

# 配置将内部BGP路由重分布到IGP

## 目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[背景信息](#)

[配置](#)

[网络图](#)

[在R2和R3之间配置OSPF](#)

[验证](#)

[在R2和R3之间配置EIGRP:](#)

[验证](#)

[在R2和R3之间配置RIP:](#)

[验证](#)

[故障排除](#)

## 简介

本文档介绍如何将内部边界网关协议 (BGP) 路由重新分配到开放最短路径优先 (OSPF) 进程中。

## 先决条件

### 要求

思科建议您了解基本BGP配置并了解以下路由协议：

- 调试输出中显示“BGP
- OSPF
- 增强型内部网关路由协议 (EIGRP)
- 路由信息协议 (RIP)

有关详细信息，请参阅 [BGP 案例研究和配置 BGP](#)。

### 使用的组件

本文档中的信息基于Cisco IOS®软件版本15.1(4)M5。

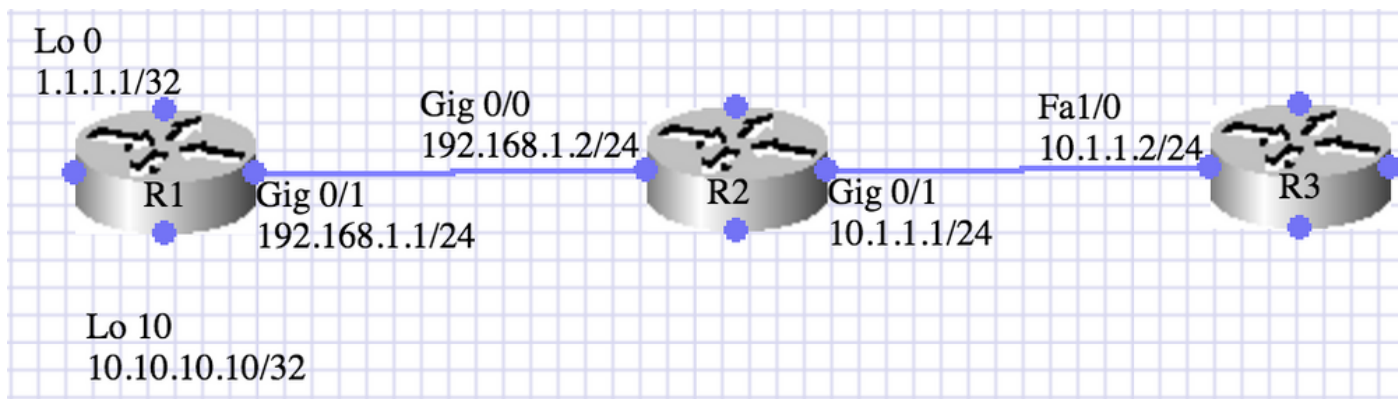
本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

## 背景信息

与其他内部网关协议(IGP)到IGP重分发一样，当内部BGP(IBGP)重分发到OSPF时，行为也不同。IBGP获知的路由不会通过redistribute命令转发到IGP路由协议。在重分发的路由器上的BGP进程下使用**bgp redistribute-internal**命令。

## 配置

### 网络图



### 在R2和R3之间配置OSPF

在此处描述的场景中，路由器R1和R2运行IBGP，路由器R2或R3运行OSPF区域0。R1通过network命令通告两条路由(1.1.1.1 /32和10.10.10.10/32)。R2将BGP重分发到OSPF区域0。需要重分发所选内部路由(10.10.10.10/32)。该任务通过使用前缀列表和路由映射来完成。

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 10.1.1.1.1 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
```

### BGP redistribute - internal之前R3中的路由表添加到路由器BGP 10的R2上 :

```
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
```

## 验证

R3:

BGP redistribute - internal之后R3的路由表添加到路由器BGP 10的R2上 :

```
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
```

## 在R2和R3之间配置EIGRP:

在此处描述的场景中，路由器R1和R2运行IBGP，路由器R2或R3运行EIGRP自治系统(AS)1。R1通过network命令通告两条路由(1.1.1.1 /32和10.10.10.10/32)。

R2将BGP重分发到EIGRP AS 1中。需要重分发所选内部路由(10.10.10.10/32)。该任务通过使用前缀列表和路由映射来完成。

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
```

BGP redistribute - internal在R2的路由器BGP 10下添加之前，R3上的show IP route的输出：

```
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
```

## 验证

在路由器BGP 10下的R2上添加BGP redistribute-internal后，R3上的show IP route的输出：

```
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
```

## 在R2和R3之间配置RIP:

在此处描述的场景中，路由器R1和R2运行IBGP，路由器R2或R3运行RIPv2。

R1通过network命令通告两条路由(1.1.1.1 /32和10.10.10.10/32)。

R2将BGP重分发到RIPv2中。需要重分发选定的内部路由(10.10.10.10/32)。该任务通过使用前缀列表和路由映射来完成。

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
```

在路由器BGP 10下启用R2上的BGP redistribute-internal之前，R3上的输出：

```
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
```

**验证**

在路由器BGP 10下启用BGP redistribute - R2上内部后，R3上的输出：

```
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

## 故障排除

当前没有可用于此配置的特定故障排除信息。