

Multi-Protocol Border Gateway Protocol and IPv6

Benoit Lourdelet

Cisco Systems

blourdel@cisco.com

Agenda

Cisco.com

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

MP-BGP

RFC2283

- **Extension to the BGP protocol**
- **Objective: carry routing information about other protocols**
 - Multicast**
 - MPLS VPN**
 - IPv6**
 - 6PE**
 - CLNS**
- **Exchange of Multi-Protocol NLRI must be negotiated at session set up**

MP-BGP Attributes

- **New non-transitive and optional 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” (RFC2283)

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

Next-hop information

- **Next-hop address must be of the same family**

MP- BGP for IPv6

RFC2545

- **IPv6 specific extensions**

Scoped addresses: Next_hop contains a global IPv6 address and/or potentially a link-local address

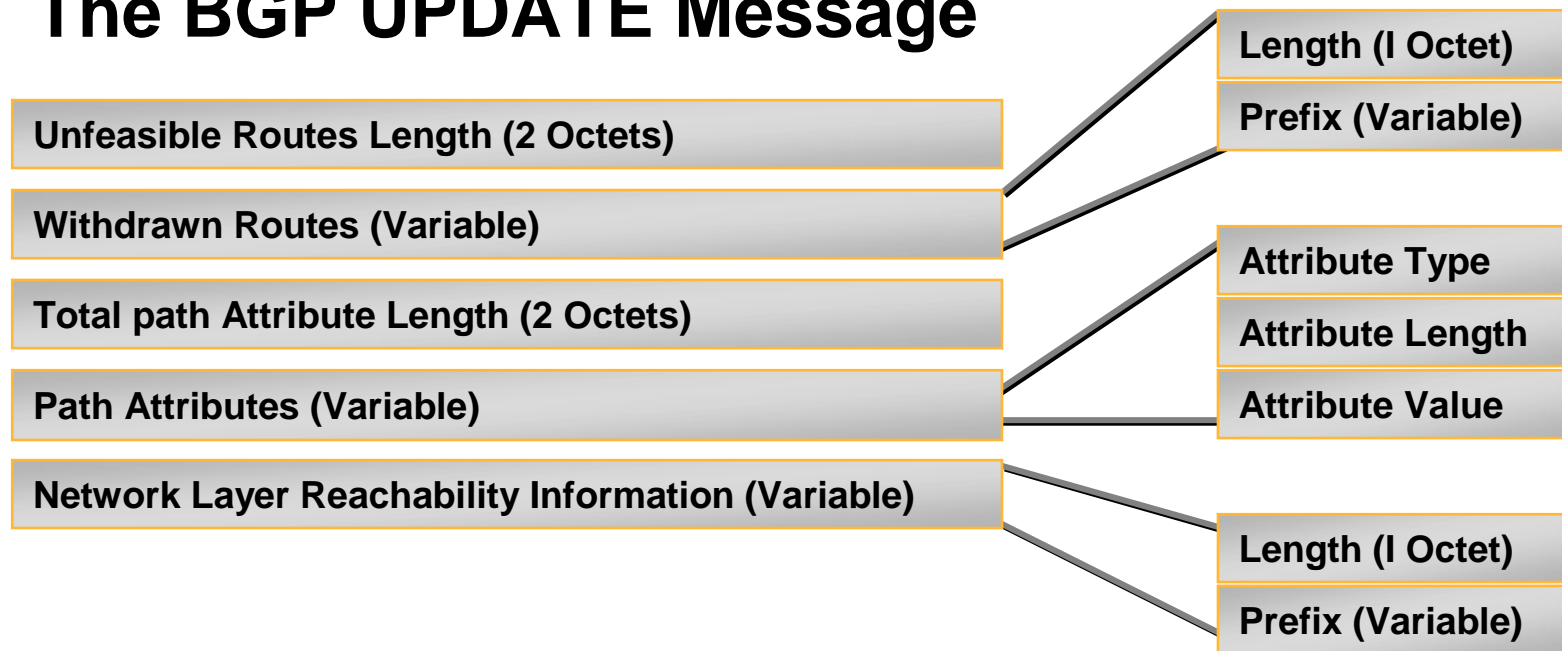
NEXT_HOP, AGGREGATOR and NLRI are expressed as IPv6 addresses and prefix

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

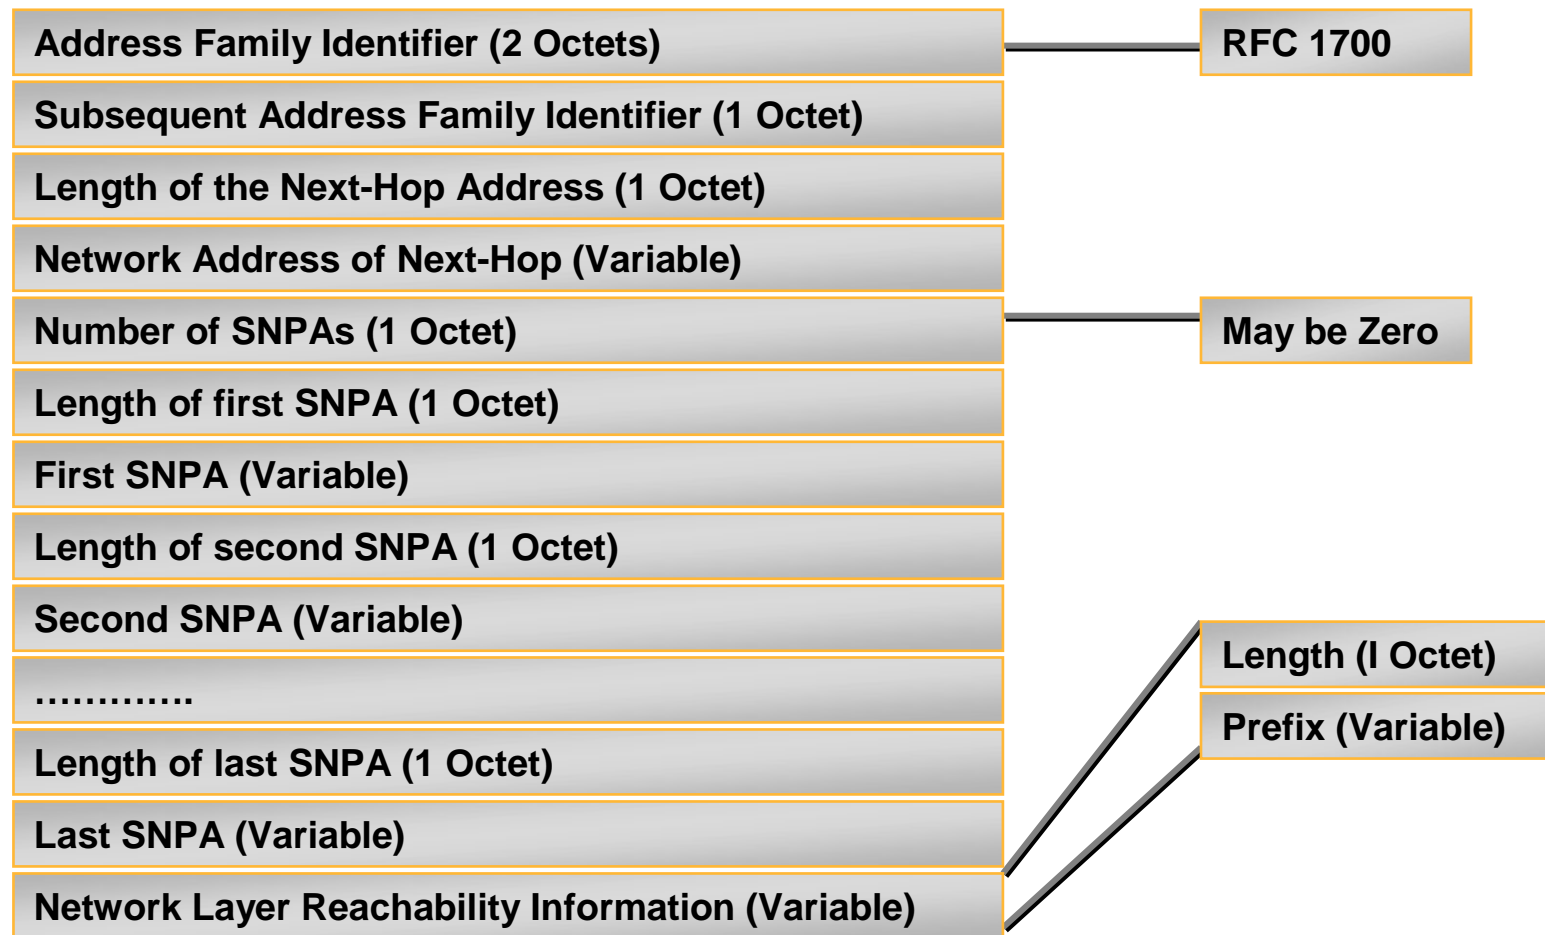
The BGP UPDATE Message



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

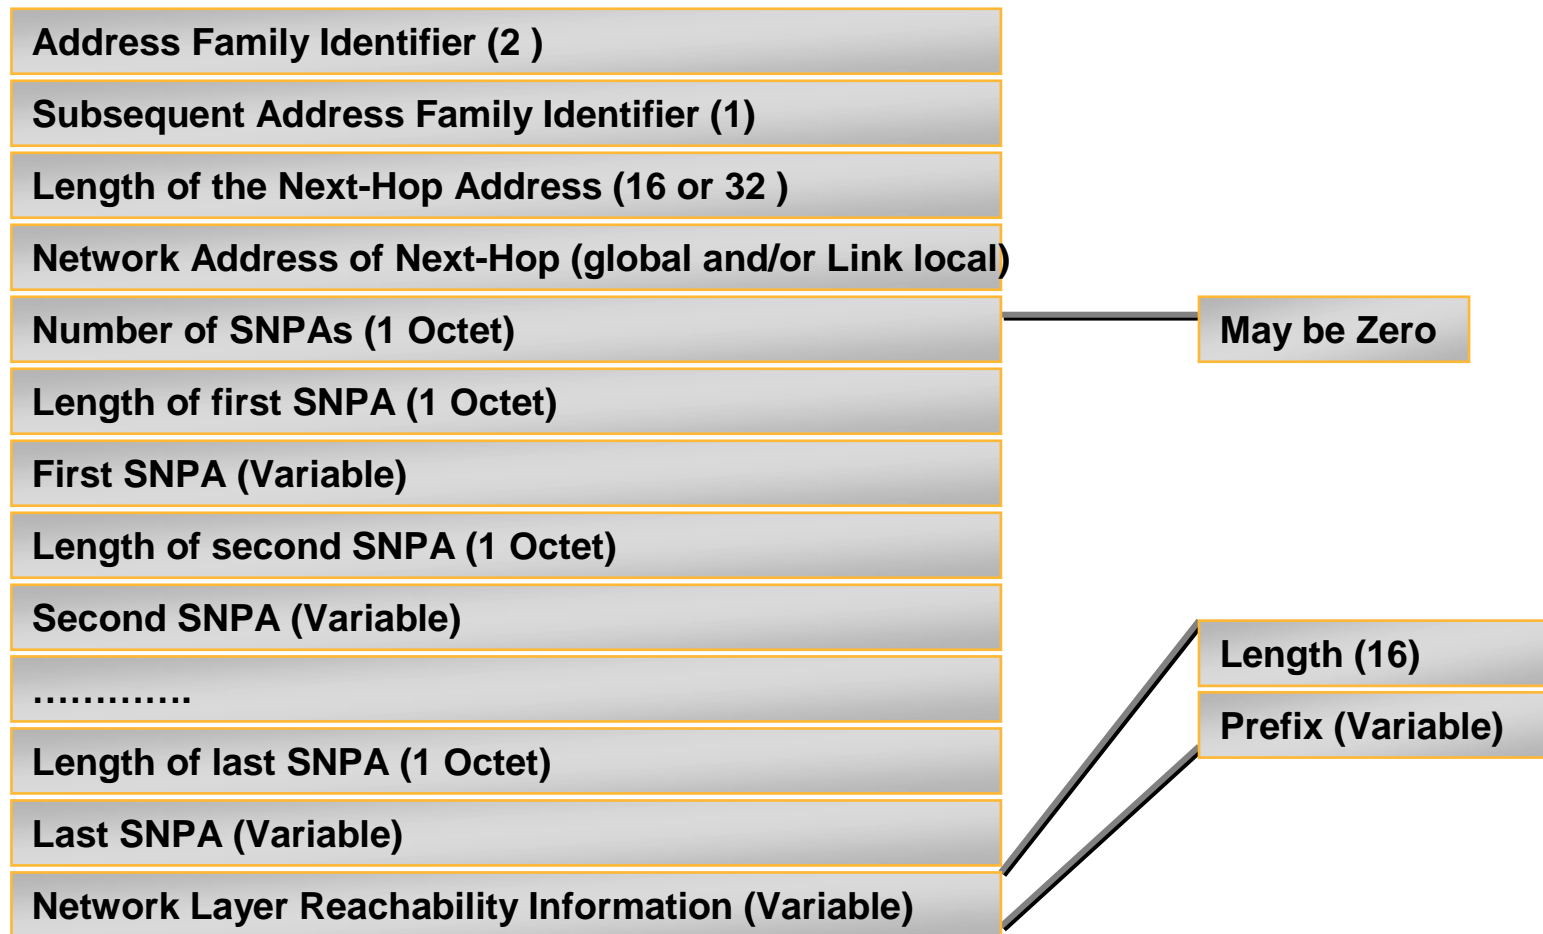
MP- BGP for IPv6 Packet Format (Cont)

MP_REACH_NLRI Attribute



Packet Format for IPv6

MP_REACH_NLRI Attribute



Packet Format for IPv6

MP_UNREACH_NLRI Attribute

Address Family Identifier (2)

Subsequent Address Family Identifier (1)

Withdrawn Routes (Variable)

Length (1 Octet)

Prefix (Variable)

Agenda

Cisco.com

- **Overview**
- **Capability negotiation**
- **MP-BGP for IPv6 configuration**

MP-BGP Capabilities Negotiation

Cisco.com

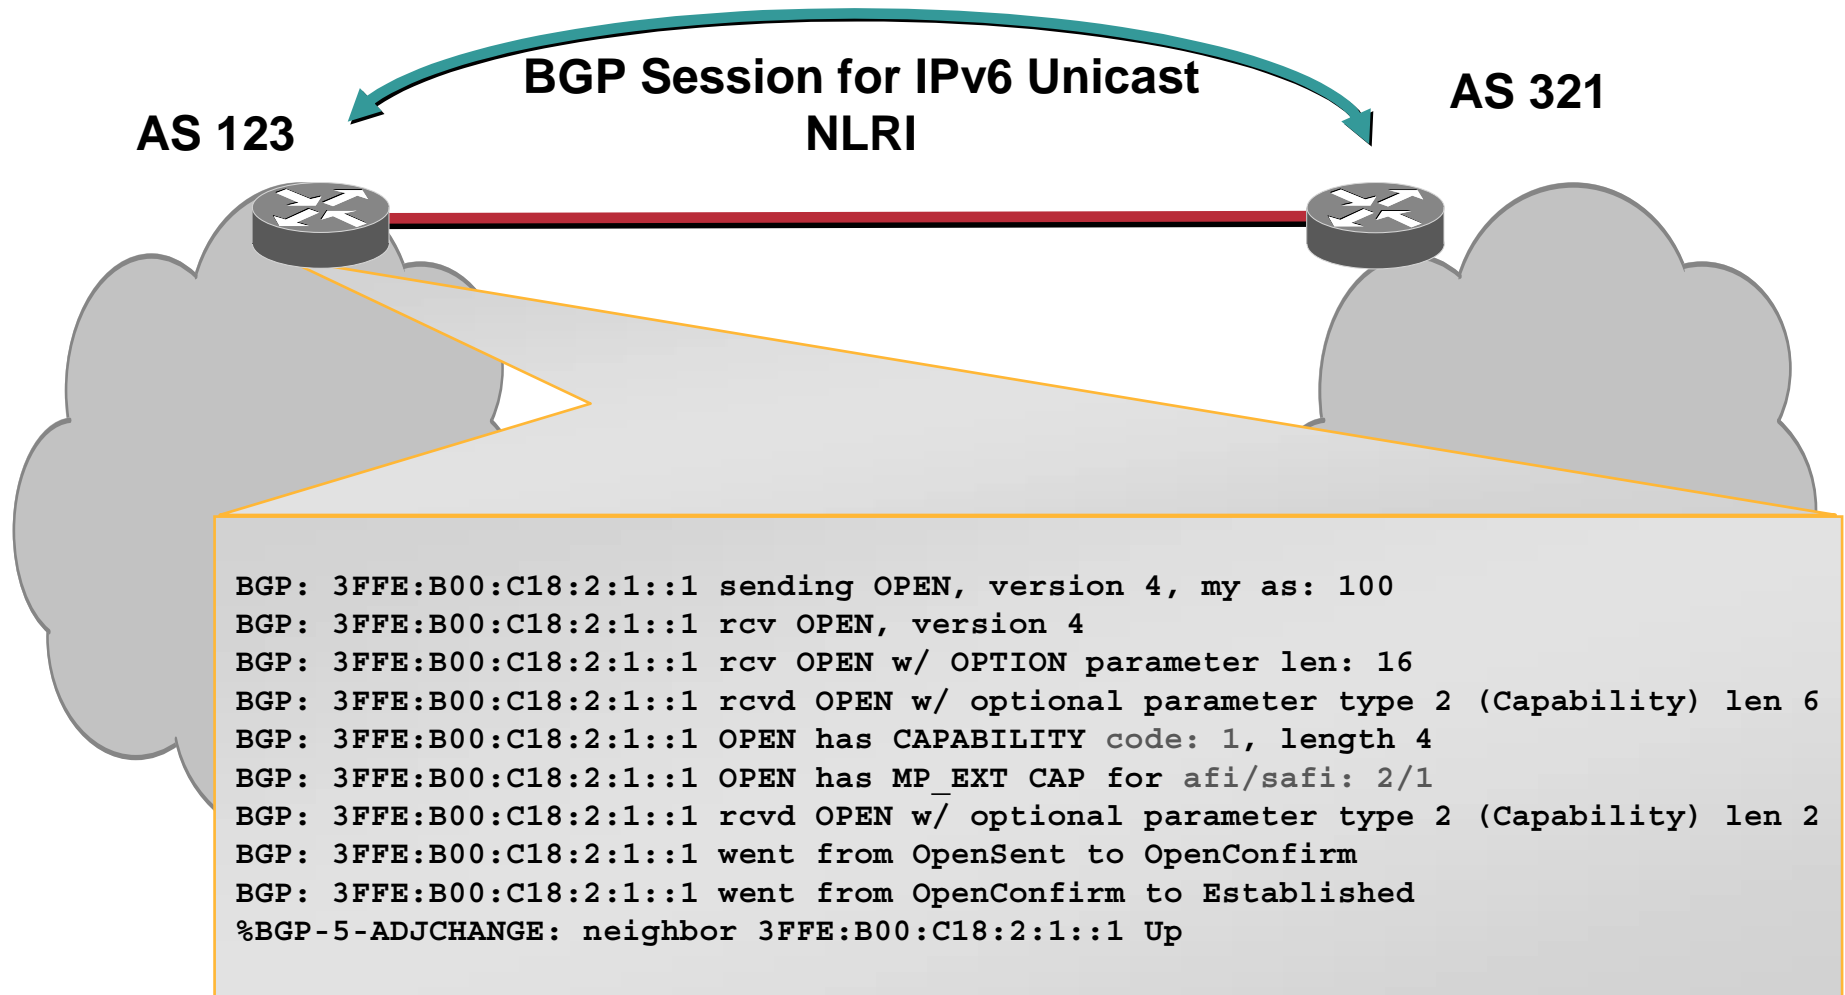
- **BGP routers establish BGP sessions through the OPEN message**
- **OPEN message contains optional parameters**
- **BGP session is terminated if OPEN parameters are not recognised**
- **A new optional parameter: CAPABILITIES**

MP-BGP Capabilities Negotiation (Cont)

Cisco.com

- **BGP router sends an OPEN message with CAPABILITIES parameter containing its capabilities**
 - Multiprotocol extension (AFI/SAFI)**
 - Route Refresh**
 - Outbound Route Filtering**

MP-BGP Session Establishment



Agenda

Cisco.com

- **Overview**
- **Capability negotiation**
- **MP-BGP for IPv6 configuration**

Availability

- **Release 12.2(2)T and above**
- **Release 12.3(1)M and above**
- **Release 12.2(14)S and above**
- **Release 12.0(22)S and above - Cisco 12000 Series Internet Routers**
- **Release 12.0(24)S and above - Cisco 10720 Internet Router**


Simple BGP Session



```
Router1#  
interface Ethernet0  
  ipv6 enable  
  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



```
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::1 FE80::205:625F:1246:5296
```

Show Commands

- **IPv6 BGP show commands take *ipv6* as argument**

`show bgp ipv6 <parameter>`

```
Router1#show bgp ipv6 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 summary

```
Router1# show bgp ipv6 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



EMPOWERING THE
INTERNET GENERATIONSM