IPv6 Routing Protocols

ISP Workshops

Last updated 5th September 2017

1

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene

- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place

- Bug fixes and improvements are welcomed
  - Please email *workshop (at) bgp4all.com*

Philip Smith

Agenda

- Enabling IPv6 Routing on Cisco IOS
- IPv6 Routing Protocols:
  - Static
  - RIPng
  - EIGRP
  - OSPFv3
  - IS-IS
  - BGP

Initial IPv6 Configuration for Cisco IOS

IPv6 Configuration on Cisco IOS

- To enable IPv6 Routing the following global commands are required:

  ```
  Router(config)# ipv6 unicast-routing
  ```

  - Also enable IPv6 CEF (not on by default):

    ```
    Router(config)# ipv6 cef
    ```

  - Also disable IPv6 Source Routing (enabled by default):

    ```
    Router(config)# no ipv6 source-route
    ```

IPv6 Configuration

- To configure a global or unique-local IPv6 address the following interface command should be entered:

   ```
   Router(config-if)# ipv6 address X:X..X:X/prefix
   ```

- To configure an EUI-64 based IPv6 address the following interface command should be entered:

   ```
   Router(config-if)# ipv6 address X:X::/prefix eui-64
   ```

   - EUI-64 is not helpful on a router and is not recommended

6

IPv6 Configuration

- If no global IPv6 address is required on an interface, yet it needs to carry IPv6 traffic:
  - Enable IPv6 on that interface using:

    ```
    Router(config-if)# ipv6 enable
    ```

  - Which will result in a link-local IPv6 address being constructed automatically
  - FE80:: is concatenated with the Interface ID to give:
    - FE80::interface-id
- Configuring an IPv6 address (whether global or unique-local) will also result in a link-local IPv6 address being created

IPv6 Configuration – Link Local

```
Router1# conf t
Router1(config)# no ipv6 source-route
Router1(config)# ipv6 unicast-routing
Router1(config)# ipv6 cef
Router1(config)# interface FastEthernet 0/0
Router1(config-int)# ipv6 enable
Router1(config-int)# ^Z

Router1#sh ipv6 interface fast 0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8B9:C0FF:FE00:F11D
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:F11D
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
```

IPv6 Configuration – EUI64

```
Router1# conf t
Router1(config)# no ipv6 source-route
Router1(config)# ipv6 unicast-routing
Router1(config)# ipv6 cef
Router1(config)# interface FastEthernet 0/0
Router1(config-int)# ipv6 address 2001:db8::/64 eui
Router1(config-int)# ^Z

Router1#sh ipv6 interface fast 0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8B9:C0FF:FE00:F11D
  Global unicast address(es):
    2001:DB8::A8B9:C0FF:FE00:F11D, subnet is 2001:DB8::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:F11D
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
```

9

IPv6 Configuration – Manual

```
Router1# conf t
Router1(config)# no ipv6 source-route
Router1(config)# ipv6 unicast-routing
Router1(config)# ipv6 cef
Router1(config)# interface FastEthernet 0/0
Router1(config-int)# ipv6 address 2001:db8::2/64
Router1(config-int)# ^Z

Router1#sh ipv6 int fast 0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8B9:C0FF:FE00:F11D
  Global unicast address(es):
    2001:DB8::2, subnet is 2001:DB8::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF08:2
    FF02::1:FF00:F11D
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
```

IPv6 Interface Configuration
Duplicate Address Detection

- If a point-to-point interface (e.g. POS, HSSI, Serial) is looped:
  - Router sees its address by Solicited Node Multicast
  - Duplicate Address Detection kicks in:

    ```
    Aug 23 09:18:41.263: %IPV6_ND-6-DUPLICATE_INFO: DAD attempt detected for
    2001:DB8:0:3:: on Serial1/1
    ```

  - Interface automatically disabled for IPv6 traffic (IPv4 traffic is unaffected)
  - This has impact for backbone links (IPv6 traffic takes alternative path) and external peering links (IPv6 peering down, IPv4 peering okay)
- Workaround?
  - Monitoring!
  - Cisco IOS "down-when-looped" configuration

Routing Protocols

Static Routing

- Syntax is:
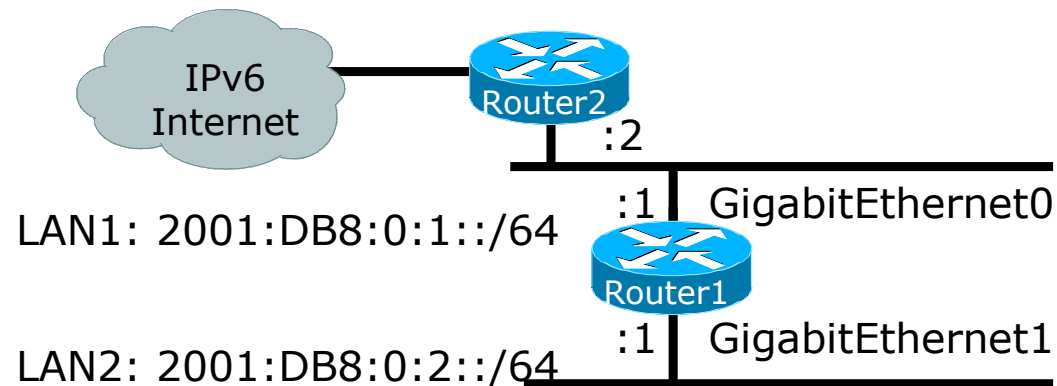 - ipv6 route ipv6-prefix/prefix-length {ipv6-address | interface-type interface-number} [administrative-distance]
- Static Route

```
ipv6 route   2001:DB8::/64   2001:DB8:2:1::1   150
```

  - Routes packets for network 2001:db8::/64 to a networking device at 2001:DB8:2:1::1 with an administrative distance of 150

Default Route Example

IPv6 Internet

Router2

:2

:1 GigabitEthernet0

LAN1: 2001:DB8:0:1::/64

Router1

:1 GigabitEthernet1

LAN2: 2001:DB8:0:2::/64

```
Router1# sh conf
<snip>
ipv6 unicast-routing
!
interface GigabitEthernet0
 ipv6 address 2001:DB8:0:1::1/64
!
interface GigabitEthernet1
 ipv6 address 2001:DB8:0:2::1/64
!
ipv6 route ::/0 2001:DB8:0:1::2
```

Default Route
to Router2

14

Dynamic Routing Protocols in IPv6

- Dynamic Routing in IPv6 is unchanged from IPv4:
  - IPv6 has 2 types of routing protocols: IGP and EGP
  - IPv6 still uses the longest-prefix match routing algorithm
- IGP
  - RIPng (RFC 2080)
  - Cisco EIGRP for IPv6
  - OSPFv3 (RFC 5340)
  - Integrated IS-ISv6 (RFC 5308)
- EGP
  - MP-BGP4 (RFC 4760 and RFC 2545)

Configuring Routing Protocols

- Dynamic routing protocols require a router-id
  - Router-id is a 32 bit integer unique to the router
  - IOS auto-generates these from loopback interface address if configured, else highest IPv4 address on the router
  - Most ISPs will deploy IPv6 dual stack – so router-id will be automatically created
- Early adopters choosing to deploy IPv6 in the total absence of any IPv4 addressing need to be aware:
  - Router-id needs to be manually configured:

```
ipv6 router ospf 100
  router-id 10.1.1.4
```

RIPng

- For the ISP industry, simply don't go here
- ISPs do not use RIP in any form unless there is absolutely no alternative
  - And there usually is
- RIPng was used in the early days of the IPv6 test network
  - Sensible routing protocols such as OSPF and BGP rapidly replaced RIPng when they became available
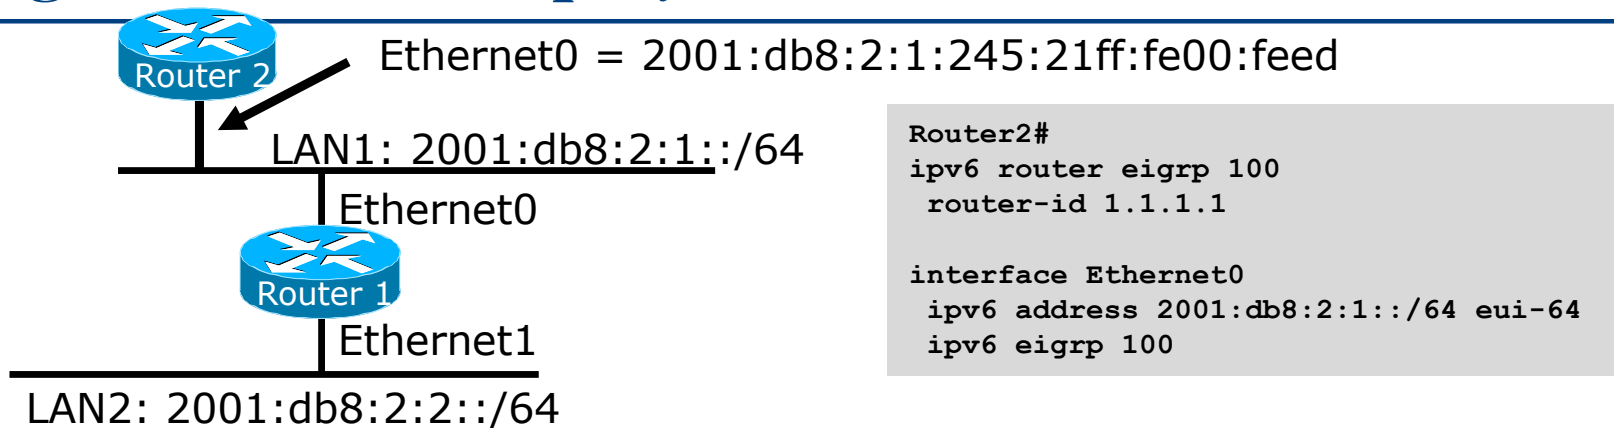
EIGRP for IPv6

- Cisco EIGRP has had IPv6 protocol support added
  - Just another protocol module (IP, IPX, AppleTalk) with three new TLVs:
    - IPv6_REQUEST_TYPE (0X0401)
    - IPv6_METRIC_TYPE (0X0402)
    - IPv6_EXTERIOR_TYPE (0X0403)
  - Router-ID is still 32-bit, protocol is still 88
- Uses similar CLI to existing IPv4 protocol support
- Easy deployment path for existing IPv4 EIGRP users
- In Cisco IOS Release 12.4 onwards

EIGRP for IPv6

- Some differences:
  - Hellos are sourced from the link-local address and destined to FF02::A (all EIGRP routers). This means that neighbors do not have to share the same global prefix (with the exception of explicitly specified neighbours where traffic is unicasted).
  - Automatic summarisation is disabled by default for IPv6 (unlike IPv4)
  - No split-horizon in the case of EIGRP for IPv6 (because IPv6 supports multiple prefixes per interface)

EIGRP for IPv6 Configuration & Display

Ethernet0 = 2001:db8:2:1:245:21ff:fe00:feed

Router 2

LAN1: 2001:db8:2:1::/64

Ethernet0

Router 1

Ethernet1

LAN2: 2001:db8:2:2::/64

```
Router2#
ipv6 router eigrp 100
 router-id 1.1.1.1

interface Ethernet0
 ipv6 address 2001:db8:2:1::/64 eui-64
 ipv6 eigrp 100
```

```
Router1#show ipv6 eigrp neighbor
IPv6-EIGRP neighbors for process 100
H Address                        Interface Hold    Uptime SRTT RTO     Q Seq
                                                   (sec) (ms)     Cnt Num
0 FE80::245:21ff:fe00:feed  E0             14 00:01:43    1 4500    0  1
```

```
Router1#show ipv6 eigrp topology all links
IPv6-EIGRP Topology Table for AS(100)/ID(1.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R -
Reply,
      r - reply Status, s - sia Status
P 2001:db8:2:1::/64, 1 successors, FD is 28160, serno 1
      via Connected, Ethernet0
      via FE80::245:21ff:fe00:feed (30720/28160), Ethernet0
```

Neighbour Identified
by Link-Local Address

20

OSPFv3 overview

- OSPFv3 is OSPF for IPv6 (RFC 5340)
- Based on OSPFv2, with enhancements
- Distributes IPv6 prefixes
- Runs directly over IPv6
- Ships-in-the-night with OSPFv2

Differences from OSPFv2

- Runs over a link, not a subnet
  - Multiple instances per link
- Topology not IPv6 specific
  - Router ID
  - Link ID
- Standard authentication mechanisms
- Uses link local addresses
- Generalized flooding scope
- Two new LSA types

OSPFv3 configuration example

```
Router1#
interface GigabitEthernet0
 ipv6 address 2001:db8:1:1::1/64
 ipv6 ospf 1 area 0
!
interface GigabitEthernet1
 ipv6 address 2001:db8:1:2::1/64
 ipv6 ospf 1 area 1
!
ipv6 router ospf 1
 router-id 192.168.0.1
```

NB: Router-id only required in absence of any IPv4 configuration

Area 0

Router2

LAN1: 2001:db8:1:1::/64

Gig0

Router1

Gig1

LAN2: 2001:db8:2:2::/64

Area 1
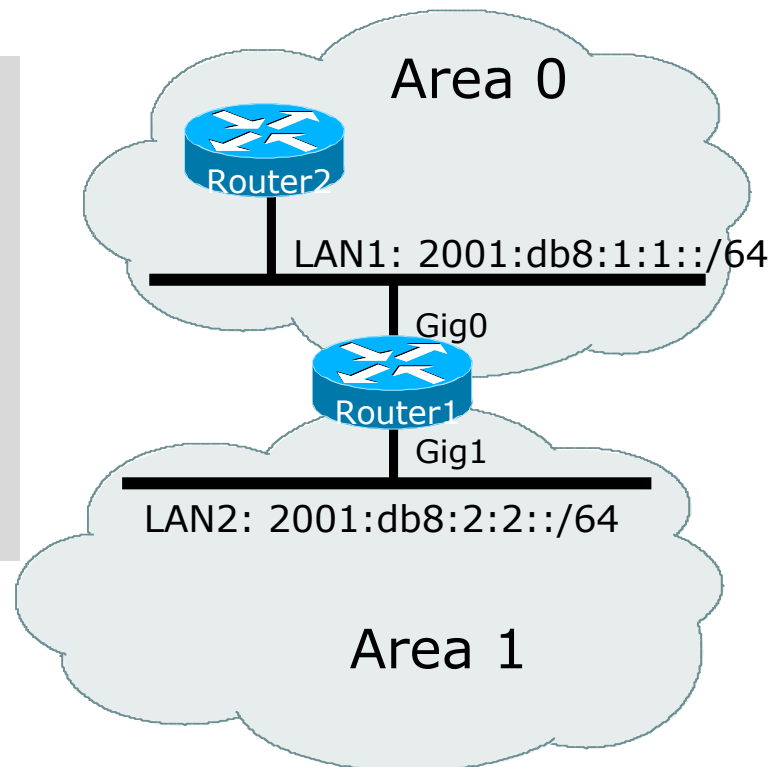
Dual stack OSPF configuration example

```
Router1#
interface GigabitEthernet0
 ip address 192.168.1.1 255.255.255.0
 ip ospf 1 area 0
 ipv6 address 2001:db8:1:1::1/64
 ipv6 ospf 1 area 0
!
interface GigabitEthernet1
 ip address 192.168.2.2 255.255.255.0
 ip ospf 1 area 1
 ipv6 address 2001:db8:1:2::1/64
 ipv6 ospf 1 area 1
!
router ospf 1
 passive-interface default
 no passive-interface GigabitEthernet0
 no passive-interface GigabitEthernet1
!
ipv6 router ospf 1
 passive-interface default
 no passive-interface GigabitEthernet0
 no passive-interface GigabitEthernet1
!
```

Area 0

Router2

LAN1: 2001:db8:1:1::/64

Gig0

Router1

Gig1

LAN2: 2001:db8:2:2::/64

Area 1

24
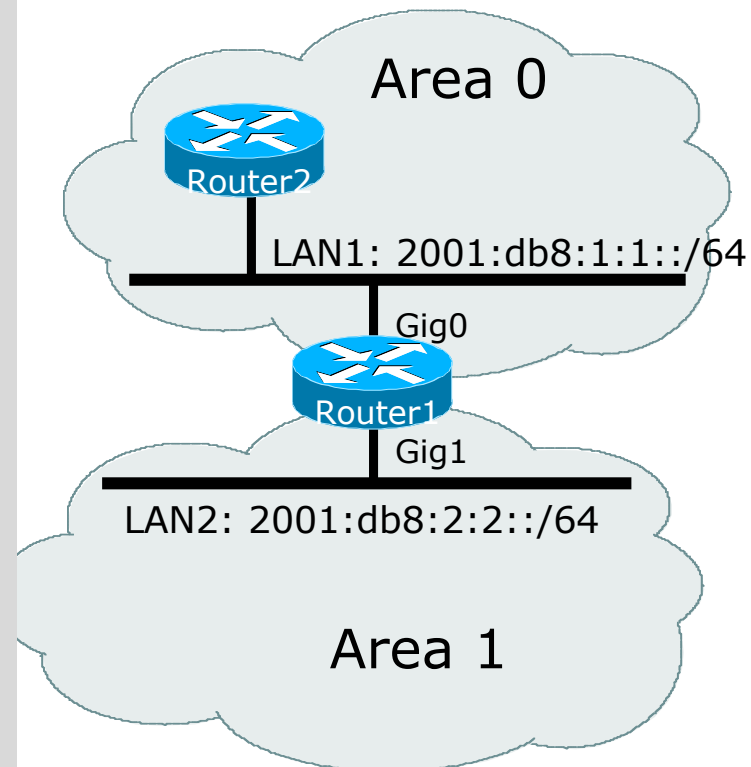
OSPFv3 entries in Routing Table

```
Router2#sh ipv6 route
IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
OI  2001:DB8:FFFF:1::/64 [110/2]
     via FE80::2D0:FFFF:FE60:DFFF, POS3/0
C   2001:DB8:1:1::/64 [0/0]
     via ::, POS3/0
L   2001:DB8:1:1::1/128 [0/0]
     via ::, POS3/0
L   FE80::/10 [0/0]
     via ::, Null0
L   FF00::/8 [0/0]
     via ::, Null0
```

IS-IS Standards History

- ISO 10589 specifies the OSI IS-IS routing protocol for CLNS traffic

- RFC 1195 added IPv4 support
  - Also known as Integrated IS-IS (I/IS-IS)
  - I/IS-IS runs on top of the Data Link Layer

- RFC5308 adds IPv6 address family support

- RFC5120 defines Multi-Topology
  - Permits IPv4 and IPv6 topologies which are not identical
 - Permits roll out of IPv6 without impacting IPv4 operations

Cisco IOS IS-IS dual stack configuration – single topology

LAN1: 2001:db8:1:1::1/64

Gig0

Router1

Gig1

LAN2: 2001:db8:1:2::1/64
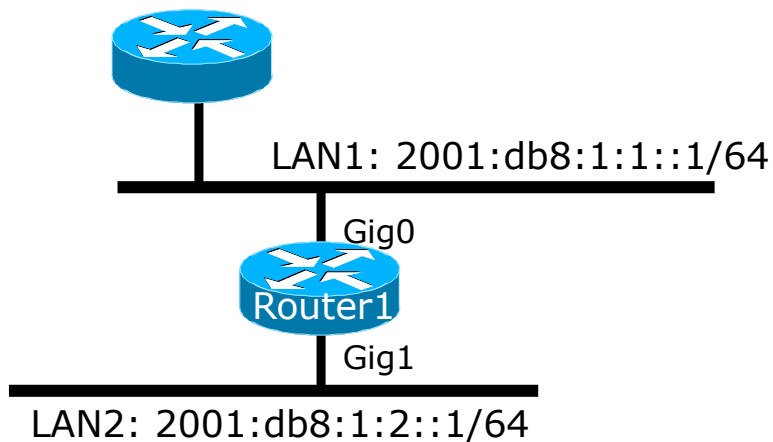
Dual stack IPv4/IPv6 single topology configuration.

```
Router1#
interface GigabitEthernet0
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 ipv6 address 2001:db8:1:1::1/64
 ipv6 router isis
!
interface GigabitEthernet1
 ip address 192.168.2.1 255.255.255.0
 ip router isis
 ipv6 address 2001:db8:1:2::1/64
 ipv6 router isis
!
router isis
 net 42.0001.1921.6800.0001.00
 metric-style wide
!
```

Cisco IOS IS-IS dual stack configuration – multi-topology

LAN1: 2001:db8:1:1::1/64

Gig0

Router1

Gig1

LAN2: 2001:db8:1:2::1/64
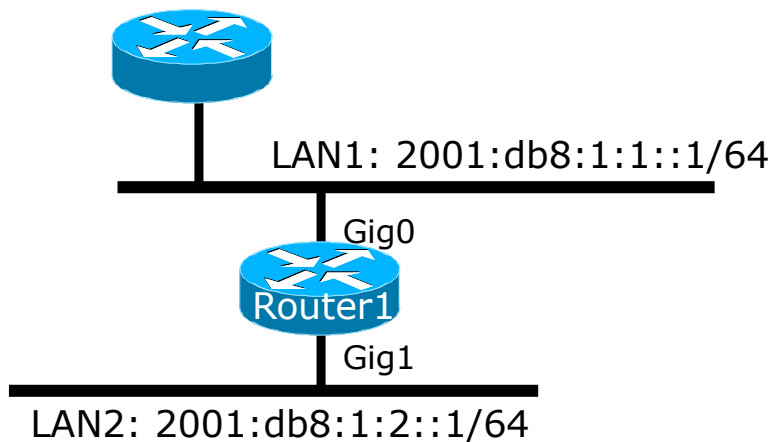
Dual stack IPv4/IPv6 multi-topology configuration.

```
Router1#
interface GigabitEthernet0
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 ipv6 address 2001:db8:1:1::1/64
 ipv6 router isis
!
interface GigabitEthernet1
 ip address 192.168.2.1 255.255.255.0
 ip router isis
 ipv6 address 2001:db8:1:2::1/64
 ipv6 router isis
!
router isis
 net 42.0001.1921.6800.0001.00
 metric-style wide
!
 address-family ipv6
  multi-topology
!
```

# IS-IS entries in Routing Table

```
router1#sh ipv6 route isis

IPv6 Routing Table - default - 46 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP

       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary

       D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, l - LISP

       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

I2   2001:DB8::2/128 [115/2]

      via FE80::C801:3CFF:FE4F:8, FastEthernet0/0

I2   2001:DB8::3/128 [115/20]

      via FE80::C802:3CFF:FE4F:0, Serial1/0

I2   2001:DB8::4/128 [115/22]

      via FE80::C801:3CFF:FE4F:8, FastEthernet0/0

I2   2001:DB8::5/128 [115/40]

      via FE80::C802:3CFF:FE4F:0, Serial1/0

I2   2001:DB8::6/128 [115/42]

      via FE80::C801:3CFF:FE4F:8, FastEthernet0/0

      via FE80::C802:3CFF:FE4F:0, Serial1/0
```

Multi-Protocol 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 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)

Simple IPv6 BGP Session

AS 70000

router1 ──────────── **router2**

AS 80000

2001:db8:2:1::0   2001:db8:2:1::1
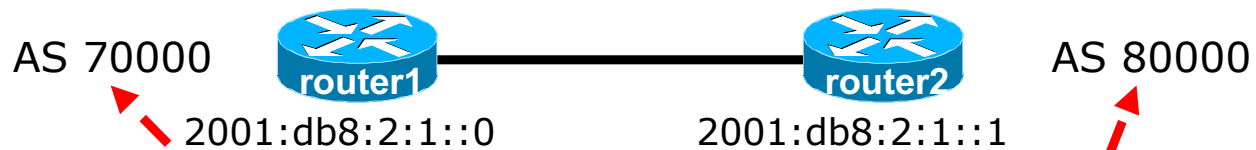
```
Router1#
interface GigabitEthernet0
 ipv6 address 2001:db8:2:1::/127
!
router bgp 70000
 bgp router-id 192.168.0.1
 no bgp default ipv4-unicast
 neighbor 2001:db8:2:1::1 remote-as 80000
 address-family ipv6
  neighbor 2001:db8:2:1::1 activate
  neighbor 2001:db8:2:1::1 prefix-list bgp80000in in
  neighbor 2001:db8:2:1::1 prefix-list bgp80000out
out
  exit-address-family
!
```

Dual Stack BGP Configuration

AS 70000   router1 ———— router2   AS 80000

2001:db8:2:1::0/172.16.1.1   2001:db8:2:1::1/172.16.1.2
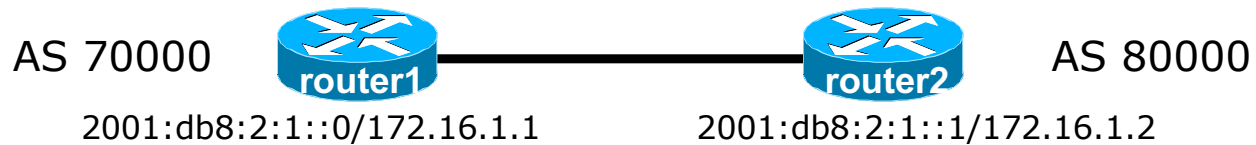
```
router bgp 70000
 no bgp default ipv4-unicast
 neighbor 2001:db8:2:1::1 remote-as 80000
 neighbor 172.16.1.2 remote-as 80000
!
 address-family ipv4
 neighbor 172.16.1.2 activate
 neighbor 172.16.1.2 prefix-list ipv4-ebgp in
 neighbor 172.16.1.2 prefix-list v4out out
 network 172.16.0.0 mask 255.255.0.0
 exit-address-family
!
 address-family ipv6
 neighbor 2001:db8:2:1::1 activate
 neighbor 2001:db8:2:1::1 prefix-list ipv6-ebgp in
 neighbor 2001:db8:2:1::1 prefix-list v6out out
 network 2001:db8::/32
 exit-address-family
!
```

Routing Protocols for IPv6
Summary

- Support for IPv6 in the major routing protocols
- More details about OSPF, IS-IS and BGP in separate presentations

IPv6 Routing Protocols

ISP Workshops