下个跳越不匹配和BGP非活动路由技术说明

目录

[简介](#)
[非活动路由和下个跳越不匹配](#)
[示例拓扑](#)
[显示输出](#)
[抑制非活动路由在BGP配置方面](#)
[添加静态路由匹配下一跳](#)
[ECMP的暗示在下一跳和非活动路由的](#)

简介

本文描述**bgp suppress-inactive**命令如何防止在路由信息库(RIB)没有安装路由的广告;它也描述非活动路由和下个跳越不匹配之间的交互作用。

RIB故障发生，当边界网关协议(BGP)设法安装bestpath前缀到RIB，但是RIB拒绝BGP路由，因为有更加好的管理距离的一个路由在路由表里已经存在。非激活BGP路由是在RIB没有安装的路由，但是安装在BGP表里作为RIB故障。

参考请[抑制非活动路由的BGP广告](#)关于其他详细信息。

非活动路由和下个跳越不匹配

当您使用**bgp suppress-inactive**命令时，非常重要的是您了解下个跳越不匹配影响。

示例拓扑

路由器1 (R1)和路由器2 (R2)有两条并行链路;一连接运行BGP AS 65535，并且另一条链路运行增强的内部网关路由选择协议(EIGRP) AS1。BGP和EIGRP通告在R1的网络10.1.1.1/32。

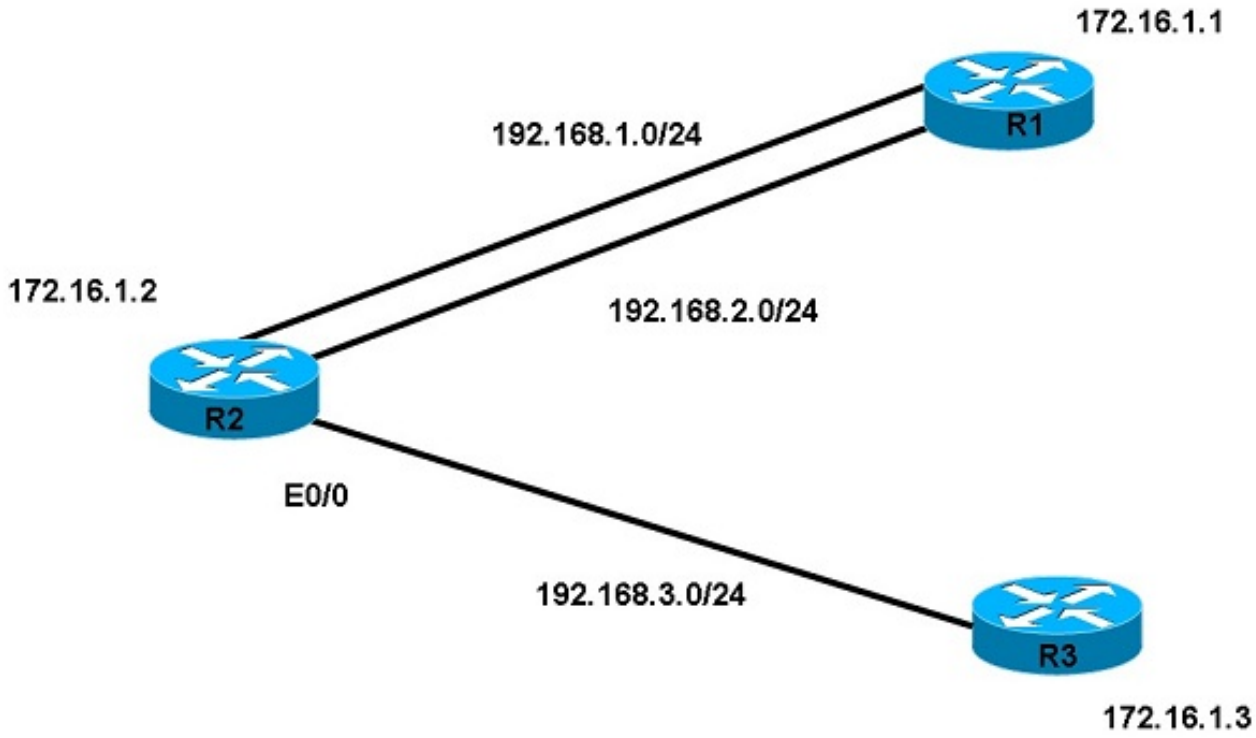

R2在路由表里得知10.1.1.1/32路由通过两个EIGRP和BGP，但是安装仅EIGRP路由由于更短的管理距离。因为BGP路由在R2路由表里没有安装，路由在R2 BGP表里出现作为RIB故障。然而，R2通告BGP路由到路由器3 (R3)不管RIB故障。

显示输出

对于R2，请输入**show ip route命令**为了确定路由表的当前状态在10.1.1.1的，并且输入**show ip bgp命令**为了显示条目在BGP路由表里：

```
Router2#show ip route 10.1.1.1
Routing entry for 10.1.1.1/32
 Known via "eigrp 1", distance 90, metric 409600, type internal
 Last update from 192.168.1.1 on Ethernet0/2, 00:07:15 ago
 Routing Descriptor Blocks:
 * 192.168.1.1, from 192.168.1.1, 00:07:15 ago, via Ethernet0/2
>>>>>>>>NEXT HOP IS LINK A
     Route metric is 409600, traffic share count is 1
     Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
     Loading 1/255, Hops 1

Router2#show ip bgp
BGP table version is 4, local router ID is 172.16.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network            Next Hop         Metric LocPrf Weight Path
r>i10.1.1.1/32        172.16.1.1             0     100      0    I
```

因为它是在R1的一环回检查递归路由下一跳：

```
Router2#show ip route 172.16.1.1
Routing entry for 172.16.1.1/32
 Known via "eigrp 1", distance 90, metric 409600, type internal
 Last update from 192.168.2.1 on Ethernet0/1, 00:07:15 ago
 Routing Descriptor Blocks:
 * 192.168.2.1, from 192.168.2.1, 00:07:15 ago, via Ethernet0/1
>>>>>>>>NEXT HOP IS LINK B
     Route metric is 409600, traffic share count is 1
     Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
     Loading 1/255, Hops 1
```

即使下一跳不匹配， R2通告路由对R3，并且R3得知路由，因为非活动路由没有被抑制：

```
Router3#show ip bgp
BGP table version is 2, local router ID is 172.16.1.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network              Next Hop            Metric   LocPrf   Weight     Path
*> 10.1.1.1/32       172.16.1.2              0                 0          I
```

在BGP配置方面抑制非活动路由

输入**bgp suppress-inactive**命令为了抑制非激活BGP路由。

```
Router2(config)#router bgp 65535
Router2(config-router)#bgp suppress-inactive
Router2(config-router)#end

Router2#show ip bgp neighbors 192.168.3.3 advertised-routes
Total number of prefixes 0
```

> **注意**：**只有当**BGP RIB故障路由的下一跳是与在路由表里，当前安装的同样路由不同下一跳**bgp suppress-inactive**命令抑制RIB失败的路由。

```
Router2#show ip bgp rib-failure
Network        Next Hop      RIB-failure            RIB-NH Matches
10.1.1.1/32    172.16.1.1    Higher admin distance     No   <<<<< No match
```

在RIB-NH匹配列，请注意RIB下一跳不配比。由于10.1.1.1/32路由的下一跳是不同的在EIGRP和BGP，您能抑制RIB失败的路由用**bgp suppress-inactive**命令。

换句话说，如果next-hop in路由表匹配BGP下一跳， **bgp suppress-inactive**命令不再抑制。那含义R3开始再接收10.1.1.1/32路由，即使它是失败的RIB。

添加静态路由匹配下一跳

添加前缀的静态路由为了匹配其与BGP:通告的下一跳的next-hop in RIB

```
Router2(config)#ip route 10.1.1.1 255.255.255.255 192.168.2.1

Router2#show ip bgp rib-failure
Network        Next Hop      RIB-failure            RIB-NH Matches
10.1.1.0/24    192.168.2.1   Higher admin distance    Yes  <<<< Next-Hop matches
```

用**bgp suppress-inactive**命令， R2仍然通告路由，并且R3仍然接收路由。

```
Router3#show ip bgp
BGP table version is 6, local router ID is 172.16.1.3
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop              Metric  LocPrf  Weight  Path
*> 10.1.1.0/24     172.16.1.2             0                 1      i
```

要汇总， **bgp suppress-inactive**命令让BGP抑制非活动路由广告对邻居，只有当路由在与更加好的管理距离的路由表里已经安装，并且，只有当比同一个路由的BGP下一跳有一不同的下一跳。

ECMP的暗示在下一跳和非活动路由的

在前一个示例中，如果在RIB安装的路由(从EIGRP)是相等代价多重通道的(ECMP)，并且，如果非活动路由被抑制，看到被抑制路由的仅零件。

运行在链路在R1之间和R2的EIGRP。R2了解一套从R1的前缀作为在两以一跳192.168.1.1和192.168.2.1之间的ECMP。例如：

```
R2#sh ip route 10.1.1.1

Routing entry for 10.1.1.1/32
 Known via "eigrp 1", distance 170, metric 40030720, type internal
     Last update from 192.168.1.1 on TenGigabitEthernet0/0/0, 2d02h ago
    Routing Descriptor Blocks:
   *192.168.1.1, from 192.168.1.1, 2d02h ago, via TenGigabitEthernet0/1/0
     Route metric is 40030720, traffic share count is 1
     Total delay is 1200 microseconds, minimum bandwidth is 64 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
      Loading 32/255, Hops 2

    192.168.2.1, from 192.168.2.1, 2d02h ago, viaTenGigabitEthernet0/0/0
     Route metric is 40030720, traffic share count is 1
     Total delay is 1200 microseconds, minimum bandwidth is 64 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
     Loading 32/255, Hops 2

R2#sh ip route 10.1.1.5

Routing entry for 10.1.1.5/32
 Known via "eigrp 1", distance 170, metric 40030720, type internal
     Last update from 192.168.1.1 on TenGigabitEthernet0/0/0, 2d02h ago
    Routing Descriptor Blocks:
  192.168.1.1, from 192.168.1.1, 2d02h ago, via TenGigabitEthernet0/1/0
     Route metric is 40030720, traffic share count is 1
     Total delay is 1200 microseconds, minimum bandwidth is 64 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
      Loading 32/255, Hops 2

 * 192.168.2.1, from 192.168.2.1, 2d02h ago, viaTenGigabitEthernet0/0/0
     Route metric is 40030720, traffic share count is 1
     Total delay is 1200 microseconds, minimum bandwidth is 64 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
     Loading 32/255, Hops 2
```

R2了解同一套从R1的前缀在BGP，并且下个跳越环回在两条链路了解。

```
Router2#show ip bgp
BGP table version is 4, local router ID is 172.16.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
  Network          Next Hop            Metric  LocPrf   Weight   Path
r>i10.1.1.1/32      172.16.1.1             0     100      0        I
r>i10.2.2.2/32      172.16.1.1             0     100       0        I
r>i10.3.3.3/32      172.16.1.1              0     100       0        I
r>i10.4.4.4/32      172.16.1.1              0     100       0        I
r>i10.5.5.5/32     172.16.1.1             0     100       0        I
r>i10.6.6.6/32     172.16.1.1             0     100       0        I
r>i10.7.7.7/32      172.16.1.1             0     100       0        I
r>i10.8.8.8/32      172.16.1.1              0     100       0        I
r>i10.9.9.9/32      172.16.1.1              0     100       0        I
r>i10.10.10.10/32   172.16.1.1              0     100      0        I
```

```
R2#sh ip route 172.16.1.1

Routing entry for 172.16.1.1/32
 Known via "eigrp 1", distance 170, metric 40030720 type internal
  Redistributing via eigrp 109
  Last update from 192.168.1.1 on TenGigabitEthernet0/0/0, 2d02h ago
 Routing Descriptor Blocks:
 * 192.168.1.1, from 192.168.1.1, 2d02h ago, via TenGigabitEthernet0/1/0
     Route metric is 40030720, traffic share count is 1
     Total delay is 1200 microseconds, minimum bandwidth is 64 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
      Loading 32/255, Hops 2

   192.168.2.1, from 192.168.2.1, 2d02h ago, viaTenGigabitEthernet0/0/0
     Route metric is 40030720, traffic share count is 1
     Total delay is 1200 microseconds, minimum bandwidth is 64 Kbit
     Reliability 255/255, minimum MTU 1500 bytes
     Loading 32/255, Hops 2
```

因为下个跳越路由是在同样两条链路的一ECMP，您预计下一跳匹配为在BGP的所有前缀和R2to通告所有对R3。 当您查看RIB-NH匹配输出的列时，一些下个跳越(NH)匹配是，并且其他是不。

```
Router2#sh ip bgp rib-failure
 Network              Next Hop              RIB-failure          RIB-NH Matches
10.1.1.1/32         172.16.1.1         Higher admin distance          Yes
10.2.2.2/32         172.16.1.1        Higher admin distance          Yes
10.3.3.3/32         172.16.1.1         Higher admin distance          Yes
10.4.4.4/32         172.16.1.1         Higher admin distance          Yes
10.5.5.5/32         172.16.1.1         Higher admin distance           No
10.6.6.6/32         172.16.1.1        Higher admin distance          No
10.7.7.7/32         172.16.1.1        Higher admin distance          No
10.8.8.8/32         172.16.1.1        Higher admin distance          No
10.9.9.9/32         172.16.1.1        Higher admin distance          No
10.10.10.10/32      172.16.1.1         Higher admin distance           No
```

有是RIB-NH匹配的所有路由通告对R3;其他被抑制。

```
R3#sh ip bgp
BGP table version is 17, local router ID is 172.16.1.3
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
*> 10.1.1.1/32       172.16.1.2             0        2        1        i
*> 10.2.2.2/32       172.16.1.2             0        2        1        i
*> 10.3.3.3/32       172.16.1.2             0        2        1        i
*> 10.4.4.4/32       172.16.1.2             0        2        1        i
```

在Cisco IOS软件方面， BGP能只选择一下一跳并且通告仅最佳路径有该下一跳的(没有添加PATH、多重通道， BGP好外部或者其它特性)。

当RIB安装目的地的EIGRP路由(请注释*在输出中)时， RIB也许选择其中一个路径作为最佳路径。如果该路径匹配那个BGP下一跳的，报告作为下个跳越匹配的是。

在本例中， RIB选择了192.168.1.1作为10.1.1.1/32网络的下一跳(请注释*在192.168.1.1在从**sh ip route 172.16.1.1**)的输出中，匹配用BGP下个跳跃的路由172.16.1.1;这报告作为在下个跳越匹配的一是。RIB选择了192.168.2.1作为10.1.1.5/32的下一跳，不用BGP下个跳跃的路由配比;这报告作为不在下个跳越不匹配。

总之，只有当抑制非活动路由，下个跳越匹配是重要;如果没有匹配，您看到在RIB-NH匹配列的一n/a标志，并且R2通告所有路由对R3。

```
R3#sh ip bgp
BGP table version is 17, local router ID is 172.16.1.3
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
*> 10.1.1.1/32       172.16.1.2             0        2        1        i
*> 10.2.2.2/32       172.16.1.2             0        2        1        i
*> 10.3.3.3/32       172.16.1.2             0        2        1        i
*> 10.4.4.4/32       172.16.1.2             0        2        1        i
```