边界网关协议(BGP)最佳路由反映

目录

简介

本文描述如何可以影响路由，当有一个或更多路由反射器(RR)时在避免在iBGP路由器之间的一全网状的网络。

背景信息

在BGP最佳路径选择算法的第8步将选有最低的IGP权值的路径对BGP下一跳。因此，如果所有步骤，在第8步是相等前，然后第8步可以是在什么的决定因素最佳路径在RR。IGP费用从RR到通告的iBGP路由器然后取决于RR的安置。默认情况下， RR只通告最佳路径给其客户端。根据放置的地方RR， IGP开销对通告路由器可以是更小或更大的。如果路径的所有IGP费用是相同的，则它是可能去结束至有的通告路由器的同分决赛最低的BGP路由器号。

网络图

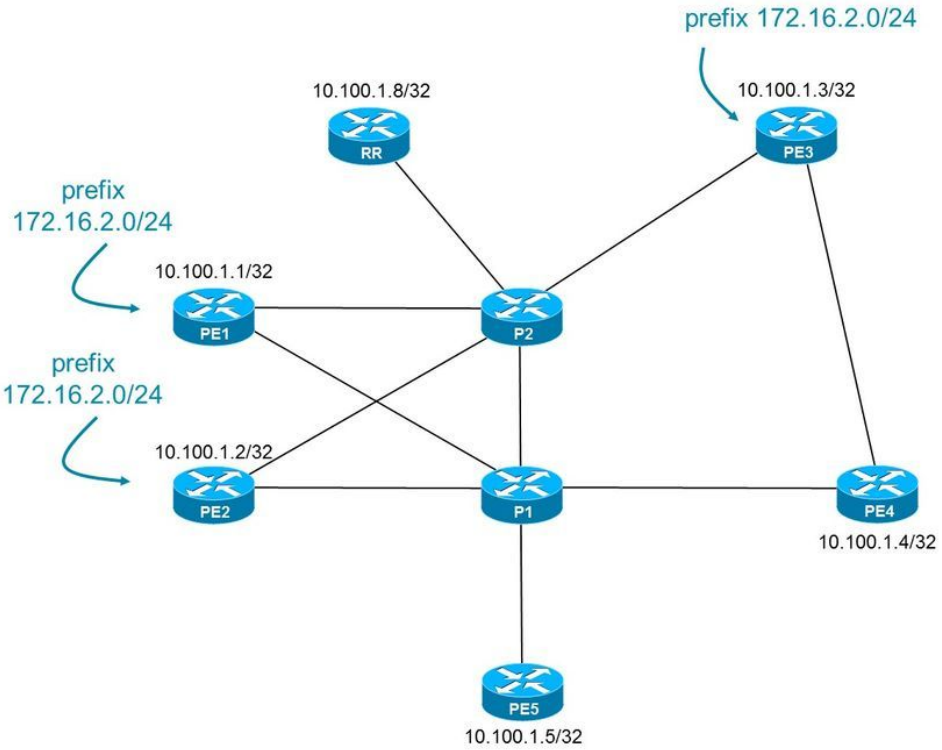

路由器PE1、PE2和PE3通告前缀172.16.2.0/24。如果链路的所有IGP费用是相同的那么RR将看到从PE1、PE2和PE3的路径以IGP费用的2。最后，因为有更低的BGP路由器号， RR选择从PE1的路径作为最佳。这是在BGP最佳路径选择算法的第11步。结果是所有PE路由器，包括PE4，将选择PE1作为前缀的172.16.2.0/24出口PE路由器。从PE4立场，更短的IGP路径向所有出口PE路由器是路径对PE3，以IGP费用的1。IGP开销对所有其他PE路由器是2。对于许多网络，传输数据流事实通过在最短的方式的穿透网络，是重要。这叫作热土豆路由。

有RR的另一个可能的来源选择了从PE1的最佳路径。如果在镜像，链路P2-PE3的内部网关Protool (IGP)费用是10，并且所有其他链路仍然有IGP费用的1，则RR不会选择从PE3的路径作为最佳，即使PE3有最低的BGP路由器号。

如果此网络的管理员要有热土豆路由，则机制一定到位，以便，当有在网络时的RR，入口路由器能仍然了解路径到iBGP网络的最接近的出口路由器。BGP功能添加路径能达到此。然而与该功能，RR和了解功能的边界路由器必须有更多最近代码。使用BGP最佳路由反映功能，这不是需求。此功能将允许RR发送最佳路径到进入BGP边界路由器，根据什么RR认为是最佳路径从该进入BGP路由器的角度。

将允许热土豆路由的另一个解决方案，当RR配置时，是RR的轴向安置。这些RR不是专用的RR，只运行BGP和IGP。这些轴向RR在转发路径和安置在网络，以便他们有他们自己的套RR客户端，因此他们能反射最佳路径对每个RR客户端，也是最佳路径从该RR客户端的方面。

如此镜像所显示， RR在网络安置，以便他们他们能服务的少数附近的RR客户端。由于网络设计，RR客户端接收是最佳路径从他们的观点的最佳路径，从RR，以便可以有在网络的棘手问题路由。

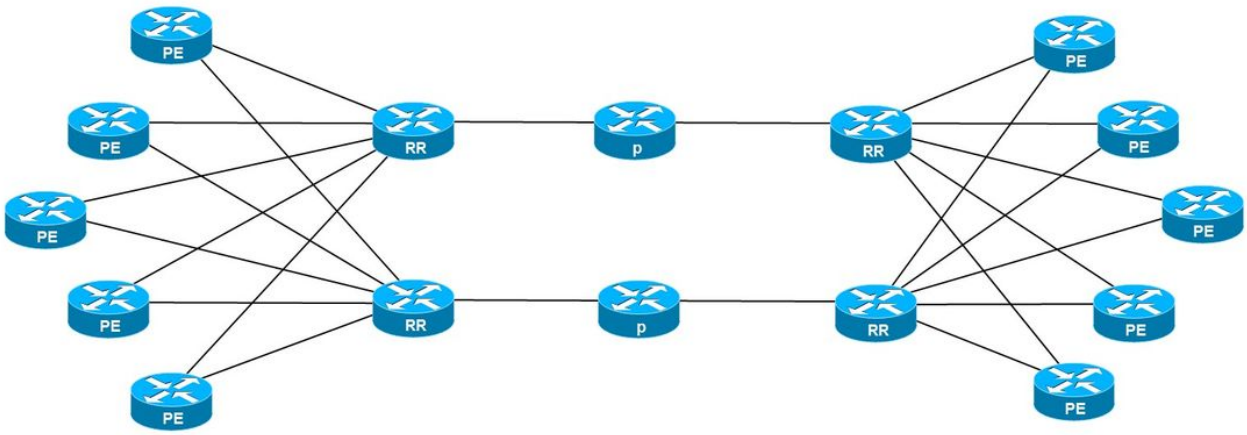

理论

BGP最佳路由反映在IETF草案草稿ietf IDR BGP最佳路由反映描述。

BGP最佳路由反映解决方案允许RR发送特定最佳路径到一个特定BGP边界路由器。RR能选择发送不同的最佳路径到另外BGP边界路由器或套边界路由器。边界路由器必须是RR的RR客户端。目标是每个进入BGP边界路由器能有一个不同的退出或出口BGP路由器同一个前缀的。如果进入边界路由器能总是转发对的数据流关闭和退出路由器，则这允许热土豆路由。

问题是RR通常只发送同一最佳路径到每个BGP边界路由器，防止热土豆路由。为了解决此，您需要RR能计算同一个前缀的不同的最佳路径依靠进入BGP边界路由器。对RR的最佳路径计算根据进入BGP边界路由器的位置进行，因此， RR从进入边界路由器的角度将执行BGP最佳路径计算。能只执行此的RR是有网络结构的完整图片从IGP方面RR和进入边界路由器查找的RR。对于此需求满足，IGP必须是链路状态路由协议。

在那种情况下， RR能运行Shortest Path First (SPF)计算以进入边界路由器作为树的根和计算费用到其他路由器。这样，从进入边界路由器的费用其他出口边界路由器的将知道。此特殊SPF计算用另一个路由器作为根，指反向SPF (rSPF)。如果RR了解从所有BGP边界路由器的所有BGP路径这可能只执行。尽管有RR客户端，可能有运行的许多rSPFs。这在RR将某种程度增加CPU负载。

解决方案允许根据BGP最佳路径选择算法的最佳路径计算，将导致选择最佳路径的RR从进入边界路由器的角度RR发送路径。这意味着最佳路径根据最短的IGP将选择开销对BGP下一跳。解决方案也允许根据若干被配置的策略将被选择的最佳路径。进入边界路由器可能选择他们的最佳路径若干被配置的策略和不根据最低的IGP费用。解决方案准许RR实现最佳路由反映在开销的IGP (网络的位置)或在若干被配置的策略或者两个。如果两个配置，则首先运用策略基于IGP的最佳路由反映在剩余的路径然后将发生。

IOS-XR实施

IOS-XR实施允许三个根节点rSPF计算。如果有许多RR客户端在一个更新组，则没有对一个rSPF计算的需要每个RR客户端，如果那些RR客户端将有同一个策略和同样IGP费用到不同的出口BGP边

界路由器。此后者通常意味着RR客户端是并行定位的(可能在同样POP)。如果那是实际情形，没有需要配置每个RR客户端作为根。IOS-XR实施准许配置三，主要，附属和第三根，每套RR客户端，为冗余目的。为了使BGP奥尔功能应用于所有RR客户端，必须配置该RR客户端是奥尔策略组的一部分。

BGP奥尔功能每个地址家族是启用的。

需要链路状态协议。它可以是OSPF或IS-IS。

IOS XP只实现根据IGP的BGP奥尔功能开销对BGP下一跳和没根据若干被配置的策略。

有同一出局策略的BGP对等体在同一个更新组安置。这通常是iBGP的论点在RR。当feature bgp奥尔是启用的，然后从不同的奥尔组的对等体用不同的更新组。这是逻辑，因为从RR发送的更新到RR客户端用不同的BGP奥尔组将是不同的，因为BGP最佳路径是不同的。

rSPF计算的结果在数据库存储。

ORRSPF是为BGP奥尔功能是需要的在IOS-XR的新的组件。ORRSPF保重：

1. 收集链路状态信息和维护链路状态数据库
2. 运行rSPFs和维护SPTs，每个策略组
3. 下载前缀从SPT到与权值的RIB

数据库获得其链路状态信息直接地从连接状态的IGP或从BGP-LS。

rSPF计算导致显示最短路径的结构从RR客户端其他路由器在区域/级。

暂停在结构的每个路由器的路由在一张特殊RIB表里存储奥尔组策略的和每AFI/SAFI。此表是由RSI创建的。表由与原始根源的rSPFs计算的路由填充作为根。如果原始根源变得未提供，则附属根是根并且填充路由在奥尔RIB表里。同样适用于第三根。

配置

需要的最小配置：

1. 奥尔需要为BGP地址族被启用，特定组的BGP邻居
2. 对于每个组BGP邻居，需要配置至少一个根。随意地，可以配置一个附属和第三根。
3. 奥尔路由的再分配从IGP的到BGP里需要被启用。

配置示例

如第一个镜像所显示， RR是有BGP奥尔功能的一个IOS-XR路由器。

所有其他路由器运行IOS。这些路由器没有BGP奥尔功能。

PE1、PE2和PE3通告在AFI/SAFI 1/1 (IPv4单播)的前缀172.16.2.0/24。RR均等地是接近PE1及PE2比与PE3。所有链路的IGP费用是1。此前缀的最佳路径是那个与R1作为下个跳跃由于最低的BGP路由器号。

```
RP/0/0/CPU0:RR#show bgp ipv4 unicast 172.16.2.0/24 bestpath-compare
BGP routing table entry for 172.16.2.0/24
```

```
Versions:
 Process           bRIB/RIB  SendTblVer
 Speaker                34          34
Last Modified: Mar  7 20:29:48.156 for 11:36:44
Paths: (3 available, best #1)
 Advertised to update-groups (with more than one peer):
   0.3
 Path #1: Received by speaker 0
 Advertised to update-groups (with more than one peer):
   0.3
 Local, (Received from a RR-client)
   10.100.1.1 (metric 3) from 10.100.1.1 (10.100.1.1)
     Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
     Received Path ID 0, Local Path ID 1, version 34
     best of local AS, Overall best
  Path #2: Received by speaker 0
 Not advertised to any peer
 Local, (Received from a RR-client)
   10.100.1.2 (metric 3) from 10.100.1.2 (10.100.1.2)
     Origin IGP, metric 0, localpref 100, valid, internal, add-path
     Received Path ID 0, Local Path ID 6, version 33
     Higher router ID than best path (path #1)
  Path #3: Received by speaker 0
 ORR bestpath for update-groups (with more than one peer):
   0.1
 Local, (Received from a RR-client)
   10.100.1.3 (metric 5) from 10.100.1.3 (10.100.1.3)
     Origin IGP, metric 0, localpref 100, valid, internal, add-path
     Received Path ID 0, Local Path ID 7, version 34
     Higher IGP metric than best path (path #1)
```

PE4将接收有PE1的路径作为下个跳跃。因此，没有PE4的热土豆路由。

如果要有热土豆路由在PE4，则PE1、PE2和PE3做通告的前缀的(例如前缀172.16.2.0/24)，则PE1应该有PE3作为出口点。这意味着在PE4的路径应该是那个与PE3作为下个跳跃。我们能做RR发送有下个跳跃的PE3路由到与此奥尔配置的PE4。

```
router ospf 1
distribute bgp-ls
 area 0
 interface Loopback0
 !
 interface GigabitEthernet0/0/0/0
  network point-to-point
 !
 !
!

router bgp 1
 address-family ipv4 unicast
 optimal-route-reflection ipv4-orr-group 10.100.1.4
 !
 address-family vpnv4 unicast
 !
 neighbor 10.100.1.1
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  route-reflector-client
 !
 !
 neighbor 10.100.1.2
```

```
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  route-reflector-client
 !
 !
 neighbor 10.100.1.3
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  route-reflector-client
 !
 !
 neighbor 10.100.1.4
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group
   route-reflector-client
 !
 !
 neighbor 10.100.1.5
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  route-reflector-client
 !
 !
!
```

如果IGP是IS-IS ：

```
router isis 1
net 49.0001.0000.0000.0008.00
 distribute bgp-ls
address-family ipv4 unicast
metric-style wide
!
interface Loopback0
address-family ipv4 unicast
!
!
interface GigabitEthernet0/0/0/0
address-family ipv4 unicast
!
!
!
```

Note:地址家族连接状态不需要被配置，全局或在BGP邻居下。

在根路由器的MPLS流量工程

RR需要查找在IGP数据库的被配置的根地址，为了运行rSPF。在ISIS中， router-ID是存在ISIS数据库。对于OSPF，没有router-ID当前在OSPF LSA。解决方案将安排根路由器通告匹配在RR的多协议标签交换(MPLS) TE router-ID被配置的根地址。

OSPF，使BGP的根路由器需要更多的配置奥尔工作。一种最小MPLS TE配置在所有根路由器必要为了通告此MPLS TE router-ID。确切的最小集命令取决于根路由器的操作系统。在根路由器的MPLS TE配置需要有被启用的MPLS TE的最小配置，以便OSPF通告在不透明区域LSA (类型10)的MPLS TE路由器ID。

一旦RR有与匹配被配置的根路由器地址的MPLS TE router-ID的不透明区域LSA， rSPF能运行，并且在RR的BGP能通告最佳路由。

为在根路由器的OSPF需要的最小配置，如果它是IOS路由器是：

```
!
interface GigabitEthernet0/2
 ip address 10.1.34.4 255.255.255.0
 ip ospf network point-to-point
mpls traffic-eng tunnels
!

router ospf 1
mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 router-id 10.200.1.155
 network 10.0.0.0 0.255.255.255 area 0
!
```

注意那：

- MPLS TE在特定OSPF区域被启用
- 配置MPLS TE router-ID匹配在RR的被配置的根地址
- MPLS TE在至少一个接口被配置
- 没有需要安排RSVP-TE被配置
- 没有需要有在其他路由器配置的MPLS TE在区域

为在根路由器的OSPF需要的最小配置，如果它是IOS-XR路由器是：

```
!
router ospf 1
 router-id 5.6.7.8
 area 0
 mpls traffic-eng
  interface Loopback0
 !
 interface GigabitEthernet0/0/0/0
  network point-to-point
 !
 !
mpls traffic-eng router-id 10.100.1.11
!
mpls traffic-eng
!
```

如果上述配置是到位在根路由器，则RR应该有在OSPF数据库的MPLS TE router-ID。

```
RP/0/0/CPU0:RR#show ospf 1 database


            OSPF Router with ID (10.100.1.99) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.12.1       10.1.12.1       1297        0x8000002b 0x006145 3
10.100.1.2      10.100.1.2      646         0x80000025 0x00fb6f 7
10.100.1.3      10.100.1.3      1693        0x80000031 0x003ba9 5
10.100.1.99     10.100.1.99     623         0x8000001e 0x00ade1 3
```

```
10.200.1.155    10.200.1.155    28          0x80000002 0x009b2e 5

               Type-10 Opaque Link Area Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Opaque ID
1.0.0.0         10.200.1.155    34          0x80000001 0x00a1ad        0

1.0.0.3         10.200.1.155    34          0x80000001 0x0057ff        3

RP/0/0/CPU0:RR#show ospf 1 database opaque-area adv-router 10.200.1.155


           OSPF Router with ID (10.100.1.99) (Process ID 1)

               Type-10 Opaque Link Area Link States (Area 0)

  LS age: 184
 Options: (No TOS-capability, DC)
 LS Type: Opaque Area Link
 Link State ID: 1.0.0.0
 Opaque Type: 1
 Opaque ID: 0
 Advertising Router: 10.200.1.155
 LS Seq Number: 80000001
 Checksum: 0xa1ad
 Length: 28

   MPLS TE router ID : 10.100.1.4

   Number of Links : 0

 LS age: 184
 Options: (No TOS-capability, DC)
 LS Type: Opaque Area Link
 Link State ID: 1.0.0.3
 Opaque Type: 1
 Opaque ID: 3
 Advertising Router: 10.200.1.155
 LS Seq Number: 80000001
 Checksum: 0x57ff
 Length: 132

   Link connected to Point-to-Point network
     Link ID : 10.100.1.3      (all bandwidths in bytes/sec)
     Interface Address : 10.1.34.4
     Neighbor Address : 10.1.34.3
     Admin Metric : 1
     Maximum bandwidth : 125000000
     Maximum reservable bandwidth global: 0
     Number of Priority : 8
     Priority 0 :                    0  Priority 1 :                    0
     Priority 2 :                    0  Priority 3 :                    0
     Priority 4 :                    0  Priority 5 :                    0
     Priority 6 :                    0  Priority 7 :                    0
     Affinity Bit : 0
     IGP Metric : 1

   Number of Links : 1
```

注意MPLS TE router-ID (10.100.1.4)和OSPF路由器ID是不同的。

PE4有PE3作为前缀的下个跳跃(与对下个跳跃的正确的IGP权值) ：

```
PE4#show bgp ipv4 unicast 172.16.2.0
BGP routing table entry for 172.16.2.0/24, version 37
Paths: (1 available, best #1, table default)
 Not advertised to any peer
 Refresh Epoch 1
 Local
   10.100.1.3 (metric 2) from 10.100.1.8 (10.100.1.8)
     Origin IGP, metric 0, localpref 100, valid, internal, best
     Originator: 10.100.1.3, Cluster list: 10.100.1.8
     rx pathid: 0, tx pathid: 0x0
```

PE5仍然有PE1作为前缀的下个跳跃(与对下个跳跃的正确的IGP权值) ：

```
PE5#show bgp ipv4 unicast 172.16.2.0/24
BGP routing table entry for 172.16.2.0/24, version 13
Paths: (1 available, best #1, table default)
 Not advertised to any peer
 Refresh Epoch 1
 Local
   10.100.1.1 (metric 3) from 10.100.1.8 (10.100.1.8)
     Origin IGP, metric 0, localpref 100, valid, internal, best
     Originator: 10.100.1.1, Cluster list: 10.100.1.8
     rx pathid: 0, tx pathid: 0x0
```

故障排除

验证在RR的前缀：

```
RP/0/0/CPU0:RR#show bgp ipv4 unicast 172.16.2.0
BGP routing table entry for 172.16.2.0/24
Versions:
 Process           bRIB/RIB  SendTblVer
 Speaker                19          19
Last Modified: Mar  7 16:41:20.156 for 03:07:40
Paths: (3 available, best #1)
 Advertised to update-groups (with more than one peer):
   0.3
 Path #1: Received by speaker 0
 Advertised to update-groups (with more than one peer):
   0.3
 Local, (Received from a RR-client)
   10.100.1.1 (metric 3) from 10.100.1.1 (10.100.1.1)
     Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
     Received Path ID 0, Local Path ID 1, version 14
 Path #2: Received by speaker 0
 Not advertised to any peer
 Local, (Received from a RR-client)
   10.100.1.2 (metric 3) from 10.100.1.2 (10.100.1.2)
     Origin IGP, metric 0, localpref 100, valid, internal, add-path
     Received Path ID 0, Local Path ID 4, version 14
 Path #3: Received by speaker 0
 ORR bestpath for update-groups (with more than one peer):
   0.1
 Local, (Received from a RR-client)
   10.100.1.3 (metric 5) from 10.100.1.3 (10.100.1.3)
     Origin IGP, metric 0, localpref 100, valid, internal, add-path
     Received Path ID 0, Local Path ID 5, version 19
```

注意添加路径被添加了到其他最非好的路径，因此他们可以做通告，除最佳路径以外。添加路径功能没有使用在RR和其客户端之间：路径没有用路径标识符通告。

验证路由(仍然)做通告给特定BGP邻居。

对相邻PE4，下个跳跃是前缀的172.16.2.0/24 PE3 ：

```
RP/0/0/CPU0:RR#show bgp ipv4 unicast neighbors 10.100.1.4 advertised-routes
Network            Next Hop        From            AS Path
172.16.1.0/24      10.100.1.5      10.100.1.5      i
172.16.2.0/24      10.100.1.3      10.100.1.3      i

Processed 2 prefixes, 2 paths
```

对相邻PE5，下个跳跃是前缀的172.16.2.0/24 PE1 ：

```
RP/0/0/CPU0:RR#show bgp ipv4 unicast neighbors 10.100.1.5 advertised-routes
Network            Next Hop        From            AS Path
172.16.1.0/24      10.100.1.8      10.100.1.5      i
172.16.2.0/24      10.100.1.1      10.100.1.1      i
```

邻接10.100.1.4到位在其自己的更新组由于奥尔策略：

```
RP/0/0/CPU0:RR#show bgp ipv4 unicast update-group
Update group for IPv4 Unicast, index 0.1:
 Attributes:
   Neighbor sessions are IPv4
   Internal
   Common admin
   First neighbor AS: 1
   Send communities
   Send GSHUT community if originated
   Send extended communities
   Route Reflector Client
   ORR root (configured): ipv4-orr-group; Index: 0
   4-byte AS capable
   Non-labeled address-family capable
   Send AIGP
   Send multicast attributes
   Minimum advertisement interval: 0 secs
 Update group desynchronized: 0
 Sub-groups merged: 0
 Number of refresh subgroups: 0
 Messages formatted: 8, replicated: 8
 All neighbors are assigned to sub-group(s)
   Neighbors in sub-group: 0.1, Filter-Groups num:1
    Neighbors in filter-group: 0.3(RT num: 0)
     10.100.1.4

Update group for IPv4 Unicast, index 0.3:
 Attributes:
   Neighbor sessions are IPv4
   Internal
   Common admin
   First neighbor AS: 1
   Send communities
   Send GSHUT community if originated
   Send extended communities
   Route Reflector Client
   4-byte AS capable
```

```
   Non-labeled address-family capable
   Send AIGP
   Send multicast attributes
   Minimum advertisement interval: 0 secs
 Update group desynchronized: 0
 Sub-groups merged: 1
 Number of refresh subgroups: 0
 Messages formatted: 12, replicated: 42
 All neighbors are assigned to sub-group(s)
   Neighbors in sub-group: 0.3, Filter-Groups num:1
    Neighbors in filter-group: 0.1(RT num: 0)
     10.100.1.1               10.100.1.2               10.100.1.3
10.100.1.5
```

database命令显示的orrspf显示ORR组其根，

```
RP/0/0/CPU0:RR#show orrspf database

ORR policy: ipv4-orr-group, IPv4, RIB tableid: 0xe0000012
Configured root: primary: 10.100.1.4, secondary: NULL, tertiary: NULL
Actual Root: 10.100.1.4


Number of mapping entries: 1
```

与详细资料关键字的同一个命令彼此提供rSPF路由器/前缀的根的费用在同样OSPF区域：

```
RP/0/0/CPU0:RR#show orrspf database detail

ORR policy: ipv4-orr-group, IPv4, RIB tableid: 0xe0000012
Configured root: primary: 10.100.1.4, secondary: NULL, tertiary: NULL
Actual Root: 10.100.1.4

Prefix                                        Cost
10.100.1.6                                    2
10.100.1.1                                    3
10.100.1.2                                    3
10.100.1.3                                    2
10.100.1.4                                    0
10.100.1.5                                    3
10.100.1.7                                    3
10.100.1.8                                    4

Number of mapping entries: 9
```

表id由RSI分配为奥尔组和为AFI/SAFI ：

```
RP/0/0/CPU0:RR#show rsi table-id 0xe0000012

TBL_NAME=ipv4-orr-group, AFI=IPv4, SAFI=Ucast TBL_ID=0xe0000012 in VRF=default/0x60000000 in
VR=default/0x20000000
Refcnt=1
VRF Index=4 TCM Index=1
Flags=0x0 LST Flags=(0x0) NULL

RP/0/0/CPU0:RR#show rib tables

Codes: N - Prefix Limit Notified, F - Forward Referenced
      D - Table Deleted, C - Table Reached Convergence

VRF/Table             SAFI  Table ID     PrfxLmt   PrfxCnt TblVersion  N F D C
```

```
default/default        uni   0xe0000000   5000000        22        128  N N N Y
**nVSatellite/default  uni   0xe0000010   5000000         2          4  N N N Y
default/ipv4-orr-grou  uni   0xe0000012   5000000         9         27  N N N Y
default/default        multi 0xe0100000   5000000         0          0  N N N Y
```

根(R4/10.100.1.4)的费用rSPF彼此路由器是相同的象在PE4的**show ip route ospf**看到的费用：

```
PE4#show ip route ospf
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
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 20 subnets, 2 masks
O        10.100.1.1/32 [110/3] via 10.1.7.6, 4d05h, GigabitEthernet0/1
O        10.100.1.2/32 [110/3] via 10.1.7.6, 4d05h, GigabitEthernet0/1
O        10.100.1.3/32 [110/2] via 10.1.8.3, 4d06h, GigabitEthernet0/2
O        10.100.1.5/32 [110/3] via 10.1.7.6, 4d05h, GigabitEthernet0/1
O        10.100.1.6/32 [110/2] via 10.1.7.6, 4d05h, GigabitEthernet0/1
O        10.100.1.7/32 [110/3] via 10.1.8.3, 4d06h, GigabitEthernet0/2
                       [110/3] via 10.1.7.6, 4d05h, GigabitEthernet0/1
O        10.100.1.8/32 [110/4] via 10.1.8.3, 4d05h, GigabitEthernet0/2
                       [110/4] via 10.1.7.6, 4d05h, GigabitEthernet0/1
```

BGP奥尔组的RIB ：

```
RP/0/0/CPU0:RR#show route afi-all safi-all topology ipv4-orr-group

IPv4 Unicast Topology ipv4-orr-group:
-------------------------------------

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
       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 - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR, l - LISP
       A - access/subscriber, a - Application route
       M - mobile route, r - RPL, (!) - FRR Backup path

Gateway of last resort is not set

o    10.100.1.1/32 [255/3] via 0.0.0.0, 14:14:52, Unknown
o    10.100.1.2/32 [255/3] via 0.0.0.0, 14:14:52, Unknown
o    10.100.1.3/32 [255/2] via 0.0.0.0, 00:04:53, Unknown
o    10.100.1.4/32 [255/0] via 0.0.0.0, 14:14:52, Unknown
o    10.100.1.5/32 [255/3] via 0.0.0.0, 14:14:52, Unknown
o    10.100.1.6/32 [255/2] via 0.0.0.0, 14:14:52, Unknown
o    10.100.1.7/32 [255/3] via 0.0.0.0, 14:14:52, Unknown
o    10.100.1.8/32 [255/4] via 0.0.0.0, 14:14:52, Unknown

RP/0/0/CPU0:RR#show rsi table name ipv4-orr-group

VR=default:
```

```
TBL_NAME=ipv4-orr-group, AFI=IPv4, SAFI=Ucast TBL_ID=0xe0000012 in VRF=default/0x60000000 in
VR=default/0x20000000
Refcnt=1
VRF Index=4 TCM Index=1
Flags=0x0 LST Flags=(0x0) NULL
```

neighbor命令的show bgp显示对等体是否是ORR根：

```
RP/0/0/CPU0:RR#show bgp neighbor 10.100.1.4

BGP neighbor is 10.100.1.4
 Remote AS 1, local AS 1, internal link
 Remote router ID 10.100.1.4
 Cluster ID 10.100.1.8
 BGP state = Established, up for 01:17:41
 NSR State: None
 Last read 00:00:52, Last read before reset 01:18:30
 Hold time is 180, keepalive interval is 60 seconds
 Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
 Last write 00:00:34, attempted 19, written 19
 Second last write 00:01:34, attempted 19, written 19
 Last write before reset 01:17:43, attempted 19, written 19
 Second last write before reset 01:18:43, attempted 19, written 19
 Last write pulse rcvd  Mar  8 10:20:13.779 last full not set pulse count 12091
 Last write pulse rcvd before reset 01:17:42
 Socket not armed for io, armed for read, armed for write
 Last write thread event before reset 01:17:42, second last 01:17:42
 Last KA expiry before reset 01:17:43, second last 01:18:43
 Last KA error before reset 00:00:00, KA not sent 00:00:00
 Last KA start before reset 01:17:43, second last 01:18:43
 Precedence: internet
 Non-stop routing is enabled
 Multi-protocol capability received
 Neighbor capabilities:
   Route refresh: advertised (old + new) and received (old + new)
   4-byte AS: advertised and received
   Address family IPv4 Unicast: advertised and received
 Received 6322 messages, 0 notifications, 0 in queue
 Sent 5782 messages, 4 notifications, 0 in queue
 Minimum time between advertisement runs is 0 secs
 Inbound message logging enabled, 3 messages buffered
 Outbound message logging enabled, 3 messages buffered

 For Address Family: IPv4 Unicast
 BGP neighbor version 41
 Update group: 0.1 Filter-group: 0.1  No Refresh request being processed
 Route-Reflector Client
 ORR root (configured): ipv4-orr-group; Index: 0
  Route refresh request: received 0, sent 0
 0 accepted prefixes, 0 are bestpaths
 Cumulative no. of prefixes denied: 0.
 Prefix advertised 2, suppressed 0, withdrawn 0
 Maximum prefixes allowed 1048576
 Threshold for warning message 75%, restart interval 0 min
 AIGP is enabled
 An EoR was received during read-only mode
 Last ack version 41, Last synced ack version 0
 Outstanding version objects: current 0, max 2
 Additional-paths operation: None
 Send Multicast Attributes
 Advertise VPNv4 routes enabled with  option
 Advertise VPNv6 routes is enabled with Local with stitching-RT option
```

```
 Connections established 6; dropped 5
 Local host: 10.100.1.8, Local port: 25176, IF Handle: 0x00000000
 Foreign host: 10.100.1.4, Foreign port: 179
 Last reset 01:17:42, due to User clear requested (CEASE notification sent - administrative
reset)
 Time since last notification sent to neighbor: 01:17:42
 Error Code: administrative reset
 Notification data sent:
   None
```

如此镜像所显示，多个套RR客户端被配置

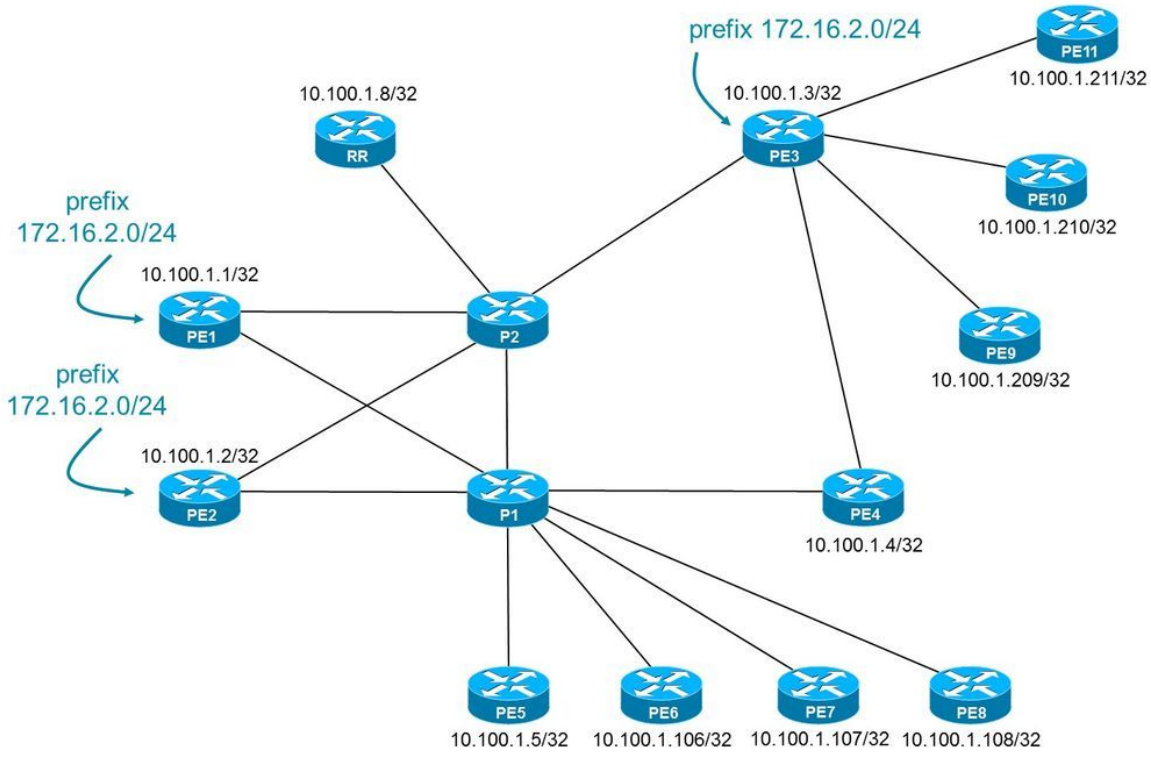

有一套RR客户端连接对PE3，并且别的设置已连接对P1。每集的每个RR客户端是在相等的距离到所有出口BGP边界路由器。

```
router bgp 1
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-1 10.100.1.4 10.100.1.209 10.100.1.210
  optimal-route-reflection ipv4-orr-group-2 10.100.1.5 10.100.1.106 10.100.1.107
 !
…
neighbor 10.100.1.4
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-1
   route-reflector-client
 !
 !
 neighbor 10.100.1.5
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-2
   route-reflector-client
 !
```

```
 !
neighbor 10.100.1.106
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-2
   route-reflector-client
 !
 !
 neighbor 10.100.1.107
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-2
   route-reflector-client
 !
 !
 neighbor 10.100.1.108
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-2
   route-reflector-client
 !
 !
 neighbor 10.100.1.209
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-1
   route-reflector-client
 !
 !
 neighbor 10.100.1.210
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-1   route-reflector-client
 !
 !
 neighbor 10.100.1.211
 remote-as 1
 update-source Loopback0
 address-family ipv4 unicast
  optimal-route-reflection ipv4-orr-group-1
   route-reflector-client
 !
 !
!
```

两个组的orrspf数据库：

```
RP/0/0/CPU0:RR#show orrspf database detail

ORR policy: ipv4-orr-group-1, IPv4, RIB tableid: 0xe0000012
Configured root: primary: 10.100.1.4, secondary: 10.100.1.209, tertiary: 10.100.1.210
Actual Root: 10.100.1.4

Prefix                                        Cost
10.100.1.1                                    3
10.100.1.2                                    3
10.100.1.3                                    2
10.100.1.4                                    0
```

```
10.100.1.5                                3
10.100.1.6                                2
10.100.1.7                                3
10.100.1.8                                4
10.100.1.106                              3
10.100.1.107                              3
10.100.1.108                              3
10.100.1.209                              3
10.100.1.210                              3
10.100.1.211                              3
ORR policy: ipv4-orr-group-2, IPv4, RIB tableid: 0xe0000013
Configured root: primary: 10.100.1.5, secondary: 10.100.1.106, tertiary: 10.100.1.107
Actual Root: 10.100.1.5

Prefix                                    Cost
10.100.1.1                                3
10.100.1.2                                3
10.100.1.3                                4
10.100.1.4                                3
10.100.1.5                                0
10.100.1.6                                2
10.100.1.7                                3
10.100.1.8                                4
10.100.1.106                              3
10.100.1.107                              3
10.100.1.108                              3
10.100.1.209                              5
10.100.1.210                              5
10.100.1.211                              5

Number of mapping entries: 30
```

如果为组原始根源发生故障或不可得到，则附属根将是使用的实际根。在本例中，组ipv4-orr-group-1原始根源是不可得到的。附属根成为实际根：

```
RP/0/0/CPU0:RR#show orrspf database ipv4-orr-group-1

ORR policy: ipv4-orr-group-1, IPv4, RIB tableid: 0xe0000012
Configured root: primary: 10.100.1.4, secondary: 10.100.1.209, tertiary: 10.100.1.210
Actual Root: 10.100.1.209

Prefix                                    Cost
10.100.1.1                                4
10.100.1.2                                5
10.100.1.3                                2
10.100.1.5                                5
10.100.1.6                                4
10.100.1.7                                3
10.100.1.8                                4
10.100.1.106                              5
10.100.1.107                              5
10.100.1.108                              5
10.100.1.209                              0
10.100.1.210                              3
10.100.1.211                              3

Number of mapping entries: 14
```

结论

BGP最佳路由反映(奥尔)是允许在iBGP网络的热土豆路由的功能，当路由反射器存在时，没有对更新的操作系统软件的需要在边界路由器。前提是IGP是链路状态路由协议。