

IPv6 Routing Protocols



AfNOG 2012 AR-E Workshop

Initial IPv6 Configuration for IOS



IPv6 Configuration on Cisco IOS

- To enable IPv6 the following global command should be entered:

`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-routing`

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
```

- This is not useful on a router and is not recommended

IPv6 Configuration

- Link-local address:
 - Enabling IPv6 on an interface using:
`Router(config-if)# ipv6 enable`
 - 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

```
Router1# conf t
Router1(config)# ipv6 unicast-routing
Router1(config)# ipv6 cef
Router1(config)# int fast 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#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
ND DAD is enabled, number of DAD attempts: 1
ND reachable time is 30000 milliseconds
ND advertised reachable time is 0 milliseconds
ND advertised retransmit interval is 0 milliseconds
ND router advertisements are sent every 200 seconds
ND router advertisements live for 1800 seconds
Hosts use stateless autoconfig for addresses.
```

IPv6 Configuration – Static

```
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:FE00:F11D
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
```


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:DB8:0:ABCD::1 150
```

- Routes packets for network 2001:db8::/64 to a networking device at 2001:DB8:0:ABCD::1 with an administrative distance of 150

Default Routing Example



```
ipv6 unicast-routing
!  
interface Ethernet0  
  ipv6 address 2001:db8:2:1::a/64  
!  
interface Ethernet1  
  ipv6 address 2001:db8:2:2::a/64  
!  
ipv6 route ::/0 2001:db8:2:1::e
```

Default Route
to Router2

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 router-id
 - Router-id is a 32 bit integer
 - 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



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

interface Ethernet0
ipv6 address 2001:db8:2:1::/64 eui-64
ipv6 enable
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

OSPFv3



OSPFv3 overview

- ❑ OSPF for IPv6
- ❑ Based on OSPFv2, with enhancements
- ❑ Distributes IPv6 unicast prefixes
- ❑ Runs directly over IPv6
- ❑ Ships-in-the-night with OSPFv2
- ❑ OSPFv3 does **not** carry IPv4 prefixes
 - RFC5838 proposes an extension which adds address family support

OSPFv3 / OSPFv2 Similarities

- Basic packet types
 - Hello, DBD, LSR, LSU, LSA
- Mechanisms for neighbor discovery and adjacency formation
- Interface types
 - P2P, P2MP, Broadcast, NBMA, Virtual
- LSA flooding and aging
- Nearly identical LSA types

V2, V3 Differences

OSPFv3 runs on a Link instead of per IP Subnet

- ❑ A link by definition is a medium over which two nodes can communicate at link layer
- ❑ In IPv6 multiple IP subnet can be assigned to a link and two nodes in different subnet can communicate at link layer therefore OSPFv3 is running per link instead of per IP subnet
- ❑ An Interface connect to a link and multiple interface can be connected to a link

V2, V3 Differences (Cont.)

Support of Multiple Instance per Link

- ❑ New field (instance) in OSPF packet header allow running multiple instance per link
- ❑ Instance ID should match before packet being accepted
- ❑ Useful for traffic separation, multiple areas per link and address families (RFC5838)

V2, V3 Differences (Cont.)

Address Semantic Change in LSA

- ❑ Router and Network LSA carry only topology information
- ❑ Router LSA can be split across multiple LSAs; Link State ID in LSA header is a fragment ID
- ❑ Intra area prefix are carried in a new LSA payload called intra-area-prefix-LSAs
- ❑ Prefix are carried in payload of inter-area and external LSA

V2, V3 Differences (Cont.)

Generalisation of Flooding Scope

- ❑ In OSPFv3 there are three flooding scope for LSAs (link-local scope, area scope, AS scope) and they are coded in LS type explicitly
- ❑ In OSPFv2 initially only area and AS wide flooding was defined; later opaque LSAs introduced link local scope as well

V2, V3 Differences (Cont.)

Explicit Handling of Unknown LSA

- ❑ The handling of unknown LSA is coded via U-bit in LS type
- ❑ When U bit is set, the LSA is flooded with the corresponding flooding scope, as if it was understood
- ❑ When U bit is clear, the LSA is flooded with link local scope
- ❑ In v2 unknown LSA were discarded

V2, V3 Differences (Cont.)

Authentication is Removed from OSPF

- ❑ Authentication in OSPFv3 has been removed
- ❑ OSPFv3 relies now on the IPv6 authentication header since OSPFv3 run over IPv6
- ❑ Autype and Authentication field in the OSPF packet header therefore have been suppressed

V2, V3 Differences (Cont.)

OSPF Packet format has been changed

- ❑ The mask field has been removed from Hello packet
- ❑ IPv6 prefix are only present in payload of Link State update packet

V2, V3 Differences (Cont.)

Two New LSAs Have Been Introduced

- Link-LSA has a link local flooding scope and has three purposes:
 - The router link local address
 - List all IPv6 prefixes attached to the link
 - Assert a collection of option bit for the Router-LSA
- Intra-area-prefix-LSA
 - Used to advertise router's IPv6 address within the area

Configuring OSPFv3 in Cisco IOS

- Similar to OSPFv2
 - Prefixing existing Interface and Exec mode commands with `ipv6`
- Interfaces configured directly
 - Replaces `network` command
 - (Also available in OSPFv2 from IOS 12.4)
- “Native” IPv6 router mode
 - Not a sub-mode of `router ospf`

Configuring OSPFv3

- ❑ Setting up the OSPFv3 process:

```
[no] ipv6 router ospf <process ID>
```

- ❑ Applying the OSPFv3 process to an interface:

```
interface <router-int-name>
```

```
  [no] ipv6 ospf <process ID> area <area ID>
```

- ❑ Configuring summarisation:

```
ipv6 router ospf <process ID>
```

```
  [no] area <area ID> range <prefix>/<length>
```

OSPFv3 exec mode commands

- ❑ Exec mode commands:

 - `show ipv6 ospf [<process ID>]`

 - `clear ipv6 ospf [<process ID>]`

- ❑ Showing new LSA:

 - `show ipv6 ospf [<process ID>] database link`

 - `show ipv6 ospf [<process ID>] database prefix`

OSPFv3 Authentication

- ❑ Configuring authentication per area:

- SPI value has to be unique per area:

```
ipv6 router ospf <process ID>
```

```
area 0 authentication ipsec spi 256 md5 <passwd>
```

- ❑ Disabling authentication on a specific link when area authentication is activated:

```
interface fastethernet 0/0
```

```
ipv6 ospf authentication null
```

- ❑ Configuring authentication per interface:


- SPI value has to be unique per link:

```
interface fastethernet 0/0
```

```
ipv6 ospf authentication ipsec spi 256 md5 <passwd>
```


OSPFv3 Configuration Example

```
Router1#  
interface POS1/1  
  ipv6 address 2001:db8:FFFF:1::1/64  
  ipv6 ospf 100 area 0  
!  
interface POS2/0  
  ipv6 address 2001:db8:1:1::2/64  
  ipv6 ospf 100 area 1  
!  
  ipv6 router ospf 100  
  log-adjacency-changes  
!  
  
Router2#  
interface POS3/0  
  ipv6 address 2001:db8:1:1::1/64  
  ipv6 ospf 100 area 1  
!  
  ipv6 router ospf 100  
  log-adjacency-changes
```



OSPFv3 Interface Status

```
Router2#sh ipv6 ospf int pos 3/0
POS3/0 is up, line protocol is up
  Link Local Address FE80::290:86FF:FE5D:A000, Interface ID 7
  Area 1, Process ID 100, Instance ID 0, Router ID 10.1.1.4
  Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 3, maximum is 3
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.3
  Suppress hello for 0 neighbor(s)
```

OSPFv3 Neighbour Status

```
Router2#sh ipv6 ospf neighbor detail
```

```
Neighbor 10.1.1.3
```

```
  In the area 1 via interface POS3/0
```

```
Neighbor: interface-id 8, link-local address FE80::2D0:FFFF:FE60:DFFF
```

```
Neighbor priority is 1, State is FULL, 12 state changes
```

```
Options is 0x630C34B9
```

```
Dead timer due in 00:00:33
```

```
Neighbor is up for 00:49:32
```

```
Index 1/1/1, retransmission queue length 0, number of retransmission 1
```

```
First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
```

```
Last retransmission scan length is 2, maximum is 2
```

```
Last retransmission scan time is 0 msec, maximum is 0 msec
```

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
```

OSPFv3 link troubleshooting

- ❑ Next router address in OSPFv3 is a link-local address

```
OI 2001:db8:FFFF:1::/64 [110/2]
   via FE80::2D0:FFFF:FE60:DFFF, POS3/0
```

- ❑ How to troubleshoot??
 - SSH to neighbouring router needs extended SSH command, for example:
`ssh FE80::2D0:FFFF:FE60:DFFF /source-int POS3/0`
 - Source interface has to be specified – a router with multiple interfaces has no idea which interface the remote link local address is attached to

Conclusion

- ❑ Based on existing OSPFv2 implementation
- ❑ Similar CLI and functionality

ISIS



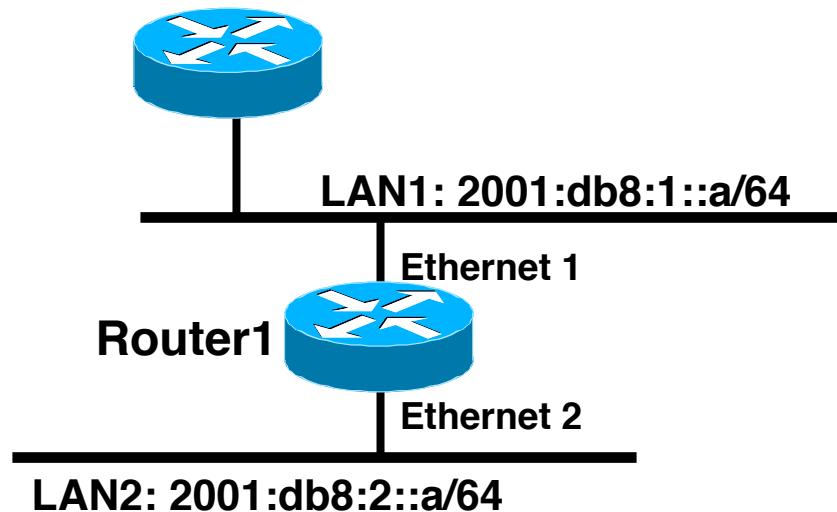
IS-IS Standards History

- ❑ IETF ISIS for Internets Working Group
- ❑ ISO 10589 specifies OSI IS-IS routing protocol for CLNS traffic
 - Tag/Length/Value (TLV) options to enhance the protocol
 - A Link State protocol with a 2 level hierarchical architecture.
- ❑ RFC 1195 added IP support, also known as Integrated IS-IS (I/IS-IS)
 - I/IS-IS runs on top of the Data Link Layer
 - Requires CLNP to be configured
- ❑ RFC5308 adds IPv6 address family support to IS-IS
- ❑ RFC5120 defines Multi-Topology concept for IS-IS
 - Permits IPv4 and IPv6 topologies which are not identical
 - Allows gradual roll out of IPv6 across backbone without⁴⁰ impacting IPv4

IS-IS for IPv6

- ❑ 2 Tag/Length/Values added to introduce IPv6 routing
- ❑ IPv6 Reachability TLV (0xEC)
 - External bit
 - Equivalent to IP Internal/External Reachability TLV's
- ❑ IPv6 Interface Address TLV (0xE8)
 - For Hello PDUs, must contain the Link-Local address
 - For LSP, must only contain the non-Link Local address
- ❑ IPv6 NLPID (0x8E) is advertised by IPv6 enabled routers

Cisco IOS IS-IS dual stack configuration




**Dual IPv4/IPv6 configuration.
Redistributing both IPv6 static routes
and IPv4 static routes.**

```
Router1#  
interface ethernet 1  
  ip address 10.1.1.1 255.255.255.0  
  ipv6 address 2001:db8:1::a/64  
  ip router isis  
  ipv6 router isis  
  
interface ethernet 2  
  ip address 10.2.1.1 255.255.255.0  
  ipv6 address 2001:db8:2::a/64  
  ip router isis  
  ipv6 router isis  
  
router isis  
  address-family ipv6  
    redistribute static  
  exit-address-family  
  net 42.0001.0000.0000.072c.00  
  redistribute static
```

Multi-Topology IS-IS extensions

- New TLVs attributes for Multi-Topology extensions.
 - Multi-topology TLV: contains one or more multi-topology ID in which the router participates. It is theoretically possible to advertise an infinite number of topologies. This TLV is included in IIH and the first fragment of a LSP.
 - MT Intermediate Systems TLV: this TLV appears as many times as the number of topologies a node supports. A MT ID is added to the extended IS reachability TLV type 22.
 - Multi-Topology Reachable IPv4 Prefixes TLV: this TLV appears as many times as the number of IPv4 announced by an IS for a given MT ID. Its structure is aligned with the extended IS Reachability TLV Type 236 and add a MT ID.
 - Multi-Topology Reachable IPv6 Prefixes TLV: this TLV appears as many times as the number of IPv6 announced by an IS for a given MT ID. Its structure is aligned with the extended IS Reachability TLV Type 236 and add a MT ID.
- Multi-Topology ID Values
 - Multi-Topology ID (MT ID) standardized and in use in Cisco IOS:
 - MT ID #0 – “standard” topology for IPv4/CLNS
 - MT ID #2 – IPv6 Routing Topology.

Multi-Topology ISIS configuration example



- ❑ The optional keyword **transition** may be used for transitioning existing IS-IS IPv6 single SPF mode to MT IS-IS
- ❑ Wide metric is mandated for Multi-Topology to work

```
Router1#
interface ethernet-1
 ip address 10.1.1.1 255.255.255.0
 ipv6 address 2001:db8:1::1/64
 ip router isis
 ipv6 router isis
 isis ipv6 metric 20

interface ethernet-2
 ip address 10.2.1.1 255.255.255.0
 ipv6 address 2001:db8:2::1/64
 ip router isis
 ipv6 router isis
 isis ipv6 metric 20

router isis
 net 49.0000.0100.0000.0000.0500
 metric-style wide
 !
 address-family ipv6
 multi-topology
 exit-address-family
```

BGP



Adding IPv6 to BGP...

□ RFC4760

- Defines Multi-protocol Extensions for BGP4
- Enables BGP to carry routing information of protocols other than IPv4
 - e.g. MPLS, IPv6, Multicast etc
- Exchange of multiprotocol NLRI must be negotiated at session startup

□ RFC2545

- Use of BGP Multiprotocol Extensions for IPv6 Inter-Domain Routing

RFC4760

- ❑ New optional and non-transitive BGP attributes:
 - MP_REACH_NLRI (Attribute code: 14)
 - ❑ Carry the set of reachable destinations together with the next-hop information to be used for forwarding to these destinations (RFC2858)
 - MP_UNREACH_NLRI (Attribute code: 15)
 - ❑ Carry the set of unreachable destinations
- ❑ Attribute contains one or more Triples:
 - AFI Address Family Information
 - Next-Hop Information
(must be of the same address family)
 - NLRI Network Layer Reachability Information

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)

BGP Considerations


- Rules for constructing the NEXTHOP attribute:
 - When two peers share a common subnet, the NEXTHOP information is formed by a global address and a link local address
 - Redirects in IPv6 are restricted to the usage of link local addresses

Routing Information

- Independent operation
 - One RIB per protocol
 - e.g. IPv6 has its own BGP table
 - Distinct policies per protocol
- Peering sessions can be shared when the topology is congruent

BGP next-hop attribute

- ❑ Next-hop contains a global IPv6 address (or potentially a link local address)
- ❑ Link local address is set as a next-hop only if the BGP peer shares the subnet with both routers (advertising and advertised)



More BGP considerations

- TCP Interaction
 - BGP runs on top of TCP
 - This connection could be set up either over IPv4 or IPv6
- Router ID
 - When no IPv4 is configured, an explicit bgp router-id needs to be configured
 - BGP identifier is a 32 bit integer currently generated from the router identifier – which is generated from an IPv4 address on the router
 - This is needed as a BGP identifier, this is used as a tie breaker, and is sent within the OPEN message

BGP Configuration

- Two options for configuring BGP peering
- Using link local addressing
 - ISP uses FE80:: addressing for BGP neighbours
 - **NOT RECOMMENDED**
 - There are plenty of IPv6 addresses
 - Unnecessary configuration complexity
- Using global unicast addresses
 - As with IPv4
 - **RECOMMENDED**

BGP Configuration

- ❑ Cisco IOS assumes that all BGP neighbours will be IPv4 unicast neighbours
 - We need to remove this assumption

```
router bgp 100
  no bgp default ipv4-unicast
```


- ❑ Failing to do this will result in all neighbours being defined as IPv4 unicast neighbours
 - Non-IPv4 neighbours will have no specific unicast IPv4 configuration
 - Cluttered configuration, confusing troubleshooting and diagnosis

BGP Configurations

Regular Peering

Router A

```
router bgp 1
  no bgp default ipv4-unicast
  bgp router-id 1.1.1.1
  neighbor 2001:db8:0:2::2 remote-as 2
  !
  address-family ipv6
  neighbor 2001:db8:0:2::2 activate
  network 2001:db8:2::/48
  network 2001:db8:3::/48
  !
```




`activate` means that the BGP peering is activated for this particular address family

BGP Configurations

Link Local Peering

Router A

```
interface fastethernet 0/0
  ipv6 address 2001:db8:ffc0:1::1/64
  !
router bgp 1
  no bgp default ipv4-unicast
  bgp router-id 1.1.1.1
  neighbor fe80::260:3eff:c043:1143 remote-as 2
  !
  address-family ipv6
    neighbor fe80::260:3eff:c043:1143 activate
    neighbor fe80::260:3eff:c043:1143 route-map next-hop in
  !
  route-map next-hop permit 5
    set ipv6 next-hop 2001:db8:ffc0:1::1
  !
```



BGP Configuration

IPv4 and IPv6

```
router bgp 10
  no bgp default ipv4-unicast
  neighbor 2001:db8:1:1019::1 remote-as 20
  neighbor 172.16.1.2 remote-as 30
  !
  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
  exit-address-family
  !
  address-family ipv6
    neighbor 2001:db8:1:1019::1 activate
    neighbor 2001:db8:1:1019::1 prefix-list ipv6-ebgp in
    neighbor 2001:db8:1:1019::1 prefix-list v6out out
    network 2001:db8::/32
  exit-address-family
  !
  ! Continued -->
```

BGP Configuration

IPv4 and IPv6

```
ip prefix-list ipv4-ebgp permit 0.0.0.0/0 le 32
!  
ip prefix-list v4out permit 172.16.0.0/16
!  
ipv6 prefix-list ipv6-ebgp permit ::/0 le 128
!  
ipv6 prefix-list v6out permit 2001:db8::/32
!
```

- Compare IPv4 prefix filters with IPv6 prefix filters

```
ip prefix-list <name> permit|deny <ipv4 address>
```

```
ipv6 prefix-list <name> permit|deny <ipv6 address>
```

BGP Configuration

IPv4 and IPv6

- ❑ When configuring the router, recommendation is:
 - Put all IPv6 configuration directly into IPv6 address family
 - Put all IPv4 configuration directly into IPv4 address family
- ❑ Router will sort generic from specific address family configuration when the configuration is saved to NVRAM or displayed on the console
- ❑ Example follows...
 - Notice how activate is added by the router to indicate that the peering is activated for the particular address family

BGP Address Families

Applied Configuration

```
router bgp 10
  no bgp default ipv4-unicast
  !
  address family ipv4
    neighbor 172.16.1.2 remote-as 30
    neighbor 172.16.1.2 prefix-list ipv4-ebgp in
    neighbor 172.16.1.2 prefix-list v4out out
    network 172.16.0.0
  !
  address-family ipv6
    neighbor 2001:db8:1:1019::1 remote-as 20
    neighbor 2001:db8:1:1019::1 prefix-list ipv6-ebgp in
    neighbor 2001:db8:1:1019::1 prefix-list v6out out
    network 2001:db8::/32
  !
  ip prefix-list ipv4-ebgp permit 0.0.0.0/0 le 32
  ip prefix-list v4out permit 172.16.0.0/16
  ipv6 prefix-list ipv6-ebgp permit ::/0 le 128
  ipv6 prefix-list v6out permit 2001:db8::/32
```

Generic Configuration

Specific Configuration

BGP Address Families

End result

```
router bgp 10
  no bgp default ipv4-unicast
  neighbor 2001:db8:1:1019::1 remote-as 20
  neighbor 172.16.1.2 remote-as 30
!
  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
  exit-address-family
!
  address-family ipv6
  neighbor 2001:db8:1:1019::1 activate
  neighbor 2001:db8:1:1019::1 prefix-list ipv6-ebgp in
  neighbor 2001:db8:1:1019::1 prefix-list v6out out
  network 2001:db8::/32
  exit-address-family
!
ip prefix-list ipv4-ebgp permit 0.0.0.0/0 le 32
ip prefix-list v4out permit 172.16.0.0/16
ipv6 prefix-list ipv6-ebgp permit ::/0 le 128
ipv6 prefix-list v6out permit 2001:db8::/32
```

Generic Configuration

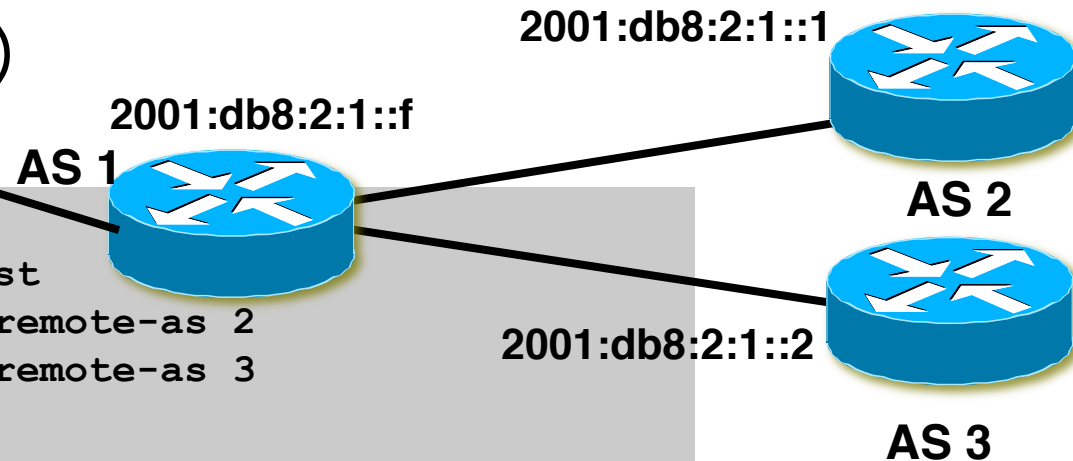
Specific Configuration

BGP Configuration

Manipulating Attributes

- ❑ Prefer routes from AS 2 (local preference)

```
router bgp 1
no bgp default ipv4-unicast
neighbor 2001:db8:2:1::1 remote-as 2
neighbor 2001:db8:2:1::2 remote-as 3
!
address-family ipv6
neighbor 2001:db8:2:1::1 activate
neighbor 2001:db8:2:1::1 prefix-list in-filter in
neighbor 2001:db8:2:1::1 route-map fromAS2 in
neighbor 2001:db8:2:1::2 activate
neighbor 2001:db8:2:1::2 prefix-list in-filter in
network 2001:db8::/32
exit-address-family
!
route-map fromAS2 permit 10
set local-preference 120
```



BGP Configuration

Carrying IPv4 inside IPv6 peering

- IPv4 prefixes can be carried inside an IPv6 peering
 - Note that the next-hop for received prefixes needs to be “fixed”
- Example

```
router bgp 1
  neighbor 2001:db8:0:2::2 remote-as 2
  !
address-family ipv4
  neighbor 2001:db8:0:2::2 activate
  neighbor 2001:db8:0:2::2 route-map ipv4 in
  !
route-map ipv4 permit 10
  set ip next-hop 131.108.1.1
```

BGP Status Commands

- ❑ IPv6 BGP show commands take ipv6 as argument
`show bgp ipv6 unicast <parameter>`

```
Router1#show bgp ipv6 unicast 2017::/32
BGP routing table entry for 2017::/32, version 11
Paths: (1 available, best #1)
Local
2001:db8:c18:2:1::1 from 2001:db8:c18:2:1::1
(10.10.20.2)
Origin incomplete, localpref 100, valid, internal, best
```

- ❑ IPv4 BGP show commands can also use this format:
`show bgp ipv4 unicast <parameter>`

BGP Status Commands

- Display summary information regarding the state of the BGP neighbours

```
show bgp ipv6 unicast summary
```

```
BGP router identifier 128.107.240.254, local AS number 109
BGP table version is 400386, main routing table version 400386
585 network entries using 78390 bytes of memory
9365 path entries using 674280 bytes of memory
16604 BGP path attribute entries using 930384 bytes of memory
8238 BGP AS-PATH entries using 228072 bytes of memory
42 BGP community entries using 1008 bytes of memory
9451 BGP route-map cache entries using 302432 bytes of memory
584 BGP filter-list cache entries using 7008 bytes of memory
BGP using 2221574 total bytes of memory
2 received paths for inbound soft reconfiguration
BGP activity 63094/62437 prefixes, 1887496/1878059 paths, scan interval 60secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
2001:1458:C000::64B:4:1									
	4	513	1294728	460213	400386	0	0	3d11h	498

↑
Neighbour Information

↑
BGP Messages Activity

Conclusion

- ❑ BGP extended to support multiple protocols
 - IPv6 is but one more address family
- ❑ Operators experienced with IPv4 BGP should have no trouble adapting
 - Configuration concepts and CLI is familiar format

Routing Protocols for IPv6

Summary

- ❑ Support for IPv6 in the major routing protocols
- ❑ We will talk about ISIS in more detail later in the workshop

IPv6 Routing Protocols



AfNOG 2012 AR-E Workshop