



IPv6 Addressing & Routing Protocols

ISP/IXP Workshops

Topics

- Addressing plans for IPv6
- Configuring IPv6 in IOS
- IPv6 Routing Protocols



Addressing

Where to get IPv6 addresses

- Your upstream ISP
- Africa
 - AfriNIC – <http://www.afrinic.net>
- Asia and the Pacific
 - APNIC – <http://www.apnic.net>
- North America
 - ARIN – <http://www.arin.net>
- Latin America and the Caribbean
 - LACNIC – <http://www.lacnic.net>
- Europe and Middle East
 - RIPE NCC – <http://www.ripe.net/info/ncc>

Internet Registry Regions



Getting IPv6 address space

- Become a member of your Regional Internet Registry and get your own allocation

Require a plan for a year ahead

General allocation policies are outlined in RFC2050, more specific details for IPv6 are on the individual RIR website

or

- Take part of upstream ISP's PA space

or

- Use 6to4

- There is **plenty** of IPv6 address space

The RIRs require high quality documentation

Getting IPv6 address space

- From the RIR

 - Receive a /32 (or larger if you have more than 65k /48 assignments)

- From your upstream ISP

 - Get one /48 from your upstream ISP

 - More than one /48 if you have more than 65k subnets

- Use 6to4

 - Take a single public IPv4 /32 address

 - 2002:<ipv4 /32 address>::/48 becomes your IPv6 address block, giving 65k subnets

 - Requires a 6to4 gateway

Addressing Plans – ISP Infrastructure

- ISPs should receive /32 from their RIR
- Address block for router loop-back interfaces
 - Generally number all loopbacks out of **one** /64
 - /128 per loopback
- Address block for infrastructure
 - /48 allows 65k subnets
 - /48 per PoP or region (for large networks)
 - /48 for whole backbone (for small to medium networks)
 - Summarise between sites if it makes sense

Addressing Plans – ISP Infrastructure

- What about LANs?

 - /64 per LAN

- What about Point-to-Point links?

 - Expectation is that /64 is used

 - People have used /126s

 - Mobile IPv6 Home Agent discovery won't work

 - People have used /112s

 - Leaves final 16 bits free for node IDs

 - See RFC3627 for more discussion

Addressing Plans – Customer

- Customers get **one** /48

Unless they have more than 65k subnets in which case they get a second /48 (and so on)

(Still on going RIR policy discussion about giving “small” customers a /56 and single LAN end-sites a /64)

- Should not be reserved or assigned on a per PoP basis


ISP iBGP carries customer nets

Aggregation within the iBGP not required and usually not desirable

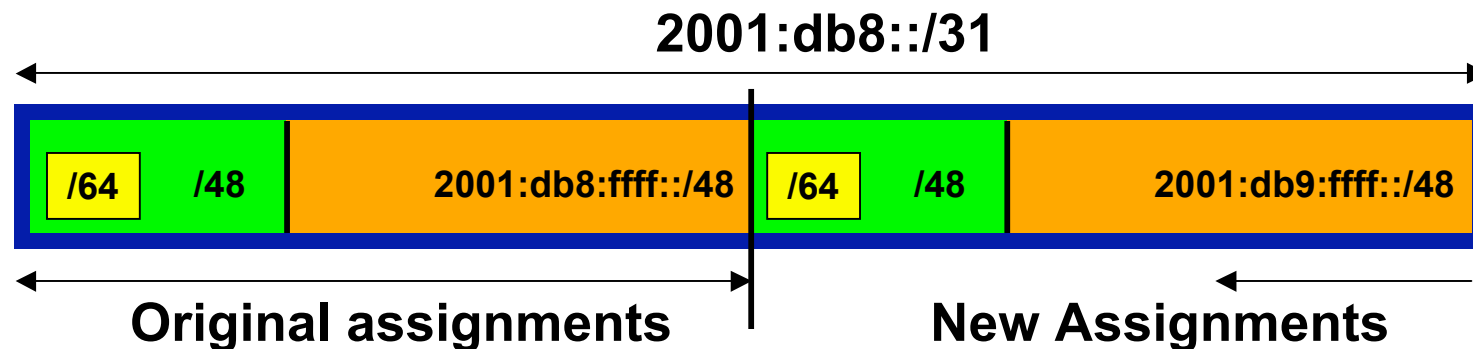
Aggregation in eBGP is very necessary

Addressing Plans – ISP Infrastructure

- Phase One



- Phase Two – second /32



Addressing Plans Planning

- Registries will usually allocate the next block to be contiguous with the first allocation
 - Minimum allocation is /32
 - Very likely that subsequent allocation will make this up to a /31
 - So plan accordingly

Addressing Plans (contd)

- Document infrastructure allocation
 - Eases operation, debugging and management
- Document customer allocation
 - Customers get /48 each
 - Prefix contained in iBGP
 - Eases operation, debugging and management
 - Submit network object to RIR Database



Initial IPv6 Configuration for IOS

IPv6 Configuration

- 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

```
rtr# conf t
rtr(config)# ipv6 unicast-routing
rtr(config)# ipv6 cef
rtr(config)# int eth 0/0
rtr(config-int)# ipv6 enable
rtr(config-int)# ^Z
```

```
rtr#sh ipv6 interface eth 0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:1E00
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1E00
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
```

IPv6 Configuration – EUI64

```
rtr#sh ipv6 interface eth 0/0
Ethernet0/0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:1E00
Global unicast address(es):
  2001:DB8::A8BB:CCFF:FE00:1E00, subnet is 2001:DB8::/64 [EUI]
Joined group address(es):
  FF02::1
  FF02::2
  FF02::1:FF00:1E00
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

```
rtr#sh ipv6 int fast 0/1
FastEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::21D:70FF:FEBD:4CA1
  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:FFBD:4CA1
  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:

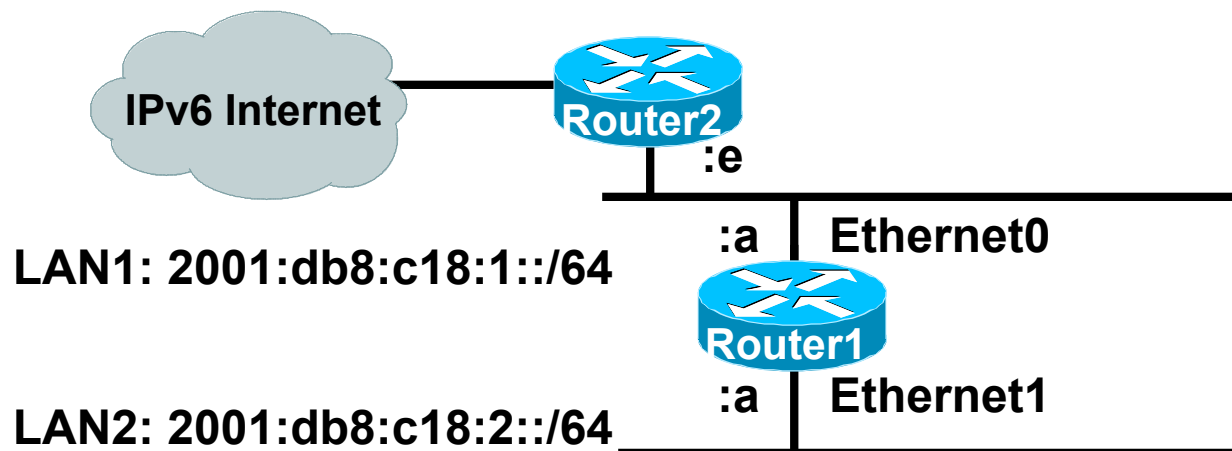
```
ipv6 route ipv6-prefix/prefix-length {ipv6-address | interface-type  
interface-number} [administrative-distance]
```

- Static Route

```
ipv6 route 2001:db8::/64 2001:db8:0:CC00::1 150
```

Routes packets for network 2001:db8::/64 to a networking device at 2001:db8:0:CC00::1 with an administrative distance of 150

Default Routing Example



```
ipv6 unicast-routing
!  
interface Ethernet0  
  ipv6 address 2001:db8:c18:1::a/64  
!  
interface Ethernet1  
  ipv6 address 2001:db8:c18:2::a/64  
!  
ipv6 route ::/0 2001:db8:c18:1::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 IOS Release 12.4 onwards

EIGRP for IPv6

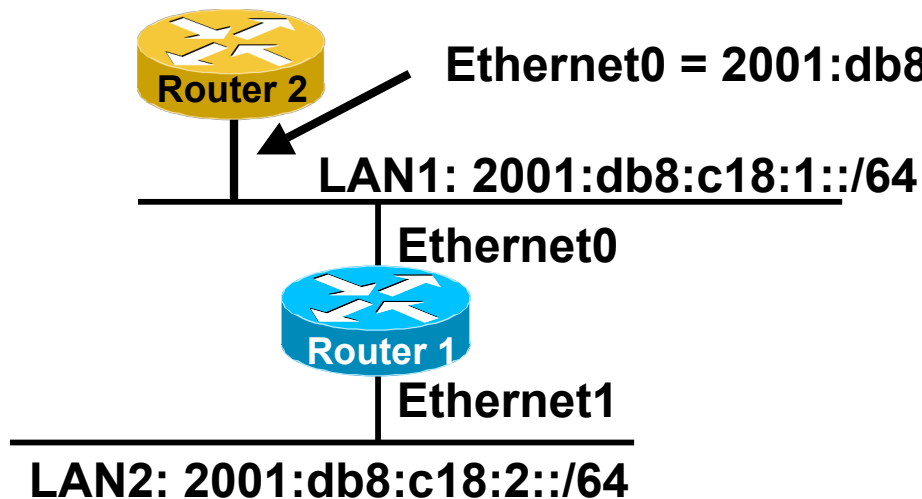
- Some differences:

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 10.10.10.1

interface Ethernet0
 ipv6 address 2001:db8:c18:1::/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::260:3eff:fe47:1530 E0 14 00:01:43 1 4500 0 1
```

```
Router1#show ipv6 eigrp topology all links
IPv6-EIGRP Topology Table for AS(100)/ID(10.10.10.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status
P 2001:db8:c18:1::/64, 1 successors, FD is 28160, serno 1
via Connected, Ethernet0
via FE80::260:3eff:fe47:1530 (30720/28160), Ethernet0
```

Neighbour Identified by
Link-Local Address

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 Ethernet0  
  ipv6 address 2001:db8:1:1::1/64  
  ipv6 ospf 1 area 0  
  
interface Ethernet1  
  ipv6 address 2001:db8:2:2::2/64  
  ipv6 ospf 1 area 1  
  
ipv6 router ospf 1  
  router-id 1.1.1.1
```




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

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 IP 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::45c/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::45a/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::45c/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::45a/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
```

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)

A Simple MP-BGP Session



```
Router1#  
interface Ethernet0  
  ipv6 address 2001:db8:c18:2:1::F/64  
!  
router bgp 65001  
  bgp router-id 10.10.10.1  
  no bgp default ipv4-unicast  
  neighbor 2001:db8:c18:2:1::1 remote-as 65002  
  address-family ipv6  
    neighbor 2001:db8:c18:2:1::1 activate  
    neighbor 2001:db8:c18:2:1::1 prefix-list bgp65002in in  
    neighbor 2001:db8:c18:2:1::1 prefix-list bgp65002out out  
  exit-address-family
```

Routing Protocols for IPv6 Summary

- Support for IPv6 in the major routing protocols
- More details for OSPF, ISIS and BGP in following presentations



IPv6 Addressing & Routing Protocols

ISP/IXP Workshops