了解在单/多宿主环境中使用BGP进行负载共享

目录

简介

先决条件

- 要求
- 使用的组件

背景信息

使用环回地址作为 BGP 邻居的负载共享

- 网络图
- 配置
- 验证
- 故障排除

通过一个本地路由器双重归属到一个网络 服务 提供商 (ISP)时的负载均衡

- 网络图
- 配置
- 验证
- 故障排除

通过多个本地路由器双重归属到一个ISP时的负载均衡

- 网络图
- 配置
- 验证
  - AS11 和 AS10 之间的两条链路都启动时验证
  - 出站流量验证
  - 从 AS 10 到 AS 11 的入站流量验证
  - 当R101-R103链路发生故障时进行验证
- 故障排除

通过单个本地路由器多重归属于两个 ISP 时的负载分摊

- 网络图
- 配置
- 验证
- 故障排除

通过多个本地路由器多宿主到两个 ISP 时的负载平衡

- 网络图
- 配置
- 验证
- 故障排除

相关信息

简介

本文档介绍负载共享，它允许路由器在多个路径之间分配传出和传入流量。

先决条件

要求

在尝试此配置之前，请确保满足以下要求：

- BGP 最佳路径选择算法的知识
- 配置 BGP 的知识

使用的组件

本文档不限于特定的软件和硬件版本。

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

传入和传出流量路径是静态派生的，也可以使用动态协议派生，例如：

- 路由信息协议 (RIP)
- 增强型内部网关路由协议 (EIGRP)
- 开放最短路径优先 (OSPF) 协议

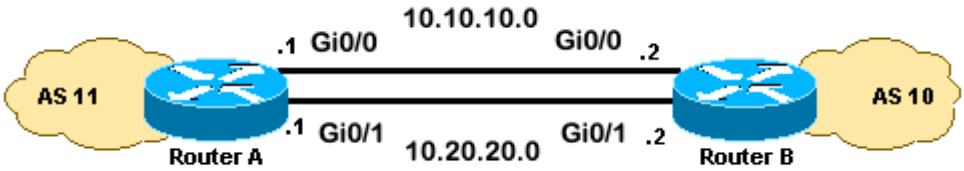
在缺省情况下，边界网关协议（BGP）只选择一个最佳路径而不进行负载均衡。本文档介绍如何使用BGP在不同的场景中执行负载共享。有关负载均衡的其他信息，请参阅负载均衡如何工作。

使用环回地址作为 BGP 邻居的负载共享

此方案显示有多个（最多六个）等成本链路时如何实现负载共享。链路在本地自治系统(AS)的一个路由器中终止，在单宿主BGP环境中的远程AS的另一个路由器中终止。以网络图为例。

网络图

本部分使用以下网络设置：



配置

本部分使用以下配置：

- 路由器A
- 路由器B

路由器A

<#root>

```
interface loopback 0
 ip address 192.168.1.1 255.255.255.255

interface GigabitEthernet0/1
 ip address 10.20.20.1 255.255.255.0

interface GigabitEthernet0/0
 ip address 10.10.10.1 255.255.255.0

router bgp 11
 neighbor 192.168.2.2 remote-as 10
 neighbor 192.168.2.2
```

update-source loopback 0

```
!--- Use the IP address of the loopback interface for TCP connections.

neighbor 192.168.2.2

ebgp-multihop


!--- You must configure ebgp-multihop whenever the external BGP (eBGP) connections are not on the same

router eigrp 12
 network 192.168.1.1 0.0.0.0
 network 10.0.0.0
 no auto-summary
```

路由器B

<#root>

```
interface loopback 0
 ip address 192.168.2.2 255.255.255.255

interface GigabitEthernet0/1
 ip address 10.20.20.2 255.255.255.0

interface GigabitEthernet0/0
 ip address 10.10.10.2 255.255.255.0

router bgp 10
 neighbor 192.168.1.1 remote-as 11
 neighbor 192.168.1.1

update-source loopback 0


!--- Use the IP address of the loopback interface for TCP connections.

neighbor 192.168.1.1

ebgp-multihop


!--- You must configure ebgp-multihop whenever the eBGP connections are not on the same network address

router eigrp 12
 network 192.168.2.2 0.0.0.0
 network 10.0.0.0
 no auto-summary
```

注意：您可以使用静态路由代替路由协议来引入两条到达目的地的等价路径。在本例中，我们使用EIGRP来共享环回信息。

验证

使用本部分可确定配置能否正常运行。

思科 CLI 分析器（仅适用于注册客户）支持某些 show 命令。要查看对 show 命令输出的分析，请使用思科 CLI 分析器。

show ip route命令的输出显示了到192.168.2.2网络的两条路径（通过EIGRP获取）。show ip bgp summary命令显示BGP邻居是使用远程路由器的环回接口构建的。traceroute命令的输出表明负载分布在两个串行链路之间。在此方案中，负载共享是针对每个数据包逐一进行的。可以在串行接口发出 ip route-cache 命令，针对每个目标逐一执行负载共享。您还可以使用 Cisco 快速转发功能来配置基于包和目的地的负载均衡。有关如何配置 Cisco 快速转发的更多信息，请参阅配置 Cisco 快速转发。

<#root>

RouterA#

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
       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, 4 subnets, 2 masks
C        10.10.10.0/24 is directly connected, GigabitEthernet0/0
L        10.10.10.1/32 is directly connected, GigabitEthernet0/0
C        10.20.20.0/24 is directly connected, GigabitEthernet0/1
L        10.20.20.1/32 is directly connected, GigabitEthernet0/1
      192.168.1.0/32 is subnetted, 1 subnets
C        192.168.1.1 is directly connected, Loopback0
      192.168.2.0/32 is subnetted, 1 subnets

D 192.168.2.2 [90/130816] via 10.20.20.2, 00:02:01, GigabitEthernet0/1 [90/130816] via 10.10.10.2, 00:02

RouterA#
```

RouterA#

show ip bgp summary

```
BGP router identifier 192.168.1.1, local AS number 11
```

```
BGP table version is 1, main routing table version 1

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.2.2     4           10      20      20        1    0    0 00:15:05        0
```

RouterA#

traceroute 192.168.2.2

```
Type escape sequence to abort.
Tracing the route to 192.168.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.10.2 2 msec
    10.20.20.2 2 msec
    10.10.10.2 2 msec
RouterA#
```

故障排除

当前没有故障排除此配置的特定可用资料。

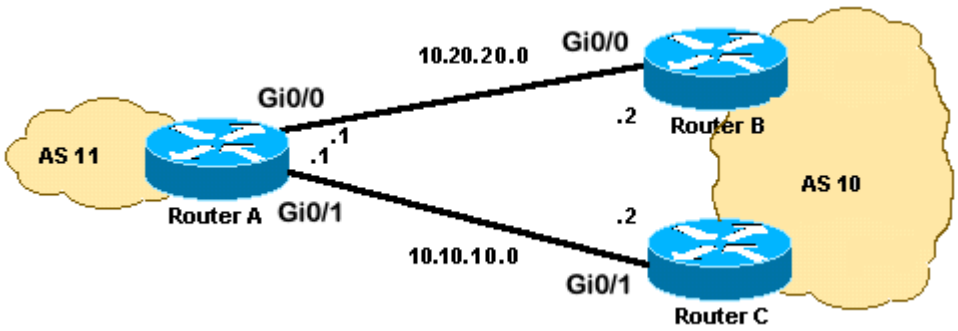
通过一个本地路由器双重归属到一个网络 服务 提供商 (ISP)时的负载均衡

此方案显示当远程 AS 与本地 AS 之间存在多条链路时，如何实现负载共享。这些链路在本地 AS 的一个路由器中终止，并在单宿主 BGP 环境中远程 AS 的多个路由器中终止。我们使用下列网络图为例。

本示例配置采用 maximum-paths 命令。默认情况下，BGP 在通过一个 AS 获知的可能的等成本路径中选择一个最佳路径。但是，您可以更改允许的并行等成本路径的最大数量。为进行此更改，请在 BGP 配置中包括 maximum-paths paths 命令。对于 paths 参数，可使用 1 和 6 之间的数字。

网络图

本部分使用以下网络设置：



配置

本部分使用以下配置：

- 路由器A
- 路由器B
- RouterC

路由器A

<#root>

```
interface Loopback0
 ip address 192.168.1.1 255.255.255.255
!
interface GigabitEthernet0/0
 ip address 10.20.20.1 255.255.255.0
!
!
interface GigabitEthernet0/1
 ip address 10.10.10.1 255.255.255.0
!
!
router bgp 11
 neighbor 10.20.20.2 remote-as 10
 neighbor 10.10.10.2 remote-as 10
 network 192.168.1.1 mask 255.255.255.255
```

maximum-paths 2

```
!--- This command specifies the maximum number of paths to install in the routing table for a specific
```

路由器B

```
interface GigabitEthernet0/2
 ip address 172.16.2.1 255.255.255.0
!
interface GigabitEthernet0/0
 ip address 10.20.20.2 255.255.255.0
!
!
router bgp 10
 neighbor 10.20.20.1 remote-as 11
 network 172.16.2.0 mask 255.255.255.0
```

RouterC

```
interface GigabitEthernet0/2
 ip address 172.16.2.2 255.255.255.0
!
interface GigabitEthernet0/1
 ip address 10.10.10.2 255.255.255.0
!
!
router bgp 10
neighbor 10.10.10.1 remote-as 11
network 172.16.2.0 mask 255.255.255.0
```

验证

使用本部分可确定配置能否正常运行。

思科 CLI 分析器（仅适用于注册客户）支持某些 show 命令。要查看对 show 命令输出的分析，请使用思科 CLI 分析器。

show ip route命令的输出显示，到172.16.2.0网络的两条路径都是通过BGP获取的。traceroute命令的输出表明负载分布在两个串行链路之间。在此方案中，负载共享是针对每个目标逐一进行的。show ip bgp 命令可生成用于 172.16.2.0 网络的有效条目。

<#root>

RouterA#

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
       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, 4 subnets, 2 masks
C        10.10.10.0/24 is directly connected, GigabitEthernet0/1
L        10.10.10.1/32 is directly connected, GigabitEthernet0/1
C        10.20.20.0/24 is directly connected, GigabitEthernet0/0
L        10.20.20.1/32 is directly connected, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets

B 172.16.2.0 [20/0] via 10.20.20.2, 00:08:51 [20/0] via 10.10.10.2, 00:08:51


      192.168.1.0/32 is subnetted, 1 subnets
C        192.168.1.1 is directly connected, Loopback0
```

RouterA#

```
traceroute 172.16.2.2 source loopback0


Type escape sequence to abort.
Tracing the route to 172.16.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.10.2 3 msec
    10.20.20.2 3 msec
    10.10.10.2 3 msec
RouterA#
```

RouterA#

```
show ip bgp


BGP table version is 4, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path

 *m 172.16.2.0/24 10.10.10.2 0 0 10 i *> 10.20.20.2 0 0 10 i

 *>   192.168.1.1/32   0.0.0.0                  0         32768 i
```

故障排除

当前没有故障排除此配置的特定可用资料。

通过多个本地路由器双重归属到一个ISP时的负载均衡

本节演示在通过多个本地路由器与同一个 ISP 有多个连接时如何实现负载平衡。两个 eBGP 对等体在两个单独的本地路由器终止。由于BGP选择从eBGP和内部BGP(iBGP)获知的网络中的单一最佳路径，因此无法实现两个链路上的负载均衡。 在指向 AS 10 的多个路径中进行负载共享是次佳选项。采用这种类型的负载共享，流向特定网络的流量将根据预定义的策略通过两条链路进行传输。此外，如果一条链路发生故障，另一条将充当其备份。

为简单起见，我们假定AS 11的BGP路由策略如下：

- AS 11 接受来自 AS 10 的本地路由以及缺省情况下的其余互联网路由。
- 出局数据流策略如下：
    - 从R101发往Internet的所有流量都会通过R101-R103链路传出。
    - 如果R101-R103链路发生故障，则从R101到Internet的所有流量都将通过R102到AS 10。
    - 同样，从R102发往Internet的所有流量都会通过R102-R104链路。
    - 如果R102-R104链路发生故障，则从R102到Internet的所有流量都将通过R101到AS 10。
- 入站数据流策略如下：
    - 从Internet发往网络192.168.11.0/24的流量来自R103-R101链路。
    - 从Internet发往网络192.168.12.0/24的流量来自R104-R102链路。
    - 如果到AS 10的一条链路发生故障，另一条链路会将发往所有网络的流量从Internet路由回AS 11。

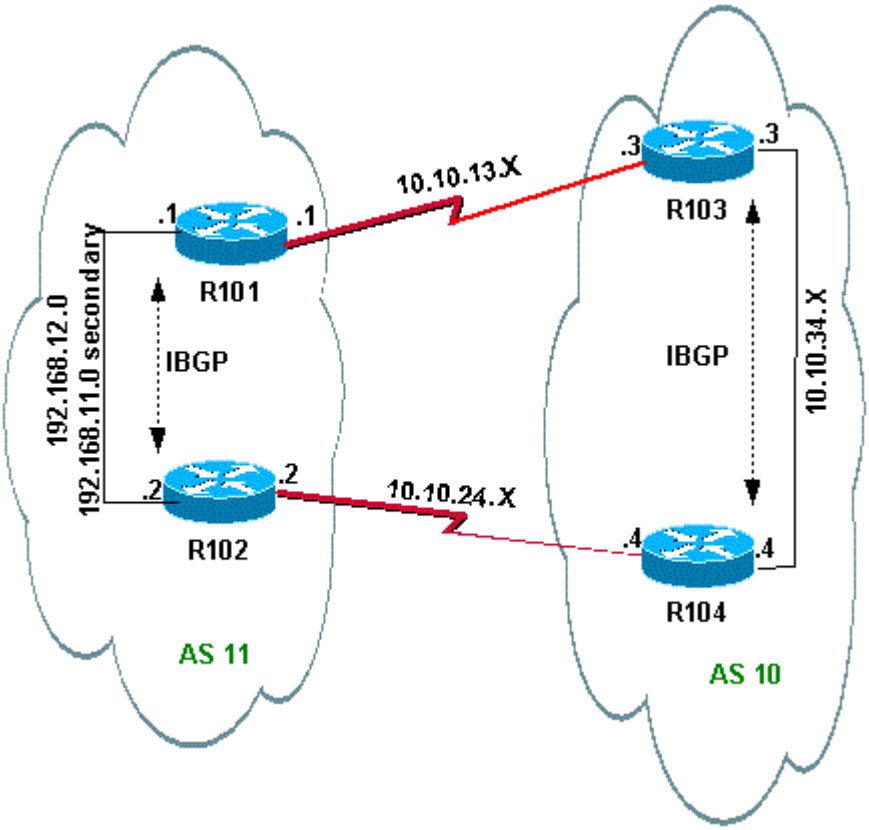
为此，从R101到R103通告了192.168.11.0，其AS_PATH比从R102到R104通告的短。AS 10通过R103-R101链路找到最佳路径。同样，192.168.12.0 通过 R102-R104 链路宣布的路径较短。AS 10优先使用R104-R102链路传输到AS 11中的192.168.12.0的流量。

对于出站流量，BGP根据通过eBGP获知的路由确定最佳路径。与通过 iBGP 获知的路由相比，这些路由更为可取。因此，R101 通过 eBGP 从 R103 获知 10.10.34.0，通过 iBGP 从 R102 获知 10.10.34.0。选择外部路径而不是内部路径。如果您查看R101配置中的BGP表，指向10.10.34.0的路由将通过R101-R103链路，下一跳为10.10.13.3。在R102上，指向10.10.34.0的路由通过R102-

R104链路，下一跳10.10.24.4将实现此目的。为发往10.10.34.0的流量分担负载。类似原因适用于R101和R102上的默认路由。有关BGP路径选择标准的详细信息，请参阅BGP最佳路径选择算法。

网络图

本部分使用以下网络设置：



配置

本部分使用以下配置：

- R101
- R102
- R103
- R104

R101

```
hostname R101
!
interface Ethernet0/0
 ip address 192.168.11.1 255.255.255.0 secondary
 ip address 192.168.12.1 255.255.255.0
!
interface Serial8/0
 ip address 10.10.13.1 255.255.255.0
!
router bgp 11
 no synchronization
 bgp log-neighbor-changes
 network 192.168.11.0
 network 192.168.12.0
 neighbor 10.10.13.3 remote-as 10
 neighbor 10.10.13.3 route-map R101-103-MAP out

!--- The AS_PATH is increased for 192.168.12.0.

 neighbor 192.168.12.2 remote-as 11
 neighbor 192.168.12.2 next-hop-self
 maximum-paths 2
 no auto-summary
!
access-list 1 permit 192.168.12.0
access-list 2 permit 192.168.11.0
route-map R101-103-MAP permit 10
 match ip address 1
 set as-path prepend 11 11 11
!
route-map R101-103-MAP permit 20
 match ip address 2
```

R102

```
hostname R102
!
interface Ethernet0/0
 ip address 192.168.11.2 255.255.255.0 secondary
 ip address 192.168.12.2 255.255.255.0
!
interface Serial8/0
 ip address 10.10.24.2 255.255.255.0
!
router bgp 11
 no synchronization
 bgp log-neighbor-changes
 network 192.168.11.0
 network 192.168.12.0
 neighbor 10.10.24.4 remote-as 10
 neighbor 10.10.24.4 route-map R102-104-MAP out

!--- The AS_PATH is increased for 192.168.11.0.

 neighbor 192.168.12.1 remote-as 11
 neighbor 192.168.12.1 next-hop-self
```

```
 no auto-summary
!
access-list 1 permit 192.168.11.0
access-list 2 permit 192.168.12.0
route-map R102-104-MAP permit 10
 match ip address 1
 set as-path prepend 11 11 11
!
route-map R102-104-MAP permit 20
 match ip address 2
```

R103

```
hostname R103
!
interface Ethernet0/0
 ip address 10.10.34.3 255.255.255.0
!
interface Serial8/0
 ip address 10.10.13.3 255.255.255.0
!
router bgp 10
no synchronization
 bgp log-neighbor-changes
 network 10.10.34.0 mask 255.255.255.0
 neighbor 10.10.13.1 remote-as 11
 neighbor 10.10.13.1 default-originate
 neighbor 10.10.34.4 remote-as 10
 neighbor 10.10.34.4 next-hop-self
 no auto-summary
```

R104

```
hostname R104
!
interface Ethernet0/0
 ip address 10.10.34.4 255.255.255.0
!
interface Serial8/0
 ip address 10.10.24.4 255.255.255.0
!
router bgp 10
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.10.24.2 remote-as 11
 neighbor 10.10.24.2 default-originate
 neighbor 10.10.34.3 remote-as 10
 neighbor 10.10.34.3 next-hop-self
 no auto-summary
```

验证

本部分提供的信息可帮助您确认您的配置是否可正常运行。

思科 CLI 分析器支持某些 show 命令（仅限于注册客户），您可以通过该命令查看对 show 命令输出的分析。

AS11 和 AS10 之间的两条链路都启动时验证

出站流量验证

注意：show ip bgp command output 中的大于号 (>) 表示可能的路径中用于该网络的最佳路径。有关详细信息，请参阅 BGP 最佳路径选择算法。

R101中的BGP表显示通过R101-R103链路到达互联网的所有传出流量的最佳路径。show ip route 命令输出将确认路由表中的路由。

<#root>

```
R101#

show ip bgp


BGP table version is 5, local router ID is 192.168.12.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i0.0.0.0          192.168.12.2                 100      0 10 i

*>                  10.10.13.3                            0 10 i


!--- This is the next hop of R103.

* i10.10.34.0/24    192.168.12.2                 100      0 10 i

*>                  10.10.13.3               0             0 10 i


!--- This is the next hop of R103.
* i192.168.11.0     192.168.12.2             0    100      0 i
```

```
*>                  0.0.0.0                  0         32768 i
* i192.168.12.0     192.168.12.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i

R101#
```

show ip route

```
!--- Output suppressed.

Gateway of last resort is 10.10.13.3 to network 0.0.0.0

C    192.168.12.0/24 is directly connected, Ethernet0/0
C    192.168.11.0/24 is directly connected, Ethernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.10.13.0 is directly connected, Serial8/0
```

B   10.10.34.0 [20/0] via 10.10.13.3, 00:08:53

```
!--- This is the next hop of R103.
```

B*   0.0.0.0/0 [20/0] via 10.10.13.3, 00:08:53

```
!--- This is the next hop of R103.
```

以下是R102的BGP和路由表。根据该策略，R102通过R102-R104链路将所有流量路由到AS 10:

<#root>

```
R102#
```

show ip bgp

```
BGP table version is 7, local router ID is 192.168.12.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
```

***> 0.0.0.0   10.10.24.4   0 10 i**

```
!--- This is the next hop of R104.

* i                 192.168.12.1                 100      0 10 i
```

***> 10.10.34.0/24   10.10.24.4   0 10 i**

```
!--- This is the next hop of R104.

* i                 192.168.12.1             0    100      0 10 i
* i192.168.11.0     192.168.12.1             0    100      0 i
*>                  0.0.0.0                  0         32768 i
* i192.168.12.0     192.168.12.1             0    100      0 i
*>                  0.0.0.0                  0         32768 i

R102#
```

show ip route

```
!--- Output suppressed.

Gateway of last resort is 10.10.24.4 to network 0.0.0.0

C    192.168.12.0/24 is directly connected, Ethernet0/0
C    192.168.11.0/24 is directly connected, Ethernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.10.24.0 is directly connected, Serial8/0
```

B   10.10.34.0 [20/0] via 10.10.24.4, 00:11:21

```
!--- This is the next hop of R104.
```

B*   0.0.0.0/0 [20/0] via 10.10.24.4, 00:11:21

```
!--- This is the next hop of R104.
```

从 AS 10 到 AS 11 的入站流量验证

网络192.168.11.0和192.168.12.0属于AS 11。根据该策略，AS 11优先使用R103-R101链路来传输发往网络192.168.11.0的流量，而使用R104-R102链路来传输发往网络192.168.12.0的流量。

<#root>

```
R103#
```

show ip bgp

```
BGP table version is 4, local router ID is 10.10.34.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.34.0/24    0.0.0.0                  0         32768 i
```

```
*> 192.168.11.0     10.10.13.1               0             0 11 i


!--- The next hop is R101.
*  192.168.12.0     10.10.13.1               0             0 11 11 11 11 i

*>i                 10.10.34.4               0    100      0 11 i


!--- The next hop is R104.


R103#

show ip route


!--- Output suppressed.
Gateway of last resort is not set


B    192.168.12.0/24 [200/0] via 10.10.34.4, 00:04:46


!--- The next hop is R104.


B    192.168.11.0/24 [20/0] via 10.10.13.1, 00:04:46


!--- The next hop is R101.

     10.0.0.0/24 is subnetted, 2 subnets
C       10.10.13.0 is directly connected, Serial8/0
C       10.10.34.0 is directly connected, Ethernet0/0
```

网络 192.168.11.0 在 R103 上的最佳路径是通过 R103-R101 链路，而网络 192.168.12.0 的最佳路径是通过 R104 至 AS 11。在这个案例中，最短路径的长度决定了最佳路径。

同样，在 R104 中，BGP 和路由表与此类似：

<#root>

```
R104#

show ip bgp


BGP table version is 13, local router ID is 10.10.34.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i10.10.34.0/24    10.10.34.3               0    100      0 i

*>i192.168.11.0     10.10.34.3               0    100      0 11 i

*                   10.10.24.2               0             0 11 11 11 11 i

*> 192.168.12.0     10.10.24.2               0             0 11 i


R104#s

how ip route


!--- Output suppressed.
Gateway of last resort is not set


B    192.168.12.0/24 [20/0] via 10.10.24.2, 00:49:06


!--- The next hop is R102.

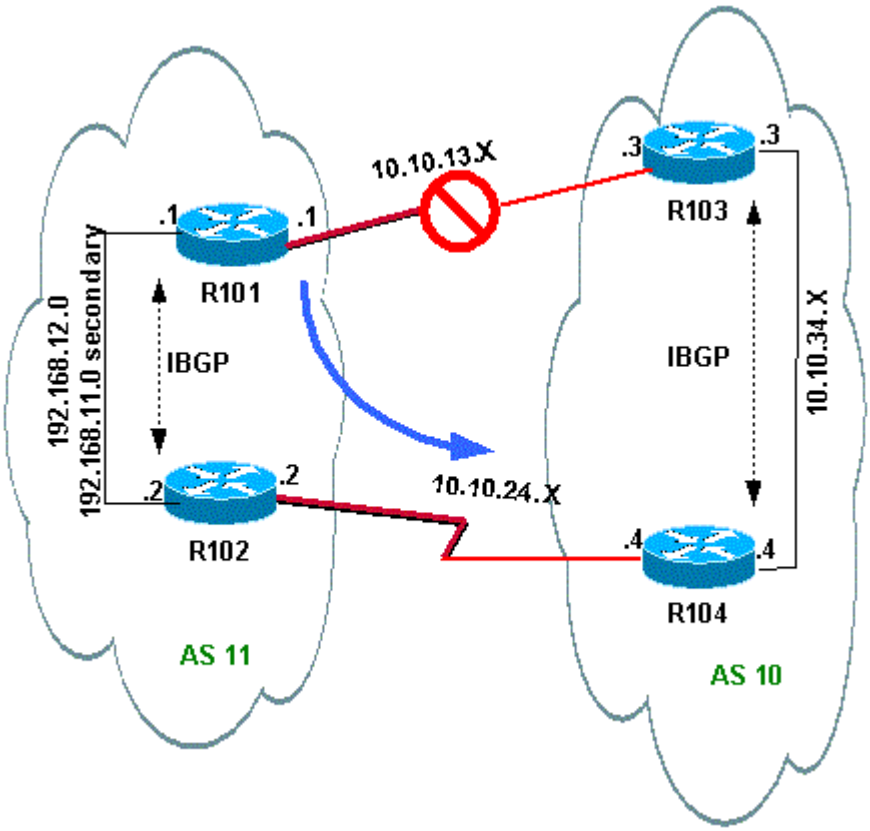

B    192.168.11.0/24 [200/0] via 10.10.34.3, 00:07:36


!--- The next hop is R103.

     10.0.0.0/24 is subnetted, 2 subnets
C       10.10.24.0 is directly connected, Serial8/0
C       10.10.34.0 is directly connected, Ethernet0/0
```

当R101-R103链路发生故障时进行验证

当R101-R103链路发生故障时，所有流量都必须通过R102重新路由。下图说明了此更改：



关闭 R103 上的 R103-R101 链路以模拟这种情况。

<#root>

```
R103(config)#
```

interface serial 8/0

```
R103(config-if)#
```

shutdown

```
*May 1 00:52:33.379: %BGP-5-ADJCHANGE: neighbor 10.10.13.1 Down Interface flap
*May 1 00:52:35.311: %LINK-5-CHANGED: Interface Serial8/0, changed state to administratively down
*May 1 00:52:36.127: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial8/0, changed state to down
```

校验到 AS 10 的出局路由。

<#root>

```
R101#
```

show ip bgp

```
BGP table version is 17, local router ID is 192.168.12.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path

*>i0.0.0.0          192.168.12.2                 100      0 10 i


!--- This is the next hop of R102.

*>i10.10.34.0/24    192.168.12.2                 100      0 10 i


!--- This is the next hop of R102.

* i192.168.11.0     192.168.12.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i
* i192.168.12.0     192.168.12.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i

R101#

show ip route


!--- Output suppressed.

Gateway of last resort is 192.168.12.2 to network 0.0.0.0

C    192.168.12.0/24 is directly connected, Ethernet0/0
C    192.168.11.0/24 is directly connected, Ethernet0/0
     10.0.0.0/24 is subnetted, 1 subnets

B       10.10.34.0 [200/0] via 192.168.12.2, 00:01:34


B*   0.0.0.0/0 [200/0] via 192.168.12.2, 00:01:34


!--- All outbound traffic goes through R102.

R102#

show ip route


!--- Output suppressed.

Gateway of last resort is 10.10.24.4 to network 0.0.0.0
```

```
C    192.168.12.0/24 is directly connected, Ethernet0/0
 C    192.168.11.0/24 is directly connected, Ethernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.10.24.0 is directly connected, Serial8/0

B       10.10.34.0 [20/0] via 10.10.24.4, 00:13:22



B*   0.0.0.0/0 [20/0] via 10.10.24.4, 00:55:22



!--- All outbound traffic on R102 goes through R104.
```

在 R101-R103 关闭时校验入站数据流路由。

<#root>

```
R103#s

how ip bgp


BGP table version is 6, local router ID is 10.10.34.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path

*> 10.10.34.0/24    0.0.0.0                  0         32768 i


*>i192.168.11.0     10.10.34.4               0    100      0 11 11 11 11 i

*>i192.168.12.0     10.10.34.4               0    100      0 11 i

R103#

show ip route


!--- Output suppressed.
Gateway of last resort is not set


B    192.168.12.0/24 [200/0] via 10.10.34.4, 00:14:55


!--- The next hop is R104.
```

```
B    192.168.11.0/24 [200/0] via 10.10.34.4, 00:05:46


!--- The next hop is R104.
     10.0.0.0/24 is subnetted, 1 subnets
C       10.10.34.0 is directly connected, Ethernet0/0
```

在 R104 上，发往 192.168.11.0 和 192.168.12.0 的业务通过 R104-R102 链路传输。

```
<#root>

R104#

show ip route


!--- Output suppressed.

Gateway of last resort is not set


B    192.168.12.0/24 [20/0] via 10.10.24.2, 00:58:35


!--- The next hop is R102.


B    192.168.11.0/24 [20/0] via 10.10.24.2, 00:07:57


!--- The next hop is R102.

     10.0.0.0/24 is subnetted, 2 subnets
C       10.10.24.0 is directly connected, Serial8/0
C       10.10.34.0 is directly connected, Ethernet0/0
```

故障排除

当前没有故障排除此配置的特定可用资料。

通过单个本地路由器多重归属于两个 ISP 时的负载分摊

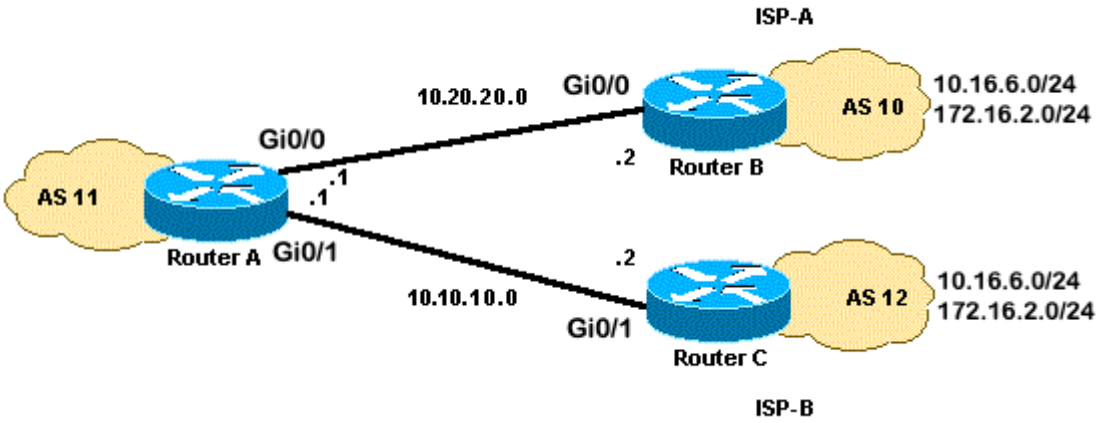
在此方案中，负载均衡不适用于多宿主环境，因此只能进行负载共享。您无法执行负载均衡，因为BGP在从不同的AS获取的BGP路由中仅选择一条到达目的地的最佳路径。其思路是为1.0.0.x到

128.0.0.x范围内从ISP(A)获知的路由设置更好的度量，为从ISP(B)获知的其余路由设置更好的度量。 以下面的网络图为例：

有关其他信息，请参阅两个不同服务提供商（多宿主）间的 BGP 的配置示例。

网络图

本部分使用以下网络设置：



配置

本部分使用以下配置：

- 路由器A
- 路由器B
- RouterC

路由器A

```
interface GigabitEthernet0/0
 ip address 10.20.20.1 255.255.255.0

interface GigabitEthernet0/1
 ip address 10.10.10.1 255.255.255.0

router bgp 11
 neighbor 10.20.20.2 remote-as 10
```

```
 neighbor 10.20.20.2 route-map UPDATES-1 in

!--- This allows only the networks up to 128.0.0.x.

neighbor 10.10.10.2 remote-as 12
neighbor 10.10.10.2 route-map UPDATES-2 in

!--- This allows anything above the 128.0.0.x network.

route-map UPDATES-1 permit 10
 match ip address 1
 set weight 100

route-map UPDATES-1 permit 20
 match ip address 2

route-map UPDATES-2 permit 10
 match ip address 1

route-map UPDATES-2 permit 20
 match ip address 2
 set weight 100

access-list 1 permit 0.0.0.0 127.255.255.255
access-list 2 deny 0.0.0.0 127.255.255.255
access-list 2 permit any
```

路由器B

```
interface GigabitEthernet0/2
 ip address 172.16.2.1 255.255.255.0

interface GigabitEthernet0/3
 ip address 10.16.6.1 255.255.255.0

interface GigabitEthernet0/0
 ip address 10.20.20.2 255.255.255.0

router bgp 10
 neighbor 10.20.20.1 remote-as 11
 network 172.16.2.0 mask 255.255.255.0
 network 10.16.6.0 mask 255.255.255.0
```

RouterC

```
interface GigabitEthernet0/3
 ip address 10.16.6.2 255.255.255.0

interface GigabitEthernet0/2
 ip address 172.16.2.2 255.255.255.0

interface GigabitEthernet0/1
 ip address 10.10.10.2 255.255.255.0
```

```
router bgp 12
 neighbor 10.10.10.1 remote-as 11
 network 172.16.2.0 mask 255.255.255.0
 network 10.16.6.0 mask 255.255.255.0
```

验证

使用本部分可确定配置能否正常运行。

思科 CLI 分析器（仅适用于注册客户）支持某些 show 命令。要查看对 show 命令输出的分析，请使用思科 CLI 分析器。

show ip route命令的输出和traceroute命令的输出显示任何低于128.0.0.x的网络都会通过10.20.20.2退出RouterA。此路由是从Serial 0接口发出的下一跳。其余网络通过 10.10.10.2 退出，这是跳出 serial 1 接口的下一跳。

<#root>

RouterA#

show ip route

!--- Output suppressed.

Gateway of last resort is not set

B   172.16.2.0/24 [20/0] via 10.10.10.2, 00:13:16

!--- This is the next hop out through GigabitEthernet0/0.

B   10.16.6.0/24 [20/0] via 10.20.20.2, 00:13:16

!--- This is the next hop out through GigabitEthernet0/1.

!--- Output suppressed.

RouterA#

show ip cef 172.16.2.0

172.16.2.0/24

nexthop 10.10.10.2 GigabitEthernet0/1

RouterA#

show ip cef 10.16.6.0

10.16.6.0/24

nexthop 10.20.20.2 GigabitEthernet0/0

RouterA#

show ip bgp

```
BGP table version is 10, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
```

***> 10.16.6.0/24 10.20.20.2 0 100 10 i * 10.10.10.2 0 0 12 i * 172.16.2.0/24 10.20.20.2 0 0 10 i *> 10.1**

*> 192.168.1.1/32 0.0.0.0 0 32768 i

RouterA

#traceroute 172.16.2.1 source loopback0

```
Type escape sequence to abort.
Tracing the route to 172.16.2.1
VRF info: (vrf in name/id, vrf out name/id)
```

1 10.10.10.2 2 msec 3 msec 2 msec

2 172.16.2.1 [AS 12] 3 msec 3 msec *

RouterA#

traceroute 10.16.6.1 source loopback0

```
Type escape sequence to abort.
Tracing the route to 10.16.6.1
VRF info: (vrf in name/id, vrf out name/id)
```

1 10.20.20.2 3 msec 2 msec *

```
RouterA#
```

故障排除

当前没有故障排除此配置的特定可用资料。

通过多个本地路由器多宿主到两个 ISP 时的负载平衡

在具有两个 ISP 的多宿主环境中，不可能实现负载均衡。BGP只从从不同AS获取的BGP路径中选择通往目的地的单一最佳路径，因此不可能进行负载均衡。但是，负载共享可以在这种多宿主BGP网络中实现。根据预先确定的策略，使用不同的 BGP 属性控制流量。

本部分讨论使用最频繁的多宿主配置。配置显示如何实现负载共享。请参阅网络图，其中 AS 100 的多宿主实现了可靠性和负载共享。

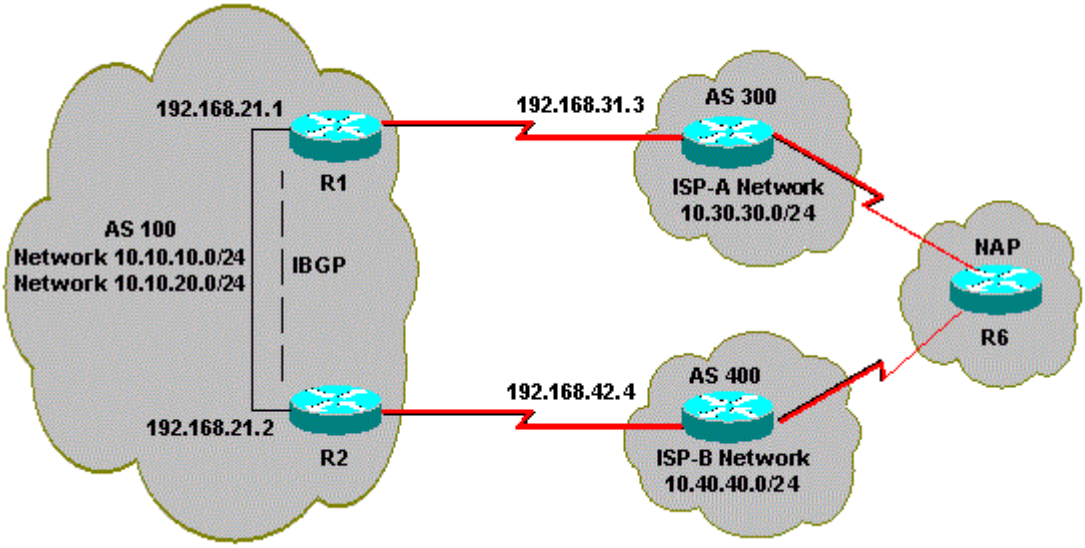
注意：本示例中的IP地址符合RFC 1918专用地址空间标准。

为了简单起见，假定AS 100的BGP路由策略如下：

- AS 100接受来自两个提供商的本地路由，以及其他Internet路由的默认路由。
- 出局数据流策略如下：
    - 到AS 300的业务经过R1 - ISP(A) 链路。
    - 到AS 400的业务经过R2 - ISP(B)链路。
    - 所有其他流量首选通过R1-ISP(A)链路的默认路由0.0.0.0。
    - 如果R1-ISP(A)链路发生故障，所有流量都会通过R2-ISP(B)链路。
- 入站数据流策略如下：
    - 从Internet发往网络10.10.10.0/24的流量来自ISP(A)-R1链路。
    - 从Internet发往网络10.10.20.0/24的流量来自ISP(B)-R2链路。
    - 如果一个ISP发生故障，另一个ISP会将所有网络的流量从Internet路由回AS 100。

网络图

本部分使用以下网络设置：



配置

本部分使用以下配置：

- R2
- R1

R2

```
interface Ethernet0
 ip address 192.168.21.2 255.255.255.0
 !
interface Serial0
  ip address 192.168.42.2 255.255.255.0
router bgp 100
 no synchronization
 bgp log-neighbor-changes

!--- The next two lines announce the networks to BGP peers.

 network 10.10.10.0 mask 255.255.255.0
 network 10.10.20.0 mask 255.255.255.0

!--- The next line configures iBGP on R1.
```

```
  neighbor 192.168.21.1 remote-as 100
  neighbor 192.168.21.1 next-hop-self

!--- The next line configures eBGP with ISP(B).

 neighbor 192.168.42.4 remote-as 400

!--- This is the incoming policy route map for the application of attributes to specific routes.

 neighbor 192.168.42.4 route-map AS-400-INCOMING in

!--- This is the outgoing policy route map for the application of attributes to specific routes.

 neighbor 192.168.42.4 route-map AS-400-OUTGOING out
 no auto-summary

!--- This line sets the AS path access list, it permits all routes within the routing domain of the pro

ip as-path access-list 1 permit ^400$

!--- These two lines set the access list.

access-list 10 permit 10.10.10.0 0.0.0.255
access-list 20 permit 10.10.20.0 0.0.0.255

!--- The next three lines configure LOCAL_PREF for routes that match AS path access list 1.

route-map AS-400-INCOMING permit 10
 match as-path 1
 set local-preference 150

!--- Here, the route map prepends AS 100 to BGP updates for networks that are permitted by access list

route-map AS-400-OUTGOING permit 10
 match ip address 10
 set as-path prepend 100

!--- This line announces the network that is permitted by access list 20 without any changes in BGP att

route-map AS-400-OUTGOING permit 20
 match ip address 20
```

R1

```
interface Serial0/0
 ip address 192.168.31.1 255.255.255.0
!
interface Ethernet1/0
 ip address 192.168.21.1 255.255.255.0
!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 network 10.10.10.0 mask 255.255.255.0
 network 10.10.20.0 mask 255.255.255.0

!--- IBGP peering with R2
```

```
neighbor 192.168.21.2 remote-as 100
neighbor 192.168.21.2 next-hop-self

!--- This line sets eBGP peering with ISP(A).

neighbor 192.168.31.3 remote-as 300

!--- This is the incoming policy route map for the application of attributes to specific routes.

 neighbor 192.168.31.3 route-map AS-300-INCOMING in

!--- This is the outgoing policy route map for the application of attributes to specific routes.

neighbor 192.168.31.3 route-map AS-300-OUTGOING out
 no auto-summary

!--- This line sets the AS path access list, it permits all routes within the routing domain of the pro

ip as-path access-list 1 permit ^300$

!--- These two lines set the IP access list.

access-list 10 permit 10.10.20.0 0.0.0.255
access-list 20 permit 10.10.10.0 0.0.0.255

!--- The next three lines configure LOCAL_PREF for routes that match AS path access list 1.

route-map AS-300-INCOMING permit 10
 match as-path 1
 set local-preference 200

!--- Here, the route map prepends AS 100 to BGP updates for networks that are permitted by access list

route-map AS-300-OUTGOING permit 10
 match ip address 10
 set as-path prepend 100

!--- This line announces the network that is permitted by access list 20 without any changes in BGP att

route-map AS-300-OUTGOING permit 20
 match ip address 20
!
```

验证

使用本部分可确定配置能否正常运行。

思科 CLI 分析器（仅适用于注册客户）支持某些 show 命令。要查看对 show 命令输出的分析，请使用思科 CLI 分析器。

发出 show ip bgp 命令以验证出站/入站策略是否有效。



注意：show ip bgp output 中的大于号 (>) 表示可能的路径中用于该网络的最佳路径。有关详

细信息，请参阅 BGP 最佳路径选择算法。

<#root>

R1#

show ip bgp

```
BGP table version is 6, local router ID is 192.168.31.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

BGP table version is 6, local router ID is 192.168.31.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network          Next Hop            Metric LocPrf Weight Path

*> 0.0.0.0          192.168.31.3                  200      0 300 i


!--- This line shows that the default route 0.0.0.0/0 is preferred through AS 300, ISP(A).

* i10.10.10.0/24    192.168.21.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i
* i10.20.20.0/24    192.168.21.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i

*> 10.30.30.0/24    192.168.31.3             0    200      0 300 i




*>i10.40.40.0/24    192.168.21.2             0    150      0 400 i



!--- The route to network 10.30.30.0/24 (AS 300) is preferred through the R1-ISP(A) link.
!--- The route to network 10.40.40.0/24 (AS 400) is preferred through the R2-ISP(B) link.
```

现在，请看 R2 上的 show ip bgp output：

<#root>

R2#

show ip bgp

```
BGP table version is 8, local router ID is 192.168.42.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  0.0.0.0          192.168.42.4                 150      0 400 i

*>i                 192.168.21.1                 200      0 300 i


!--- This line shows that the default route 0.0.0.0/0 is preferred through AS 300, through the R2-ISP(B

*> 10.10.10.0/24    0.0.0.0                  0         32768 i
* i                 192.168.21.1             0    100      0 i
*> 10.10.20.0/24    0.0.0.0                  0         32768 i
* i                 192.168.21.1             0    100      0 i

*>i10.30.30.0/24    192.168.21.1             0    200      0 300 i




*> 10.40.40.0/24    192.168.42.4             0    150      0 400 i


!--- The route to network 10.30.30.0/24 (AS 300) is preferred through the R1-ISP(A) link.
!--- The route to network 10.40.40.0/24 (AS 400) is preferred through the R2-ISP(B) link.
```

在Router 6上发出show ip bgp命令，以观察网络10.10.10.0/24和10.10.20.0/24的入站策略：

```
<#root>

R6#

show ip bgp


BGP table version is 15, local router ID is 192.168.64.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path

*> 10.10.10.0/24    192.168.63.3                           0 300 100 100 i


!--- This line shows that network 10.10.10.0/24 is routed through AS 300
!--- with the ISP(A)-R1 link.

*                   192.168.64.4                           0 400 100 100 100 i
*  10.10.20.0/24    192.168.63.3                           0 300 100 100 i

*>                  192.168.64.4                           0 400 100 i
```

```
!--- This line shows that network 10.10.20.0/24 is routed through AS 400
!--- with the ISP(B)-R2 link.

*> 10.30.30.0/24    192.168.63.3             0             0 300 i
*> 10.40.40.0/24    192.168.64.4             0             0 400 i
```

关闭 R1 上的 R1-ISP(A) 链路并观察 BGP 表。所有流向Internet的流量都通过R2-ISP(B)链路路由：

<#root>

```
R1(config)#

interface serial 0/0

R1(config-if)#

shutdown



*May  2 19:00:47.377: %BGP-5-ADJCHANGE: neighbor 192.168.31.3 Down Interface flap
*May  2 19:00:48.277: %LINK-5-CHANGED: Interface Serial0/0, changed state to
 administratively  down
*May 23 12:00:51.255: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed
 state to down

R1#

show ip bgp



BGP table version is 12, local router ID is 192.168.31.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i0.0.0.0          192.168.21.2                 150      0 400 i

!--- The best default path is now through the R2-ISP(B) link.

* i10.10.10.0/24    192.168.21.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i
* i10.10.20.0/24    192.168.21.2             0    100      0 i
*>                  0.0.0.0                  0         32768 i
*>i10.40.40.0/24    192.168.21.2             0    150      0 400 i

R2#

show ip bgp
```

```
BGP table version is 14, local router ID is 192.168.42.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path

*> 0.0.0.0          192.168.42.4                  150      0 400 i



!--- The best default route is now through ISP(B) with a local preference of 150.

* i10.10.10.0/24    192.168.21.1             0    100      0 i
*>                  0.0.0.0                  0         32768 i
* i10.10.20.0/24    192.168.21.1             0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 10.40.40.0/24    192.168.42.4             0    150      0 400 i
```

请看 Router 6 中网络 10.10.10.0/24 的路由：

```
<#root>

R6#

show ip bgp



BGP table version is 14, local router ID is 192.168.64.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.10.0/24    192.168.64.4                           0 400 100 100 i

!--- Network 10.10.10.0 is reachable through ISP(B), which announced the network with AS path prepend.

*> 10.10.20.0/24    192.168.64.4                           0 400 100 i

*> 10.30.30.0/24    192.168.63.3             0             0 300 i
*> 10.40.40.0/24    192.168.64.4             0             0 400 i
```

故障排除

当前没有故障排除此配置的特定可用资料。

相关信息

- BGP 多宿主：设计和故障排除 - 网络直播视频
- BGP 多宿主：设计和故障排除 — 网络直播中的问答
- 负载均衡如何工作
- 两个不同服务提供商（多宿）间的 BGP 的示例配置
- BGP 路由器如何使用多出口标识符进行最佳路径选择
- IP 路由技术支持页
- 技术支持和文档 - Cisco Systems

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。