

目录

简介

[BGP路由反射的说明](#)

[路由反射配置示例](#)

[与默认设置的单个团星](#)

[与禁用的客户端到客户端的反射的单个团星](#)

[两集群、内部站点和站点之间路由反映](#)

[两集群，没有客户端到客户端的反射](#)

[团星列表和环路预防](#)

[在客户端和无委托人之间的反射](#)

[簇内部反射](#)

[集群间反射](#)

[MCIDs和环路预防](#)

[参考](#)

[相关的思科支持社区讨论](#)

简介

此条款描述边界网关协议(BGP)路由反射不同的多个集群ID方案和使用情况。BGP特别是概念集群和路由反映前期知识假设。

BGP路由反射的说明

BGP扬声器是BGP可用的路由器。默认情况下BGP扬声器不通告IBGP学习的前缀给iBGP对等体-这执行维护环路预防。RFC4456介绍取消全网状需要在iBGP扬声器之间的路由反映功能。当路由反射器反射前缀时，创建/修改添加呼叫的CLUSTER_LIST一个可选非传递属性其自己的集群ID到它。此属性使用环路预防：当路由器CLUSTER_LIST包含路由器的自己的集群ID的接收更新，此更新丢弃。


默认情况下集群ID设置为BGP路由器号值，但是可以设置为一个任意32位值。多个集群ID(MCID)功能提供分配每邻居集群ID。因此，有路由反映方案的3种类型。

1. 在客户端和无委托人之间
2. 在同一集群的客户端之间(簇内部)
3. 在不同的集群的客户端之间(集群间)

路由反映配置示例

Following是一些路由器反射方案和各自配置示例。

与默认设置的单个团星



在这种配置是完成在作为路由反射器的路由器R1和R2是RR的客户端。在常规设计，无委托人路由器将是路由器的路由反射器在下层级水平，但是在本例中另一个PE使用的为了简化。

```
#-neighbors C2C-rfl-CFG C2C-rfl-USERR#show ip bgp 172.16.1.1BGP routing table entry for
172.16.1.1/32, version 2Paths: (1 available, best #1, table default) Advertised to update-
groups:      1          2 Refresh Epoch 2 Local, (Received from a RR-client) 10.0.10.2 from
10.0.10.2 (172.16.1.1) Origin IGP, metric 0, localpref 100, valid, internal, best rx
pathid: 0, tx pathid: 0x0RR#show ip bgp update-group 1BGP version 4 update-group 1, internal,
Address Family: IPv4 Unicast BGP Update version : 4/0, messages 0 Topology: global, highest
version: 4, tail marker: 4 Format state: Current working (OK, last not in list)
Refresh blocked (not in list, last not in list) Update messages formatted 2, replicated 2,
current 0, refresh 0, limit 1000 Number of NLRI in the update sent: max 1, min 0 Minimum time
between advertisement runs is 0 seconds Has 1 member: 10.0.40.2 RR#show ip bgp update-group
2BGP version 4 update-group 2, internal, Address Family: IPv4 Unicast BGP Update version : 4/0,
messages 0 Route-Reflector Client Topology: global, highest version: 4, tail marker: 4 Format
state: Current working (OK, last not in list) Refresh blocked (not in list, last
not in list) Update messages formatted 3, replicated 6, current 0, refresh 0, limit 1000
Number of NLRI in the update sent: max 1, min 0 Minimum time between advertisement runs is 0
seconds Has 2 members: 10.0.10.2 10.0.20.2
```

这些输出显示RR接收从S1PE1的172.16.1.1/32前缀并且反射它给客户端S1PE2和无委托人S2PE1。在这个特定情况下，更新也被退还的对S1PE1，但是发生，因为S1PE1和S1PE2有同一项路由策略，并且，因此，建立同一更新组。

与禁用的客户端到客户端的反射的单个团星

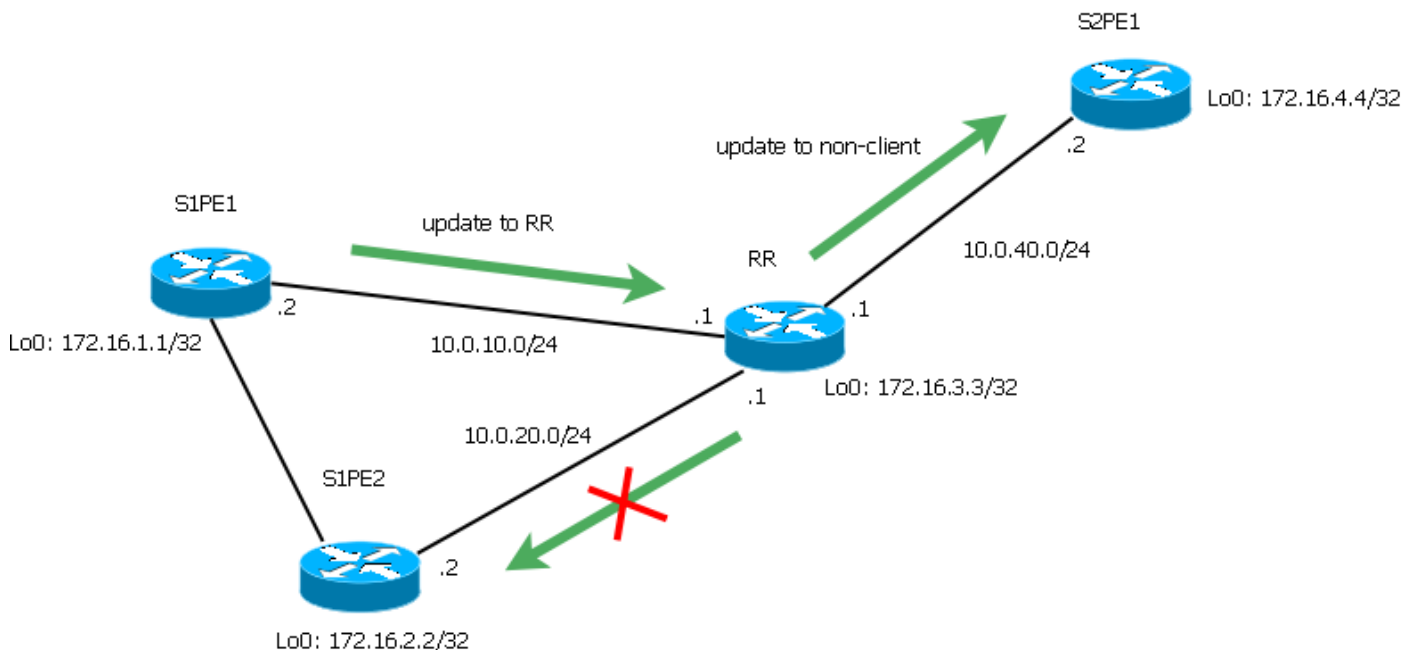


图 2

跟随的配置是完成在作为路由反射器的路由器RR。

```
RR#show run | sec bgprouter bgp 1 no bgp client-to-client reflection bgp log-neighbor-changes
neighbor 10.0.10.2 remote-as 1 neighbor 10.0.10.2 route-reflector-client neighbor 10.0.20.2
remote-as 1 neighbor 10.0.20.2 route-reflector-client neighbor 10.0.40.2 remote-as 1
```

假设AS1部分地网状连接：S1PE1和S1PE2表iBGP相邻(例如，他们在同一个站点查找，并且我们要优化网络处理更新)的方式。在这种情况下RR有禁用的客户端到客户端的反射，并且反射仅来自S1PE1的172.16.1.1/32到无委托人S2PE1。

```
RR#show ip bgp cluster-idsGlobal cluster-id: 172.16.3.3 (configured: 0.0.0.0)BGP client-to-
client reflection:      Configured      Used all (inter-cluster and intra-cluster): DISABLED
intra-cluster:          ENABLED          DISABLEDList of cluster-ids:Cluster-id
#-neighbors C2C-rfl-CFG C2C-rfl-USERR#show ip bgp 172.16.1.1BGP routing table entry for
172.16.1.1/32, version 5Paths: (1 available, best #1, table default, RIB-failure(17))
Advertised to update-groups:      1 Refresh Epoch 2 Local, (Received from a RR-client)
```

```

10.0.10.2 from 10.0.10.2 (172.16.1.1)      Origin IGP, metric 0, localpref 100, valid, internal,
best      rx pathid: 0, tx pathid: 0x0
RR#show ip bgp update-group 1BGP version 4 update-group 1, internal, Address Family: IPv4
Unicast BGP Update version : 7/0, messages 0 Topology: global, highest version: 7, tail marker:
7 Format state: Current working (OK, last not in list) Refresh blocked (not in list, last not in
list) Update messages formatted 4, replicated 4, current 0, refresh 0, limit 1000 Number of
NLRIs in the update sent: max 1, min 0 Minimum time between advertisement runs is 0 seconds Has
1 member: 10.0.40.2

```

两集群、内部站点和站点之间路由反映

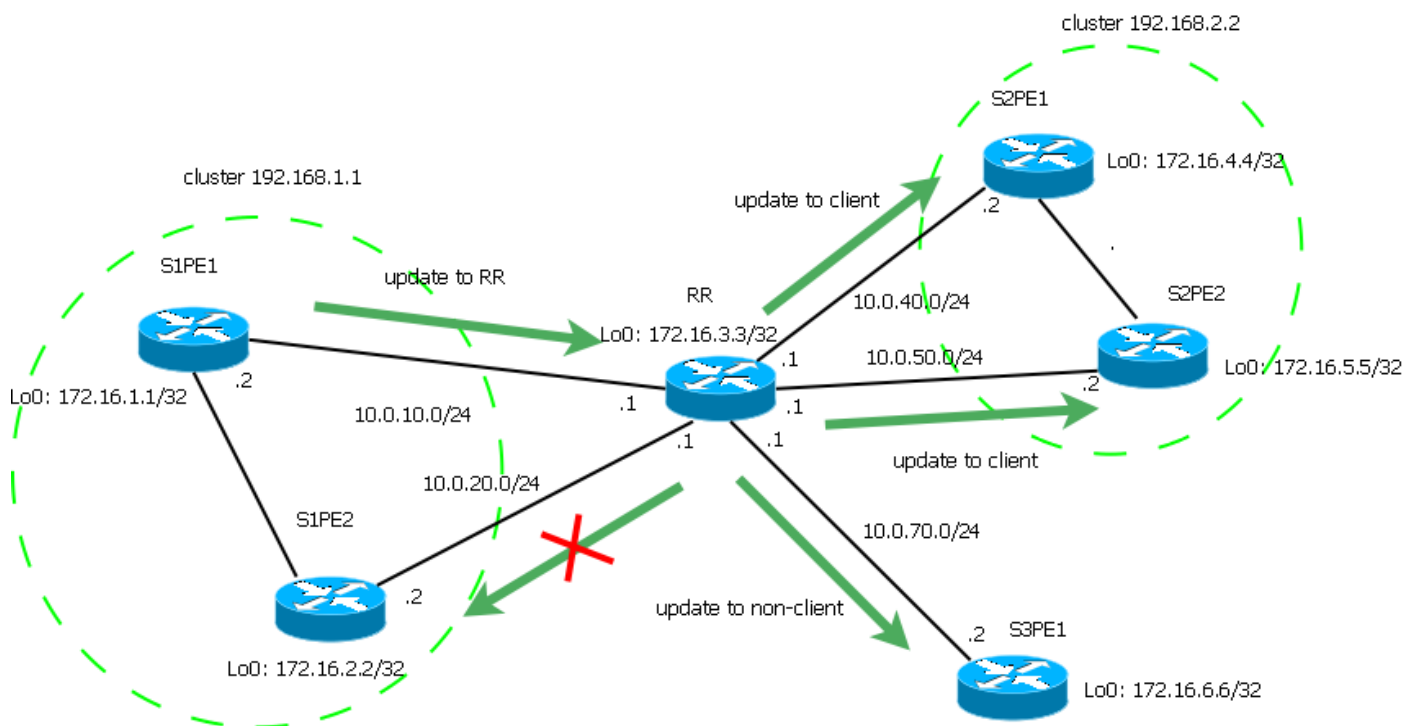


图 3

跟随的配置是完成在作为路由反射器的路由器RR。

```

RR#sh run | sec bgprouter bgp 1 no bgp client-to-client reflection intra-cluster cluster-id
192.168.1.1 bgp log-neighbor-changes neighbor 10.0.10.2 remote-as 1 neighbor 10.0.10.2 cluster-
id 192.168.1.1 neighbor 10.0.10.2 route-reflector-client neighbor 10.0.20.2 remote-as 1 neighbor
10.0.20.2 cluster-id 192.168.1.1 neighbor 10.0.20.2 route-reflector-client neighbor 10.0.40.2
remote-as 1 neighbor 10.0.40.2 cluster-id 192.168.2.2 neighbor 10.0.40.2 route-reflector-client
neighbor 10.0.50.2 remote-as 1 neighbor 10.0.50.2 cluster-id 192.168.2.2 neighbor 10.0.50.2
route-reflector-client neighbor 10.0.70.2 remote-as 1

```

在这种情况下两在站点1表的观点扫描器集群192.168.1.1，当两在站点2表的观点扫描器集群192.168.2.2时。S3PE1是无委托人。在站点1的观点扫描器有直接iBGP会话，簇内部反射为集群192.168.1.1禁用，但是为集群192.168.2.2仍然启用。集群间反射启用。

```

RR#show ip bgp cluster-idsGlobal cluster-id: 172.16.3.3 (configured: 0.0.0.0)BGP client-to-
client reflection:      Configured      Used all (inter-cluster and intra-cluster): ENABLED
intra-cluster:          ENABLED          ENABLEDList of cluster-ids:Cluster-id
#-neighbors C2C-rfl-CFG C2C-rfl-USE192.168.1.1      2 DISABLED      DISABLED192.168.2.2
2 ENABLED      ENABLEDRR#show ip bgp 172.16.1.1BGP routing table entry for 172.16.1.1/32, version
5Paths: (1 available, best #1, table default, RIB-failure(17)) Advertised to update-groups:
3      5 Refresh Epoch 9 Local, (Received from a RR-client) 10.0.10.2 from 10.0.10.2
(172.16.1.1)      Origin IGP, metric 0, localpref 100, valid, internal, best      rx pathid: 0,
tx pathid: 0x0
RR#show ip bgp update-group 3BGP version 4 update-group 3, internal, Address Family: IPv4
Unicast BGP Update version : 11/0, messages 0 Topology: global, highest version: 11, tail
marker: 11 Format state: Current working (OK, last not in list) Refresh blocked (not in list,

```

last not in list) Update messages formatted 20, replicated 20, current 0, refresh 0, limit 1000
Number of NLRI in the update sent: max 1, min 0 Minimum time between advertisement runs is 0
seconds Has 1 member: 10.0.70.2

```
RR#show ip bgp update-group 5BGP version 4 update-group 5, internal, Address Family: IPv4
Unicast BGP Update version : 11/0, messages 0 Route-Reflector Client Configured with cluster-id
192.168.2.2 Topology: global, highest version: 11, tail marker: 11 Format state: Current working
(OK, last not in list) Refresh blocked (not in list, last not in list) Update messages formatted
22, replicated 34, current 0, refresh 0, limit 1000 Number of NLRI in the update sent: max 3,
min 0 Minimum time between advertisement runs is 0 seconds Has 2 members: 10.0.40.2 10.0.50.2
```

前缀172.16.1.1/32已接收从S1PE1反射给集群的192.168.2.2客户端和对无委托人。同时，请加前缀
172.16.4.4/32已接收从S2PE1反射给所有客户端和无委托人。

```
RR#show ip bgp cluster-idsGlobal cluster-id: 172.16.3.3 (configured: 0.0.0.0)BGP client-to-
client reflection:          Configured      Used all (inter-cluster and intra-cluster): ENABLED
intra-cluster:              ENABLED        ENABLEDList of cluster-ids:Cluster-id
#-neighbors C2C-rfl-CFG C2C-rfl-USE192.168.1.1      2 DISABLED      DISABLED192.168.2.2
2 ENABLED      ENABLEDRR#show ip bgp 172.16.1.1BGP routing table entry for 172.16.1.1/32, version
5Paths: (1 available, best #1, table default, RIB-failure(17)) Advertised to update-groups:
3          5 Refresh Epoch 9 Local, (Received from a RR-client) 10.0.10.2 from 10.0.10.2
(172.16.1.1) Origin IGP, metric 0, localpref 100, valid, internal, best rx pathid: 0,
tx pathid: 0x0
```

```
RR#show ip bgp update-group 3BGP version 4 update-group 3, internal, Address Family: IPv4
Unicast BGP Update version : 11/0, messages 0 Topology: global, highest version: 11, tail
marker: 11 Format state: Current working (OK, last not in list) Refresh blocked (not in list,
last not in list) Update messages formatted 20, replicated 20, current 0, refresh 0, limit 1000
Number of NLRI in the update sent: max 1, min 0 Minimum time between advertisement runs is 0
seconds Has 1 member: 10.0.70.2
```

```
RR#show ip bgp update-group 5BGP version 4 update-group 5, internal, Address Family: IPv4
Unicast BGP Update version : 11/0, messages 0 Route-Reflector Client Configured with cluster-id
192.168.2.2 Topology: global, highest version: 11, tail marker: 11 Format state: Current working
(OK, last not in list) Refresh blocked (not in list, last not in list) Update messages formatted
22, replicated 34, current 0, refresh 0, limit 1000 Number of NLRI in the update sent: max 3,
min 0 Minimum time between advertisement runs is 0 seconds Has 2 members: 10.0.40.2 10.0.50.2
```

您能禁用集群的192.168.2.2内部站点路由反映，但是该集群的客户端应该在这种情况下有iBGP会
话全网状：

```
RR(config-router)#no bgp client-to-client reflection intra-cluster cluster-id 192.168.2.2
```

```
RR#sh ip bgp cluster-ids
Global cluster-id: 172.16.3.3 (configured: 0.0.0.0)
BGP client-to-client reflection:          Configured      Used
  all (inter-cluster and intra-cluster): ENABLED
  intra-cluster:              ENABLED        ENABLED
```

```
List of cluster-ids:
Cluster-id      #-neighbors C2C-rfl-CFG C2C-rfl-USE
192.168.1.1      2 DISABLED      DISABLED
192.168.2.2      2 DISABLED      DISABLED
```

内部站点反射可以为所有集群也禁用：

```
RR(config-router)#no bgp client-to-client reflection intra-cluster cluster-id 192.168.2.2
```

```
RR#sh ip bgp cluster-ids
Global cluster-id: 172.16.3.3 (configured: 0.0.0.0)
BGP client-to-client reflection:          Configured      Used
  all (inter-cluster and intra-cluster): ENABLED
  intra-cluster:              ENABLED        ENABLED
```

```
List of cluster-ids:
Cluster-id      #-neighbors C2C-rfl-CFG C2C-rfl-USE
192.168.1.1      2 DISABLED      DISABLED
192.168.2.2      2 DISABLED      DISABLED
```

两集群，没有客户端到客户端的反射

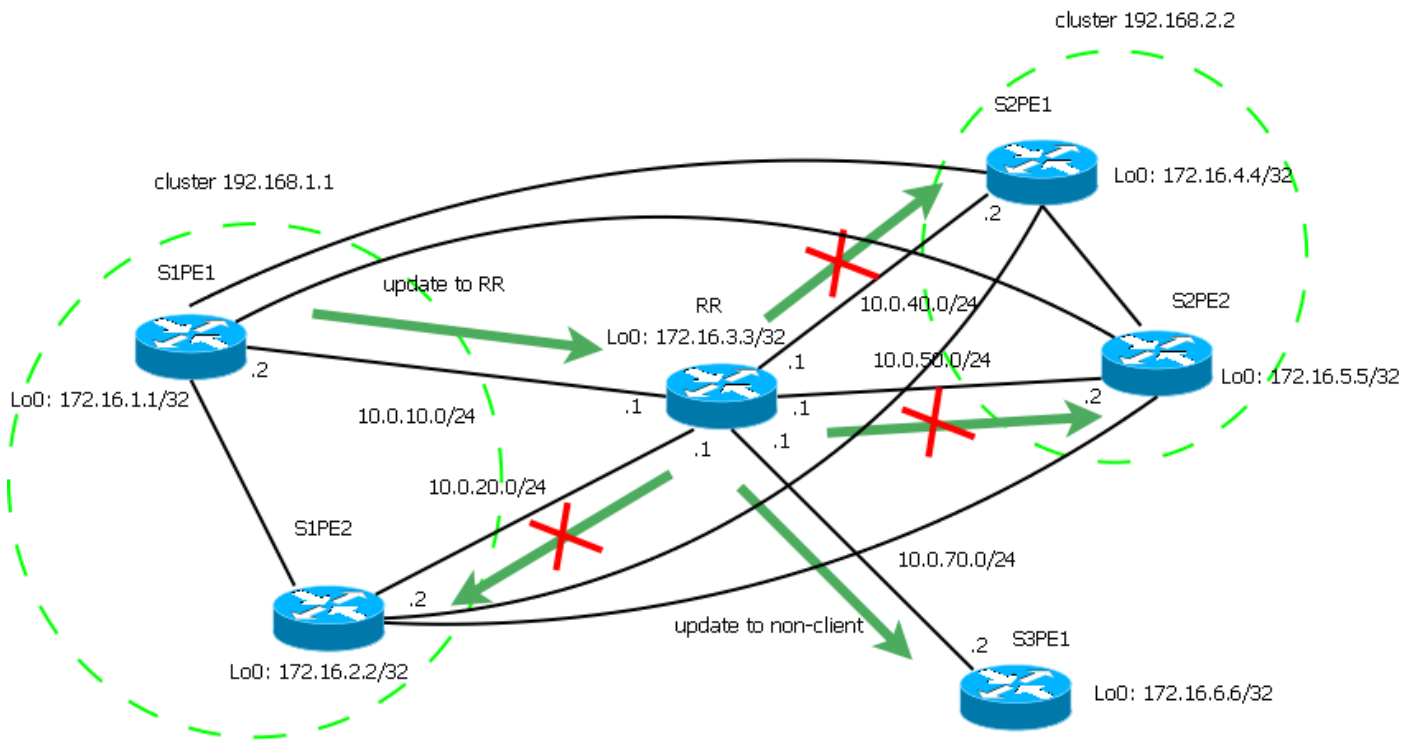


图 4

跟随的配置是完成在作为路由反射器的路由器RR。

```
RR#show run | sec bgprouter bgp 1 no bgp client-to-client reflection bgp log-neighbor-changes
neighbor 10.0.10.2 remote-as 1 neighbor 10.0.10.2 cluster-id 192.168.1.1 neighbor 10.0.10.2
route-reflector-client neighbor 10.0.20.2 remote-as 1 neighbor 10.0.20.2 cluster-id 192.168.1.1
neighbor 10.0.20.2 route-reflector-client neighbor 10.0.40.2 remote-as 1 neighbor 10.0.40.2
cluster-id 192.168.2.2 neighbor 10.0.40.2 route-reflector-client neighbor 10.0.50.2 remote-as 1
neighbor 10.0.50.2 cluster-id 192.168.2.2 neighbor 10.0.50.2 route-reflector-client neighbor
10.0.70.2 remote-as 1
```

禁用簇内部和集群间反射是可能的。在这种情况下，在客户端之间的仅反射和无委托人将执行。

```
RR#show ip bgp cluster-idsGlobal cluster-id: 172.16.3.3 (configured: 0.0.0.0)BGP client-to-
client reflection:          Configured   Used all (inter-cluster and intra-cluster): DISABLED
intra-cluster:              ENABLED      DISABLEDList of cluster-ids:Cluster-id
#-neighbors C2C-rfl-CFG C2C-rfl-USE192.168.1.1          2 ENABLED      DISABLED192.168.2.2
2 ENABLED      DISABLED
```

```
RR#show ip bgp 172.16.1.1
BGP routing table entry for 172.16.1.1/32, version 5
Paths: (1 available, best #1, table default, RIB-failure(17))
Advertised to update-groups:
 3
Refresh Epoch 9
Local, (Received from a RR-client)
 10.0.10.2 from 10.0.10.2 (172.16.1.1)
  Origin IGP, metric 0, localpref 100, valid, internal, best
  rx pathid: 0, tx pathid: 0x0
```

```
RR#show ip bgp 172.16.4.4
BGP routing table entry for 172.16.4.4/32, version 9
Paths: (1 available, best #1, table default, RIB-failure(17))
```

Advertised to update-groups:

3

Refresh Epoch 6

Local, (Received from a RR-client)

10.0.40.2 from 10.0.40.2 (172.16.4.4)

Origin IGP, metric 0, localpref 100, valid, internal, best

rx pathid: 0, tx pathid: 0x0

RR#show ip bgp update-group 3

BGP version 4 update-group 3, internal, Address Family: IPv4 Unicast

BGP Update version : 11/0, messages 0

Topology: global, highest version: 11, tail marker: 11

Format state: Current working (OK, last not in list)

Refresh blocked (not in list, last not in list)

Update messages formatted 20, replicated 20, current 0, refresh 0, limit 1000

Number of NLRI in the update sent: max 1, min 0

Minimum time between advertisement runs is 0 seconds

Has 1 member:

10.0.70.2

前缀172.16.1.1/32和172.16.4.4/32由集群192.168.1.1和192.168.2.2产生，分别。这两个前缀仅反射对无委托人S3PE1。在这种情况下，所有客户端必须充分地网状连接。通常，在此特定的方案MCIDs确实没有意义(同一种行为可能用单个集群完成)，但是可以仍然使用他们，如果要有路由的不同的集群列表从不同的邻居。

注意：是不可能的启用簇内部反射(特定集群或所有集群的)，当集群间反射禁用时。

团星列表和环路预防

当RR反射前缀时，添加集群ID到可选非传递属性CLUSTER_LIST。并且它设置可选非传递属性ORIGINATOR_ID为对等体的路由器ID，那通告前缀对RR。

当使用时MCIDs，并且RR反射前缀，使用配置的集群ID通告该前缀对RR的对等体。如果该对等体不安排特定集群ID配置，使用全局集群ID。

请参见一些示例。RR有启用的路由反映所有表。全局集群ID是172.16.3.3，集群ID 192.168.1.1和192.168.2.2分别设置为在站点1和站点2的观点扫描器(以上参考的拓扑图)。

```
RR#show ip bgp cluster-idsGlobal cluster-id: 172.16.3.3 (configured: 0.0.0.0)BGP client-to-
client reflection:          Configured      Used all (inter-cluster and intra-cluster): ENABLED
intra-cluster:              ENABLED          ENABLEDList of cluster-ids:Cluster-id
#-neighbors C2C-rfl-CFG C2C-rfl-USE192.168.1.1          2 ENABLED          ENABLED192.168.2.2
2 ENABLED          ENABLED
```

在客户端和无委托人之间的反射

```
S2PE3#show ip bgp 172.16.1.1BGP routing table entry for 172.16.1.1/32, version 2Paths: (1
available, best #1, table default, RIB-failure(17)) Not advertised to any peer Refresh Epoch 1
Local 10.0.10.2 (metric 20) from 10.0.70.1 (172.16.3.3) Origin IGP, metric 0, localpref
100, valid, internal, best Originator: 172.16.1.1, Cluster list: 192.168.1.1 rx
pathid: 0, tx pathid: 0x0
```

```
S2PE3#show ip bgp 172.16.4.4BGP routing table entry for 172.16.4.4/32, version 4Paths: (1
available, best #1, table default, RIB-failure(17)) Not advertised to any peer Refresh Epoch 1
Local 10.0.40.2 (metric 20) from 10.0.70.1 (172.16.3.3) Origin IGP, metric 0, localpref 100,
valid, internal, best Originator: 172.16.4.4, Cluster list: 192.168.2.2 rx pathid: 0, tx
pathid: 0x0
```

无委托人S2PE3集群172.16.1.1/32产生的recevies前缀192.168.1.1 -集群ID 192.168.1.1被添加到集

群列表。它也接收集群172.16.4.4/32产生的前缀192.168.2.2 -集群ID 192.168.2.2被添加到集群列表

。

```
S1PE1#show ip bgp 172.16.6.6BGP routing table entry for 172.16.6.6/32, version 5Paths: (1
available, best #1, table default, RIB-failure(17)) Not advertised to any peer Refresh Epoch 1
Local 10.0.70.2 (metric 20) from 10.0.10.1 (172.16.3.3) Origin IGP, metric 0, localpref
100, valid, internal, best Originator: 172.16.6.6, Cluster list: 172.16.3.3 rx pathid:
0, tx pathid: 0x0
```

客户端S1PE1无委托人172.16.6.6/32产生的接收前缀-全局集群ID 172.16.3.3被添加到集群列表。

簇内部反射

```
S1PE1#show ip bgp 172.16.6.6BGP routing table entry for 172.16.6.6/32, version 5Paths: (1
available, best #1, table default, RIB-failure(17)) Not advertised to any peer Refresh Epoch 1
Local 10.0.70.2 (metric 20) from 10.0.10.1 (172.16.3.3) Origin IGP, metric 0, localpref
100, valid, internal, best Originator: 172.16.6.6, Cluster list: 172.16.3.3 rx pathid:
0, tx pathid: 0x0
```

S1PE2属于集群192.168.1.1，并且接收加前缀也属于集群192.168.1.1的S1PE1产生的172.16.1.1/32。团星ID 192.168.1.1被添加到集群列表。

集群间反射

```
S2PE1#show ip bgp 172.16.1.1/32BGP routing table entry for 172.16.1.1/32, version 4Paths: (1
available, best #1, table default, RIB-failure(17)) Not advertised to any peer Refresh Epoch 1
Local 10.0.10.2 (metric 20) from 10.0.40.1 (172.16.3.3) Origin IGP, metric 0, localpref
100, valid, internal, best Originator: 172.16.1.1, Cluster list: 192.168.1.1 rx
pathid: 0, tx pathid: 0x0 S1PE1#sh ip bgp 172.16.4.4/32BGP routing table entry for
172.16.4.4/32, version 4Paths: (1 available, best #1, table default, RIB-failure(17)) Not
advertised to any peer Refresh Epoch 1 Local 10.0.40.2 (metric 20) from 10.0.10.1
(172.16.3.3) Origin IGP, metric 0, localpref 100, valid, internal, best Originator:
172.16.4.4, Cluster list: 192.168.2.2 rx pathid: 0, tx pathid: 0x0
```

S2PE1属于集群192.168.2.2和集群172.16.1.1/32产生的接收前缀192.168.1.1 -集群ID设置为192.168.1.1。

S1PE1属于集群192.168.1.1和集群172.16.4.4/32产生的接收前缀192.168.2.2 -集群ID设置为192.168.2.2。

MCIDs和环路预防

如果路由器接收为集群列表的前缀的更新包含路由器的自己的集群ID，更新丢弃。如果包含其中任一已配置的集群ID使用的MCIDs，请更新(全局或每相邻)将丢弃。

参考

[RFC 4456](#)

[BGP多个团星ID](#)