IPv6 Routing Protocols

Texas IPv6 Task Force Summit 2012

Faraz Shamim - Technical Leader
Cisco Systems Inc

Agenda

- IPv6 Routing Protocols
  - RIPng
  - EIGRPv6
  - ISISv6
  - OSPFv3
  - BGP4+
- IPv6 challenges & IGP Selection
- Co-existence & Convergence of Routing protocols

Routing in IPv6

- As in IPv4, IPv6 has 2 catagories of routing protocols: IGP and EGP, and still uses the longest-prefix match routing algorithm
- IGP
  - RIPng (RFC 2080)
  - Cisco EIGRP for IPv6
  - Integrated IS-ISv6 (RFC 5308)
  - OSPFv3 (RFC 5340)
- EGP : MP-BGP4 (RFC 2858 and RFC 2545)
- Cisco IOS supports all of them
  - Pick one that meets your objectives

Learn. Connect. Collaborate. together.

IPv6 Default & Static Routing

Default & Static Routing

- Similar to IPv4. Need to define the next hop / interface.
- Default route denoted as ::/0

*ipv6 route ipv6-prefix/**prefix-length*** **{*ipv6-address* | *interface-type interface-number* [*ipv6-address*]} [*administrative-distance*] [*administrative-multicast-distance* |** *unicast* **|** *multicast***]** **[***tag* ***tag*****]**

Examples:

- Forward packets for network 2001:DB8::0/32 through 2001:DB8:1:1::1 with an administrative distance of 10

```
Router(config)# ipv6 route 2001:DB8::0/32 2001:DB8:1:1::1 10
```

- Default route to 2001:DB8:1:1::1

```
Router(config)# ipv6 route ::/0 2001:DB8:1:1::1
```

Learn. Connect.
Collaborate. together.

RIPng for IPv6

Enhanced Routing Protocol Support
RIPng Overview RFC 2080

command	version	must be zero
Address Family Identifier		Route Tag
IPv4 Address		
Subnet Mask		
Next Hop		
Metric		

command	version	must be zero
IPv6 prefix		
route tag	prefix len	metric

- Similar characteristics as IPv4
  - Distance-vector, hop limit of 15, split-horizon, multicast based (FF02::9), UDP port (521) etc.
- Updated features for IPv6
  - IPv6 prefix & prefix len.
- Special Handling for the NH
  - Route tag and prefix len for NH is all 0. Metric will have 0xFF. NH must be link local

EIGRP for IPv6

EIGRP for IPv6 Overview

- Just another protocol module (IP, IPX, AppleTalk) with three new TLVs:
  - -IPv6_REQUEST_TYPE (0X0401)
  - -IPv6_METRIC_TYPE (0X0402)
  - -IPv6_EXTERIOR_TYPE (0X0403)
- Other similarities
  - -Same protocol number 88
  - -Router ID stays 32 bits (must be configured explicitly if there is no IPv4 interface on the router)
  - -Uses MD5 like for IPv4
(IPSec authentication will be available soon)
  - -Same metrics as EIGRP for IPv4

EIGRP for IPv6 Specific Features

Several IPv6 Specific Differences with Respect to IPv4:

- 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 neighbors where traffic is unicasted).
- Automatic summarization 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)
- By default EIGRP starts in shutdown mode & needs no shut cmd.

Learn. Connect.
Collaborate. together.

ISIS for IPv6

Integrated IS-IS for IPv6—Overview

- RFC 5308
- Two TLVs added to support IPv6:

  -IPv6 Reachability TLV (0xEC)—Describes network reachability (IPv6 routing prefix, metric information and option bits). The option bits indicate the advertisement of IPv6 prefix from a higher level, redistribution from other routing protocols. Equivalent to IP Internal/External Reachability TLVs described in RFC1195.

  -IPv6 Interface Address TLV (0xE8)—Contains 128-bit address. Hello PDUs, must contain the link-local address but for LSP, must only contain the non-link-local address.

- A new Network Layer Protocol Identifier (NLPID)—Allows IS-IS routers with IPv6 support to advertise IPv6 prefix payload using 0x8E value (IPv4 and OSI uses different values)

Single Topology IS-IS

- Single topology (default for all protocols supported). Potentially beneficial in saving resources (same topology and same SPF):
  - All routers must support the same address families (dual-stack, topologically congruent network). Adjacency checking should be disabled during migration.
  - Interface metrics apply to both IPv4 and IPv6

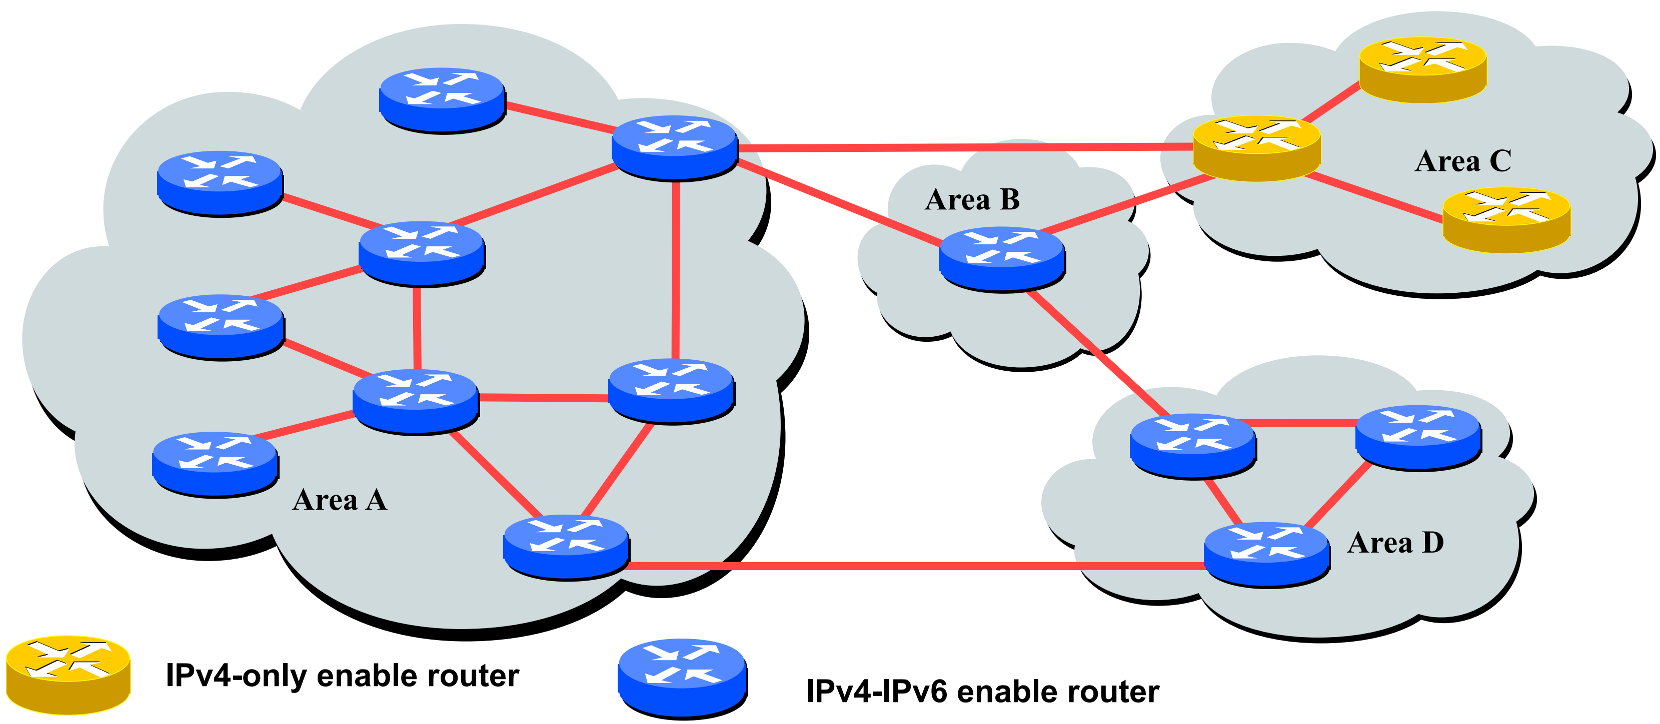

Multi-Topology IS-IS

- Multi-topology (RFC 5120)
  - Independent IPv4 and IPv6 topologies
  - Independent interface metrics
- Transition mode available—both types of TLVs are advertised

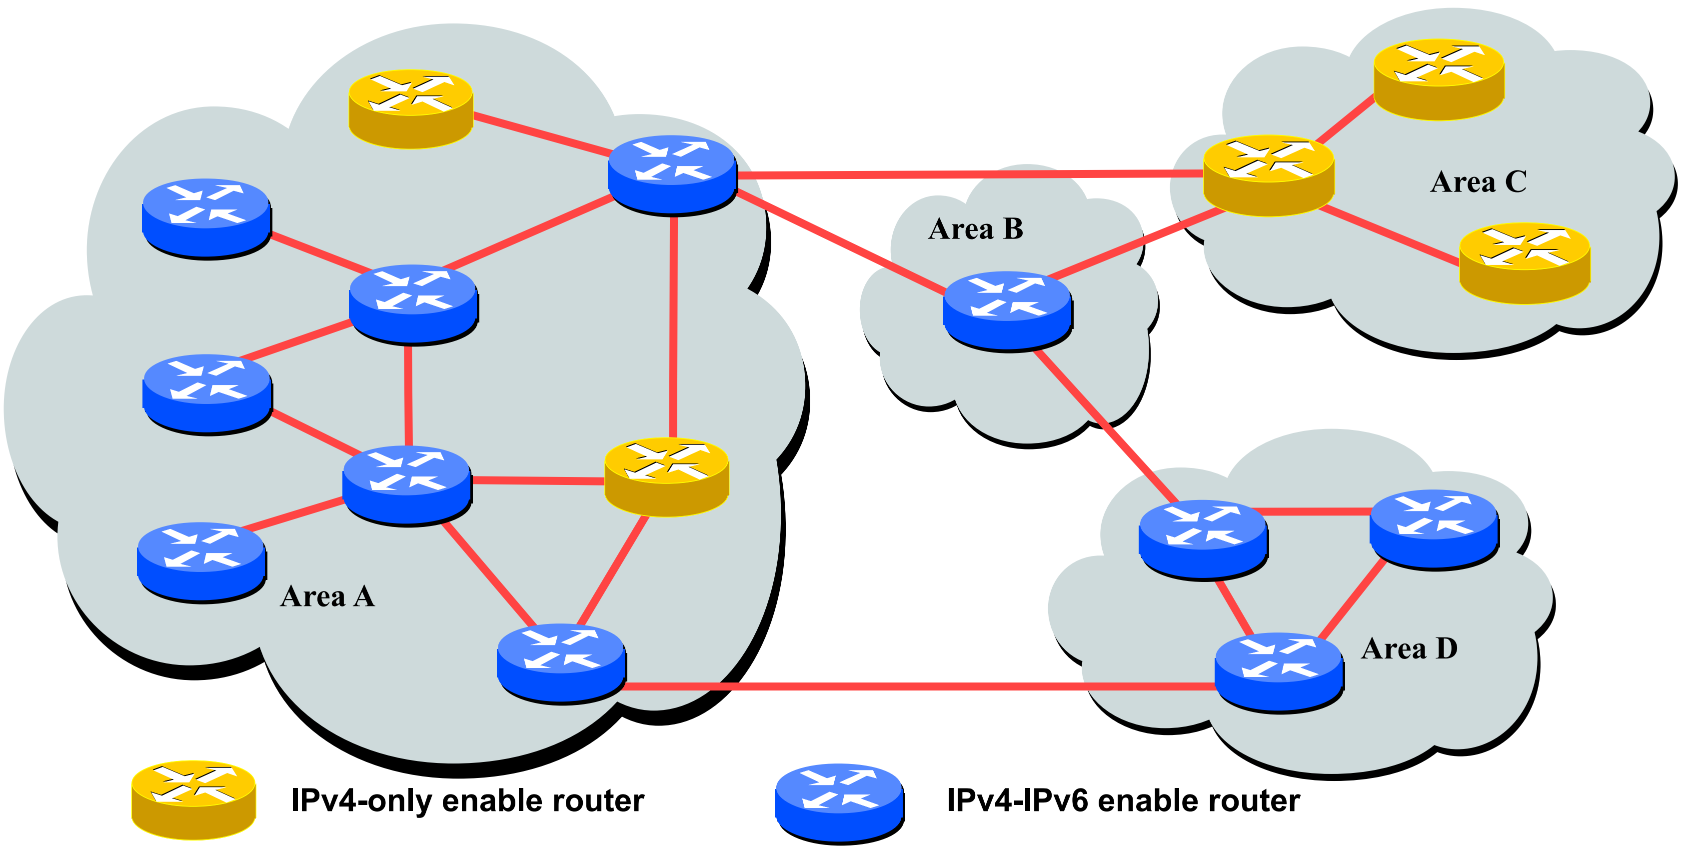

OSPFv3 – RFC 5340

Similarities with OSPFv2

- OSPFv3 is based on OSPFv2:
  - -Runs directly over IPv6 (port 89)
  - -Uses the same basic packet types
  - -Neighbor discovery and adjacency formation mechanisms are identical (all OSPF routers FF02::5, all OSPF DRs FF02::6)
  - -LSA flooding and aging mechanisms are identical
  - -Same interface types (P2P, P2MP, broadcast, NBMA, virtual)
- OSPFv3 and OSPFv2 are independent processes and run as ships in the night

V2, V3 Differences

OSPFv3 Is Running per Link Instead of per Node (and IP Subnet)

- A link by definition is a medium over which two nodes can communicate at link layer
- Regardless of assigned prefixes, two devices can communicate using link-local addresses therefore OSPFv3 is running per link instead of per IP prefix
- Multiple IPv6 prefixes can be assigned to the same link

V2, V3 Differences (Cont.)

Authentication Is Removed from OSPF

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

Version	Type	Packet Length	
Router ID			
Area ID			
Checksum		Autype	
Authentication			
Authentication			

Version	Type	Packet Length	
Router ID			
Area ID			
Checksum		Instance ID	0

V2, V3 Differences (Cont.)

Support of Multiple Instances per Link

- New field (instance) in OSPF packet header allows running multiple instances per link
- Instance ID should match before packet is being accepted
- Useful for traffic separation, multiple areas per link

Version	Type	Packet Length	
Router ID			
Area ID			
Checksum		Instance ID	0

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

Network Mask		
HelloInterval	Options	Rtr Pri
RouterDeadInterval		
Designated Router		
Backup Designated Router		
Neighbor ID		

Interface ID	
Rtr Pri	Options
HelloInterval	RouterDeadInterval
Designated Router	
Backup Designated Router	
Neighbor ID	

V2, V3 Differences (Cont.)

Address Semantic Changes 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 prefixes are carried in a new LSA payload called intra-area-prefix-LSAs
- Prefixes are carried in the payload of inter-area and external LSA

V2, V3 Differences (Cont.)

Generalization of Flooding Scope

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

LS age	U	S2	S1	LSA Function Code

S2	S1	Flooding scope
0	0	Link-Local flooding scope
0	1	Area flodding scope
1	0	AS flooding scope
1	1	Reserved

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 within the corresponding flooding scope, as if it was understood
- When U bit is not set, the LSA is flooded within the link local scope
- In v2 unknown LSA were discarded

U-bit	LSA Handling
0	Treat this LSA as if it has link-local Scope
1	Store and flood this LSA as if type understood

V2, V3 Differences (Cont.)

Two New LSAs Have Been Introduced

- Link-LSA has a link local flooding scope and has three purposes

  -Carry IPv6 link local address used for NH calculation

  -Advertise IPv6 global address to other routers on the link (used for multi-access link)

  -Convey router options to DR on the link

- Intra-area-prefix-LSA to advertise router's IPv6 address within the area

LSA Types

	LSA Function Code	LSA Type
Router-LSA	1	Ox2001
Network-LSA	2	Ox2002
Inter-Area-Prefix-LXA	3	Ox2003
Inter-Area-Router-LSA	4	Ox2004
AS-External-LSA	5	Ox4005
Group-Membership-LSA	6	Ox2006
Type-7-LSA	7	Ox2007
Link-LSA	8	Ox0008
Intra-Area-Prefix-LSA	9	Ox2009

New

Learn. Connect.
Collaborate. together.

BGP for IPv6

BGP-4 Extensions for IPv6

BGP-4 Carries Only 3 Pieces of Information Which Are Truly IPv4 Specific:

- NLRI in the UPDATE message contains an IPv4 prefix
- NEXT_HOP path attribute in the UPDATE message contains an IPv4 address
- BGP Identifier is in the OPEN message and AGGREGATOR attribute

BGP-4 Extensions for IPv6

To Make BGP-4 Available for Other Network Layer Protocols, RFC 2858 (Obsoletes RFC 2283) Defines Multiprotocol Extensions for BGP-4:

- Enables BGP-4 to carry information of other protocols (MPLS, IPv6, etc.)
- New BGP-4 optional and non-transitive attributes
  - -MP_REACH_NLRI
  - -MP_UNREACH_NLRI
- Protocol independent NEXT_HOP attribute
- Protocol independent NLRI attribute

BGP-4 Extensions for IPv6

- 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 14 and 15 contains one or more triples:

    Address Family Information (AFI)

    Next-Hop Information
    (must be of the same address family)

    NLRI

BGP-4 Extensions for IPv6

Address Family Information (AFI) for IPv6

- AFI = 2 (RFC 1700)
- Sub-AFI = 1 unicast
- Sub-AFI = 2 (multicast for RPF check)
- Sub-AFI = 3 for both unicast and multicast
- Sub-AFI = 4 label
- Sub-AFI= 128 VPN

BGP-4 Extensions for IPv6

- TCP Interaction
  - -BGP-4 runs on top of TCP
  - -This connection could be setup either over IPv4 or IPv6
- Router ID
  - -When no IPv4 is configured, an explicit bgp router-id needs to be configured
  - -This is needed as a BGP Identifier, this is used as a tie breaker, and is sent within the OPEN message

BGP-4 Extensions for IPv6

- Next-hop contains a global IPv6 address or potentially a link local (for iBGP update this has to be changed to global IPv6 address with route-map in old codes)
- The value of the length of the next hop field on MP_REACH_NLRI attribute is set to 16 when only global is present and is set to 32 if link local is present as well
- Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)

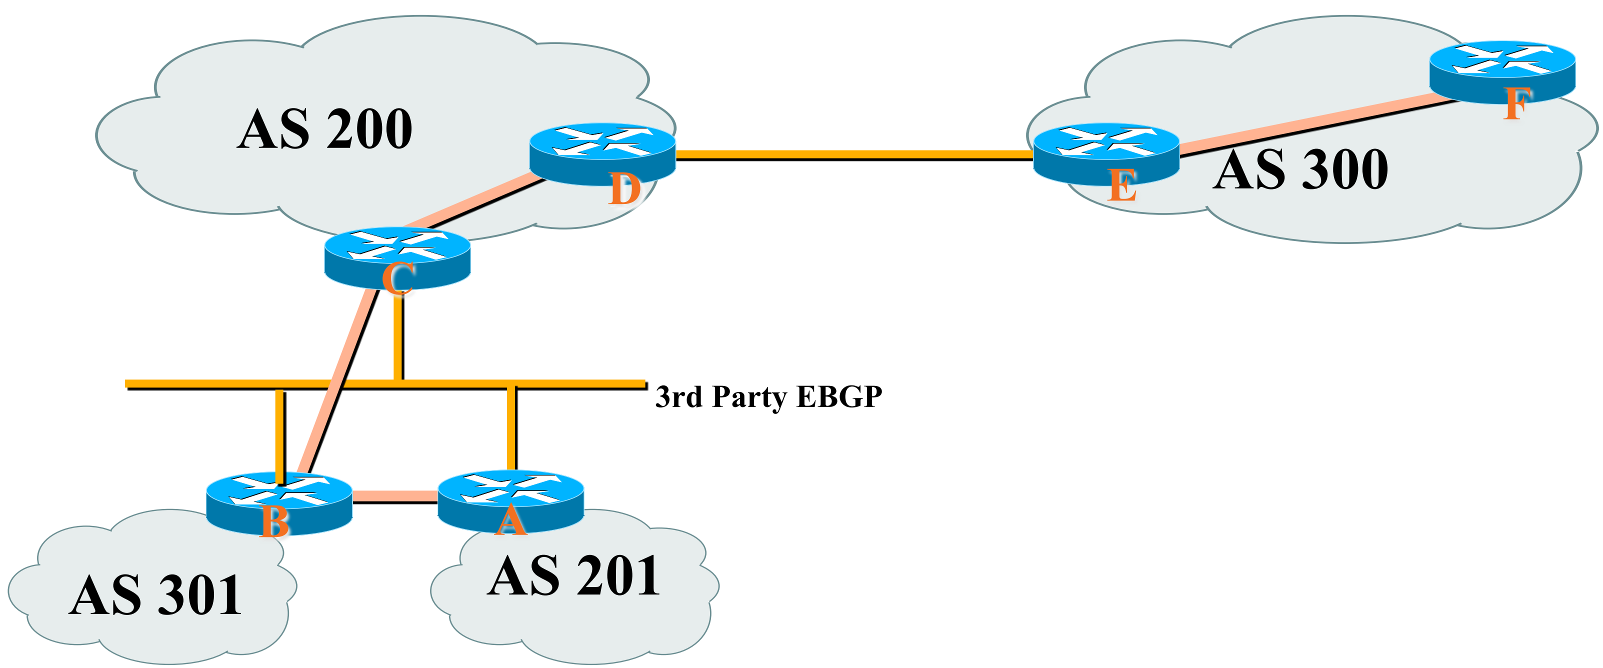

IPv6 Challenges & IGP Selection

IPv6 Challenges to Router Performance

Addressing Driven

- Forwarding challenges—lookup not impacted as much as originally thought, different size prefixes typically see little difference in forwarding performance
- Control plane challenges—routing table sizes:
  - -IPv6 supports multiple addresses per interface (not the most significant concern at this time but it could be in the future)
  - -IPv6 can have a lot more prefixes due to a significantly larger address space

The Questions Are the Same as for IPv4… Almost

- Is one routing protocol better than any other routing protocol?
- Define “Better” →

- Converges faster?
- Uses less resources?
- Easier to troubleshoot?
- Easier to configure?
- Scales to a larger number of routers, routes, or neighbors?
- More flexible?
- Degrades more gracefully?
- And so on

IPv6 IGP Selection—In Theory

In Theory:

- The similarity between the IPv6 and IPv4 routing protocols leads to similar behavior and expectations
- To select the IPv6 IGP, start by using the IPv4 IGP rules of thumb

IPv6 IGP Selection—In Practice

- In practice:
  - -The IPv6 IGP implementations might not be fully optimized yet so there is a bit more uncertainty
  - -Not all knobs for Fast Convergence might be available
  - -The average convergence time is 100% larger than IPv4, as IPv6 converges after IPv4
  - -No significant operational experience with large scale IPv6 networks

Routing Protocols Coexistence & Convergence

The Questions Are **Almost** the Same as for IPv4

- Most likely the IPv6 IGP will not be deployed in a brand new network and just by itself
- Most likely the IPv4 services are more important at first since they are generating most of the revenue
- Redefine "better"

- What is the impact on the convergence of IPv4?
- Are the resources optimally shared?
- Are the topologies going to be congruent?
- Etc.

Co-existence—Convergence Considerations

At first, the IPv6 IGP Convergence might be less important than the impact of IPv6 on the Convergence of the existent IPv4 infrastructure

- What IGPs coexist better?
- What IPv6 IGP impacts IPv4 the least (hopefully not at all)?

Nothing is for Free

- Resources **will be shared** between the two IGPs and they **will compete** for processor cycles in a way that reflects their relative configuration
- This has implications on:
  - Expected convergence behavior
  - Single process/topology vs Multi process/topology selection
  - Resources (Memory, CPU) planning

Coexistence—Resources Considerations

- With the exception of ISIS single topology, the IPv4 and IPv6 routing processes claim their own memory and processing resources for maintaining adjacencies, databases and related calculations
- It is important to define the IPv6 network design in order to understand the new resource requirements (memory) and the new operational parameters (max CPU) for the network devices

Coexistence—Topology Considerations

- The IPv4 and IPv6 topologies can be:
  - -Congruent
    - Dual-stack deployment
  - -Non-Congruent
    - Not all network devices are supporting the necessary IPv6 features so they must be avoided during migration
- Non-congruent is not necessarily bad, even though it might be more difficult to manage and troubleshoot. Strive for congruent topologies.

Convergence Considerations

The IGPs Will Compete over Processor Cycles Based on Their Relative Tuning

- If you configure the IPv4 and IPv6 IGPs the same way (aggressively tuned for fast convergence), naturally expect a doubling of their stand alone operation convergence time
- If the IPv6 IGP is operating under default settings, the convergence time for the optimally tuned IPv4 IGP is not significantly affected

Summary

- In summary we learned:
- IPv6 Routing Protocols
- IPv6 challenges & IGP Selection
- Co-existence & Convergence of Routing protocols