

# Configuración de Redistribución de Rutas BGP Internas en IGP

## Contenido

[Introducción](#)

[Prerequisites](#)

[Requirements](#)

[Componentes Utilizados](#)

[Antecedentes](#)

[Configurar](#)

[Diagrama de la red](#)

[Configuración de OSPF entre R2 y R3](#)

[Verificación](#)

[Configuración de EIGRP entre R2 y R3:](#)

[Verificación](#)

[Configure RIP entre R2 y R3:](#)

[Verificación](#)

[Troubleshoot](#)

## Introducción

En este documento se describe cómo redistribuir rutas del protocolo de gateway fronterizo (BGP, Border Gateway Protocol) interno en el proceso de abrir primero la ruta más corta (OSPF, Open Shortest Path First).

## Prerequisites

### Requirements

Cisco recomienda que conozca la configuración básica de BGP y comprenda los protocolos de ruteo de:

- BGP
- OSPF
- Protocolo de routing de gateway interior mejorado (EIGRP)
- Routing Information Protocol (RIP)

Para obtener más información, consulte [Casos Prácticos de BGP](#) y [Configuración de BGP](#).

## Componentes Utilizados

La información de este documento se basa en la versión 15.1(4)M5 del software del IOS® de Cisco.

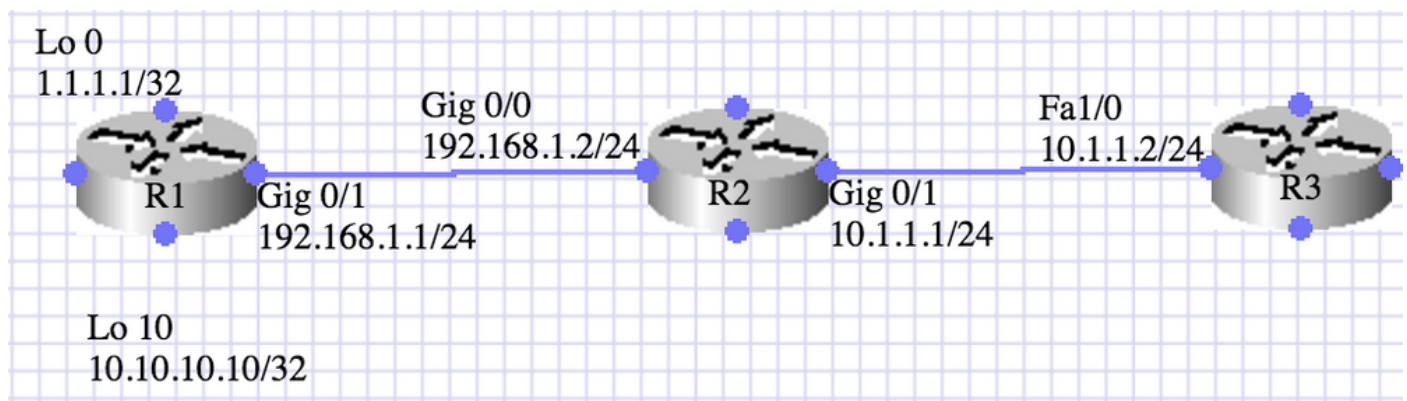
The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. Si tiene una red en vivo, asegúrese de entender el posible impacto de cualquier comando.

## Antecedentes

Al igual que en otros protocolos de gateway interior (IGP) para la redistribución IGP, el comportamiento es diferente cuando el BGP interno (IBGP) se redistribuye en OSPF. Las rutas aprendidas de IBGP no se reenvían a un protocolo de ruteo IGP a través del comando `redistribute`. Utilice el comando [bgp redistribute-internal](#) en el proceso BGP en el router que redistribuye.

## Configurar

### Diagrama de la red



### Configuración de OSPF entre R2 y R3

En el escenario representado aquí, los routers R1 y R2 ejecutan IBGP, y los routers R2 o R3 ejecutan OSPF Area 0. R1 anuncia dos rutas (1.1.1.1 /32 y 10.10.10.10/32) a través del comando `network`.

R2 redistribuye BGP en el Área 0 OSPF. Es necesario redistribuir las rutas internas seleccionadas (10.10.10.10/32).

La tarea se logra con el uso de `prefix-list` y `route-map`.

R1:

```
interface Loopback0
ip address 1.1.1.1 255.255.255.255
!
interface Loopback10
ip address 10.10.10.10 255.255.255.255
!
interface GigabitEthernet0/1
ip address 192.168.1.1 255.255.255.0
duplex auto
speed auto
!
```

```
router bgp 10
no synchronization
bgp router-id 1.1.1.1
bgp log-neighbor-changes
network 1.1.1.1 mask 255.255.255.255
network 10.10.10.10 mask 255.255.255.255
neighbor 192.168.1.2 remote-as 100
no auto-summary
```

```
R1#show ip bgp summary
BGP router identifier 10.10.10.10, local AS number 10
BGP table version is 3, main routing table version 3
2 network entries using 296 bytes of memory
2 path entries using 128 bytes of memory
1/1 BGP path/bestpath attribute entries using 136 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 560 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
192.168.1.2 4 10 6 7 3 0 0 00:03:10 0
```

## R2:

```
interface Loopback0
ip address 2.2.2.2 255.255.255.255
!
interface GigabitEthernet0/0
ip address 192.168.1.2 255.255.255.0
duplex auto
speed auto
!
interface GigabitEthernet0/1
ip address 10.1.1.1 255.255.255.0
duplex auto
speed auto
!
```

```
router ospf 1
router-id 2.2.2.2
log-adjacency-changes
redistribute bgp 100 metric 100 metric-type 1 subnets route-map BGP-To_OSPF
network 10.1.1.1 0.0.0.0 area 0
```

```
R2#show ip ospf neighbor
Neighbor ID Pri State Dead Time Address Interface
3.3.3.3 1 FULL/BDR 00:00:38 10.1.1.2 GigabitEthernet0/1
```

```
router bgp 10
no synchronization
bgp router-id 2.2.2.2
bgp log-neighbor-changes
bgp redistribute-internal
neighbor 192.168.12.1 remote-as 10
no auto-summary
!
ip prefix-list BGP-to-ospf seq 5 permit 172.16.0.0/16
!
route-map BGP-To_OSPF permit 10
match ip address prefix-list BGP-to-ospf
```

```

R2#show ip bgp summary
BGP router identifier 192.168.1.2, local AS number 10
BGP table version is 3, main routing table version 3
2 network entries using 272 bytes of memory
2 path entries using 112 bytes of memory
1/1 BGP path/bestpath attribute entries using 128 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 512 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
192.168.1.1 4 10 8 7 3 0 0 00:03:52 2
R2#show ip bgp
BGP table version is 3, local router ID is 192.168.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
*>i1.1.1.1/32 192.168.1.1 0 100 0 i
*>i10.10.10.10/32 192.168.1.1 0 100 0 i

R2#show ip route 1.1.1.1
Routing entry for 1.1.1.1/32
Known via "bgp 10", distance 200, metric 0, type internal
Last update from 192.168.1.1 00:04:53 ago
Routing Descriptor Blocks:
* 192.168.1.1, from 192.168.1.1, 00:04:53 ago
Route metric is 0, traffic share count is 1
AS Hops 0
MPLS label: none

R2#show ip route 10.10.10.10
Routing entry for 10.10.10.10/32
Known via "bgp 10", distance 200, metric 0, type internal
Last update from 192.168.1.1 00:04:56 ago
Routing Descriptor Blocks:
* 192.168.1.1, from 192.168.1.1, 00:04:56 ago
Route metric is 0, traffic share count is 1
AS Hops 0
MPLS label: none

```

### R3:

```

interface FastEthernet1/0
ip address 10.1.1.2 255.255.255.0
duplex auto
speed auto

```

```

router ospf 1
log-adjacency-changes
network 10.1.1.2 0.0.0.0 area 0

```

```

R3#show ip ospf neighbor
Neighbor ID Pri State Dead Time Address Interface
192.168.1.2 1 FULL/DR 00:00:36 10.1.1.1 GigabitEthernet0/1

```

La tabla de ruteo en R3 antes de la redistribución de BGP - interna se agrega en R2 bajo el router BGP 10:

```
R3#show 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 not set
3.0.0.0/32 is subnetted, 1 subnets
C 3.3.3.3 is directly connected, Loopback0
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.1.1.0/24 is directly connected, GigabitEthernet0/1
L 10.1.1.2/32 is directly connected, GigabitEthernet0/1
```

R2:

```
router bgp 10
bgp redistribute-internal
```

## Verificación

R3:

La tabla de ruteo para R3 después de la redistribución BGP - internal se agrega en R2 bajo el router BGP 10:

```
R3#show 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 not set
3.0.0.0/32 is subnetted, 1 subnets
C 3.3.3.3 is directly connected, Loopback0
10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.1.1.0/24 is directly connected, GigabitEthernet0/1
L 10.1.1.2/32 is directly connected, GigabitEthernet0/1
O E1 10.10.10.10/32 [110/11] via 10.1.1.1, 00:00:06, GigabitEthernet0/1
```

## Configuración de EIGRP entre R2 y R3:

En el escenario representado aquí, los routers R1 y R2 ejecutan IBGP, y los routers R2 o R3 ejecutan EIGRP Autonomous System (AS) 1. R1 anuncia dos rutas (1.1.1.1 /32 y 10.10.10.10/32) a través del comando network.

R2 redistribuye BGP en EIGRP AS 1. Es necesario redistribuir las rutas internas seleccionadas (10.10.10.10/32).

La tarea se logra con el uso de prefix-list y route-map.

R2:

```
router eigrp 1
network 10.0.0.0
redistribute bgp 10 metric 1544 10 255 1 1500 route-map BGP_To_EIGRP
eigrp router-id 2.2.2.2
```

```
route-map BGP_To_EIGRP, permit, sequence 10
Match clauses:
ip address prefix-lists: BGP-to-eigrp
Set clauses:
Policy routing matches: 0 packets, 0 bytes
```

```
ip prefix-list BGP-to-eigrp: 1 entries
seq 1 permit 10.10.10.10/32
```

**R3:**

```
router eigrp 1
network 10.0.0.0
eigrp router-id 3.3.3.3
```

El resultado del comando show IP route en R3 antes de la redistribución de BGP - internal se agrega en R2 bajo el router BGP 10:

```
R3#show 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 not set
3.0.0.0/32 is subnetted, 1 subnets
C 3.3.3.3 is directly connected, Loopback0
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.1.1.0/24 is directly connected, GigabitEthernet0/1
L 10.1.1.2/32 is directly connected, GigabitEthernet0/1
```

**R2:**

```
router bgp 10
bgp redistribute-internal
```

## Verificación

El resultado del comando show IP route en R3 después de que BGP redistribute-internal se agrega en R2 bajo el router BGP 10:

```
R3#show 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 not set  
3.0.0.0/32 is subnetted, 1 subnets  
C 3.3.3.3 is directly connected, Loopback0  
10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks  
C 10.1.1.0/24 is directly connected, GigabitEthernet0/1  
L 10.1.1.2/32 is directly connected, GigabitEthernet0/1  
**D EX 10.10.10.10/32**  
**[170/1660672] via 10.1.1.1, 00:00:04, GigabitEthernet0/1**

## Configure RIP entre R2 y R3:

En el escenario representado aquí, los routers R1 y R2 ejecutan IBGP, y los routers R2 o R3 ejecutan RIPv2.

R1 anuncia dos rutas (1.1.1.1 /32 y 10.10.10.10/32) a través del comando network.  
R2 redistribuye BGP en RIPv2. Es necesario redistribuir las rutas internas seleccionadas (10.10.10.10/32).

La tarea se logra con el uso de prefix-list y route-map.

R2:

```
router rip
version 2
redistribute bgp 10 metric 1 route-map BGP_To_RIP
network 10.0.0.0
no auto-summary
```

```
route-map BGP_To_RIP, permit, sequence 10
Match clauses:
ip address prefix-lists: BGP-to-rip
Set clauses:
Policy routing matches: 0 packets, 0 bytes
```

```
ip prefix-list BGP-to-rip: 1 entries
seq 1 permit 10.10.10.10/32
```

R3:

```
router rip
version 2
network 10.0.0.0
no auto-summary
```

Salida en R3 antes de habilitar el BGP redistribute-internal en R2 bajo el router BGP 10:

```
R3#show 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 not set
3.0.0.0/32 is subnetted, 1 subnets
C 3.3.3.3 is directly connected, Loopback0
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.1.1.0/24 is directly connected, GigabitEthernet0/1
L 10.1.1.2/32 is directly connected, GigabitEthernet0/1
```

R2:

```
router bgp 10
bgp redistribute-internal
```

## Verificación

Salida en R3 después de habilitar la redistribución BGP - interna en R2 bajo el router BGP 10:

```
R3#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 not set
3.0.0.0/32 is subnetted, 1 subnets
C 3.3.3.3 is directly connected, Loopback0
10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.1.1.0/24 is directly connected, GigabitEthernet0/1
L 10.1.1.2/32 is directly connected, GigabitEthernet0/1
R 10.10.10.10/32 [120/1] via 10.1.1.1, 00:00:09, GigabitEthernet0/1
```

## Troubleshoot

Actualmente no hay información de troubleshooting específica disponible para esta configuración.