



MULTI-PROTOCOL BORDER GATEWAY PROTOCOL AND IPV6

INTERNET TECHNOLOGIES DIVISION

OCTOBER 2004

Agenda

- **Overview**
- **Capability negotiation**
- **Applications**
- **Multi-Protocol Border Gateway Protocol (MP-BGP) configuration for IPv6**

OVERVIEW



MP-BGP (RFC2858)

- **Extension to the BGP protocol**
- **Objective is to carry routing information about other protocols:**
 - IPv6 Unicast**
 - 6PE**
 - Multicast (IPv4 and IPv6)**
 - Multi-Protocol Label Switching (MPLS) VPN (IPv4 and IPv6)**
 - Connectionless Network Service (CLNS)**
- **Exchange of Multi-Protocol Reachability Information must be negotiated at session setup time**

MP-BGP Attributes

- **New non-transitive and optional Border Gateway Protocol (BGP) attributes**

MP_REACH_NLRI

“Carry the set of reachable destinations together with the next-hop information to be used for forwarding to these destinations” (RFC2858)

MP_UNREACH_NLRI

Carry the set of unreachable destinations

MP-BGP Attributes (Cont.)

- **Attribute contains one or more triples:**

Address Family Information (AFI) with Sub-AFI

**Identifies protocol information carried in the Network Layer
Reachability Info (NRLI) field**

Next-hop information

Next-hop address must be of the same family

Reachability information

MP-BGP for IPv6 (RFC2545)

- **IPv6 specific extensions**

Scoped addresses:

Next_hop contains a global IPv6 address and potentially a link-local address

NEXT_HOP, **AGGREGATOR**, and **NLRI** are expressed as IPv6 addresses and prefix

Address Family Information (AFI) = 2 (IPv6)

Sub-AFI = 1 (NLRI is used for unicast)

Sub-AFI = 2 (NLRI is used for multicast RPF check)

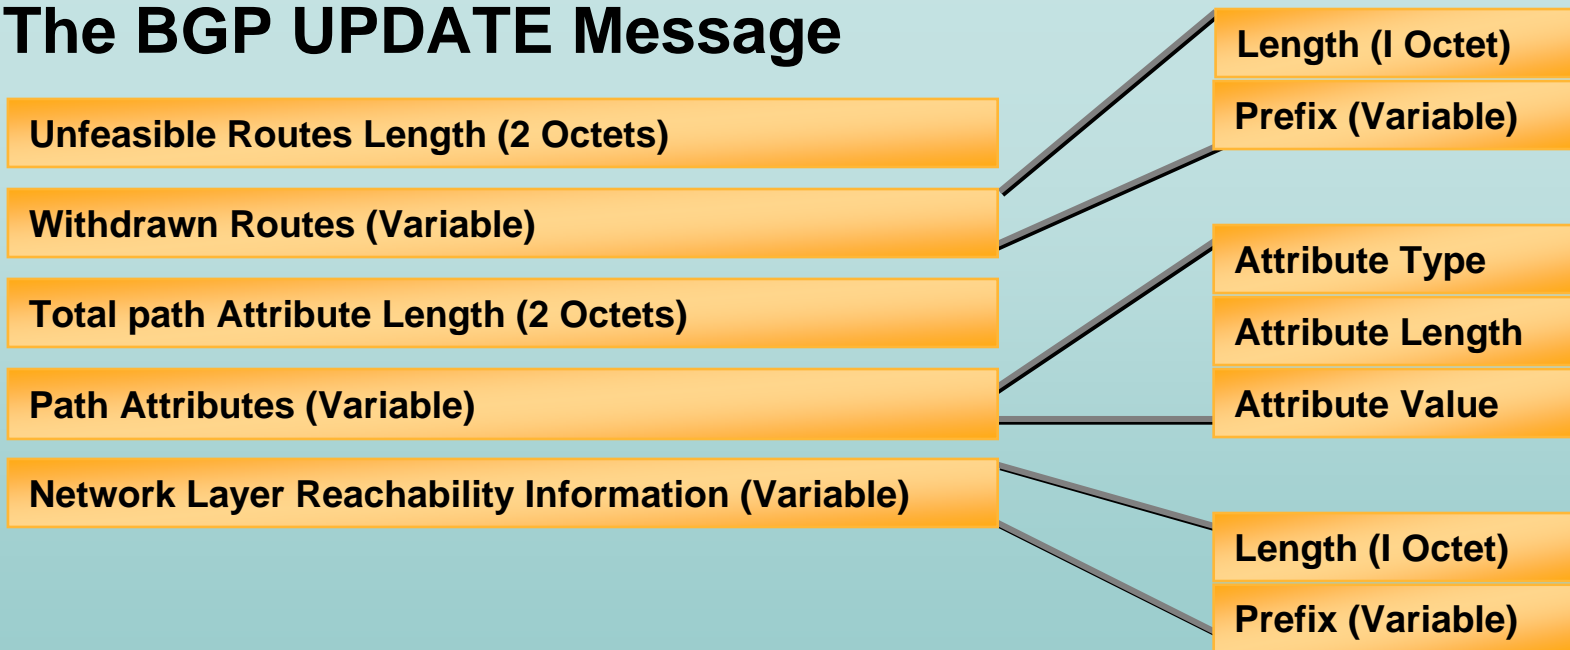
Sub-AFI = 3 (NLRI is used for both unicast and multicast RPF check)

Sub-AFI = 4 (label)

Sub-AFI = 128 (VPN)

MP-BGP for IPv6 Packet Format

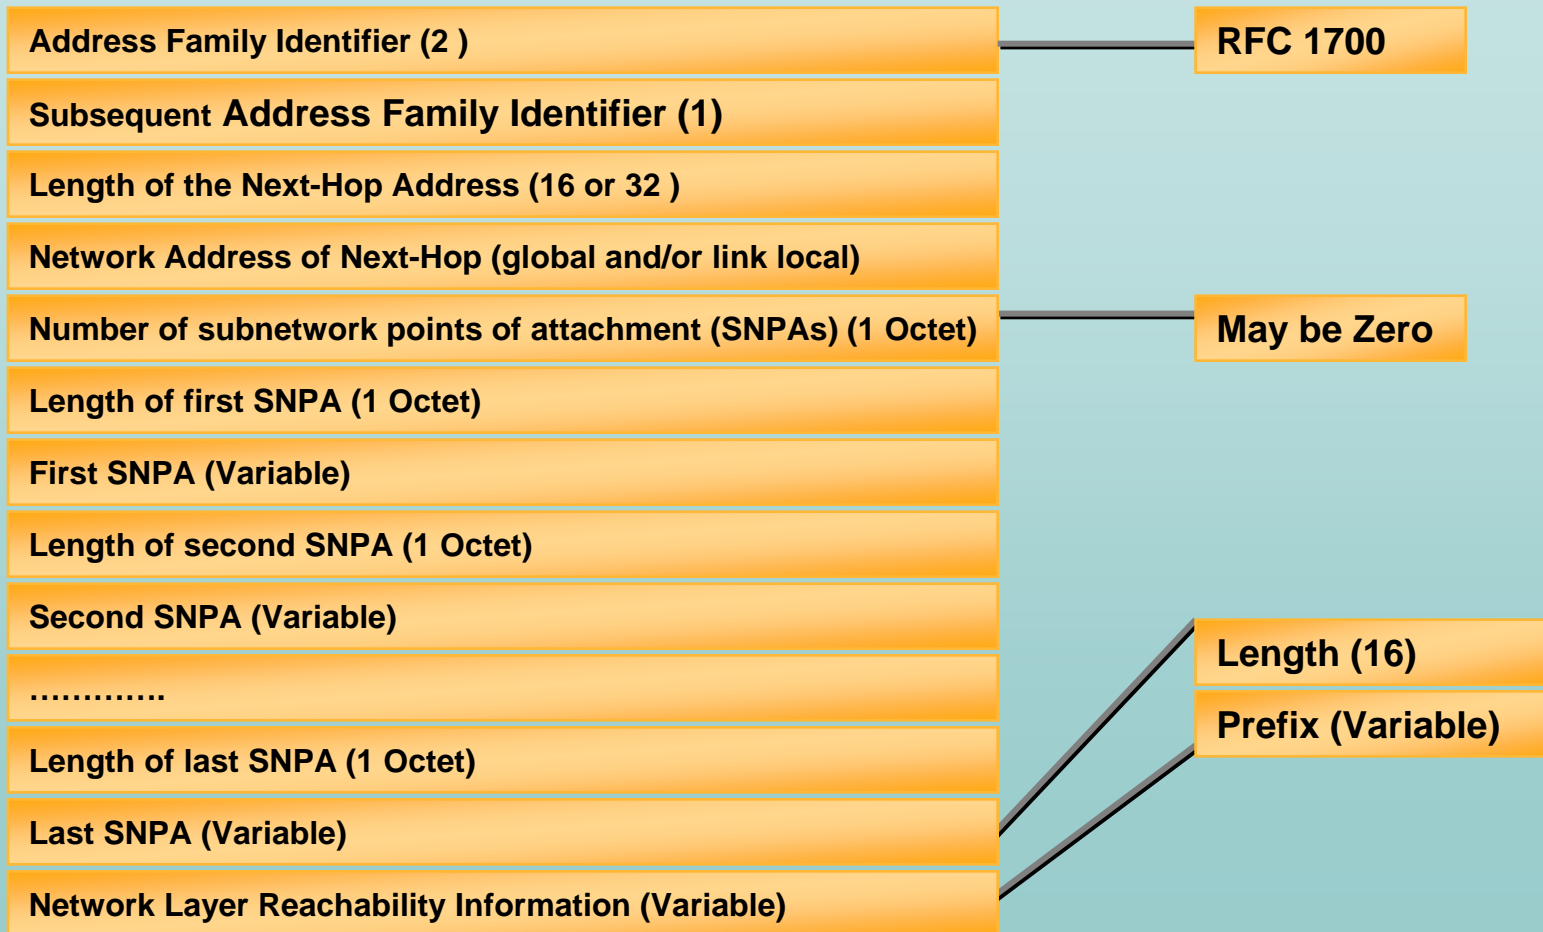
The BGP UPDATE Message



- A BGP update is used to advertise feasible routes to a peer or to withdraw multiple unfeasible routes
- Each update message contains attributes like: origin, AS-Path, Next-Hop,
MP_REACH_NLRI

MP-BGP for IPv6 Unicast Packet Format

MP_REACH_NLRI Attribute



Packet Format for IPv6 Unicast

MP_UNREACH_NLRI Attribute



CAPABILITY NEGOTIATION



MP-BGP Capabilities Negotiation

- **BGP routers establish BGP sessions through the OPEN message**
- **OPEN message contains optional parameters**
- **If OPEN parameters are not recognized, BGP session is terminated**
- **A new optional parameter:
CAPABILITIES**

MP-BGP Capabilities Negotiation (Cont.)


- **BGP router sends an OPEN message with CAPABILITIES parameter containing its capabilities :**

Multi-Protocol extensions (AFI/SAFI)

Route Refresh

Outbound Route Filtering

MP-BGP Session Establishment



```
BGP: 3FFE:B00:C18:2:1::1 sending OPEN, version 4, my as: 100
BGP: 3FFE:B00:C18:2:1::1 rcv OPEN, version 4
BGP: 3FFE:B00:C18:2:1::1 rcv OPEN w/ OPTION parameter len: 16
BGP: 3FFE:B00:C18:2:1::1 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
BGP: 3FFE:B00:C18:2:1::1 OPEN has CAPABILITY code: 1, length 4
BGP: 3FFE:B00:C18:2:1::1 OPEN has MP_EXT CAP for afi/safi: 2/1
BGP: 3FFE:B00:C18:2:1::1 rcvd OPEN w/ optional parameter type 2 (Capability) len 2
BGP: 3FFE:B00:C18:2:1::1 went from OpenSent to OpenConfirm
BGP: 3FFE:B00:C18:2:1::1 went from OpenConfirm to Established
%BGP-5-ADJCHANGE: neighbor 3FFE:B00:C18:2:1::1 Up
```

APPLICATIONS

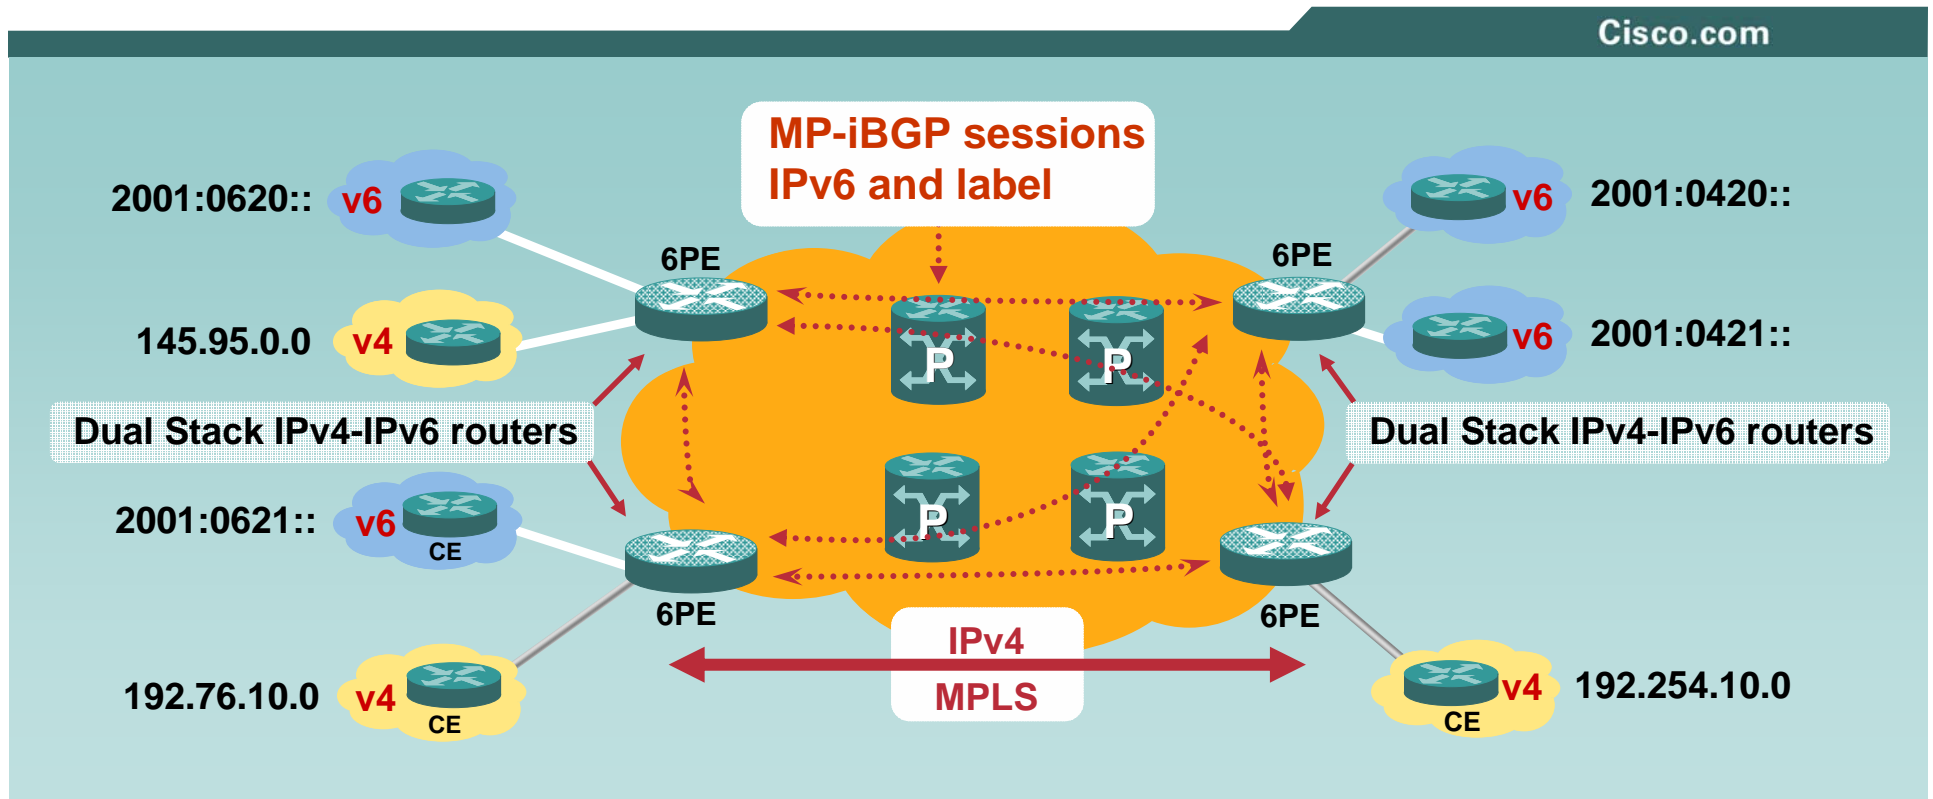


MP-BGP Value Added Features for IPv6

- **IPv6 Unicast**
MP-BGP enables the creation of IPv6 Internet with an exterior gateway protocol (EGP) as reliable as for IPv4
- **IPv6 Multicast**
Unicast prefixes for Reverse Path Forwarding (RPF) checking
RPF information is disseminated between autonomous systems
Compatible with single domain Rendezvous Points or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM)
Topology can be congruent or non-congruent with the unicast one
- **IPv6 and label (6PE)**
IPv6 packet is transported over an IPv4 MPLS backbone
- **IPv6 VPN (6VPE)**
Multiple IPv6 VPNs are created over an IPv4 MPLS backbone

IPv6 Provider Edge Router (6PE) over MPLS – IPv6 and Label

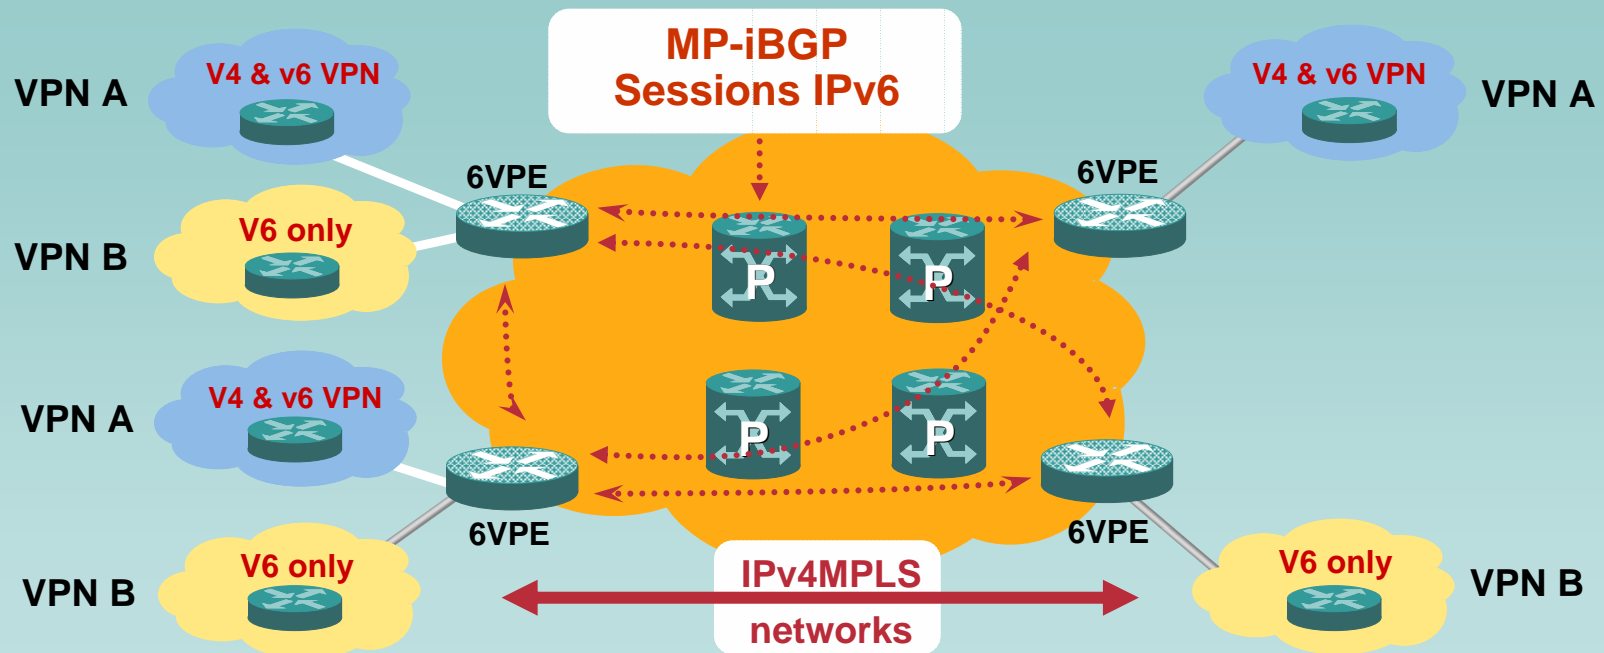
Cisco.com



- **MPLS Core Infrastructure is IPv6-unaware**
- **IPv6 reachability is exchanged among 6PEs via iBGP (MP-BGP)**
- **Co-existence with any other address families is exchanged**
- **IPv6 packets are transported from 6PE to 6PE inside MPLS**

IPv6 VPN over MPLS (6VPE) - IPv6 VPN

Cisco.com



- **MPLS Core Infrastructure is IPv6-unaware**
- **IPv6 VPN routes are exchanged via iBGP (MP-BGP)**
- **IPv6 VPN can co-exist with IPv4 VPN**

MP-BGP FOR IPV6 CONFIGURATION



MP-BGP Availability

- **IPv6 Unicast**
- **IPv6 Multicast**
- **IPv6 Unicast and label : 6PE**
- **IPv6 VPN: 6VPE (Early Field Trial)**
- **Detailed availability can be found in:**

Start Here: Cisco IOS® Software Release Specifics for IPv6 Features:

www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

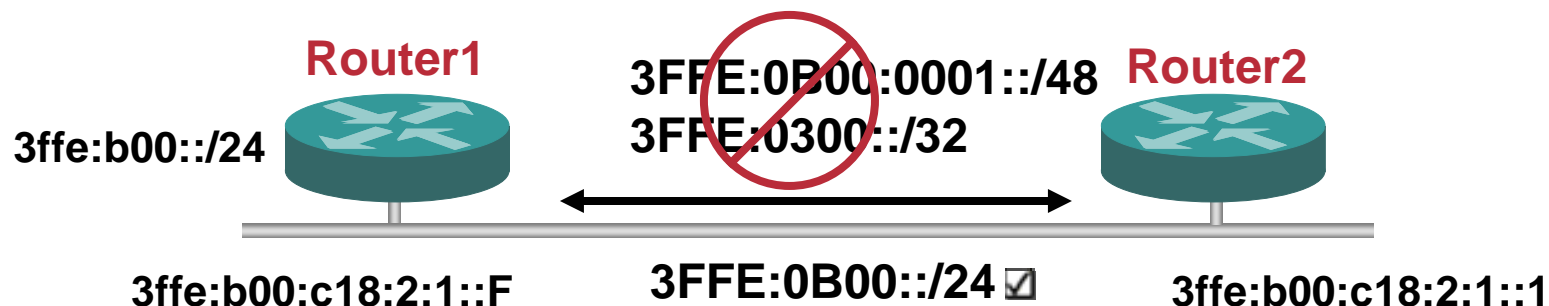
A Simple BGP Session



```
Router1#  
interface Ethernet0  
  ipv6 address 3FFE:B00:C18:2:1::F/64  
!  
router bgp 65001  
  bgp router-id 10.10.10.1  
  no bgp default ipv4-unicast  
  neighbor 3FFE:B00:C18:2:1::1 remote-as 65002  
  address-family ipv6  
    neighbor 3FFE:B00:C18:2:1::1 activate  
    neighbor 3FFE:B00:C18:2:1::1 prefix-list bgp65002in in  
    neighbor 3FFE:B00:C18:2:1::1 prefix-list bgp65002out out  
  exit-address-family
```

IPv6 Route-Map / Prefix-List


Filtering BGP routing updates



```
Router1#
router bgp 65001
  no bgp default ipv4-unicast
  neighbor 3FFE:B00:C18:2:1::1 remote-as 65002
  address-family ipv6
    neighbor 3FFE:B00:C18:2:1::1 activate
    neighbor 3FFE:B00:C18:2:1::1 prefix-list bgp65002in in
    neighbor 3FFE:B00:C18:2:1::1 prefix-list bgp65002out out
  network 3FFE:B00::/24
  exit-address-family
!
ipv6 prefix-list bgp65002in seq 5 permit 3FFE::/16 le 24
ipv6 prefix-list bgp65002out seq 5 permit 3FFE::/16 le 24
```

Adjusting BGP Attributes

- Prefer routes from AS 65002 (local preference)




```
router bgp 65001
no bgp default ipv4-unicast
neighbor 3FFE:B00:C18:2:1::1 remote-as 65002
neighbor 3FFE:B00:C18:2:1::2 remote-as 65003
address-family ipv6
neighbor 3FFE:B00:C18:2:1::1 activate
neighbor 3FFE:B00:C18:2:1::2 activate
neighbor 3FFE:B00:C18:2:1::1 prefix-list Legal in
neighbor 3FFE:B00:C18:2:1::1 route-map From_65002 in
neighbor 3FFE:B00:C18:2:1::2 prefix-list Legal in
network 3FFE:B00::/24
exit-address-family
!
route-map From_65002 permit 10
set local-preference 120
```


Peering with Link Local Addresses

Cisco.com


Router 1



Router 2



Router 3



```
Router1#
router bgp 65001
  no bgp default ipv4-unicast
  neighbor FE80::205:32FF:FE46:9601 remote-as 65002
  neighbor FE80::205:32FF:FEC3:40F0 remote-as 65003
!
address-family ipv6
  neighbor FE80::205:32FF:FE46:9601 activate
  neighbor FE80::205:32FF:FE46:9601 update-source FastEthernet0/0
  neighbor FE80::205:32FF:FE46:9601 route-map nh6 out
  neighbor FE80::205:32FF:FEC3:40F0 activate
  neighbor FE80::205:32FF:FEC3:40F0 update-source FastEthernet0/0
  neighbor FE80::205:32FF:FEC3:40F0 route-map nh6 out

  network 3FFE:B00::/24
  exit-address-family
!
route-map nh6 permit 10
  set ipv6 next-hop 3ffe:b00:c18:2:1::F FE80::205:625F:1246:5296
```

MP-BGP with IPv6 Multicast Support

Cisco.com



```
Router1#  
interface Ethernet0  
  ipv6 address 3FFE:B00:C18:2:1::F/64  
!  
router bgp 65001  
  bgp router-id 10.10.10.1  
  no bgp default ipv4-unicast  
  neighbor 3FFE:B00:C18:2:1::1 remote-as 65002  
  address-family ipv6 multicast  
    neighbor 3FFE:B00:C18:2:1::1 activate  
  redistribute ospf 1  
  exit-address-family
```

Show Commands

IPv6 BGP show commands take *ipv6* as an argument

Show bgp ipv6 unicast <parameter>

```
Router1#show bgp ipv6 unicast 2017::/96
BGP routing table entry for 2017::/96, version 11
Paths: (1 available, best #1)
Local
  3FFE:B00:C18:2:1::1 from 3FFE:B00:C18:2:1::1 (10.10.20.2)
    Origin incomplete, localpref 100, valid, internal, best
```

Show Commands (Cont.)

Displays summary information

Show bgp ipv6 unicast summary

```
Router1# show bgp ipv6 unicast summary
GP router identifier 10.10.20.1, local AS number 100
BGP table version is 18, main routing table version 18
5 network entries and 8 paths using 1261 bytes of memory
2 BGP path attribute entries using 120 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 8/71 prefixes, 15/7 paths, scan interval 15 secs
```

} resource utilization
by the BGP process

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3FFE:B00:C18:2:1::1	4	100	7	6	18	0	0	00:01:10	2

↑
Neighbor information

↑
BGP messages activity

CISCO SYSTEMS

