

Contents

Introduction	1
Prerequisites	1
Example: Configuring route selection based on the AS_PATH attribute	1
Network configuration	1
Analysis	2
Applicable hardware and software versions.....	2
Procedures	4
Configuring IP addresses for interfaces	4
Configuring BGP connections	4
Configuring routing policies	6
Verifying the configuration	6
Configuration files	7
Example: Configuring route selection based on the MED attribute	11
Network configuration	11
Analysis	11
Applicable hardware and software versions.....	11
Procedures	13
Configuring basic BGP	13
Configuring a routing policy	15
Verifying the configuration	15
Configuration files	15

Introduction

This document provides examples for configuring BGP route selection based on route attributes.

Prerequisites

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

The configuration examples in this document were created and verified in a lab environment, and all the devices were started with the factory default configuration. When you are working on a live network, make sure you understand the potential impact of every command on your network.

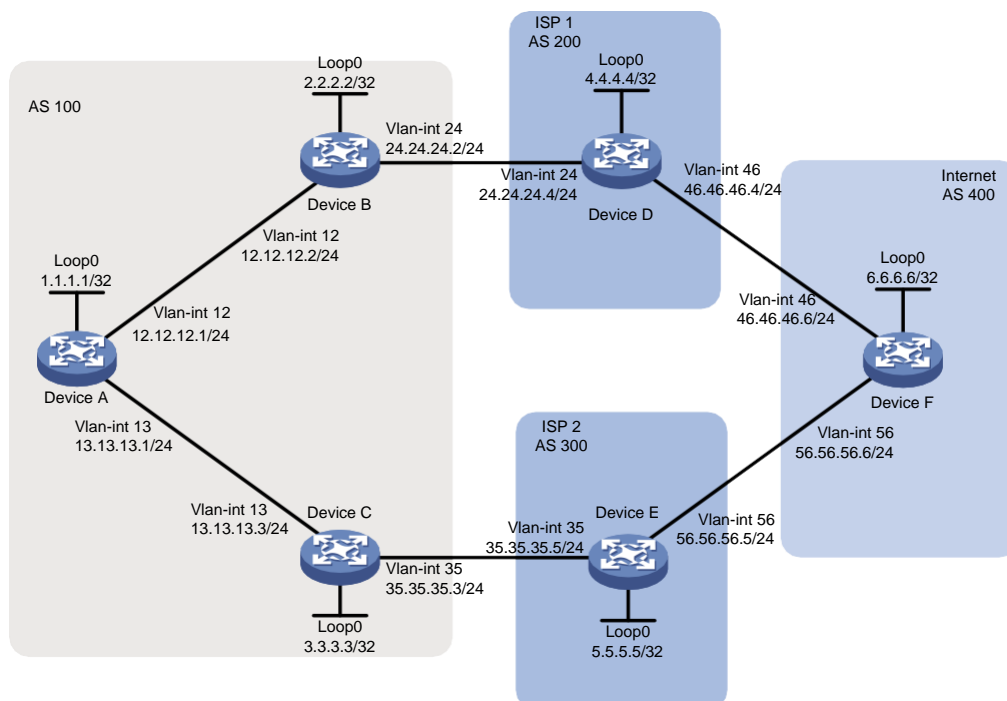
This document assumes that you have basic knowledge of BGP and routing policy.

Example: Configuring route selection based on the AS_PATH attribute

Network configuration

As shown in [Figure 1](#), all devices run BGP. Configure a routing policy on Device B and Device C so that traffic from AS 100 to AS 400 is preferentially forwarded by Device D.

Figure 1 Network diagram



Analysis

For devices in AS 100 to select the optimal route based on AS numbers, increase the local preference for routes whose AS_PATH attributes end with the specified AS number. Configure a routing policy on Device C to set the local preference to 300 for routes whose AS_PATH attributes end with AS number 400.

To filter routes based on AS numbers, use an AS path list.

Applicable hardware and software versions

The following matrix shows the hardware and software versions to which this configuration example is applicable:

Hardware	Software version
SC 3570 switch series	Release 11xx
SC 5525 switch series	Release 63xx, Release 65xx, Release 6615Pxx, Release 6628Pxx
SC 5520 switch series	Release 63xx, Release 65xx, Release 6615Pxx, Release 6628Pxx
SC 3170 switch series	Not supported
SC 3130 switch series	Not supported

Procedures

Configuring IP addresses for interfaces

Configure an IP address for the interface VLAN-interface 12 on Device A.

```
<DeviceA> system-view
[DeviceA] interface Vlan-interface 12
[DeviceA-Vlan-interface12] ip address 12.12.12.1 24
[DeviceA-Vlan-interface12] quit
```

Configure IP addresses for other interfaces as shown in [Figure 1](#). (Details not shown.)

Configuring BGP connections

On Device A, enable the default BGP instance, set the AS number to 100, and specify 12.12.12.2 and 13.13.13.3 as BGP peers.

```
[DeviceA] bgp 100
[DeviceA-bgp-default] router-id 1.1.1.1
[DeviceA-bgp-default] peer 12.12.12.2 as-number 100
[DeviceA-bgp-default] peer 13.13.13.3 as-number 100
[DeviceA-bgp-default] address-family ipv4 unicast
[DeviceA-bgp-default-ipv4] peer 12.12.12.2 enable
[DeviceA-bgp-default-ipv4] peer 13.13.13.3 enable
[DeviceA-bgp-default-ipv4] quit
[DeviceA-bgp-default] quit
```

On Device B, enable the default BGP instance, set the AS number to 100, specify 12.12.12.1 and 24.24.24.4 as BGP peers, and redistribute direct routes.

```
[DeviceB] bgp 100
[DeviceB-bgp-default] router-id 2.2.2.2
[DeviceB-bgp-default] peer 12.12.12.1 as-number 100
[DeviceB-bgp-default] peer 24.24.24.4 as-number 200
[DeviceB-bgp-default] address-family ipv4 unicast
[DeviceB-bgp-default-ipv4] peer 12.12.12.1 enable
[DeviceB-bgp-default-ipv4] peer 24.24.24.4 enable
[DeviceB-bgp-default-ipv4] import-route direct
[DeviceB-bgp-default-ipv4] quit
[DeviceB-bgp-default] quit
```

On Device C, enable the default BGP instance, set the AS number to 100, specify 13.13.13.1 and 35.35.35.5 as BGP peers, and redistribute direct routes.

```
[DeviceC] bgp 100
[DeviceC-bgp-default] router-id 3.3.3.3
```

```
[DeviceC-bgp-default] peer 13.13.13.1 as-number 100
[DeviceC-bgp-default] peer 35.35.35.5 as-number 300
[DeviceC-bgp-default] address-family ipv4 unicast
[DeviceC-bgp-default-ipv4] peer 13.13.13.1 enable
[DeviceC-bgp-default-ipv4] peer 35.35.35.5 enable
[DeviceC-bgp-default-ipv4] import-route direct
[DeviceC-bgp-default-ipv4] quit
[DeviceC-bgp-default] quit
```

On Device D, enable the default BGP instance, set the AS number to 200, specify 24.24.24.2 and 46.46.46.6 as BGP peers, and advertise the route 4.4.4.4/32.

```
[DeviceD] bgp 200
[DeviceD-bgp-default] router-id 4.4.4.4
[DeviceD-bgp-default] peer 24.24.24.2 as-number 100
[DeviceD-bgp-default] peer 46.46.46.6 as-number 400
[DeviceD-bgp-default] address-family ipv4 unicast
[DeviceD-bgp-default-ipv4] peer 24.24.24.2 enable
[DeviceD-bgp-default-ipv4] peer 46.46.46.6 enable
[DeviceD-bgp-default-ipv4] network 4.4.4.4 32
[DeviceD-bgp-default-ipv4] quit
[DeviceD-bgp-default] quit
```

On Device E, enable the default BGP instance, set the AS number to 300, specify 35.35.35.3 and 56.56.56.6 as BGP peers, and advertise the route 5.5.5.5/32.

```
[DeviceE] bgp 300
[DeviceE-bgp-default] router-id 5.5.5.5
[DeviceE-bgp-default] peer 35.35.35.3 as-number 100
[DeviceE-bgp-default] peer 56.56.56.6 as-number 400
[DeviceE-bgp-default] address-family ipv4 unicast
[DeviceE-bgp-default-ipv4] peer 35.35.35.3 enable
[DeviceE-bgp-default-ipv4] peer 56.56.56.6 enable
[DeviceE-bgp-default-ipv4] network 5.5.5.5 32
[DeviceE-bgp-default-ipv4] quit
[DeviceE-bgp-default] quit
```

On Device F, enable the default BGP instance, set the AS number to 400, specify 46.46.46.4 and 56.56.56.5 as BGP peers, and advertise the route 6.6.6.6/32.

```
[DeviceF] bgp 400
[DeviceF-bgp-default] router-id 6.6.6.6
[DeviceF-bgp-default] peer 46.46.46.4 as-number 200
[DeviceF-bgp-default] peer 56.56.56.5 as-number 300
[DeviceF-bgp-default] address-family ipv4 unicast
[DeviceF-bgp-default-ipv4] peer 46.46.46.4 enable
[DeviceF-bgp-default-ipv4] peer 56.56.56.5 enable
[DeviceF-bgp-default-ipv4] network 6.6.6.6 32
[DeviceF-bgp-default-ipv4] quit
[DeviceF-bgp-default] quit
```

Display the BGP routing table on Device A. The output shows the routes advertised by Device D, Device E, and Device F, and the AS_PATH attributes of the routes.

```
[DeviceA] display bgp routing-table ipv4
```

Total number of routes: 12

BGP local router ID is 1.1.1.1

Status codes: * - valid, > - best, d - dampened, h - history

s - suppressed, S - stale, i - internal, e - external

a - additional-path

Origin: i - IGP, e - EGP, ? - incomplete

Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
* >i 2.2.2.2/32	12.12.12.2	0	100	0	?
* >i 3.3.3.3/32	13.13.13.3	0	100	0	?
* >i 4.4.4.4/32	24.24.24.4	0	100	0	200i
* i	35.35.35.5		100	0	300 400i
					200i
* >i 5.5.5.5/32	35.35.35.5	0	100	0	300i
* i	24.24.24.4		100	0	200 400i
					300i
* >i 6.6.6.6/32	24.24.24.4		100	0	200 400i
* i	35.35.35.5		100	0	300 400i
* >i 12.12.12.0/24	12.12.12.2	0	100	0	?
* >i 13.13.13.0/24	13.13.13.3	0	100	0	?
* >i 24.24.24.0/24	12.12.12.2	0	100	0	?
* >i 35.35.35.0/24	13.13.13.3	0	100	0	?

Configuring routing policies

Create routing policy **aspath** on Device C, and set the local preference to 300 for routes whose AS_PATH attributes end with AS number 400.

```
[DeviceC] ip as-path 1 permit 400$
[DeviceC] route-policy aspath permit node 20
[DeviceC-route-policy-aspath-20] if-match as-path 1
[DeviceC-route-policy-aspath-20] apply local-preference 300
[DeviceC-route-policy-aspath-20] quit
[DeviceC] route-policy aspath permit node 25
```

Apply routing policy **aspath** to routes from the peer 35.35.35.5.

```
[DeviceC] bgp 100
[DeviceC-bgp-default] address-family ipv4
[DeviceC-bgp-default-ipv4] peer 35.35.35.5 route-policy aspath import
```

Verifying the configuration

Display the BGP routing table on Device A. The output shows that the next hop has changed for the route to AS 400.

```
[DeviceA] display bgp routing-table ipv4
```

Total number of routes: 11

BGP local router ID is 1.1.1.1

Status codes: * - valid, > - best, d - dampened, h - history

s - suppressed, S - stale, i - internal, e - external

a - additional-path

Origin: i - IGP, e - EGP, ? - incomplete

Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
* >i 2.2.2.2/32	12.12.12.2	0	100	0	?
* >i 3.3.3.3/32	13.13.13.3	0	100	0	?
* >i 4.4.4.4/32	24.24.24.4	0	100	0	200i
* i	35.35.35.5		100	0	300 400 200i
* >i 5.5.5.5/32	35.35.35.5	0	100	0	300i
* i	24.24.24.4		100	0	200 400 300i
* >i 6.6.6.6/32	35.35.35.5		300	0	300 400i
* i	24.24.24.4		100	0	200 400i
* >i 12.12.12.0/24	12.12.12.2	0	100	0	?
* >i 13.13.13.0/24	13.13.13.3	0	100	0	?
* >i 24.24.24.0/24	12.12.12.2	0	100	0	?
* >i 35.35.35.0/24	13.13.13.3	0	100	0	?

Verify that packets from Device A to 6.6.6.6 are forwarded by Device D.

[DeviceA] tracert 6.6.6.6

traceroute to 6.6.6.6 (6.6.6.6), 30 hops at most, 52 bytes each packet, press CT
RL+C to break

```
 1  12.12.12.2 (12.12.12.2)  2.417 ms  1.887 ms  1.773 ms
 2  35.35.35.5 (35.35.35.5)  4.057 ms  2.293 ms  2.739 ms
 3  6.6.6.6 (6.6.6.6)  5.145 ms  4.205 ms  4.402 ms
```

Configuration files

- Device A:

```
#
vlan 12
#
vlan 13
#
interface LoopBack0
 ip address 1.1.1.1 255.255.255.255
#
interface Vlan-interface12
 ip address 12.12.12.1 255.255.255.0
#
interface Vlan-interface13
 ip address 13.13.13.1 255.255.255.0
#
bgp 100
```

```

router-id 1.1.1.1
peer 12.12.12.2 as-number 100
peer 13.13.13.3 as-number 100
#
address-family ipv4 unicast
  peer 12.12.12.2 enable
  peer 13.13.13.3 enable
#

```

- **Device B:**

```

#
vlan 12
#
vlan 24
#
interface LoopBack0
  ip address 2.2.2.2 255.255.255.255
#
interface Vlan-interface12
  ip address 12.12.12.2 255.255.255.0
#
interface Vlan-interface24
  ip address 24.24.24.2 255.255.255.0
#
bgp 100
  router-id 2.2.2.2
  peer 12.12.12.1 as-number 100
  peer 24.24.24.4 as-number 200
#
  address-family ipv4 unicast
    import-route direct
    peer 12.12.12.1 enable
    peer 24.24.24.4 enable
#

```

- **Device C:**

```

#
vlan 13
#
vlan 35
#
interface LoopBack0
  ip address 3.3.3.3 255.255.255.255
#
interface Vlan-interface13
  ip address 13.13.13.3 255.255.255.0
#
interface Vlan-interface35
  ip address 35.35.35.3 255.255.255.0
#

```



```

bgp 100
  router-id 3.3.3.3
  peer 13.13.13.1 as-number 100
  peer 35.35.35.5 as-number 300
  #
  address-family ipv4 unicast
    import-route direct
    peer 13.13.13.1 enable
    peer 35.35.35.5 enable
    peer 35.35.35.5 route-policy aspath import
  #
  route-policy aspath permit node 20
    if-match as-path 1
    apply local-preference 300
  route-policy aspath permit node 25
  #
  ip as-path 1 permit 400$
  #

```

- **Device D:**

```

#
vlan 24
#
vlan 46
#
interface LoopBack0
  ip address 4.4.4.4 255.255.255.255
#
interface Vlan-interface24
  ip address 24.24.24.4 255.255.255.0
#
interface Vlan-interface46
  ip address 46.46.46.4 255.255.255.0
#
bgp 200
  router-id 4.4.4.4
  peer 24.24.24.2 as-number 100
  peer 46.46.46.6 as-number 400
  #
  address-family ipv4 unicast
    network 4.4.4.4 255.255.255.255
    peer 24.24.24.2 enable
    peer 46.46.46.6 enable
  #

```

- **Device E:**

```

#
vlan 35
#
vlan 56

```

```
#
interface LoopBack0
 ip address 5.5.5.5 255.255.255.255
#
interface Vlan-interface35
 ip address 35.35.35.5 255.255.255.0
#
interface Vlan-interface56
 ip address 56.56.56.5 255.255.255.0
#
bgp 300
 router-id 5.5.5.5
 peer 35.35.35.3 as-number 100
 peer 56.56.56.6 as-number 400
#
 address-family ipv4 unicast
  network 5.5.5.5 255.255.255.255
  peer 35.35.35.3 enable
  peer 56.56.56.6 enable
#
```

- **Device F:**

```
#
vlan 46
#
vlan 56
#
interface LoopBack0
 ip address 6.6.6.6 255.255.255.255
#
interface Vlan-interface46
 ip address 46.46.46.6 255.255.255.0
#
interface Vlan-interface56
 ip address 56.56.56.6 255.255.255.0
#
bgp 400
 router-id 6.6.6.6
 peer 46.46.46.4 as-number 200
 peer 56.56.56.5 as-number 300
#
 address-family ipv4 unicast
  network 6.6.6.6 255.255.255.255
  peer 46.46.46.4 enable
  peer 56.56.56.5 enable
#
```

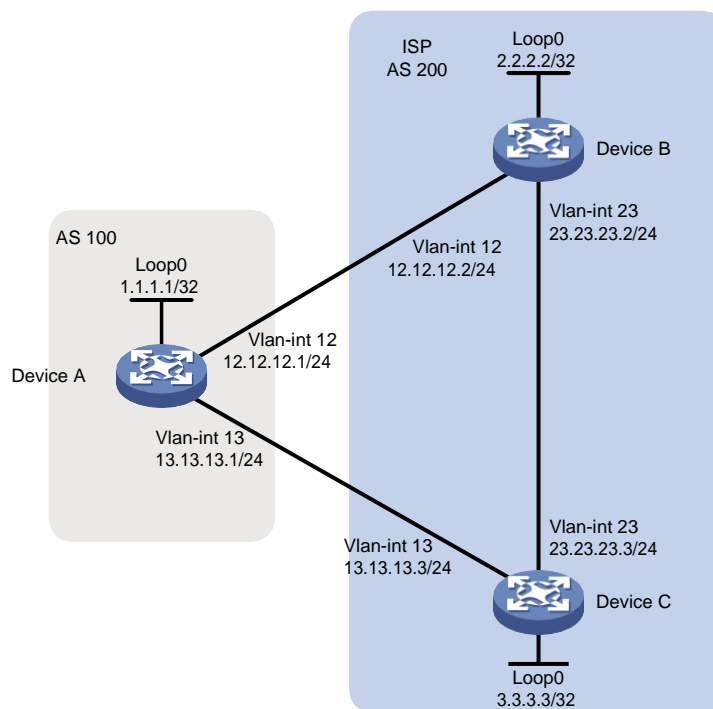
Example: Configuring route selection based on the MED attribute

Network configuration

As shown in [Figure 2](#), all devices run BGP. EBGP runs between Device A and Device B, and between Device A and Device C. IBGP runs between Device B and Device C.

Configure a routing policy to ensure that traffic from AS 100 to AS 200 is preferentially forwarded by Device C. Before you configure the routing policy, the traffic is preferentially forwarded by Device B.

Figure 2 Network diagram



Analysis

To ensure that the traffic is preferentially forwarded by Device C, configure a routing policy on Device B to change the MED value for the route to Device A. Make sure the MED value is not the default MED value 0.

Applicable hardware and software versions

The following matrix shows the hardware and software versions to which this configuration example is applicable:

Hardware	Software version
SC 3570 switch series	Release 11xx
SC 5525 switch series	Release 63xx, Release 65xx, Release 6615Pxx, Release 6628Pxx
SC 5520 switch series	Release 63xx, Release 65xx, Release 6615Pxx, Release 6628Pxx
SC 3170 switch series	Not supported
SC 3130 switch series	Not supported

Procedures

Configuring basic BGP

Configure an IP address for the interface VLAN-interface 12 on Device A.

```
<DeviceA> system-view
[DeviceA] interface Vlan-interface 12
[DeviceA-Vlan-interface12] ip address 12.12.12.1 24
[DeviceA-Vlan-interface12] quit
```

Configure IP addresses for other interfaces as shown in [Figure 2](#). (Details not shown.)

On Device A, enable the default BGP instance, set the AS number to 100, and specify 12.12.12.2 and 13.13.13.3 as BGP peers.

```
[DeviceA] bgp 100
[DeviceA-bgp-default] router-id 1.1.1.1
[DeviceA-bgp-default] peer 12.12.12.2 as-number 200
[DeviceA-bgp-default] peer 13.13.13.3 as-number 200
[DeviceA-bgp-default] address-family ipv4 unicast
[DeviceA-bgp-default-ipv4] peer 12.12.12.2 enable
[DeviceA-bgp-default-ipv4] peer 13.13.13.3 enable
[DeviceA-bgp-default-ipv4] quit
[DeviceA-bgp-default] quit
```

On Device B, enable the default BGP instance, set the AS number to 200, and specify 12.12.12.1 and 3.3.3.3 as BGP peers.

```
[DeviceB] bgp 200
[DeviceB-bgp-default] router-id 2.2.2.2
[DeviceB-bgp-default] peer 12.12.12.1 as-number 100
[DeviceB-bgp-default] peer 3.3.3.3 as-number 200
[DeviceB-bgp-default] peer 3.3.3.3 connect-interface LoopBack0
[DeviceB-bgp-default] address-family ipv4 unicast
[DeviceB-bgp-default-ipv4] peer 12.12.12.1 enable
[DeviceB-bgp-default-ipv4] peer 3.3.3.3 enable
[DeviceB-bgp-default-ipv4] network 23.23.23.0 24
[DeviceB-bgp-default-ipv4] quit
[DeviceB-bgp-default] quit
```

Configure a static route to 3.3.3.3/32 on Device B.

```
[DeviceB] ip route-static 3.3.3.3 32 23.23.23.3
```

On Device C, enable the default BGP instance, set the AS number to 200, and specify 13.13.13.1 and 2.2.2.2 as BGP peers.

```
[DeviceC] bgp 200
[DeviceC-bgp-default] router-id 3.3.3.3
[DeviceC-bgp-default] peer 13.13.13.1 as-number 100
[DeviceC-bgp-default] peer 2.2.2.2 as-number 200
[DeviceC-bgp-default] peer 2.2.2.2 connect-interface LoopBack0
[DeviceC-bgp-default] address-family ipv4 unicast
[DeviceC-bgp-default-ipv4] peer 13.13.13.1 enable
[DeviceC-bgp-default-ipv4] peer 2.2.2.2 enable
[DeviceC-bgp-default-ipv4] network 23.23.23.0 24
[DeviceC-bgp-default-ipv4] quit
[DeviceC-bgp-default] quit
```

Configure a static route to 2.2.2.2/32 on Device C.

```
[DeviceC] ip route-static 2.2.2.2 32 23.23.23.2
```

Display the BGP routing table on Device A. The output shows that the route with the next hop 12.12.12.2 becomes the optimal route to the network 23.23.23.0/24.

```
[DeviceA] display bgp routing-table ipv4
```

```
Total number of routes: 2
```

```
BGP local router ID is 1.1.1.1
```

```
Status codes: * - valid, > - best, d - dampened, h - history
               s - suppressed, S - stale, i - internal, e - external
               a - additional-path
Origin: i - IGP, e - EGP, ? - incomplete
```

	Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
* >	e 23.23.23.0/24	12.12.12.2	0		0	200i
* e		13.13.13.3	0		0	200i

Configuring a routing policy

Create routing policy 10 on Device B and set the cost to 100.

```
[DeviceB] route-policy 10 permit node 10
[DeviceB-route-policy-10-10] apply cost 100
[DeviceB-route-policy-10-10] quit
```

Apply routing policy 10 to routes to the peer 12.12.12.1.

```
[DeviceB] bgp 200
[DeviceB-bgp-default] address-family ipv4 unicast
[DeviceB-bgp-default-ipv4] peer 12.12.12.1 route-policy 10 export
[DeviceB-bgp-default-ipv4] quit
[DeviceB-bgp-default] quit
```

Verifying the configuration

Display the BGP routing table on Device A. The output shows that the MED value for the route with the next hop 12.12.12.2 changes to 100, and the route with the next hop 13.13.13.3 becomes the optimal route.

```
[DeviceA] display bgp routing-table ipv4
```

```
Total number of routes: 2
```

```
BGP local router ID is 1.1.1.1
```

```
Status codes: * - valid, >- best, d - dampened, h - history
               s - suppressed, S - stale, i - internal, e - external
               a - additional-path
Origin: i - IGP, e - EGP, ? - incomplete
```

	Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
* >	e 23.23.23.0/24	13.13.13.3	0		0	200i
* e		12.12.12.2	100		0	200i

Configuration files

- Device A:
#

```

vlan 12
#
vlan 13
#
interface LoopBack0
 ip address 1.1.1.1 255.255.255.255
#
interface Vlan-interface12
 ip address 12.12.12.1 255.255.255.0
#
interface Vlan-interface13
 ip address 13.13.13.1 255.255.255.0
#
bgp 100
 router-id 1.1.1.1
 peer 12.12.12.2 as-number 200
 peer 13.13.13.3 as-number 200
#
 address-family ipv4 unicast
  peer 12.12.12.2 enable
  peer 13.13.13.3 enable
#

```

- **Device B:**

```

#
vlan 12
#
vlan 23
#
interface LoopBack0
 ip address 2.2.2.2 255.255.255.255
#
interface Vlan-interface12
 ip address 12.12.12.2 255.255.255.0
#
interface Vlan-interface23
 ip address 23.23.23.2 255.255.255.0
#
bgp 200
 router-id 2.2.2.2
 peer 3.3.3.3 as-number 200
 peer 3.3.3.3 connect-interface LoopBack0
 peer 12.12.12.1 as-number 100
#
 address-family ipv4 unicast
  network 23.23.23.0 255.255.255.0
  peer 3.3.3.3 enable
  peer 12.12.12.1 enable
  peer 12.12.12.1 route-policy 10 export

```

```

#
route-policy 10 permit node 10
  apply cost 100
#
ip route-static 3.3.3.3 32 23.23.23.3
#
• Device C:
#
vlan 13
#
vlan 23
#
interface LoopBack0
  ip address 3.3.3.3 255.255.255.255
#
interface Vlan-interface13
  ip address 13.13.13.3 255.255.255.0
#
interface Vlan-interface23
  ip address 23.23.23.3 255.255.255.0
#
bgp 200
  router-id 3.3.3.3
  peer 2.2.2.2 as-number 200
  peer 2.2.2.2 connect-interface LoopBack0
  peer 13.13.13.1 as-number 100
#
address-family ipv4 unicast
  network 23.23.23.0 255.255.255.0
  peer 2.2.2.2 enable
  peer 13.13.13.1 enable
#
ip route-static 2.2.2.2 32 23.23.23.2
#

```