配置將內部BGP路由重分發到IGP

目錄

簡介

本文說明如何將內部邊界閘道通訊協定(BGP)路由重分佈到開放最短路徑優先(OSPF)進程中。

必要條件

需求

思科建議您瞭解BGP的基本配置，並瞭解以下路由協定：

- BGP
- OSPF
- 增強型內部閘道路由通訊協定(EIGRP)
- 路由資訊通訊協定(RIP)

如需詳細資訊，請參閱BGP個案研究和設定BGP。

採用元件

本檔案中的資訊是根據Cisco IOS®軟體版本15.1(4)M5。

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除（預設）的組態來啟動。如果您的網路運作中，請確保您瞭解任何指令可能造成的影響。

背景資訊

如同在其他內部閘道通訊協定(IGP)到IGP的重新分佈中，將內部BGP(IBGP)重新分佈到OSPF中時，行為會有所不同。IBGP獲知的路由不會通過redistribute命令轉發到IGP路由協定。在重新分配的路由器上的BGP進程下使用命令bgp redistribute-internal。

設定

網路圖表

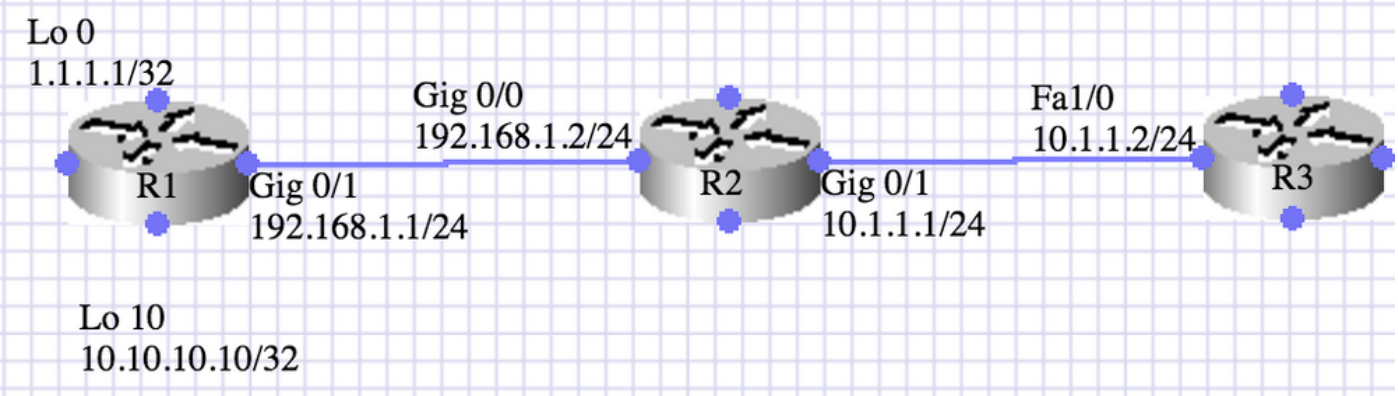

在R2和R3之間配置OSPF

此處描述的場景中，路由器R1和R2運行IBGP，路由器R2或R3運行OSPF區域0。R1通過network命令通告兩條路由(1.1.1.1 /32和10.10.10.10/32)。
R2將BGP重分佈到OSPF區域0。需要重分佈選定的內部路由(10.10.10.10/32)。
該任務通過使用字首清單和路由對映來實現。

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
```

在R2的BGP 10下新增了R3中BGP redistribute - internal之前的路由表：

```
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
```

驗證

R3:

在R2的路由器BGP 10下新增了BGP redistribute - internal後R3的路由表：

```
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與R3之間配置EIGRP:

在此所示場景中，路由器R1和R2運行IBGP，路由器R2或R3運行EIGRP Autonomous System(AS)1。R1通過network命令通告兩條路由(1.1.1.1 /32和10.10.10.10/32)。R2將BGP重分佈到EIGRP AS 1中。需要重分佈選定的內部路由(10.10.10.10/32)。該任務通過使用字首清單和路由對映來實現。

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
```

R3上BGP redistribute - internal之前的show IP route輸出已新增到R2的路由器BGP 10下：

```
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
```

驗證

在R2的路由器BGP 10下新增BGP redistribute-internal後R3上show IP route的輸出：

```
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之間配置RIP:

在此所示場景中，路由器R1和R2運行IBGP，路由器R2或R3運行RIPv2。

R1通過network命令通告兩條路由(1.1.1.1 /32和10.10.10.10/32)。
R2將BGP重分佈到RIPv2中。需要重分佈選定的內部路由(10.10.10.10/32)。
該任務通過使用字首清單和路由對映來實現。

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下在R2上啟用BGP redistribute-internal之前，R3上的輸出：

```
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
```

驗證

在路由器BGP 10下啟用BGP redistribute - R2上的內部後R3上的輸出：

```
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
```

疑難排解

目前尚無適用於此組態的具體疑難排解資訊。