Understand Load Sharing with BGP in Single/Multihomed Environments

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Background Information

Load Sharing with the Loopback Address as a BGP Neighbor

- Network Diagram
- Configurations
- Verify
- Troubleshoot

Load Sharing When Dual-Homed to One Internet Service Provider (ISP) Through a Single Local Router

- Network Diagram
- Configurations
- Verify
- Troubleshoot

Load Sharing When Dual-Homed to One ISP Through Multiple Local Routers

- Network Diagram
- Configurations
- Verify
  - Verification When Both Links Between AS 11 and AS 10 Are Up
  - Verification of Outbound Traffic
  - Verification of Inbound Traffic from AS 10 to AS 11
  - Verification When the R101-R103 Link Fails
- Troubleshoot

Load Sharing When Multihomed to Two ISPs Through a Single Local Router

- Network Diagram
- Configurations
- Verify
- Troubleshoot

Load Sharing When Multihomed to Two ISPs Through Multiple Local Routers

- Network Diagram
- Configurations
- Verify
- Troubleshoot

Related Information

Introduction

This document describes load sharing, which allows a router to distribute the outgoing and incoming traffic among multiple paths.

Prerequisites

Requirements

Ensure you meet these requirements before you attempt this configuration:

- Knowledge of BGP Best Path Selection Algorithm
- Knowledge of Configuring BGP

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

The incoming and outgoing traffic paths are derived either statically or with dynamic protocols such as:

- Routing Information Protocol (RIP)
- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Open Shortest Path First (OSPF) Protocol

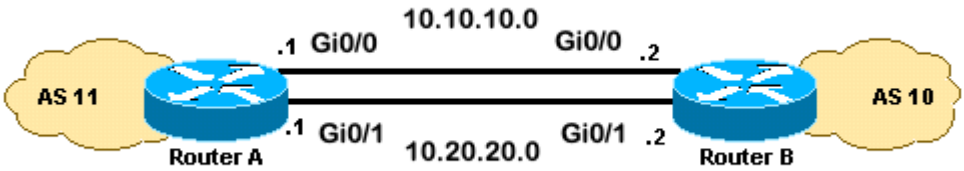
By default, Border Gateway Protocol (BGP) selects only a single best path and does not perform load balancing. This document describes how to perform load sharing in different scenarios with the use of BGP. For additional information about load balance, refer to How Does Load Balance Work.

Load Sharing with the Loopback Address as a BGP Neighbor

This scenario shows how to achieve load sharing when there are multiple (up to a maximum of six), equal-cost links. The links are terminated in one router at a local autonomous system (AS) and in another router at a remote AS in a single-homed BGP environment. The Network Diagram serves as an example.

Network Diagram

This section uses this network setup:



Configurations

This section uses these configurations:

- RouterA
- RouterB

RouterA

<#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
```

RouterB

<#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
```

Note: You can use static routes instead of a routing protocol to introduce two equal-cost paths to reach the destination. But in this example, EIGRP was used to share the loopback information.

Verify

Use this section to confirm your configuration works properly.

The Cisco CLI Analyzer (registered customers only) supports certain **show** commands. Use the Cisco CLI Analyzer to view an analysis of **show** command output.

The output of the **show ip route** command shows both paths to the 192.168.2.2 network, learned by EIGRP. The **show ip bgp summary** command shows that the BGP neighbor was built with the Loopback of the remote router. The output of the **traceroute** command indicates the load is distributed between two serial links. In this scenario, load sharing occurs on a per-packet basis. You can issue the **ip route-cache** command on the serial interfaces to do load sharing on a per-destination basis. You can also configure per-packet and per-destination load balancing with Cisco Express Forwarding. For more information on how to configure Cisco Express Forwarding, refer to Configuring Cisco Express Forwarding.

<#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#
```

Troubleshoot

There is currently no specific information available to troubleshoot this configuration.

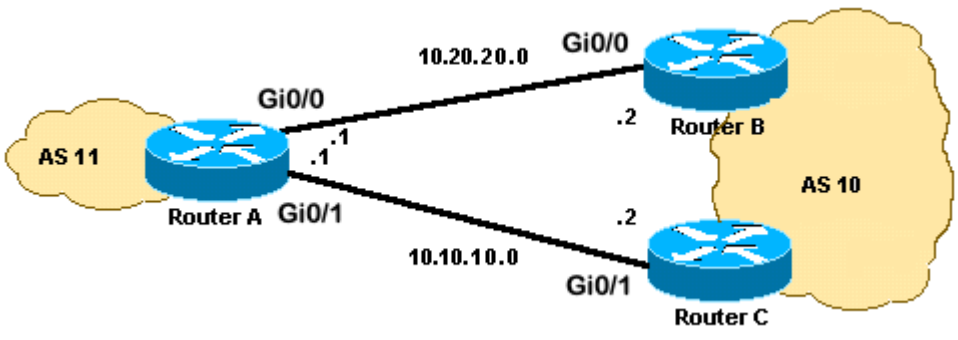
Load Sharing When Dual-Homed to One Internet Service Provider (ISP) Through a Single Local Router

This scenario shows how to achieve load sharing when multiple links exist between a remote AS and a local AS. These links are terminated in one router at the local AS and on multiple routers at remote ASs in a single-homed BGP environment. The Network Diagram is an example of such a network.

This sample configuration uses the **maximum-paths** command. By default, BGP chooses one best path among the possible equal-cost paths that are learned from one AS. However, you can change the maximum number of parallel equal-cost paths that are allowed. In order to make this change, include the **maximum-paths paths** command under the BGP configuration. Use a number between 1 and 6 for the paths argument.

Network Diagram

This section uses this network setup:



Configurations

This section uses these configurations:

- RouterA
- RouterB
- RouterC

RouterA

<#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
```

RouterB

```
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
```

Verify

Use this section to confirm your configuration works properly.

The Cisco CLI Analyzer (registered customers only) supports certain **show** commands. Use the Cisco CLI Analyzer to view an analysis of **show** command output.

The output of the **show ip route** command shows both paths to the 172.16.2.0 network are learned by BGP. The output of the **traceroute** command indicates the load is distributed between two serial links. In this scenario, load sharing occurs on a per-destination basis. The **show ip bgp** command gives the valid entries for the 172.16.2.0 network.

<#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
```

Troubleshoot

There is currently no specific information available to troubleshoot this configuration.

Load Sharing When Dual-Homed to One ISP Through Multiple Local Routers

This scenario shows how to achieve load sharing when there are multiple connections to the same ISP through multiple local routers. The two eBGP peers are terminated on two separate local routers. Load balancing on the two links is not possible because BGP chooses the single best path among the networks learned from eBGP and internal BGP (iBGP). Load sharing among the multiple paths to AS 10 is the next-best option. With this type of load sharing, traffic to specific networks, on the basis of predefined policies, travels through both links. Additionally, each link acts as a backup to the other link, in case one link fails.

For simplicity, assume the BGP routing policy for AS 11 is:

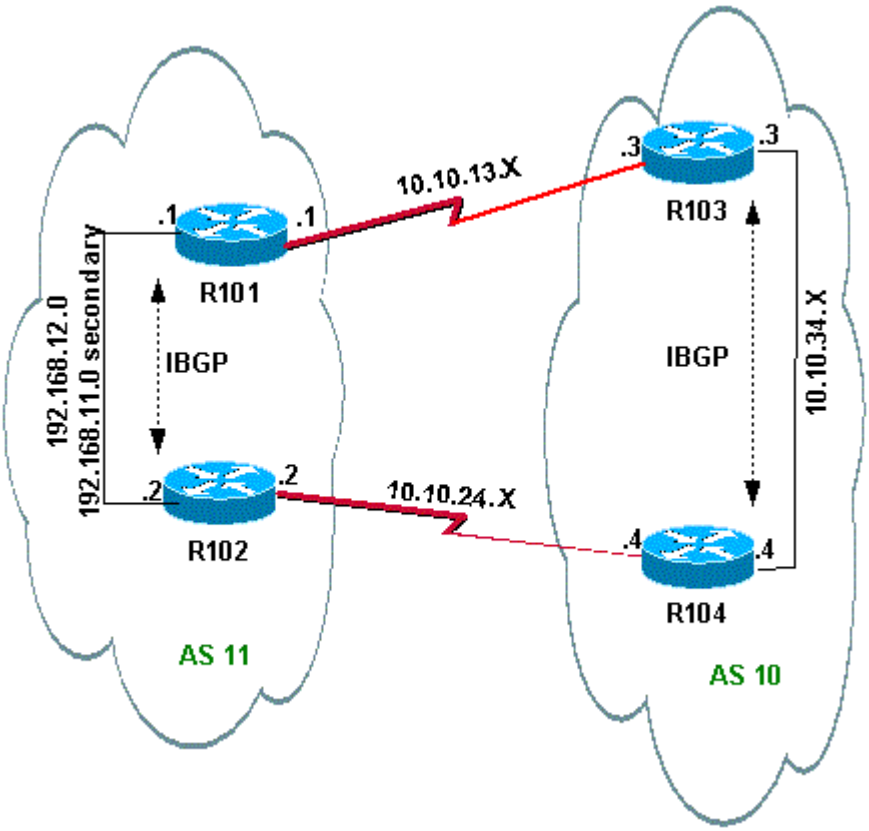
- AS 11 accepts the local routes from AS 10, along with a default for the rest of the Internet routes.
- The outbound traffic policy is:
    - All traffic destined to the Internet from R101 goes out through the R101-R103 link.
    - If the R101-R103 link fails, then all traffic to the Internet from R101 goes through R102 to AS 10.
    - Similarly, all traffic destined to the Internet from R102 goes through the R102-R104 link.
    - If the R102-R104 link fails, then all traffic to the Internet from R102 goes through R101 to AS 10.
- The inbound traffic policy is:
    - Traffic destined for network 192.168.11.0/24 from the Internet comes from the R103-R101 link.
    - Traffic destined for network 192.168.12.0/24 from the Internet comes from the R104-R102 link.
    - If one link to AS 10 fails, then the other link routes the traffic destined for all networks back to AS 11 from the Internet.

To achieve this, 192.168.11.0 is announced from R101 to R103 with a shorter AS_PATH than is announced from R102 to R104. AS 10 finds the best path through the R103-R101 link. Similarly, 192.168.12.0 is announced with a shorter path through the R102-R104 link. AS 10 prefers the R104-R102 link for traffic bound to 192.168.12.0 in AS 11.

For outbound traffic, BGP determines the best path on the basis of routes learned through eBGP. These routes are preferable to routes learned through iBGP. So R101 learns 10.10.34.0 from R103 through eBGP and from R102 through iBGP. The external path is selected over the internal path. If you look at the BGP table in the R101 configuration, the route toward 10.10.34.0 would be through the R101-R103 link, with the next hop 10.10.13.3. On R102, the route toward 10.10.34.0 is through the R102-R104 link, with the next hop 10.10.24.4. This achieves load sharing for traffic destined to 10.10.34.0. Similar reason applies to the default routes on R101 and R102. For more information on BGP path selection criteria, refer to BGP Best Path Selection Algorithm.

Network Diagram

This section uses this network setup:



Configurations

This section uses these configurations:

- 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
```

Verify

This section provides information you can use to confirm your configuration works properly.

Certain **show** commands are supported by the Cisco CLI Analyzer (registered customers only) , which allows you to view an analysis of **show** command output.

Verification When Both Links Between AS 11 and AS 10 Are Up

Verification of Outbound Traffic

Note: The greater-than sign (>) in the **show ip bgp** command output represents the best path to use for that network among the possible paths. Refer to BGP Best Path Selection Algorithm for more information.

The BGP table in R101 shows the best path for all outgoing traffic to the Internet is through the R101-R103 link. The **show ip route** command output confirms the routes in the routing table.

```
<#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.
```

Here are the BGP and routing tables for R102. Based on the policy, R102 routes all the traffic to AS 10 through the R102-R104 link:

```
<#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.
```

Verification of Inbound Traffic from AS 10 to AS 11

The networks 192.168.11.0 and 192.168.12.0 belong to AS 11. Based on the policy, AS 11 prefers the R103-R101 link for traffic destined to network 192.168.11.0 and the R104-R102 link for traffic destined to network 192.168.12.0.

<#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
```

The best path for network 192.168.11.0 on R103 is through the R103-R101 link, and the best path for network 192.168.12.0 is through R104 to AS 11. In this case, the shortest path length determines the best path.

Similarly, on R104, the BGP and routing table resembles this:

<#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.

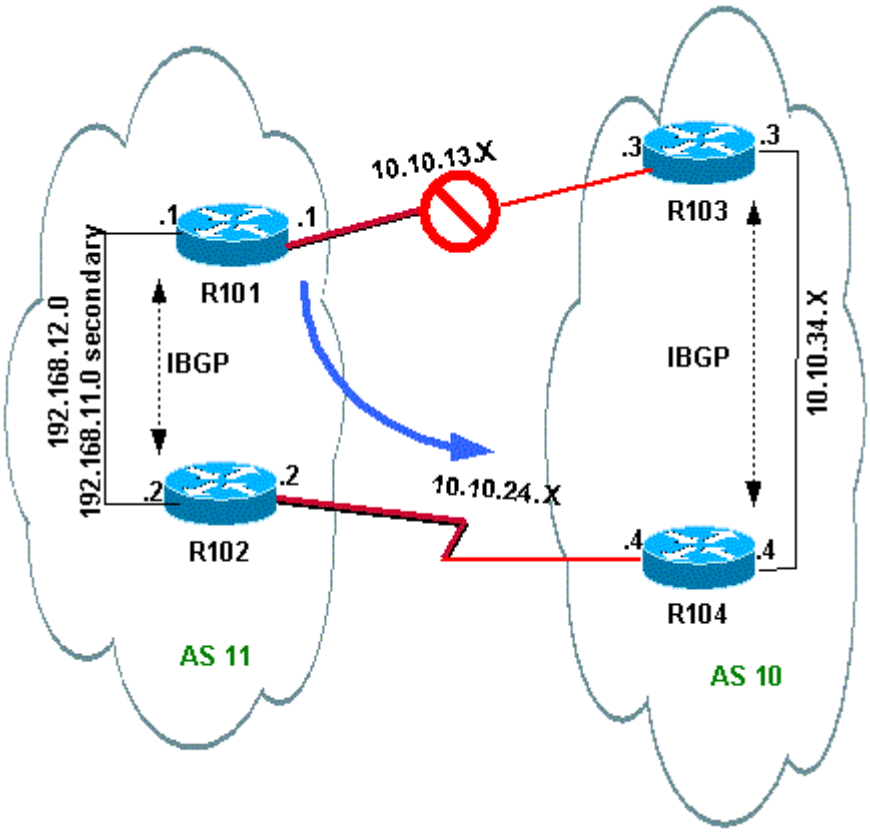

B    192.168.11.0/24 [200/0] via 10.10.34.3, 00:07:36


!--- 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
```

Verification When the R101-R103 Link Fails

When the R101-R103 link fails, all traffic must reroute through R102. This diagram illustrates this change:



Shut down the R103-R101 link on R103 in order to simulate this situation.

<#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
```

Verify the outbound route to AS 10.

```
<#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.
```

Verify the inbound traffic route when R101-R103 is down.

<#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
```

On R104, the traffic for 192.168.11.0 and 192.168.12.0 goes through the R104-R102 link.

<#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
```

Troubleshoot

There is currently no specific information available to troubleshoot this configuration.

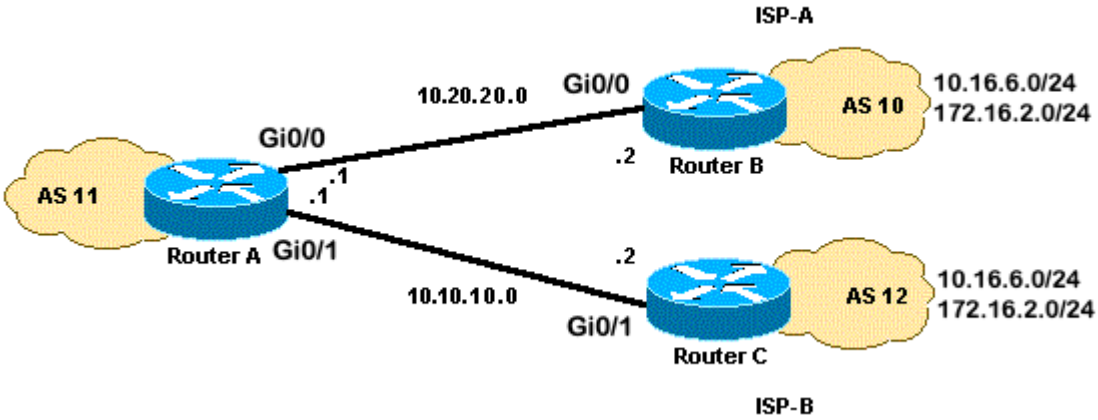
Load Sharing When Multihomed to Two ISPs Through a Single Local Router

In this scenario, load balancing is not an option in a multihomed environment, so you can only do load sharing. You cannot do load balancing because BGP selects only a single best path to a destination among the BGP routes learned from the different ASs. The idea is to set a better metric for the routes in the range 1.0.0.x to 128.0.0.x that are learned from ISP(A) and a better metric for the rest of the routes learned from ISP(B). The Network Diagram is an example.

Refer to Sample Configuration for BGP with Two Different Service Providers (Multihoming) for additional information.

Network Diagram

This section uses this network setup:



Configurations

This section uses these configurations:

- RouterA
- RouterB
- RouterC

RouterA

```
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
```

RouterB

```
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
```

Verify

Use this section to confirm your configuration works properly.

The Cisco CLI Analyzer (registered customers only) supports certain **show** commands. Use the Cisco CLI Analyzer to view an analysis of **show** command output.

The output from the **show ip route** command and the output from the **traceroute** command show any network lower than 128.0.0.x exits RouterA through 10.20.20.2. This route is the next hop out of the serial 0 interface. The rest of the networks exit through 10.10.10.2, which is the next hop out of the serial 1 interface.

<#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#
```

Troubleshoot

There is currently no specific information available to troubleshoot this configuration.

Load Sharing When Multihomed to Two ISPs Through Multiple Local Routers

Load balancing is not possible in a multihomed environment with two ISPs. BGP selects only the single best path to a destination among the BGP paths learned from different ASs, which makes load balancing impossible. But, load sharing is possible in such multihomed BGP networks. On the basis of predetermined policies, traffic flow is controlled with different BGP attributes.

This section discusses the multihomed configurations that has the most frequent use. The configuration shows how to achieve load sharing. See the Network Diagram, in which the multihome of AS 100 achieves reliability and load sharing.

Note: The IP addresses in this example conform to RFC 1918 standards for Private Address Space.

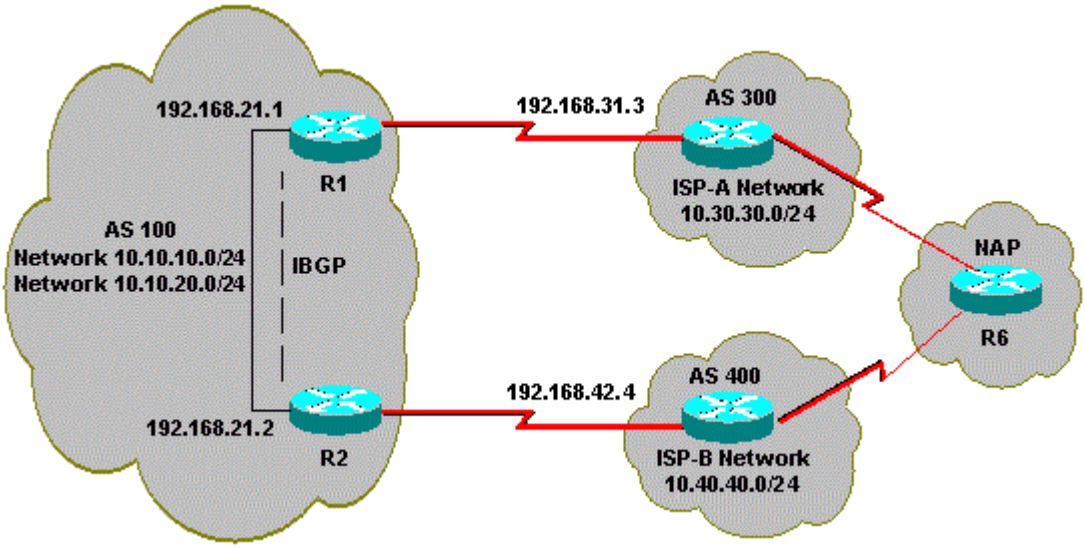
For simplicity, assume the BGP routing policy for AS 100 is:

- AS 100 accepts the local routes from both providers, along with a default for the other Internet routes.

- The outbound traffic policy is:
  - Traffic destined to AS 300 goes through the R1-ISP(A) link.
  - Traffic destined to AS 400 goes through the R2-ISP(B) link.
  - All other traffic prefers default route 0.0.0.0 through the R1-ISP(A) link.
  - If the R1-ISP(A) link fails, all traffic goes through the R2-ISP(B) link.
- The inbound traffic policy is:
  - Traffic destined for network 10.10.10.0/24 from the Internet comes from the ISP(A)-R1 link.
  - Traffic destined for network 10.10.20.0/24 from the Internet comes from the ISP(B)-R2 link.
  - If one ISP fails, the other ISP routes traffic back to AS 100 from the Internet for all the networks.

Network Diagram

This section uses this network setup:



Configurations

This section uses these configurations:

- 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
!
```

Verify

Use this section to confirm your configuration works properly.

The Cisco CLI Analyzer (registered customers only) supports certain **show** commands. Use the Cisco CLI Analyzer to view an analysis of **show** command output.

Issue the **show ip bgp** command in order to verify that the outbound/inbound policy works.

Note: The greater-than sign (>) in the **show ip bgp** output represents the best path to use for that network among the possible paths. Refer to BGP Best Path Selection Algorithm for more information.

<#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
* i10.10.20.0/24    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.
```

Now, look at the show ip bgp output on R2:

<#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.
```

Issue the **show ip bgp** command on Router 6 to observe the inbound policy for networks 10.10.10.0/24 and 10.10.20.0/24:

<#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
```

Shut down the R1-ISP(A) link on R1 and observe the BGP table. Expect all traffic to the Internet to be routed through the R2-ISP(B) link:

<#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
```

Look at the route for network 10.10.10.0/24 in Router 6:

<#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
```

Troubleshoot

There is currently no specific information available to troubleshoot this configuration.

Related Information

- BGP Multi-homing: Design and Troubleshooting - Video from live Webcast
- BGP Multi-homing: Design and Troubleshooting - Questions and Answers from Live Webcast
- How Does Load Balancing Work
- Sample Configuration for BGP with Two Different Service Providers (Multihoming)
- How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection
- IP Routing Technology Support Page
- Technical Support & Documentation - Cisco Systems