Deploying IPv6 Routing Protocols

Specifics and Considerations

Peter Palúch - Cisco Training Bootcamps, CCIE #23527
@Peter_Paluch, ppaluch@cisco.com

BRKIPV-2418

What this session is about (abstract)

Even though all major IPv6 routing protocols are rooted in their well-known IPv4 counterparts, they all come with their own unique set of idiosyncrasies. And when faced with the task of deploying an IPv6 routing protocol along with its IPv4 variant, there are often multiple ways of doing it, but the best choice is not always obvious.

The goal of this session is to discuss the specifics of IPv6 support in OSPFv3, IS-IS, EIGRP, and BGP in IOS XE. For each of these protocols, we will do a quick IPv6-focused refresher, and then focus on the protocol's unique IPv6 traits (and quirks). We will also look at how the protocol supports the co-existence of IPv4 and IPv6, and if there are multiple options to choose from, we will discuss the pros and cons of each of them.

Agenda

- Considerations for OSPF
- Considerations for IS-IS
- Considerations for EIGRP
- Considerations for BGP
- Considerations for FHRPs
- Considerations for link-local only addressing in networks

Timeline of IPv6 support

IPv6 and Support in Routing Protocols

- Work on the next-generation IP protocol started in 1992
- Basic IPv6 as we know it emerged first in 1994 as an Internet Draft proposal, and in December 1995 as RFC 1883
- Formal routing protocol specifications followed suit
  - RIPng – RFC 2080, January 1997 (first draft in February 1996)
  - OSPFv3 – RFC 2740, December 1999 (first draft in February 1996)
  - BGP – RFC 2545, March 1999 (first draft in February 1997)
  - IS-IS – RFC 5308, October 2008 (first draft in January 2000)

IPv6 Routing Protocol Support in IOS and IOS-XR

Protocol	In IOS since	In IOS-XR since
RIPng	12.2(2)T	7.5.2
OSPFv3	12.0(24)S / 12.2(15)T	3.7.2
OSPFv3 with Address Family Support	15.1(3)S / 15.2(1)T	N/A
IS-IS	12.2(8)T / 12.0(22)S	3.9.0
BGP	12.2(2)T	3.7.2
EIGRP	12.4(6)T	3.7.2

Considerations for OSPF

IPv6 Support in Open Shortest Path First

- OSPFv1 (RFC 1131) and OSPFv2 (RFC 2328) only support IPv4
- OSPFv3 (RFC 5340) is a substantial rework to support IPv6
  - Initially, OSPFv3 supported only IPv6
  - With RFC 5838, OSPFv3 supports address families to advertise both IPv6 and IPv4 prefixes
- Routing both IPv4 and IPv6 with OSPF always requires two entirely independent processes with all their packet exchanges
  - Either OSPFv2 for IPv4 and OSPFv3 for IPv6,
  - Or two OSPFv3 instances, one for IPv4 AF, the other for IPv6 AF

Brush Up on OSPFv3 Basic Traits

- OSPFv3 messages are carried in IPv6 packets
  - Sender address: link-local address of the interface
  - Destination address: multicast ff02::5, ff02::6, or the link-local address of the directly connected neighbor
  - Global addresses are only used for virtual and sham links
- The 5 basic message types (Hello, DBD, LSR, LSU, LSAck) stay
- Link State Advertisements have been reworked
- Authentication was first left to IPsec, later brought back as a protocol extension in RFC 7166

Reasons for Reworking LSAs in OSPFv3

- OSPFv2 “smartly” combined topology and addressing information, using 4-byte fields for both IDs and addresses
  - It was impossible to insert IPv6 addresses into existing LSA formats
  - It was impossible to tell apart a topology change from an address change
- OSPFv3 brings several changes to LSAs
  - Addressing information moved out from Router and Network LSAs into new LSA types: Link LSA and Intra-Area-Prefix LSA
  - Topology and addressing information are diligently separated
  - IDs remain 4 bytes long

Comparing OSPFv2 and OSPFv3 LSA Types

OSPFv2	OSPFv3
Router-LSA	Router-LSA
Network-LSA	Network-LSA
Summary-LSA (Network)	Inter-Area-Prefix-LSA
Summary-LSA (ASBR)	Inter-Area-Router-LSA
AS-External-LSA	AS-External-LSA
NSSA-LSA	NSSA-LSA
	Link-LSA
	Intra-Area-Prefix-LSA

Address Information

Will they establish an adjacency? (1)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ipv6 ospf 1 area 0
!
ipv6 router ospf 1
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ipv6 ospf 1 area 0
!
ipv6 router ospf 1
```

Yes – this is basic legacy style (non-AF) OSPFv3 configuration

Will they establish an adjacency? (2)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ipv6 ospf 1 area 0
!
ipv6 router ospf 1
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ip ospf 1 area 0
!
router ospf 1
```

No – OSPFv3 does not interoperate with OSPFv2

Will they establish an adjacency? (3)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ipv6 ospf 2 area 0
!
ipv6 router ospf 2
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ospfv3 1 ipv6 area 0
!
router ospfv3 1
 address-family ipv6 unicast
```

Yes – non-AF and AF-enabled OSPFv3 for IPv6 interoperate

Will they establish an adjacency? (4)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ip ospf 1 area 0
!
router ospf 1
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ospfv3 1 ipv4 area 0
!
router ospfv3 1
 address-family ipv4 unicast
```

No – OSPFv3 does not interoperate with OSPFv2 even in IPv4 AF

OSPFv3 Adjacency Formation Overview

R1 Configuration	R2 Configuration	Adjacency?
router ospfv3 / IPv6 AF	ipv6 router ospf	Yes
router ospfv3 / IPv6 AF	router ospf	No
router ospfv3 / IPv6 AF	router ospfv3 / IPv4 AF	No
router ospfv3 / IPv4 AF	ipv6 router ospf	No
router ospfv3 / IPv4 AF	router ospf	No
ipv6 router ospf	router ospf	No

Identical configurations on both routers omitted for brevity

IPv4 with OSPFv2 vs. OSPFv3 IPv4 AF

- When considering the choice of OSPFv2 vs. OSPFv3 IPv4 AF, there are some compelling arguments in favor of OSPFv3
  - If also running IPv6, using a single protocol means greater consistency
  - Potential better use of Partial SPF as opposed to Full SPF
  - Prefix suppression works more efficiently
  - IPv4 addresses on inter-router links do not matter much (can be from discontiguous networks or unnumbered if allowed)
  - Stub router functionality through the R-bit
  - IPsec authentication and even encryption of OSPFv3 packets if desired

CLI gotchas with OSPFv3 IPv4 AF

```
r2# show ip route
[ ... snip ... ]
 10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C        10.0.12.0/24 is directly connected, Ethernet0/1
L        10.0.12.2/32 is directly connected, Ethernet0/1
C        10.0.23.0/24 is directly connected, Ethernet0/0
L        10.0.23.2/32 is directly connected, Ethernet0/0
O        10.255.255.1/32 [110/10] via 10.0.12.1, 00:03:12, Ethernet0/1
C        10.255.255.2/32 is directly connected, Loopback0
O        10.255.255.3/32 [110/10] via 10.0.23.3, 00:03:12, Ethernet0/0
O     192.0.2.0/24 [110/11] via 10.0.12.1, 00:03:12, Ethernet0/1
O     203.0.113.0/24 [110/11] via 10.0.23.3, 00:03:12, Ethernet0/0

r2# show ip route ospf
r2#
```

What's going on here?

OSPFv3 CLI

```
r2# show ip route ospf?
Hostname or A.B.C.D  ospf  ospfv3

r2# show ip route ospfv3
[ ... snip ... ]
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O        10.255.255.1/32 [110/10] via 10.0.12.1, 00:08:15, Ethernet0/1
O        10.255.255.3/32 [110/10] via 10.0.23.3, 00:08:15, Ethernet0/0
O     192.0.2.0/24 [110/11] via 10.0.12.1, 00:08:15, Ethernet0/1
O     203.0.113.0/24 [110/11] via 10.0.23.3, 00:08:15, Ethernet0/0
r2#
```

OSPFv3 Authentication

- Initially, OSPFv3 did not have its own authentication mechanism
  - The task of authenticating OSPFv3 packets was relegated to IPsec
- Real life showed little adoption
  - As a protocol suite, IPsec is complex
  - IPsec support is far from ubiquitous
  - Configuration is purely manual and rather tedious
- RFC 7166 brings back the authentication trailer function to OSPFv3
  - Easier to configure through key chains

OSPFv3 Authentication Trailer Configuration


```
key chain keys
 key 1
  key-string CiscoLive
   cryptographic-algorithm hmac-sha-512
!
interface Gi1
 [ ... snip ... ]
 ospfv3 1 ipv6 authentication key-chain keys
```

"Keep-in-Mind" Facts About OSPFv3

- Despite providing services to IPv4, OSPFv3 for IPv4 continues to be encapsulated in IPv6
  - Important for ACLs, QoS, monitoring
- Virtual links are not supported in OSPFv3 IPv4 AFI
  - This is because OSPFv3 uses IPv6 but there is neither a guarantee that the virtual link endpoints have global IPv6 addresses, nor can they be advertised in an IPv4 AFI
- In pure IPv6 environment, manual RID configuration is required

The more you improve OSPF, the more you get IS-IS.

Considerations for IS-IS

IPv6 Support in IS-IS

- IS-IS by its very nature invites multiprotocol capability
  - Runs directly over Layer2 frames
  - Information is encoded as Type-Length-Value (TLV) records
  - Link State PDUs carry TLVs of various types within their flooding scope
- RFC 1195 brought IPv4 support to IS-IS
  - Integrated IS-IS – a single IS-IS instance handling both OSI and IPv4
- RFC 5308 brings IPv6 support to IS-IS
  - Continues the integrated approach – a single IS-IS instance possibly handling OSI, IPv4 and IPv6 routing at once

RFC 5308 IPv6 Extensions to IS-IS

- IPv6 Reachability TLV 236 (0xEC)
  - Carries a single IPv6 prefix with its metric and other attributes
- IPv6 Interface Address TLV 232 (0xE8)
  - Carries an interface's IPv6 address
- IPv6 Network Layer Protocol ID value of 142 (0x8E)
  - Indicates support of IPv6

Consequences of integrated AF handling in IS-IS

- IS-IS handles all enabled routed protocols in a single instance
  - On wire, there is only a single exchange of PDUs for all address families
  - If an adjacency is torn down or cannot be established, all address families are impacted
  - Churn in one address family still causes flooding new LSPs with information about all address families
  - By default, on IOS, IOS XE and NX-OS, all address families are forced to share the same topology, the same link costs, and hence the same best paths

Will they establish an adjacency? (1)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ipv6 router isis
!
router isis
 net 49.0001.1111.1111.1111.00
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ipv6 router isis
!
router isis
 net 49.0001.2222.2222.2222.00
```

Yes – this is basic configuration of IS-IS for IPv6

Will they establish an adjacency? (2)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ip router isis
 ipv6 router isis
!
router isis
 net 49.0001.1111.1111.1111.00
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ip router isis
!
router isis
 net 49.0001.2222.2222.2222.00
```

No – R1 is missing the IPv6 Interface Address TLV in IIH from R2

Will they establish an adjacency? (3)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ip router isis
 ipv6 router isis
!
router isis
 net 49.0001.1111.1111.1111.00
 address-family ipv6
  no adjacency-check
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ip router isis
!
router isis
 net 49.0001.2222.2222.2222.00
```

Will they establish an adjacency? (4)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ip router isis
 ipv6 router isis
!
router isis
 net 49.0001.1111.1111.1111.00
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ipv6 router isis
!
router isis
 net 49.0001.2222.2222.2222.00
```

No – R1 is missing the IPv4 Interface Address TLV in IIH from R2

Will they establish an adjacency? (5)

Solution


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ip router isis
 ipv6 router isis
!
router isis
 net 49.0001.1111.1111.1111.00
 no adjacency-check
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ipv6 router isis
!
router isis
 net 49.0001.2222.2222.2222.00
```

Enabling IPv4 and IPv6 in an IS-IS instance

- If running IS-IS for one IP protocol already, enabling IS-IS for the other IP protocol may cause the existing adjacency to drop
  - This is because one router requires the neighbor to pass adjacency sanity checks which will fail if both protocols are not enabled at the same time
  - The workaround is to temporarily disable the adjacency checks
- By default, once IPv4 and IPv6 are enabled, they will both be forced onto a single topology (IOS, IOS XE, NX-OS)
  - No ability to define independent link costs for IPv4 and IPv6
  - The best paths for IPv4 and IPv6 will be always identical

IS-IS Multi Topology Extensions

- RFC 5120 brings Multi Topology (MT) Extensions to IS-IS
  - Ability to treat IPv4 and IPv6 topologies independently
- MT IS-IS has tremendous advantages
  - Enabling MT and IPv6 on a router does not cause adjacency drops with neighbors that do not run MT and/or IPv6
  - IPv4 and IPv6 topologies can have independent link costs
  - Even with MT, IS-IS still runs a single instance
- IOS-XR defaults to the use of multi-topology extensions
- The prerequisite for using MT Extensions is to run wide metrics

Configuring Multi Topology Extensions in IS-IS


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
 ip router isis
 ipv6 router isis
 isis ipv6 metric 1234
!
router isis
 net 49.0001.1111.1111.1111.00
 metric-style wide
 address-family ipv6
  multi-topology
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
 ip router isis
!
router isis
 net 49.0001.2222.2222.2222.00
 metric-style wide
```

Considerations for EIGRP

EIGRP Support for IPv6 (1)

- EIGRP architecture is build from ground up to support multiple AFs
  - The core protocol engine stays the same
  - Protocol Dependent Modules provide the adaptation services for the particular address family (IPX, IPv4, IPv6…)
  - One protocol instance (process) handles one address family
- EIGRP supports IPv6 since IOS 12.4T
- All EIGRP configuration should be done using the named mode
  - Classic configuration mode (**`router eigrp`** *`asn`* / **`ipv6 router eigrp`** *`asn`*) is obsolete and should not be used anymore

EIGRP Support for IPv6 (2)

- IPv4 and IPv6 EIGRP processes are independent
  - Properties like timers, authentication, passive interfaces, split horizon, next-hop-self, distribute / offset lists, stub are local to every process
- Despite running a separate protocol instance for IPv4 and IPv6, it is not possible to have independent EIGRP interface metrics
  - Bandwidth, delay, reliability, load – these are protocol-agnostic
  - At most, it is possible to use different K-values in different processes
- In IOS, IPv6 EIGRP automatically runs on all IPv6-enabled interfaces
  - For better control, use **`shutdown`** in **`af-interface default`**

IPv6 EIGRP Addressing and Packets

- IPv6 EIGRP messages are carried in IPv6 packets
  - Sender address: link-local address of the interface
  - Destination address: multicast ff02::a or the link-local address of the directly connected neighbor
  - Global addresses are only used for static remote neighbors
- The 7 basic message types (Hello, Update, Query, Reply, Ack, SIA-Query, SIA-Reply) stay

Will they establish an adjacency? (1)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
!
router eigrp ROCKS
 address-family ipv4 auton 1
  network 10.0.0.0
 address-family ipv6 auton 1
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
!
router eigrp RULES
 address-family ipv4 auton 1
  network 10.0.0.0
 address-family ipv6 auton 1
```

Yes – in both IPv4 and IPv6 address families

Will they establish an adjacency? (2)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
!
router eigrp ROCKS
 address-family ipv4 auton 1
  network 10.0.0.0
 address-family ipv6 auton 1
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
!
router eigrp ROCKS
 address-family ipv4 auton 1
 address-family ipv6 auton 1
```

Yes – but only in the IPv6 address family

Will they establish an adjacency? (3)


```
interface Gi1
 ip address 10.0.0.1 255.255.255.0
 ipv6 address fe80::1 link-local
!
router eigrp ROCKS
 address-family ipv4 auton 1
  network 10.0.0.0
 address-family ipv6 auton 1
```

```
interface Gi1
 ip address 10.0.0.2 255.255.255.0
 ipv6 address fe80::2 link-local
!
router eigrp ROCKSv4
 address-family ipv4 auton 1
  network 10.0.0.0
!
router eigrp ROCKSv6
 address-family ipv6 auton 1
```

Yes – in both IPv4 and IPv6 address families

“Keep-in-Mind” Facts About EIGRP for IPv6

- For IPv6 AF, the “offset-list” is not implemented
  - Instead, use a route-map modifying the metric in a distribute-list
  - Use the max bandwidth value to prevent overriding the minBW
  - Override the delay component instead

```
route-map Offset permit 10
 set metric 4294967295 655360 255 1 1500
!
router eigrp ROCKS
 address-family ipv6 autonomous-system ...
  topology base
   distribute-list route-map Offset { in | out } Gi1
```

“Keep-in-Mind” Facts About EIGRP for IPv6

- Unequal Cost Multi Path (“variance”) for IPv6 is not functional
  - Unequal paths will be installed into IPv6 FIB as equal paths
  - This is due to a lack of metric-based UCMP support in IPv6 RIB
  - The issue is tracked through CSCwi91760 as an enhancement request
- In pure IPv6 environment, manual RID configuration is required

Considerations for BGP

Multi Protocol Extensions and IPv6 in BGP

- IPv6 support in BGP stands on two pillars
  - Multi Protocol extensions (MP), brought first in RFC 2283 (February 1998)
  - Use of MP for IPv6, brought first in RFC 2545 (March 1999)
- MP extensions (MP-BGP) allow advertising multiple address families in a single peering

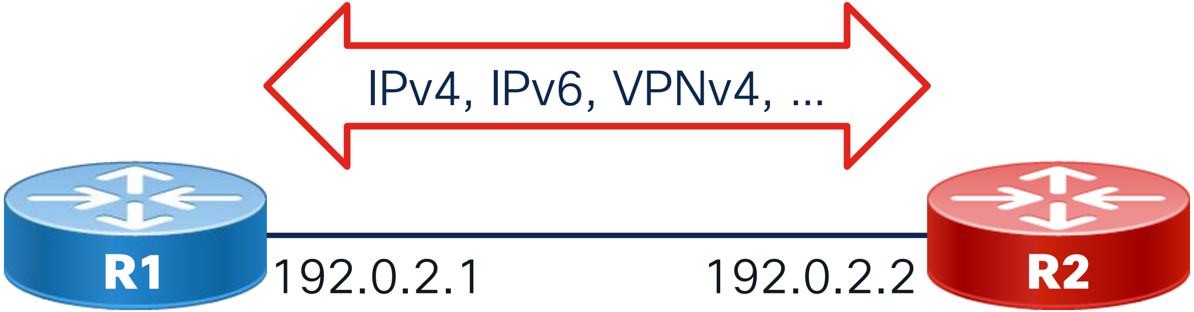

Multi Protocol (MP) BGP Considerations

- MP-BGP is only concerned with encoding of different address families and their NLRI in existing BGP messages
  - No need to be concerned with transport: BGP runs on top of TCP, and TCP already runs on top of IPv4/IPv6
- With MP-BGP, there is no longer a 1:1 relation between the neighbor's address type (IPv4 or IPv6) and the type of routes
  - Neighbor's address no longer implies the type of content to exchange
- Configuring an MP-BGP neighbor is about two distinct qualities
  - Who the neighbor is and how we talk to it (address, ASN, password, ...)
  - What content we exchange with it (IPv4 routes, IPv6 routes, ...)

MP-BGP IPv4-addressed Neighbors

```
router bgp 64512
 no bgp default ipv4-unicast
 neighbor 192.0.2.2 remote-as 64512
 neighbor 192.0.2.2 update-source Loopback0
 neighbor 192.0.2.3 remote-as 64513
 !
 address-family ipv4 unicast
  neighbor 192.0.2.2 activate
  neighbor 192.0.2.3 activate
 !
 address-family ipv6 unicast
  neighbor 192.0.2.2 activate
```

Who

What

MP-BGP IPv6-addressed Neighbors

```
router bgp 64512
 no bgp default ipv4-unicast
 neighbor 2001:db8:cafe::1 remote-as 64512
 neighbor 2001:db8:cafe::1 update-source Loopback0
 neighbor 2001:db8:f00d::2 remote-as 64513
 !
 address-family ipv4 unicast
  no neighbor 2001:db8:cafe::1 activate
  neighbor 2001:db8:f00d::2 activate
 !
 address-family ipv6 unicast
  neighbor 2001:db8:cafe::1 activate
```

Who

What

Next Hop Address Considerations with MP-BGP

- The primary consideration with MP-BGP is the next hop address
- By default, when advertising a route,
  - iBGP keeps the next hop unchanged
  - eBGP sets the next hop to the local source address of the peering
- This poses a problem when the source address of the peering session is from a different address family than the advertised route

IPv6 Next Hop Address Selection in BGP

- According to RFC 2545, BGP next hop for IPv6 routes can carry
  - Global IPv6 address only
  - Global and link local IPv6 addresses if the peers are on a common subnet

BGP Peer Configured Via...	Advertised Next Hop	Next Hop Installed in RIB
IPv6 Global Address on link (eBGP)	Global + Link Local	Link Local
IPv6 Global Address off link (eBGP)	Global	Global
IPv6 Global Address (iBGP)	Global	Global
IPv6 Link Local Address	Link Local	Link Local

Advertising IPv6 routes in IPv4-based Peerings


```
router bgp 64512
 no bgp default ipv4-unicast
 neighbor 192.0.2.2 remote-as 64513
 !
 address-family ipv6 unicast
  neighbor 192.0.2.2 activate
```

```
router bgp 64513
 no bgp default ipv4-unicast
 neighbor 192.0.2.1 remote-as 64512
 !
 address-family ipv6 unicast
  neighbor 192.0.2.1 activate
```

All IPv6 routes advertised to R2 will have their next hop set to ::ffff:192.0.2.1 and will be unreachable

All IPv6 routes advertised to R1 will have their next hop set to ::ffff:192.0.2.2 and will be unreachable

Solving the Next Hop Problem for IPv6 Routes

- There are three main approaches for proper next hop setting of IPv6 routes advertised over IPv4 peerings
  - Configure outbound route map or RPL policy to set the next hop
  - Enable the automatic IPv6 next hop address selection for IPv4 peers
  - Use multiple sessions – IPv6-based for IPv6 AF, IPv4-based for IPv4 AF
- There are pros and cons to each approach
  - Sharing a single session conserves system resources but creates a single point of failure
  - Manually setting next hops on each peering is tedious and error-prone
  - Multiple sessions are somewhat more resource hungry

Setting IPv6 Next Hops through Route Maps


```
route-map NH permit 10
 set ipv6 next-hop 2001:db8::1
!
router bgp 64512
 no bgp default ipv4-unicast
 neighbor 192.0.2.2 remote-as 64513
 !
 address-family ipv6 unicast
  neighbor 192.0.2.2 activate
  neighbor 192.0.2.2 route-map NH out
```

```
route-map NH permit 10
 set ipv6 next-hop 2001:db8::2
!
router bgp 64513
 no bgp default ipv4-unicast
 neighbor 192.0.2.1 remote-as 64512
 !
 address-family ipv6 unicast
  neighbor 192.0.2.1 activate
  neighbor 192.0.2.1 route-map NH out
```

Automatic IPv6 Next Hop Address Selection

- This feature assumes that two BGP peers communicate over IPv4 but their update source interfaces also have global IPv6 addresses
- The IPv6 next hop will then be set following these rules:
  - Use the next hop set by the outbound route-map, if any; else
  - Use the global IPv6 address of the neighbor … update-source interface, if any; else
  - Use the global and link-local IPv6 address of the local interface directly connected to the on-link peer, if any; else
  - Use the local source IPv4 address of the peering, encoded in IPv4-mapped IPv6 address format

Automatic IPv6 Next Hop Address Selection


```
router bgp 64512
 no bgp default ipv4-unicast
 no bgp default ipv6-nexthop
 neighbor 192.0.2.2 remote-as 64513
 !
 address-family ipv6 unicast
  neighbor 192.0.2.2 activate
```

```
router bgp 64513
 no bgp default ipv4-unicast
 no bgp default ipv6-nexthop
 neighbor 192.0.2.1 remote-as 64512
 !
 address-family ipv6 unicast
  neighbor 192.0.2.1 activate
```

Multiple Dedicated Sessions

- Using multiple AF-dedicated sessions is the simplest approach
  - IPv4 peer address + IPv4 AF only
  - IPv6 peer address + IPv6 AF only
- Despite mildly increased resource and configuration overhead, there are major advantages
  - Simplicity and intuitiveness of the design
  - Next hop will be handled automatically and correctly as expected
  - A hiccup in one AF won’t destabilize the other one

Multiple Sessions for IPv4 and IPv6


```
router bgp 64512
 no bgp default ipv4-unicast
 neighbor 192.0.2.2 remote-as 64513
 neighbor 2001:db8::2 remote-as 64513
 !
 address-family ipv4 unicast
  neighbor 192.0.2.2 activate
  no neighbor 2001:db8::2 activate
 !
 address-family ipv6 unicast
  neighbor 2001:db8::2 activate
```

```
router bgp 64513
 no bgp default ipv4-unicast
 neighbor 192.0.2.1 remote-as 64512
 neighbor 2001:db8::1 remote-as 64512
 !
 address-family ipv4 unicast
  neighbor 192.0.2.1 activate
  no neighbor 2001:db8::1 activate
 !
 address-family ipv6 unicast
  neighbor 2001:db8::1 activate
```

Must the Next Hop AF Match the NLRI AF?

- So far, we have assumed that routes from a particular AF must have a next hop from the same AF
- For BGP, this is long not true
  - For example, VPNv4 and VPNv6 routes use IPv4 next hops
  - The ability to decouple next hop AF from the route AF is the key to building overlay solutions over independent underlay
- RFC 8590 allows advertising IPv4 routes with an IPv6 next hop
  - IPv4 over IPv6 core, using tunneling
  - IPv4 VPN over IPv6 core, using MPLS et al.

Considerations for FHRPs

IPv6 First Hop Redundancy Protocols

- IPv6 by design offers a basic gateway redundancy
  - Router priority in RAs – High, Normal, Low
  - No rapid or seamless switchover, though
- Cisco's HSRPv2, GLBP and VRRPv3 all support IPv6
  - For VRRP, the **`fhrp version vrrp v3`** global config command must be configured before attempting to configure it on interfaces
- Group numbers may require some consideration
  - HSRPv2 and GLBP require unique group numbers for IPv4 and IPv6
  - VRRPv3 allows reusing the same group number for IPv4 and IPv6

IPv6 First Hop Redundancy Protocols

- VRRPv3, HSRPv2 and GLBP all protect the link-local address
  - VRRPv3 and HSRPv2 can also protect a non-LL address
- VRRP has a concept of "address owner"
  - The one router whose real IP matches the virtual IP
  - This router will always be the Master router
  - In IPv6, only the link-local real IP can match the virtual IP

Considerations for link-local only addressing in networks

Using link-local addresses on infrastructure links

- It is possible to set up a routed network using exclusively link-local addresses on router-to-router links
  - IGPs limit themselves to using LLAs anyway in most cases
  - For global reachability, routers can use a single loopback with a global address advertised in IGP
  - Using the same LLA on all interfaces of a router makes things very simple
- This approach is very attractive and quick to deploy
  - Similar to prefix suppression in OSPF and IS-IS
- There are, of course, pros and cons to this approach

RFC 7404

Using Only Link-Local Addressing inside an IPv6 Network

- Advantages:
  - Smaller routing tables
  - Simpler address management
  - Lower configuration complexity
  - Reduced attack surface
- Caveats:
  - Interface ping
  - Traceroute
  - Hardware dependency (MAC)
  - MPLS Traffic Engineering

Conclusion

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events mobile app
2. Click “Join the Discussion”
3. Install the Webex app or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

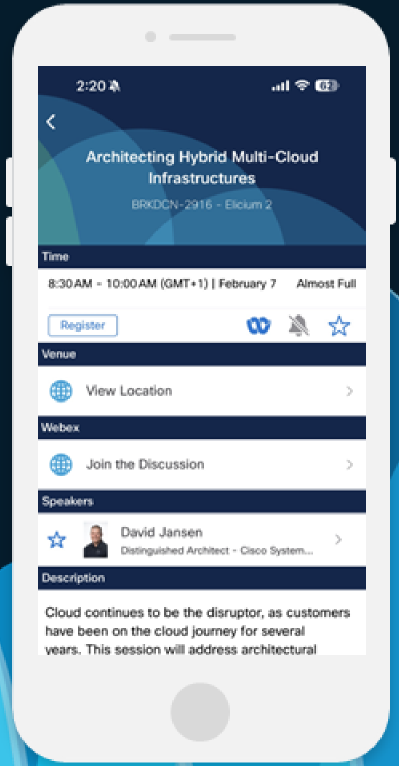

cisco Live!

Fill Out Your Session Surveys

Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)

All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the 'Participant Dashboard'

Content Catalog

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.

Contact me at: ppaluch@cisco.com

cisco Live!

Thank you

GO BEYOND