3 remote-as 200 ! end</pre>

Again, **external** is replaced with **external 1 external 2** in the configuration. This is normal unless you specify which specific external routes you want to redistribute into BGP. After performing the configuration change, Router B redistributes all OSPF routes and Router C starts to learn all routes from BGP:

```
RTB# show ip bgp
BGP table version is 30, local router ID is 192.168.1.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2.2.2.0/24	0.0.0.0	0		32768	?
*> 131.108.1.0/24	2.2.2.2	84		32768	?
*> 131.108.2.0/24	2.2.2.2	74		32768	?
*> 200.1.1.0	2.2.2.2	20		32768	?
*> 200.2.2.0	2.2.2.2	104		32768	?

```
RTB#
```

```
RTC# sh ip route
```

```
Codes: 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, 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
```

```
Gateway of last resort is not set
```

```

2.0.0.0/24 is subnetted, 1 subnets
B    2.2.2.0 [20/0] via 3.3.3.1, 00:01:24
C    3.0.0.0/8 is directly connected, Serial0/0
B    200.1.1.0/24 [20/20] via 3.3.3.1, 00:01:24
B    200.2.2.0/24 [20/104] via 3.3.3.1, 00:01:24
    131.108.0.0/24 is subnetted, 2 subnets
B    131.108.2.0 [20/74] via 3.3.3.1, 00:01:24
B    131.108.1.0 [20/84] via 3.3.3.1, 00:01:24
RTC#

```

Redistribution of OSPF NSSA-External Routes into BGP

This is a special case in which only Not-So-Stubby Area (NSSA) routes are redistributed into BGP. This case is very similar to the case described in the Redistribution of only OSPF External (Type 1 and 2) Routes into BGP section. The only difference is that BGP is now matching NSSA-external routes instead of just external routes. The routing table of Router B shows these OSPF NSSA-external routes:

```

RTB# show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

```


Gateway of last resort is not set

```

O N2 200.1.1.0/24 [110/20] via 2.2.2.1, 00:22:53, Ethernet0
O N1 200.2.2.0/24 [110/20] via 2.2.2.2, 00:22:53, Ethernet0
O IA 131.108.1.0/24 [110/20] via 2.2.2.2, 00:22:53, Ethernet0
O 131.108.2.0/24 [110/20] via 2.2.2.2, 00:22:53, Ethernet0
C    2.0.0.0/8 is directly connected, Ethernet0
C    3.0.0.0/8 is directly connected, Serial1

```

This network diagram is used for this case:



The network diagram shows that Router B receives both OSPF N1 and N2 routes. The default behavior is to redistribute both N1 and N2 routes if only the **nssa-external** keyword is used. This configuration of Router B allows us to redistribute OSPF N1 (200.1.1.0/24) and OSPF N2 (200.2.2.0/24) routes into BGP:

```

RTB
-----
hostname RTB
!
interface Ethernet0/0
 ip address 2.2.2.1 255.0.0.0
!
interface Serial1/0
 ip address 3.3.3.1 255.0.0.0
!
router ospf 1
 network 2.0.0.0 0.255.255.255 area 1
 nssa
!
router bgp 100
 redistribute ospf 1 match nssa-external 1 nssa-external 2

!--- This redistributes only OSPF NSSA-external routes
!--- Type-1 and Type-2 into BGP.

neighbor 3.3.3.3 remote-as 200
!
end
```

Note: Like the OSPF external configuration, the above configuration displays **match nssa-external 1 nssa-external 2** and the command entered was **redistribute ospf 1 match nssa-external**. This is normal because OSPF automatically appends "nssa-external 1 nssa-external 2" in the configuration. It matches both OSPF N1 and OSPF N2 routes and redistributes both routes into BGP.

After the configuration change on Router B, Router B redistributes OSPF NSSA-external routes, and Router C starts learning OSPF NSSA-external routes from BGP:

```

RTB# show ip route
Codes: 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, 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

Gateway of last resort is not set

    2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Ethernet0/0
C       3.0.0.0/8 is directly connected, Serial1/0
O N2 200.1.1.0/24 [110/94] via 2.2.2.1, 00:11:12, Ethernet0/0
O N1 200.2.2.0/24 [110/20] via 2.2.2.2, 00:12:23, Ethernet0/0
    131.108.0.0/24 is subnetted, 2 subnets
O       131.108.2.0 [110/74] via 2.2.2.2, 00:12:23, Ethernet0/0
O IA   131.108.1.0 [110/84] via 2.2.2.2, 00:12:11, Ethernet0/0
RTB#

RTB# show ip bgp
BGP table version is 21, local router ID is 3.3.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
      r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
      Network          Next Hop          Metric LocPrf Weight Path
*> 200.1.1.0          2.2.2.2             94          32768 ?
*> 200.2.2.0          2.2.2.1             20          32768 ?
RTB#
```

```
RTC# show ip route
```

```
Codes: 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, 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
```

```
Gateway of last resort is not set
```

```
C    3.0.0.0/8 is directly connected, Serial0/0
B    200.1.1.0/24 [20/94] via 3.3.3.1, 00:02:06
B    200.2.2.0/24 [20/20] via 3.3.3.1, 00:02:06
RTC#
```

In the same way as with OSPF External routes, to redistribute only OSPF N1 routes, enter the following command under router BGP 100 on Router B:

```
router bgp 100
  redistribute ospf 1 match nssa-external 1

!--- This redistributes only OSPF NSSA-external
!--- Type-1 routes into BGP.
```

To redistribute only OSPF N2 routes, enter the following command under router BGP 100 on Router B:

```
router bgp 100
  redistribute ospf 1 match nssa-external 2

!--- This redistributes only OSPF NSSA-external
!--- Type-2 routes into BGP.
```

Modify the Redistribution Option in OSPF

It is important to understand how the successive configuration change may alter your configuration. A new command with the match option does not overwrite the previous one but is added to it.. Starting with no redistribution, this example configuration command sequence shows its impact on redistribution:

```
R4# conf t
R4(config)# router bgp 100
R4(config-router)# redistribute ospf 1 match internal
R4(config-router)# ^Z

!--- Initially, we redistribute internal OSPF routes into BGP 100.

R4# sh run | i redistribute ospf
  redistribute ospf 1 match internal
R4# conf t
R4(config)# router bgp 100
R4(config-router)# redistribute ospf 1 match external
R4(config-router)# ^Z
```

```
!--- With this second command, we tell BGP  
!--- to also redistribute external OSPF routes.
```

```
R4# sh run | i redistribute ospf  
redistribute ospf 1 match internal external 1 external 2  
R4#  
R4# conf t  
R4(config)# router bgp 100  
R4(config-router)# no redistribute ospf 1 match external 2  
R4(config-router)# ^Z
```

```
!--- With this no command, we only disable the  
!--- redistribution of external type 2 into BGP.  
!--- All other types of routes previously configured remain.
```

```
R4# sh run | i redistribute ospf  
redistribute ospf 1 match internal external 1
```

```
!--- As you can see, internal and external type 1 remain.
```

```
R4# conf t  
R4(config)# router bgp 100  
R4(config-router)# no redistribute ospf 1 match internal external 1  
R4(config-router)# ^Z
```

```
!--- Now, with this no command, which includes all configured  
!--- keywords, it is important to note that we  
!--- still do not disable the redistribution fully.  
!--- We only removed the keyword. After this,  
!--- the IOS still acts as default redistributing  
!--- internal routes only.
```

```
R4# sh run | i redistribute ospf  
redistribute ospf 1  
R4# conf t  
R4(config)# router bgp 100  
R4(config-router)# no redistribute ospf 1
```

```
!--- Always use this command in order to completely  
!--- disable redistribution.
```

```
R4(config-router)# ^Z  
R4# sh run | i redistribute ospf  
R4#
```

Unable to Redistribute iBGP Learnt Routes into an IGP such as EIGRP, OSPF

Route Redistribution is used to propagate routes learned with the use of one protocol, into another routing protocol. When BGP is redistributed into an IGP, only eBGP learned routes get redistributed. The iBGP learned routes known on the router are not introduced into the IGP in order to prevent routing loops from being formed.

By default, iBGP redistribution into IGP is disabled. Issue the **bgp redistribute-internal** command in order to enable redistribution of iBGP routes into IGP. Precautions need to be taken to redistribute specific routes using route maps into IGP.

A sample configuration for redistributing iBGP routes into OSPF is shown here:

```
router bgp 65345
  bgp redistribute-internal
  !
router ospf 100
  redistribute bgp 65345 subnets
```

Note: The redistribution of internal Border Gateway Protocol (iBGP) routes into an Interior Gateway Protocol can cause routing loops within the Autonomous System (AS). This is not recommended. Route filters need to be set in order to control the information, which is imported into the IGP.

Related Information

- [OSPF Support Page](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 5242
