

ASA/PIX: BGP through ASA Configuration Example

Document ID: 6500

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Related Products
- Conventions

Configure

- Network Diagram
- Scenario 1
- Scenario 2

MD5 Authentication for BGP Neighbors through the PIX/ASA

- PIX 6.x Configuration
- PIX / ASA 7.x and Later
- Verify

Related Information

Introduction

This sample configuration demonstrates how to run Border Gateway Protocol (BGP) across a Security Appliance (PIX/ASA) and how to achieve redundancy in a multihomed BGP and PIX environment. With a network diagram as an example, this document explains how to automatically route traffic to Internet service provider B (ISP-B) when AS 64496 loses connectivity to ISP-A (or the reverse), through the use of dynamic routing protocols that run between all routers in AS 64496.

Because BGP uses unicast TCP packets on port 179 to communicate with its peers, you can configure PIX1 and PIX2 to allow unicast traffic on TCP port 179. This way, BGP peering can be established between the routers that are connected through the firewall. Redundancy and the desired routing policies can be achieved through the manipulation of the BGP attributes.

Prerequisites

Requirements

Readers of this document should be familiar with Configuring BGP and Basic Firewall Configuration.

Components Used

The example scenarios in this document are based on these software versions:

- Cisco 2600 routers with Cisco IOS® Software Release 12.2(27)
- PIX 515 with Cisco PIX Firewall Version 6.3(3) and later

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, make sure

that you understand the potential impact of any command.

Related Products

This configuration can also be used with these hardware and software versions:

- Cisco Adaptive Security Appliance (ASA) 5500 Series with 7.x version and later
- Cisco Firewall Services Module (FWSM) that runs software version 3.2 and later

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

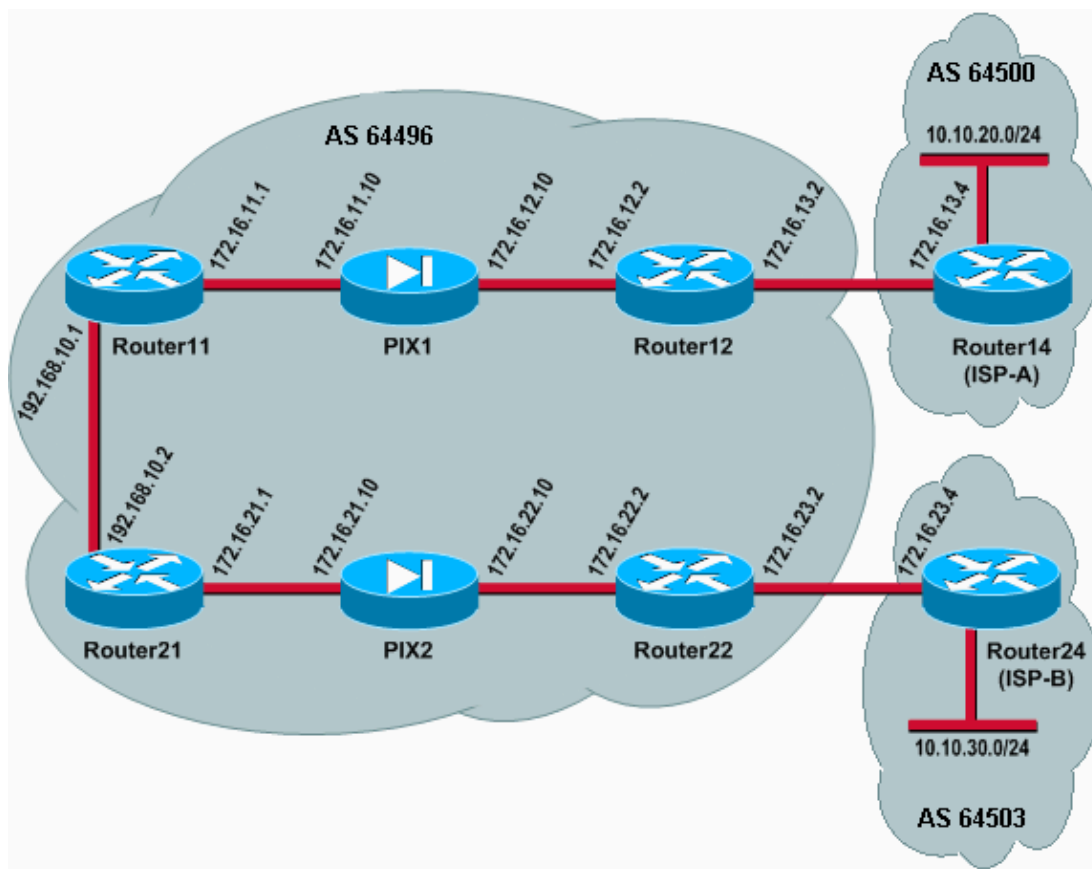
Configure

This section provides information to configure the features described in this document.

Note: To find additional information about the commands in this document, use the Command Lookup Tool (registered customers only) .

Network Diagram

This document uses this network setup:



In this network setup, Router12 and Router22 (which belong to AS 64496) are multihomed to Router14 (ISP-A) and Router24 (ISP-B) respectively, for redundancy. The internal network 192.168.10.0/24 is on the inside of the firewall. Router11 and Router21 connect to Router12 and Router22 through the firewall. PIX1

and PIX2 are not configured to perform Network Address Translation (NAT).

Scenario 1

In this scenario, Router12 in AS 64496 does external BGP (eBGP) peering with Router14 (ISP-A) in AS 64500. Router12 also does internal BGP (iBGP) peering with Router11 through PIX1. If eBGP learned routes from ISP-A are present, Router12 announces a default route 0.0.0.0/0 on iBGP to Router11. If the link to ISP-A fails, Router12 stops announcing the default route.

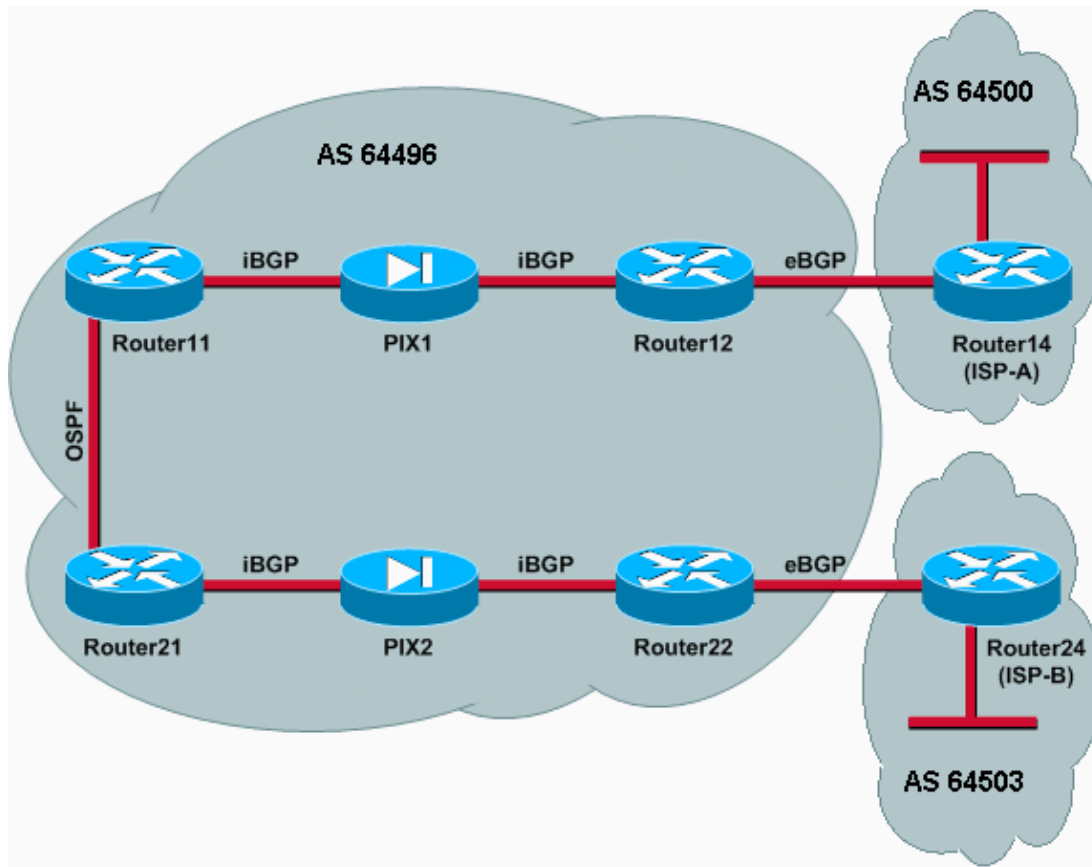
Similarly, Router22 in AS 64496 does eBGP peering with Router24 (ISP-B) in AS 64503 and announces a default route on iBGP to Router21 conditionally based on the presence of ISP-B routes in its routing table.

Through the use of an access list, PIX1 and PIX2 are configured to allow the BGP traffic (TCP, port 179) between iBGP peers. This is because PIX interfaces have an associated security level. By default, the inside interface (ethernet1) has a security level 100 and the outside interface (ethernet0) has a security level 0. Connections and traffic are normally permitted from higher to lower security level interfaces. To permit traffic from a lower security level interface to a higher security level interface, however, you must explicitly define an access list on the PIX. Also, you must configure a static NAT translation on PIX1 and PIX2, to allow routers on the outside to initiate a BGP session with routers on the inside of PIX.

Both Router11 and Router21 conditionally announce the default route into the Open Shortest Path First (OSPF) domain based on the iBGP-learned default route. Router11 announces the default route into the OSPF domain with a metric of 5, Router21 announces the default route with a metric of 30, and therefore the default route from Router11 is preferred. This configuration helps propagate only the default route 0.0.0.0/0 to Router11 and Router21, which conserves memory consumption on the inside routers and achieves optimum performance.

Thus, to summarize these conditions, this is the routing policy for AS 64496:

- AS 64496 prefers the link from Router12 to ISP-A for all outbound traffic (from 192.168.10.0/24 to the Internet).
- If connectivity to ISP-A fails, all traffic is routed via the link from Router22 to ISP-B.
- All traffic that comes from the Internet to 192.168.10.0/24 uses the link from ISP-A to Router12.
- If the link from ISP-A to Router12 fails, all inbound traffic is routed via the link from ISP-B to Router22.



Configurations

This scenario uses these configurations:

- Router11
- Router12
- Router14 (ISP-A)
- Router21
- Router22
- PIX1
- PIX2

Router11

```
hostname Router11
!
interface FastEthernet0/0
 ip address 192.168.10.1 255.255.255.0

!--- Connected to Router21.

!
interface FastEthernet0/1
 ip address 172.16.11.1 255.255.255.0

!--- Connected to PIX1.

!
router ospf 1
 log-adjacency-changes
 network 192.168.10.0 0.0.0.255 area 0
 default-information originate metric 5 route-map check-default
```

```

!--- A default route is advertised into OSPF conditionally (based on whether the link
!--- from Router12 to ISP-A is active), with a metric of 5.

router bgp 64496
  no synchronization
  bgp log-neighbor-changes
  network 192.168.10.0
  neighbor 172.16.12.2 remote-as 64496

!--- Configures Router12 as an iBGP peer
.
  distance bgp 20 105 200

!--- Administrative distance of iBGP learned routes is changed from default 200 to 105.

  no auto-summary
  !
  ip route 172.16.12.0 255.255.255.0 172.16.11.10

!--- Static route to iBGP peer, because it is not directly connected.

  !
  access-list 30 permit 0.0.0.0
  access-list 31 permit 172.16.12.2
  route-map check-default permit 10
  match ip address 30
  match ip next-hop 31

```

Router12

```

hostname Router12
!
interface FastEthernet0/0
  ip address 172.16.13.2 255.255.255.0

!--- Connected to Router14 (ISP-A).

!
interface FastEthernet0/1
  ip address 172.16.12.2 255.255.255.0

!--- Connected to PIX1.

!
router bgp 64496
  no synchronization
  neighbor 172.16.11.1 remote-as 64496
  neighbor 172.16.11.1 next-hop-self
  neighbor 172.16.11.1 default-originate route-map
  check-ispa-route

!--- A default route is advertised to Router11 conditionally (based on whether the link
!--- from Router12 to ISP-A is active).

  neighbor 172.16.11.1 distribute-list 1 out
  neighbor 172.16.13.4 remote-as 64500

!--- Configures Router14 (ISP-A) as an eBGP peer.

  neighbor 172.16.13.4 route-map adv-to-ispa out
  no auto-summary
  !
  ip route 172.16.11.0 255.255.255.0 172.16.12.10

```

```

!--- Static route to iBGP peer, because it is not directly connected.
!
access-list 1 permit 0.0.0.0
access-list 10 permit 192.168.10.0
access-list 20 permit 10.10.20.0 0.0.0.255
access-list 21 permit 172.16.13.4
!
route-map check-isp-a-route permit 10
  match ip address 20
  match ip next-hop 21
!
route-map adv-to-isp-a permit 10
  match ip address 10

```

Router14 (ISP-A)

```

hostname Router14
!
interface Ethernet0/0
  ip address 172.16.13.4 255.255.255.0
!
interface Ethernet0/1
  ip address 10.10.20.1 255.255.255.0
!
router bgp 64500
  network 10.10.20.0 mask 255.255.255.0
  neighbor 172.16.13.2 remote-as 64496

!--- Configures Router12 as an eBGP peer.
!

```

Router21

```

hostname Router21
!
interface FastEthernet0/0
  ip address 192.168.10.2 255.255.255.0

!--- Connected to Router11.
!
interface FastEthernet0/1
  ip address 172.16.21.1 255.255.255.0

!--- Connected to PIX2.
!
router ospf 1
  network 192.168.10.0 0.0.0.255 area 0
  default-information originate metric 30 route-map
  check-default

!--- A default route is advertised into OSPF conditionally (based on whether the link
!--- from Router22 to ISP-B is active), with a metric of 30.
!
router bgp 64496
  no synchronization
  network 192.168.10.0
  neighbor 172.16.22.2 remote-as 64496

!--- Configures Router22 as an iBGP peer.

```

```
!  
ip route 172.16.22.0 255.255.255.0 172.16.21.10  
  
!--- Static route to iBGP peer, because it is not directly connected.  
  
!  
access-list 30 permit 0.0.0.0  
access-list 31 permit 172.16.22.2  
route-map check-default permit 10  
  match ip address 30  
  match ip next-hop 31  
!
```

Router22

```
hostname Router22  
!  
interface FastEthernet0/0  
  ip address 172.16.23.2 255.255.255.0  
  
!--- Connected to Router24 (ISP-B).  
  
!  
interface FastEthernet0/1  
  ip address 172.16.22.2 255.255.255.0  
  
!--- Connected to PIX2.  
  
!  
router bgp 64496  
no synchronization  
  bgp log-neighbor-changes  
  neighbor 172.16.21.1 remote-as 64496  
  
!--- Configure Router21 as an iBGP peer.  
  
  neighbor 172.16.21.1 next-hop-self  
  neighbor 172.16.21.1 default-originate route-map  
  check-ispb-route  
  
!--- A default route is advertised to Router21 conditionally (based on whether the link  
!--- from Router22 to ISP-B is active).  
  
!  
  neighbor 172.16.21.1 distribute-list 1 out  
  neighbor 172.16.23.4 remote-as 64503  
  neighbor 172.16.23.4 route-map adv-to-ispb out  
!  
ip route 172.16.21.0 255.255.255.0 172.16.22.10  
  
!--- Static route to iBGP peer, because it is not directly connected.  
  
!  
access-list 1 permit 0.0.0.0  
access-list 10 permit 192.168.10.0  
access-list 20 permit 10.10.30.0 0.0.0.255  
access-list 21 permit 172.16.23.4  
!  
route-map check-ispb-route permit 10  
  match ip address 20  
  match ip next-hop 21  
!  
route-map adv-to-ispb permit 10  
  match ip address 10  
  set as-path prepend 10 10 10
```

```
!--- Route map used to change the AS path attribute of outgoing updates.
```

Router24 (ISP-B)

```
hostname Router24
!
interface Loopback0
 ip address 10.10.30.1 255.255.255.0
!
interface FastEthernet0/0
 ip address 172.16.23.4 255.255.255.0
!
router bgp 64503
 bgp log-neighbor-changes
 network 10.10.30.0 mask 255.255.255.0
 neighbor 172.16.23.2 remote-as 64496

!--- Configures Router22 as an eBGP peer.

!
```

PIX1

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 172.16.12.10 255.255.255.0
ip address inside 172.16.11.10 255.255.255.0

!--- Configures the IP addresses for the inside and outside interfaces.

access-list acl-1 permit tcp host 172.16.12.2 host 172.16.11.1 eq bgp

!--- Access list allows BGP traffic to pass from outside to inside.

access-list acl-1 permit icmp any any

!--- Allows ping to pass through for testing purposes only.

access-group acl-1 in interface outside
nat (inside) 0 0.0.0.0 0.0.0.0 0 0

!--- No NAT translation, to allow Router11 on the inside to initiate a BGP session
!--- to Router12 on the outside of PIX.

static (inside,outside) 172.16.11.1 172.16.11.1 netmask 255.255.255.255

!--- Static NAT translation, to allow Router12 on the outside to initiate a BGP session
!--- to Router11 on the inside of PIX.

route outside 0.0.0.0 0.0.0.0 172.16.12.2 1
route inside 192.168.10.0 255.255.255.0 172.16.11.1 1
```

PIX2

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 172.16.22.10 255.255.255.0
ip address inside 172.16.21.10 255.255.255.0

!--- Configures the IP addresses for the inside and outside interfaces.

access-list acl-1 permit tcp host 172.16.22.2 host 172.16.21.1 eq bgp
```



```

!--- Access list allows BGP traffic to pass from outside to inside.

access-list acl-1 permit icmp any any

!--- Allows ping to pass through for testing purposes only.

access-group acl-1 in interface outside
route outside 0.0.0.0 0.0.0.0 172.16.22.2 1
route inside 192.168.10.0 255.255.255.0 172.16.21.1 1
nat (inside) 0 0.0.0.0 0.0.0.0 0 0

!--- No NAT translation, to allow Router21 on the inside to initiate a BGP session
!--- to Router22 on the outside of PIX.

static (inside,outside) 172.16.21.1 172.16.21.1 netmask 255.255.255.255

! -- Static NAT translation, to allow Router22 on the outside to initiate a BGP session
!--- to Router21 on the inside of PIX.

```

Verify

Use this section to confirm that your configuration works properly.

The Output Interpreter Tool (registered customers only) (OIT) supports certain **show** commands. Use the OIT to view an analysis of **show** command output.

When both BGP sessions are up, you can expect all packets to be routed via ISP-A. Consider the BGP table on Router11. It learns a default route 0.0.0.0/0 from Router12 with the next hop 172.16.12.2.

```

Router11# show ip bgp

BGP table version is 14, local router ID is 192.168.10.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
  *>i0.0.0.0        172.16.12.2           100      0  i
  *> 192.168.10.0   0.0.0.0              0         32768  i

```

The 0.0.0.0/0 default route that is learned via BGP is installed in the routing table, as shown in the output of **show ip route** on Router11.

```

Router11# show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

Gateway of last resort is 172.16.12.2 to network 0.0.0.0

C    192.168.10.0/24 is directly connected, FastEthernet0/0
     172.16.0.0/24 is subnetted, 2 subnets
S      172.16.12.0 [1/0] via 172.16.11.10
C    172.16.11.0 is directly connected, FastEthernet0/1
B*   0.0.0.0/0 [105/0] via 172.16.12.2, 00:27:24

```

Now consider the BGP table on Router21. It also learns the default route via Router22.

```
Router21# show ip bgp

BGP table version is 8, local router ID is 192.168.10.2
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
*>i0.0.0.0          172.16.22.2      100      0 i
*> 192.168.10.0     0.0.0.0          0          32768
```

Now see if this BGP-learned default route gets installed in the routing table of Router21.

```
Router21# show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

Gateway of last resort is 192.168.10.1 to network 0.0.0.0

C    192.168.10.0/24 is directly connected, FastEthernet0/0
     172.16.0.0/24 is subnetted, 2 subnets
C      172.16.21.0 is directly connected, FastEthernet0/1
S      172.16.22.0 [1/0] via 172.16.21.10
O*E2 0.0.0.0/0 [110/5] via 192.168.10.1, 00:27:06, FastEthernet0/0
```

The default route in Router21 is learned via OSPF (note the o prefix on the 0.0.0.0/0 route). It is interesting to note that there is a default route learned via BGP from Router22, but the **show ip route** output shows the default route learned via OSPF.

The OSPF default route was installed in Router21 because Router21 learns the default route from two sources: Router22 via iBGP and Router11 via OSPF. The route selection process installs the route with a better administrative distance into the routing table. The administrative distance of OSPF is 110 while the administrative distance of iBGP is 200. Therefore, the OSPF-learned default route gets installed in the routing table, because 110 is less than 200. For more information on route selection, refer to Route Selection in Cisco Routers.

Troubleshoot

Use this section to troubleshoot your configuration.

Bring down the BGP session between Router12 and ISP-A.

```
Router12(config)# interface fas 0/0

Router12(config-if)# shut

1w0d: %LINK-5-CHANGED: Interface FastEthernet0/0,
      changed state to administratively down
1w0d: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0,
      changed state to down
```

Router11 does not have the default route learned via BGP from Router12.

```
Router11# show ip bgp
```

```
BGP table version is 16, local router ID is 192.168.10.1
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
*> 192.168.10.0   0.0.0.0           0
```

Check the routing table on Router11. The default route is learned via OSPF (administrative distance of 110) with a next hop of Router21.

```
Router11# show ip route
```

```
!--- Output suppressed.
```

```
Gateway of last resort is 192.168.10.2 to network 0.0.0.0

C    192.168.10.0/24 is directly connected, FastEthernet0/0
     172.16.0.0/24 is subnetted, 2 subnets
S     172.16.12.0 [1/0] via 172.16.11.10
C     172.16.11.0 is directly connected, FastEthernet0/1
O*E2 0.0.0.0/0 [110/30] via 192.168.10.2, 00:00:09, FastEthernet0/0
```

This output is expected as per the pre-defined policies. At this point, however, it is important to understand the **distance bgp 20 105 200** configuration command in Router11 and how it influences the route selection on Router11.

The default values of this command are **distance bgp 20 200 200**, where eBGP-learned routes have an administrative distance of 20, iBGP-learned routes have an administrative distance of 200, and local BGP routes have an administrative distance of 200.

When the link between Router12 and ISP-A comes up again, Router11 learns the default route via iBGP from Router12. However, because the default administrative distance of this iBGP-learned route is 200, it will not replace the OSPF-learned route (because 110 is less than 200). This forces all of the outbound traffic to the link from Router21 to Router22 to ISP-B, even though the link from Router12 to ISP-A is up again. To solve this issue, change the administrative distance of the iBGP-learned route to a value less than the Interior Gateway Protocol (IGP) used. In this example, the IGP is OSPF, so a distance of 105 was chosen (because 105 is less than 110).

For more information on the **distance bgp** command, refer to BGP Commands. For more information on multihoming with BGP, refer to Load Sharing with BGP in Single and Multihomed Environments: Sample Configurations.

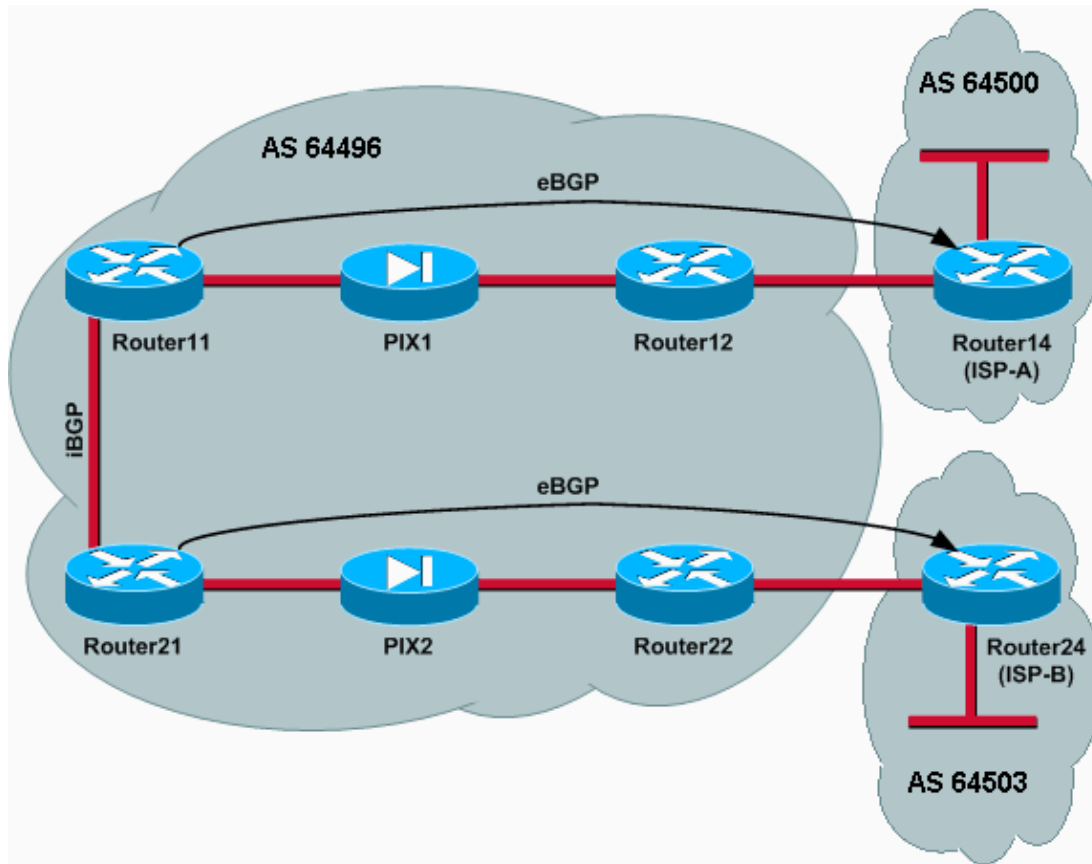
Scenario 2

In this scenario, Router11 is directly eBGP peering with Router 14 (ISP-A), and Router21 is directly eBGP peering with Router24 (ISP-B). Router12 and Router22 do not participate in BGP peering, but they do provide the IP connectivity to the ISPs. Because the eBGP peers are not directly connected neighbors, the **neighbor ebgp-multihop** command is used on the participating routers. The **neighbor ebgp-multihop** command enables BGP to override the default one hop eBGP limit because it changes the Time to Live (TTL) of eBGP packets from the default value of 1. In this scenario, the eBGP neighbor is 3 hops away, so **neighbor ebgp-multihop 3** is configured on the participating routers so that the TTL value is changed to 3. Also, static routes are configured on the routers and PIX to ensure that Router11 can ping the Router14 (ISP-A) address 172.16.13.4 and to ensure that Router21 can ping the Router24 (ISP-B) address 172.16.23.4.

By default, PIX does not allow Internet Control Message Protocol (ICMP) packets (sent when you issue the **ping** command) to pass through. To allow ICMP packets, use the **access-list** command as shown in in the next

PIX configuration. For more information on the **access-list** command, refer to the PIX Firewall A through B Commands.

The routing policy is the same as in Scenario 1: the link between Router12 and ISP-A is preferred over the link between Router22 and ISP-B, and when the ISP-A link goes down the ISP-B link is used for all inbound and outbound traffic.



Configurations

This scenario uses these configurations:

- Router11
- Router12
- Router14 (ISP-A)
- Router21
- Router22
- PIX1
- PIX2

Router11

```
hostname Router11
!
interface FastEthernet0/0
 ip address 192.168.10.1 255.255.255.0

!--- Connected to Router21.

!
interface FastEthernet0/1
```

```

ip address 172.16.11.1 255.255.255.0

!--- Connected to PIX1.

!
router bgp 64496
no synchronization
bgp log-neighbor-changes
network 192.168.10.0
neighbor 172.16.13.4 remote-as 64500
neighbor 172.16.13.4 ebgp-multihop 3

!--- To accept and attempt BGP connections to external peers that reside on networks that
!--- are not directly connected.

neighbor 172.16.13.4 route-map set-pref in

!--- Sets higher local-preference for learned routes.

neighbor 172.16.13.4 route-map adv_to_ispa out
neighbor 192.168.10.2 remote-as 64496
neighbor 192.168.10.2 next-hop-self
no auto-summary
!
ip route 172.16.12.0 255.255.255.0 172.16.11.10
ip route 172.16.13.4 255.255.255.255 172.16.11.10

!--- Static route to eBGP peer, because it is not directly connected.

!
access-list 20 permit 192.168.10.0
!
route-map set-pref permit 10
set local-preference 200
!
route-map adv_to_ispa permit 10
match ip address 20
!

```

Router12

```

hostname Router12
!
interface FastEthernet0/0
ip address 172.16.13.2 255.255.255.0

!--- Connected to ISP-A.

!
interface FastEthernet0/1
ip address 172.16.12.2 255.255.255.0

!--- Connected to PIX1.

!
ip route 172.16.11.0 255.255.255.0 172.16.12.10
ip route 192.168.10.0 255.255.255.0 172.16.12.10

```

Router14 (ISP-A)

```

hostname Router14
!
interface Ethernet0/0
ip address 172.16.13.4 255.255.255.0
!

```

```

interface Ethernet0/1
 ip address 10.10.20.1 255.255.255.0
!
router bgp 64500
no synchronization
network 10.10.20.0 mask 255.255.255.0
 neighbor 172.16.11.1 remote-as 64496
 neighbor 172.16.11.1 ebgp-multihop 3

!--- To accept and attempt BGP connections to external peers that reside on networks that
!--- are not directly connected.

 neighbor 172.16.11.1 default-originate

!--- Advertises a default route to Router11.

no auto-summary
!
ip route 172.16.11.1 255.255.255.255 172.16.13.2

!--- Static route to eBGP peers, because it is not directly connected.

```

Router21

```

hostname Router21
!
interface FastEthernet0/0
 ip address 192.168.10.2 255.255.255.0

!--- Connected to Router11.

!
interface FastEthernet0/1
 ip address 172.16.21.1 255.255.255.0

!--- Connected to PIX2.

!
router bgp 64496
 no synchronization
 network 192.168.10.0
 neighbor 172.16.23.4 remote-as 64503
 neighbor 172.16.23.4 ebgp-multihop 3

!--- To accept and attempt BGP connections to external peers that reside on networks that
!--- are not directly connected.

 neighbor 172.16.23.4 route-map adv_to_ispb out
 neighbor 192.168.10.1 remote-as 64496
 neighbor 192.168.10.1 next-hop-self
 no auto-summary
!
ip route 172.16.22.0 255.255.255.0 172.16.21.10
ip route 172.16.23.4 255.255.255.255 172.16.21.10

!--- Static routes configured to reach BGP peer.

!
access-list 20 permit 192.168.10.0
!
route-map adv_to_ispb permit 10
 match ip address 20
 set as-path prepend 10 10 10

```

Router22

```
hostname Router22
!
interface FastEthernet0/0
 ip address 172.16.23.2 255.255.255.0

!--- Connected to Router24 (ISP-B).

!
interface FastEthernet0/1
 ip address 172.16.22.2 255.255.255.0

!--- Connected to PIX2.

!
ip route 172.16.21.0 255.255.255.0 172.16.22.10
ip route 192.168.10.0 255.255.255.0 172.16.22.10
```

Router24 (ISP-B)

```
hostname Router24
!
interface Loopback0
 ip address 10.10.30.1 255.255.255.0
!
interface FastEthernet0/0
 ip address 172.16.23.4 255.255.255.0

!--- Connected to Router22.

!
router bgp 64503
 no synchronization
 bgp log-neighbor-changes
 network 10.10.30.0 mask 255.255.255.0
  neighbor 172.16.21.1 remote-as 64496
  neighbor 172.16.21.1 ebgp-multihop 3

!--- To accept and attempt BGP connections to external peers that reside on networks that
!--- are not directly connected.

neighbor 172.16.21.1 default-originate

!--- Advertises a default route to Router21.

no auto-summary
!
ip route 172.16.21.1 255.255.255.255 172.16.23.2

!--- Static route for BGP peer Router11, because it is not directly connected.
```

PIX1

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 172.16.12.10 255.255.255.0
ip address inside 172.16.11.10 255.255.255.0
access-list acl-1 permit tcp host 172.16.13.4 host 172.16.11.1 eq bgp

!-- Access list allows BGP traffic to pass from outside to inside.

access-list acl-1 permit icmp any any
```

```
!-- Allows ping to pass through for testing purposes only.

access-group acl-1 in interface outside
nat (inside) 0 0.0.0.0 0.0.0.0 0 0
static (inside,outside) 172.16.11.1 172.16.11.1 netmask 255.255.255.255
route outside 0.0.0.0 0.0.0.0 172.16.12.2 1
route inside 192.168.10.0 255.255.255.0 172.16.11.1 1
```

PIX2

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 172.16.22.10 255.255.255.0
ip address inside 172.16.21.10 255.255.255.0
access-list acl-1 permit tcp host 172.16.23.4 host 172.16.21.1 eq bgp

!-- Access list allows BGP traffic to pass from outside to inside.

access-list acl-1 permit icmp any any

!-- Allows ping to pass through for testing purposes only.

access-group acl-1 in interface outside
route outside 0.0.0.0 0.0.0.0 172.16.22.2 1
route inside 192.168.10.0 255.255.255.0 172.16.21.1 1
nat (inside) 0 0.0.0.0 0.0.0.0 0 0
static (inside,outside) 172.16.21.1 172.16.21.1 netmask 255.255.255.255
```

Verify

Begin with the situation where the links to ISP-A and ISP-B are up. The **show ip bgp summary** command output on Router11 and Router21 confirms the established BGP sessions with ISP-A and ISP-B respectively.

```
Router11# show ip bgp summary
```

```
BGP router identifier 192.168.10.1, local AS number 10
BGP table version is 13, main routing table version 13
4 network entries and 5 paths using 568 bytes of memory
7 BGP path attribute entries using 420 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 43/264 prefixes, 75/70 paths, scan interval 15 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/
172.16.13.4	4	64500	1627	1623	13	0	0	02:13:36	2
192.168.10.2	4	64496	1596	1601	13	0	0	02:08:47	2

```
Router21# show ip bgp summary
```

!--- Output suppressed.

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.23.4	4	64503	1610	1606	8	0	0	02:06:22	2
192.168.10.1	4	64496	1603	1598	8	0	0	02:10:16	3

The BGP table on Router11 shows the default route (0.0.0.0/0) toward the next hop ISP-A 172.16.13.4.

```
Router11# show ip bgp
```

```
BGP table version is 13, local router ID is 192.168.10.1
```


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	172.16.13.4		200	0	20 i
*> 10.10.20.0/24	172.16.13.4	0	200	0	64500 i
*>i10.10.30.0/24	192.168.10.2	0	100	0	64503 i
* i192.168.10.0	192.168.10.2	0	100	0	i
*>	0.0.0.0	0		32768	i

Now check the BGP table on Router21. It has two 0.0.0.0/0 routes: one learned from ISP-B with a next hop of 172.16.23.4 on eBGP, and the other learned via iBGP with a local-preference of 200. Router21 prefers iBGP-learned routes because of the higher local-preference attribute, so it installs that route in the routing table. For more information on BGP path selection, refer to BGP Best Path Selection Algorithm.

```
Router21# show ip bgp
```

```
BGP table version is 8, local router ID is 192.168.10.2  
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	172.16.23.4			0	64503 i
*>i	192.168.10.1		200	0	64500 i
*>i10.10.20.0/24	192.168.10.1	0	200	0	64500 i
*> 10.10.30.0/24	172.16.23.4	0		0	64503 i
*> 192.168.10.0	0.0.0.0	0		32768	i
* i	192.168.10.1	0	100	0	i

Troubleshoot

Bring down the Router11 and ISP-A BGP session.

```
Router11(config)# interface fas 0/1
```

```
Router11(config-if)# shut
```

```
4w2d: %LINK-5-CHANGED: Interface FastEthernet0/1,  
changed state to administratively down
```

```
4w2d: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1,  
changed state to down
```

```
4w2d: %BGP-5-ADJCHANGE: neighbor 172.16.13.4 Down BGP Notification sent
```

```
4w2d: %BGP-3-NOTIFICATION: sent to neighbor 172.16.13.4 4/0 (hold time expired)0 bytes
```

The eBGP session to ISP-A goes down when the hold-down timer (180 seconds) expires.

```
Router11# show ip bgp summary
```

```
!--- Output suppressed.
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.13.4	4	64500	1633	1632	0	0	0	00:00:58	Active
192.168.10.2	4	64496	1609	1615	21	0	0	02:18:09	

With the link to ISP-A down, Router11 installs 0.0.0.0/0 with a next hop of 192.168.10.2 (Router21), which is learned via iBGP in its routing table. This pushes all outbound traffic through Router21 and then to ISP-B, as shown in this output:

```
Router11# show ip bgp
```

```
BGP table version is 21, local router ID is 192.168.10.1
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	192.168.10.2		100	0	64503 i
*>i10.10.30.0/24	192.168.10.2	0	100	0	64503 i
* i192.168.10.0	192.168.10.2	0	100	0	i
*>	0.0.0.0	0		32768	i

```
Router21# show ip bgp
```

```
BGP table version is 14, local router ID is 192.168.10.2
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	172.16.23.4			0	64503 i
*> 10.10.30.0/24	172.16.23.4	0		0	64503 i
*> 192.168.10.0	0.0.0.0	0		32768	i
* i	192.168.10.1	0	100	0	i

MD5 Authentication for BGP Neighbors through the PIX/ASA

PIX 6.x Configuration

Just like any other routing protocol, BGP can be configured for authentication. You can configure MD5 authentication between two BGP peers, which means that each segment sent on the TCP connection between the peers is verified. MD5 authentication must be configured with the same password on both BGP peers; otherwise, the connection between them will not be made. The configuration of MD5 authentication causes Cisco IOS software to generate and check the MD5 digest of every segment sent on the TCP connection. If authentication is invoked and a segment fails authentication, an error message is generated.

When you are configuring BGP peers with MD5 authentication that pass through a PIX firewall, it is important to configure the PIX between the BGP neighbors so that the sequence numbers for the TCP flows between the BGP neighbors are not random. This is because the TCP random sequence number feature on the PIX firewall is enabled by default, and it changes the TCP sequence number of the incoming packets before it forwards them.

MD5 authentication is applied on the TCP pseudo-IP header, TCP header and data (refer to RFC 2385 [↗](#)). TCP uses this data—which includes the TCP sequence and ACK numbers—along with the BGP neighbor password to create a 128 bit hash number. The hash number is included in the packet in a TCP header option field. By default, the PIX offsets the sequence number by a random number, per TCP flow. On the sending BGP peer, TCP uses the original sequence number to create the 128 bit MD5 hash number and includes this hash number in the packet. When the receiving BGP peer gets the packet, TCP uses the PIX-modified sequence number to create a 128 bit MD5 hash number and compares it to the hash number that is included in the packet.

The hash number is different because the TCP sequence value was changed by the PIX, and TCP on the BGP neighbor drops the packet and logs an MD5 failed message similar to this one:

```
%TCP-6-BADAUTH: Invalid MD5 digest from 172.16.11.1:1778 to 172.16.12.2:179
```

Use the **norandomseq** keyword with the **static (inside,outside) 172.16.11.1 172.16.11.1 netmask 255.255.255.0 norandomseq** command to solve this problem and to stop the PIX from offsetting the TCP

sequence number. This example illustrates the use of the **norandomseq** keyword:

Router11

```
hostname Router11
!
interface FastEthernet0/0
 ip address 192.168.10.1 255.255.255.0

!--- Connected to Router21.

!
interface FastEthernet0/1
 ip address 172.16.11.1 255.255.255.0

!--- Connected to PIX1.

!
router ospf 1
 log-adjacency-changes
 network 192.168.10.0 0.0.0.255 area 0
 default-information originate metric 5 route-map
 check-default

!--- A default route is originated conditionally, with a metric of 5.

!
router bgp 64496
 no synchronization
 bgp log-neighbor-changes
 network 192.168.10.0
 neighbor 172.16.12.2 remote-as 64496
 neighbor 172.16.12.2 password 7 08345C5A001A1511110D04

!--- Configures MD5 authentication on BGP.

 distance bgp 20 105 200

!--- Administrative distance of iBGP-learned routes is changed from default 200 to 105.
!--- MD5 authentication is configured for BGP.

no auto-summary
!
ip route 172.16.12.0 255.255.255.0 172.16.11.10

!--- Static route to iBGP peer, because it is not directly connected.

!
access-list 30 permit 0.0.0.0
access-list 31 permit 172.16.12.2
route-map check-default permit 10
 match ip address 30
 match ip next-hop 31
```

Router12

```
hostname Router12
!
interface FastEthernet0/0
 ip address 172.16.13.2 255.255.255.0

!--- Connected to ISP-A.

!
```

```

interface FastEthernet0/1
 ip address 172.16.12.2 255.255.255.0

!--- Connected to PIX1.

!
router bgp 64496
no synchronization
 neighbor 172.16.11.1 remote-as 64496
 neighbor 172.16.11.1 next-hop-self
 neighbor 172.16.11.1 default-originate route-map
 neighbor 172.16.11.1 password 7 08345C5A001A1511110D04

!--- Configures MD5 authentication on BGP.

check-ispera-route

!--- Originate default to Router11 conditionally if check-ispera-route is a success.
!--- MD5 authentication is configured for BGP.

 neighbor 172.16.11.1 distribute-list 1 out
 neighbor 172.16.13.4 remote-as 64500
 neighbor 172.16.13.4 route-map adv-to-ispera out
 no auto-summary
!
ip route 172.16.11.0 255.255.255.0 172.16.12.10

!--- Static route to iBGP peer, because it is not directly connected.

!
access-list 1 permit 0.0.0.0
access-list 10 permit 192.168.10.0
access-list 20 permit 10.10.20.0 0.0.0.255
access-list 21 permit 172.16.13.4
!
route-map check-ispera-route permit 10
 match ip address 20
 match ip next-hop 21
!
route-map adv-to-ispera permit 10
 match ip address 10

```

PIX1

```

nameif ethernet0 outside security0
nameif ethernet1 inside security100
ip address outside 172.16.12.10 255.255.255.0
ip address inside 172.16.11.10 255.255.255.0
access-list acl-1 permit tcp host 172.16.13.4 host 172.16.11.1 eq bgp

!--- Access list allows BGP traffic to pass from outside to inside.

 access-list acl-1 permit icmp any any

!--- Allows ping to pass through for testing purposes only.

access-group acl-1 in interface outside
nat (inside) 0 0.0.0.0 0.0.0.0 0 0
static (inside,outside) 172.16.11.1 172.16.11.1 netmask 255.255.255.255 norandomseq


!--- Stops the PIX from offsetting the TCP sequence number.

route outside 0.0.0.0 0.0.0.0 172.16.12.2 1
route inside 192.168.10.0 255.255.255.0 172.16.11.1 1

```

PIX / ASA 7.x and Later

This section uses this Network setup.



PIX/ASA version 7.x and later introduces an additional challenge when you try to establish a BGP peering session with MD5 authentication. By default, PIX/ASA version 7.x and later rewrites any TCP MD5 option included on a TCP datagram that goes through the device and replaces the option kind, size and value with NOP option bytes. This effectively breaks BGP MD5 authentication, and results in error messages like this on each peering router:

```
000296: Apr 7 2010 15:13:22.221 EDT: %TCP-6-BADAUTH: No MD5 digest from 172.16.11.1(28894)
to 172.16.12.2(179)
```

In order for a BGP session with MD5 authentication to be successfully established, these three issues must be resolved:

- Disable TCP sequence number randomization
- Disable TCP MD5 option rewriting
- Disable NAT between peers

A class-map and an access-list are used to select the traffic between the peers that must both be exempted from the TCP sequence number randomization feature and allowed to carry an MD5 option without rewriting. A tcp-map is used to specify the option type to be allowed, in this case, option kind 19 (TCP MD5 option). The class-map and the tcp-map are both linked together through a policy-map, part of the Modular Policy Framework infrastructure. The configuration is then activated with the **service-policy** command.

Note: The need to disable NAT between the peers is handled by the **no nat-control** command.

In version 7.0 and later, the default nature of an ASA is **no nat-control**, which states that every connection through ASA, by default, need not pass the NAT test. It is assumed that ASA has a default setting of **no nat-control**. Refer to nat-control for more information. If **nat-control** is enforced, you must explicitly disable NAT for the BGP peers. This can be done with the **static** command between inside and outside interfaces.

```
static (inside, outside) 172.16.11.1 172.16.11.1 netmask 255.255.255.255
```

PIX/ASA 7.x/8.x

```
ciscoasa# sh run
: Saved
:
ASA Version 8.2(1)
!
```

```
hostname ciscoasa
domain-name example.com
enable password 8Ry2YjIyt7RRXU24 encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
names
!

!--- Configure the outside interface.

interface Ethernet0/0
 nameif outside
 security-level 0
 ip address 172.16.12.10 255.255.255.0
!

!--- Configure the inside interface.

interface Ethernet0/1
 nameif inside
 security-level 100
 ip address 172.16.11.10 255.255.255.0
!

!-- Output suppressed.

!--- Access list to allow incoming BGP sessions
!--- from the outside peer to the inside peer

access-list OUTSIDE-ACL-IN extended permit tcp host 172.16.12.2 host 172.16.11.1 eq bgp

!--- Access list to match BGP traffic.

!--- The next line matches traffic from the inside peer to the outside peer

access-list BGP-MD5-ACL extended permit tcp host 172.16.11.1 host 172.16.12.2 eq bgp

!--- The next line matches traffic from the outside peer to the inside peer

access-list BGP-MD5-ACL extended permit tcp host 172.16.12.2 host 172.16.11.1 eq bgp

!

!--- TCP-MAP to allow MD5 Authentication.

tcp-map BGP-MD5-OPTION-ALLOW
  tcp-options range 19 19 allow
!

!--- Apply the ACL that allows traffic
!--- from the outside peer to the inside peer

access-group OUTSIDE-ACL-IN in interface outside
!
asdm image disk0:/asdm-621.bin
no asdm history enable
```

```

arp timeout 14400

route outside 0.0.0.0 0.0.0.0 172.16.12.2 1
route inside 192.168.10.0 255.255.255.0 172.16.11.1 1
http server enable
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
crypto ipsec security-association lifetime seconds 28800
crypto ipsec security-association lifetime kilobytes 4608000
telnet timeout 5
ssh timeout 5
console timeout 0
threat-detection basic-threat
threat-detection statistics access-list
no threat-detection statistics tcp-intercept

!
class-map inspection_default
match default-inspection-traffic
class-map BGP-MD5-CLASSMAP
match access-list BGP-MD5-ACL
!
!
policy-map type inspect dns preset_dns_map
parameters
  message-length maximum 512
policy-map global_policy
class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
class BGP-MD5-CLASSMAP
set connection random-sequence-number disable
set connection advanced-options BGP-MD5-OPTION-ALLOW

!
service-policy global_policy global
prompt hostname context
Cryptochecksum:64ea55d7271e19eea87c8603ab3768a2
: end

```

Router11

```

Router11#sh run
hostname Router11
!
ip subnet-zero
!
interface Loopback0
  no ip address
  shutdown
!
interface Loopback1

```

```

ip address 192.168.10.1 255.255.255.0
!
interface Ethernet0
ip address 172.16.11.1 255.255.255.0
!
interface Serial0
no ip address
shutdown
no fair-queue
!
interface Serial1
no ip address
shutdown
!
interface BRI0
no ip address
encapsulation hdlc
shutdown
!
router bgp 64496
no synchronization
bgp log-neighbor-changes
network 192.168.10.0
neighbor 172.16.12.2 remote-as 64496

!--- Configures MD5 authentication on BGP.

neighbor 172.16.12.2 password 7 123456789987654321

!--- Administrative distance of iBGP-learned routes is changed from default 200 to 105.
!--- MD5 authentication is configured for BGP.

distance bgp 20 105 200
no auto-summary
!
ip classless

!--- Static route to iBGP peer, because it is not directly connected.

ip route 172.16.12.0 255.255.255.0 172.16.11.10
ip http server
!

!--- Output suppressed

```

Router12

```

Router12#sh run
hostname Router12
!
aaa new-model
!
ip subnet-zero
!
interface Ethernet0
ip address 172.16.13.2 255.255.255.0
!
interface Ethernet1
ip address 172.16.12.2 255.255.255.0
!
interface Serial0
no ip address

```



```

no fair-queue
!
interface Serial1
no ip address
shutdown
!
router bgp 64496
no synchronization
bgp log-neighbor-changes
neighbor 172.16.11.1 remote-as 64496

!--- Configures MD5 authentication on BGP.

neighbor 172.16.11.1 password 7 123456789987654321
neighbor 172.16.11.1 next-hop-self

!--- Originate default to Router11 conditionally if check-ispa-route is a success

neighbor 172.16.11.1 default-originate route-map check-ispa-route
neighbor 172.16.11.1 distribute-list 1 out
neighbor 172.16.13.4 remote-as 64500
no auto-summary
!
ip classless

!--- Static route to iBGP peer, because it is not directly connected.

ip route 172.16.11.0 255.255.255.0 172.16.12.10
ip http server
!
access-list 1 permit 0.0.0.0
access-list 10 permit 192.168.10.0
access-list 20 permit 10.10.20.0 0.0.0.255
access-list 21 permit 172.16.13.4
route-map check-ispa-route permit 10
match ip address 20
match ip next-hop 21
!
route-map adv-to-ispa permit 10
match ip address 10
!

!--- Output suppressed

```

Router14 (ISP-A)

```

Router14#sh run
hostname Router14
!
!
ip subnet-zero
!
interface Ethernet0
ip address 172.16.13.4 255.255.255.0
!
interface Ethernet1
ip address 10.10.20.1 255.255.255.0
!
interface Serial0
no ip address
shutdown

```

```

no fair-queue
!
interface Serial11
no ip address
shutdown
!
router bgp 64500
bgp log-neighbor-changes
network 10.10.20.0 mask 255.255.255.0

!--- Configures Router12 as an eBGP peer.

neighbor 172.16.13.2 remote-as 64496
!

!--- Output suppressed

ip classless

```

Verify

Output from the **show ip bgp summary** command indicates that the authentication is successful and that the BGP session is established on Router11.

```

Router11#show ip bgp summary
BGP router identifier 192.168.10.1, local AS number 64496
BGP table version is 8, main routing table version 8
3 network entries using 360 bytes of memory
3 path entries using 156 bytes of memory
2/2 BGP path/bestpath attribute entries using 248 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 764 total bytes of memory
BGP activity 25/22 prefixes, 26/23 paths, scan interval 60 secs

Neighbor          V           AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.13.2       4           64496   137    138     8     0     0 02:01:16      1
Router11#

```

Related Information

- [BGP Support Page](#)
- [BGP Best Path Selection Algorithm](#)
- [Load Sharing with BGP in Single and Multihomed Environments: Sample Configurations](#)
- [Cisco PIX Firewall Software](#)
- [Cisco Secure PIX Firewall Command References](#)
- [Configuring and Testing PIX Firewall](#)
- [Technical Support & Documentation - Cisco Systems](#)