IPv6 Essentials

Philip Smith

SANOG 23

16 January 2014

Thimphu

Last updated 13 January 2014

Presentation Slides

- Will be available on
  - http://thyme.apnic.net/ftp/seminars/SANOG23-IPv6-Essentials.pdf
  - And on the SANOG 23 website
- Feel free to ask questions any time

Agenda

- IPv6 Background
- IPv6 Protocol
- IPv6 Addressing

Early Internet History

- Late 1980s
  - Exponential growth of the Internet
- Late 1990: CLNS proposed as IP replacement
- 1991-1992
  - Running out of “class-B” network numbers
  - Explosive growth of the “default-free” routing table
  - Eventual exhaustion of 32-bit address space
- Two efforts – short-term vs. long-term
  - More at “The Long and Windy ROAD”

    http://rms46.vlsm.org/1/42.html

Early Internet History

- CIDR and Supernetting proposed in 1992-3
  - Deployment started in 1994
- IETF “ipng” solicitation – RFC1550, Dec 1993
- Proliferation of proposals:
  - TUBA – RFC1347, June 1992
  - PIP – RFC1621, RFC1622, May 1994
  - CATNIP – RFC1707, October 1994
  - SIPP – RFC1710, October 1994
  - NIMROD – RFC1753, December 1994
  - ENCAPS – RFC1955, June 1996
- Direction and technical criteria for ipng choice
  - RFC1752, January 1995

Early Internet History
→ 1996

- IPv6 Specification (RFC1883) published in December 1995
- Other activities included:
    - Development of NAT, PPP, DHCP,…
    - Some IPv4 address reclamation
    - The RIR system was introduced
- → Brakes were put on IPv4 address consumption
- IPv4 32 bit address = 4 billion hosts
    - HD Ratio (RFC3194) realistically limits IPv4 to 250 million hosts

Recent Internet History
The "boom" years → 2001

- IPv6 Development in full swing
  - Rapid IPv4 consumption
  - IPv6 specifications sorted out
  - (Many) Transition mechanisms developed
- 6bone
  - Experimental IPv6 backbone sitting on top of Internet
  - Participants from over 100 countries
- Early adopters
  - Japan, Germany, France, UK,…

Recent Internet History
The "bust" years: 2001 → 2004

- The DotCom "crash"
  - i.e. Internet became mainstream
- IPv4:
  - Consumption slowed
  - Address space pressure "reduced"
- Indifference
  - Early adopters surging onwards
  - Sceptics more sceptical
  - Yet more transition mechanisms developed

2004 → 2011

- Resurgence in demand for IPv4 address space
  - All IPv4 address space was allocated by IANA by 3rd February 2011
  - Exhaustion predictions did range from wild to conservative
  - …but by early 2011 IANA had no more!
  - …and what about the market for address space?
- Market for IPv4 addresses:
  - Creates barrier to entry
  - Condemns the less affluent to tyranny of NATs
- IPv6 offers vast address space
  - The only compelling reason for IPv6

Do we really need a larger address space?

- Internet population
  - ~630 million users end of 2002 – 10% of world pop.
  - ~1320 million users end of 2007 – 20% of world pop.
  - Doubles every 5 years (approximately)
  - Future? (World pop. ~9B in 2050)
- US uses 93.7 /8s – this is 6.4 IPv4 addresses per person
  - Repeat this the world over…
  - 6 billion population could require 26 billion IPv4 addresses
  - (7 times larger than the IPv4 address pool)

Do we really need a larger address space?

- Other Internet Economies:

China	19.7 IPv4 /8s
Japan	12.0 IPv4 /8s
UK	7.3 IPv4 /8s
Germany	7.1 IPv4 /8s
Korea	6.7 IPv4 /8s

  - Source: http://bgp.potaroo.net/iso3166/v4cc.html
- Emerging Internet economies need address space:
  - China would need more than a /4 of IPv4 address space if every student (320M) is to get an IPv4 address
  - India lives behind NATs (using only 2.1 /8s)
  - Africa lives behind NATs (using less than 1.5 /8s)

Do we really need a larger address space?

- Mobile Internet introduces new generation of Internet devices
  - PDA (~20M in 2004), Mobile Phones (~1.5B in 2003), Tablet PC
  - Enable through several technologies, eg: 3G, 802.11,…
- Transportation – Mobile Networks
  - 1B automobiles forecast for 2008 – Begin now on vertical markets
  - Internet access on planes, e.g. Connexion by Boeing
  - Internet access on trains, e.g. Narita Express
- Consumer, Home and Industrial Appliances

Do we really need a larger address space?

- RFC 1918 is not sufficient for large environments
  - Cable Operators (e.g. Comcast – NANOG37 presentation)
  - Mobile providers (fixed/mobile convergence)
  - Large enterprises
- The Policy Development process of the RIRs turned down a request to increase private address space
  - RIR community guideline is to use global addresses instead
  - This leads to an accelerated depletion of the global address space
- Some wanted 240/4 as new private address space
  - But how to back fit onto all TCP/IP stacks released since 1995?

Status in Internet Operational Community

- Service Providers get an IPv6 prefix from their regional Internet Registries
    - Very straight forward process when compared with IPv4
- Much discussion amongst operators about transition:
    - NOG experiments of 2008
        - http://www.civil-tongue.net/6and4/
    - What is really still missing from IPv6
        - http://www.nanog.org/meetings/nanog41/presentations/Bush-v6-op-reality.pdf
    - Many presentations on IPv6 deployment experiences

Service Provider Status

- Many transit ISPs have “quietly” made their backbones IPv6 capable as part of infrastructure upgrades
  - Native is common (dual stack)
  - Providers using MPLS use 6PE/6VPE
  - Tunnels still used (unfortunately)
- Today finding IPv6 transit is not as challenging as it was 5 years ago

OS, Services, Applications, Content

- Operating Systems
  - MacOS X, Linux, BSD Family, many SYS V
  - Windows: XP SP2 (hidden away), Vista, 7
  - All use IPv6 first if available
    - (MacOS 10.7 has “happy eyeballs”)
- Applications
  - Browsers
    - Firefox has “happy eyeballs”
  - E-mail clients, IM, bittorrent,…
- Services
  - DNS, Apache WebServer, E-mail gateways,…
- Content Availability
  - Needs to be on IPv4 and on IPv6

Why not use Network Address Translation?

- Private address space and Network address translation (NAT) could be used instead of IPv6
- But NAT has many serious issues:
  - Breaks the end-to-end model of IP
  - Breaks end-to-end network security
  - Serious consequences for Lawful Intercept
  - Non-NAT friendly applications means NAT has to be upgraded
  - Some applications don’t work through NATs
  - Layered NAT devices
  - Mandates that the network keeps the state of the connections
  - How to scale NAT performance for large networks??
  - Makes fast rerouting and multihoming difficult
  - How to offer content from behind a NAT?

IPv4 run-out

- Policy Development process in each RIR region has discussed and implemented many proposals relating to IPv4 run-out, for example:
  - The Last /8
    - All RIRs will receive one /8 from the IANA free pool
  - IPv4 address transfer
    - Permits LIRs to transfer address space to each other rather than returning to their RIR
  - Soft landing
    - Reduce the allocation sizes for an LIR as IPv4 pool is depleted
  - IPv4 distribution for IPv6 transition
    - Reserving a range of IPv4 address to assist with IPv6 transition (for Large Scale NATs etc)

Conclusion

- There is a need for a larger address space
    - IPv6 offers this – will eventually replace NAT
    - But NAT will be around for a while too
    - Market for IPv4 addresses looming also
- Many challenges ahead

The IPv6 Protocol

A more detailed look at IPv6 itself

So what has really changed?

- Expanded address space
  - Address length quadrupled to 16 bytes
- Header Format Simplification
  - Fixed length, optional headers are daisy-chained
  - IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- No checksum at the IP network layer
- No hop-by-hop fragmentation
  - Path MTU discovery
- 64 bits aligned
- Authentication and Privacy Capabilities
  - IPsec is mandated
- No more broadcast

IPv4 and IPv6 Header Comparison

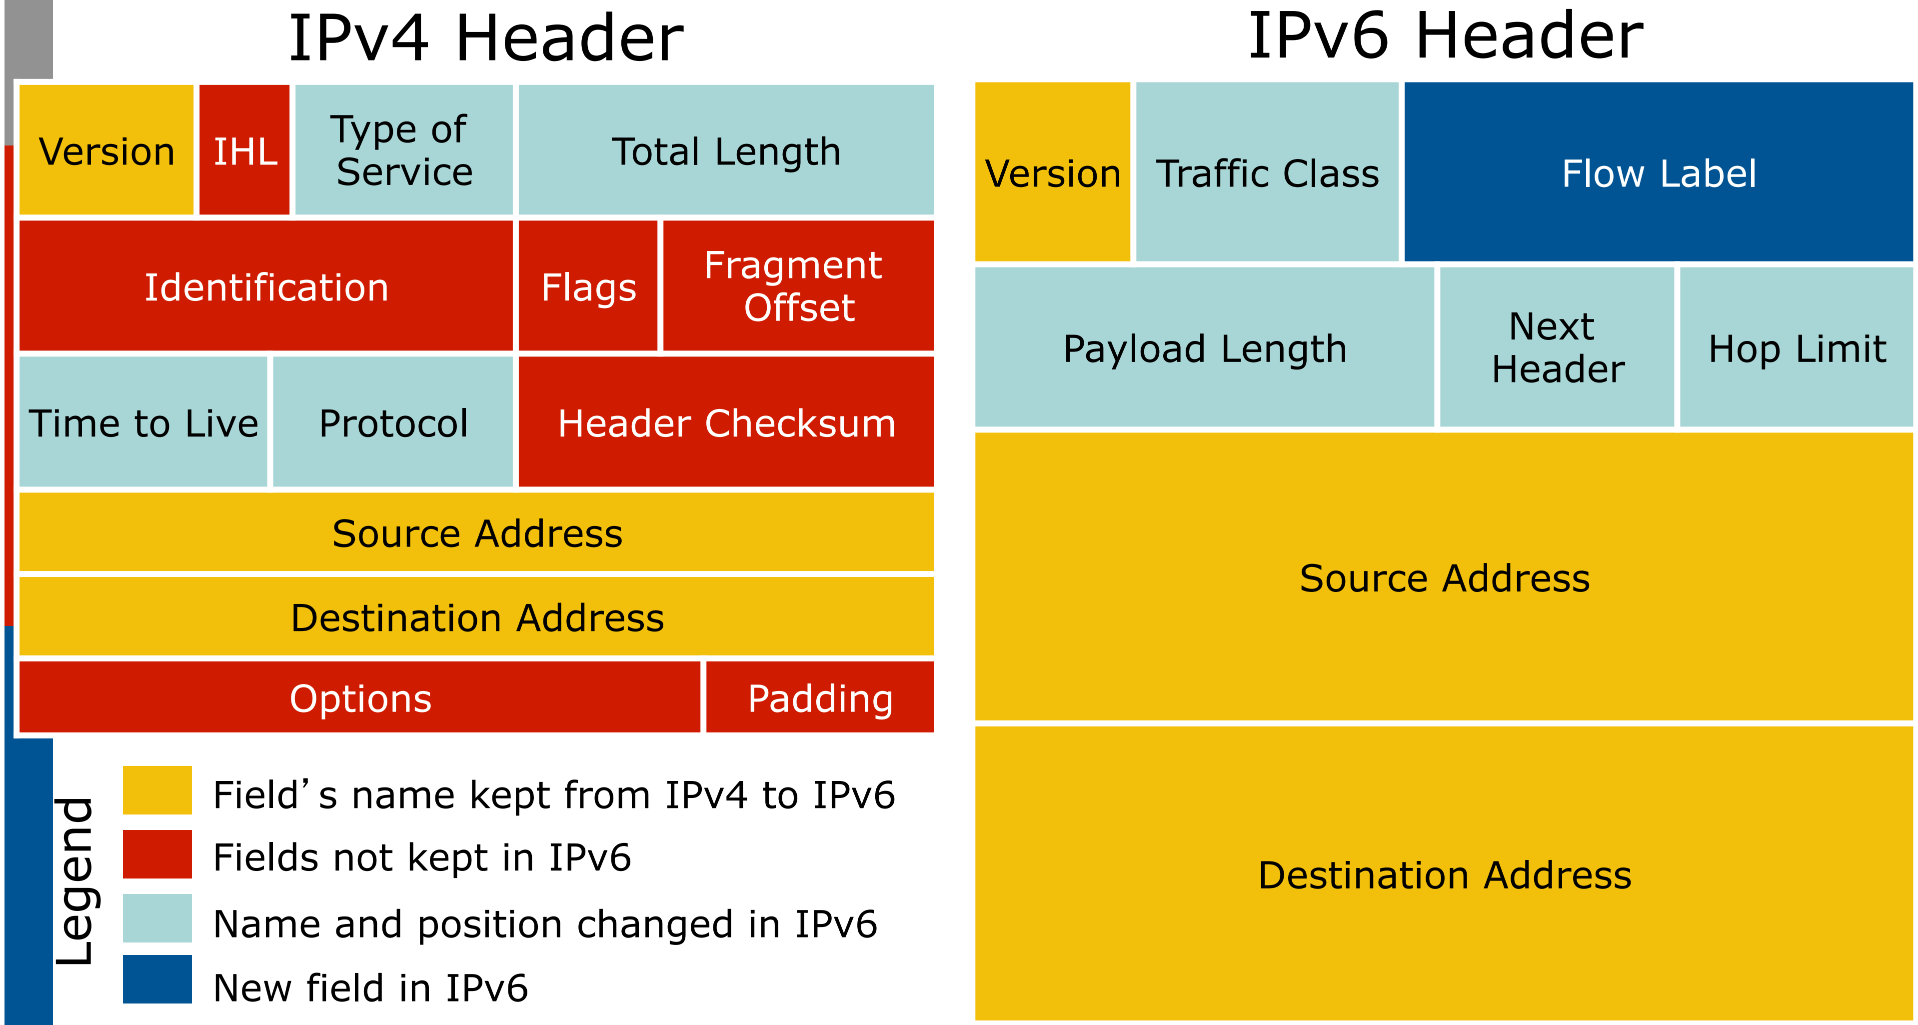

IPv4 Header

Version
IHL
Type of Service
Total Length
Identification
Flags
Fragment Offset
Time to Live
Protocol
Header Checksum
Source Address
Destination Address
Options
Padding

IPv6 Header

Version
Traffic Class
Flow Label
Payload Length
Next Header
Hop Limit
Source Address
Destination Address

Legend

🟨 Field's name kept from IPv4 to IPv6
🟥 Fields not kept in IPv6
▨ Name and position changed in IPv6
🟦 New field in IPv6

IPv6 Header

- Version = 4-bit value set to 6
- Traffic Class = 8-bit value
  - Replaces IPv4 TOS field
- Flow Label = 20-bit value
- Payload Length = 16-bit value
  - The size of the rest of the IPv6 packet following the header – replaces IPv4 Total Length
- Next Header = 8-bit value
  - Replaces IPv4 Protocol, and indicates type of next header
- Hop Limit = 8-bit value
  - Decreased by one every IPv6 hop (IPv4 TTL counter)
- Source address = 128-bit value
- Destination address = 128-bit value

Header Format Simplification

- Fixed length
  - Optional headers are daisy-chained
- 64 bits aligned
- IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- IPv4 contains 10 basic header fields
- IPv6 contains 6 basic header fields
  - No checksum at the IP network layer
  - No hop-by-hop fragmentation

Header Format – Extension Headers

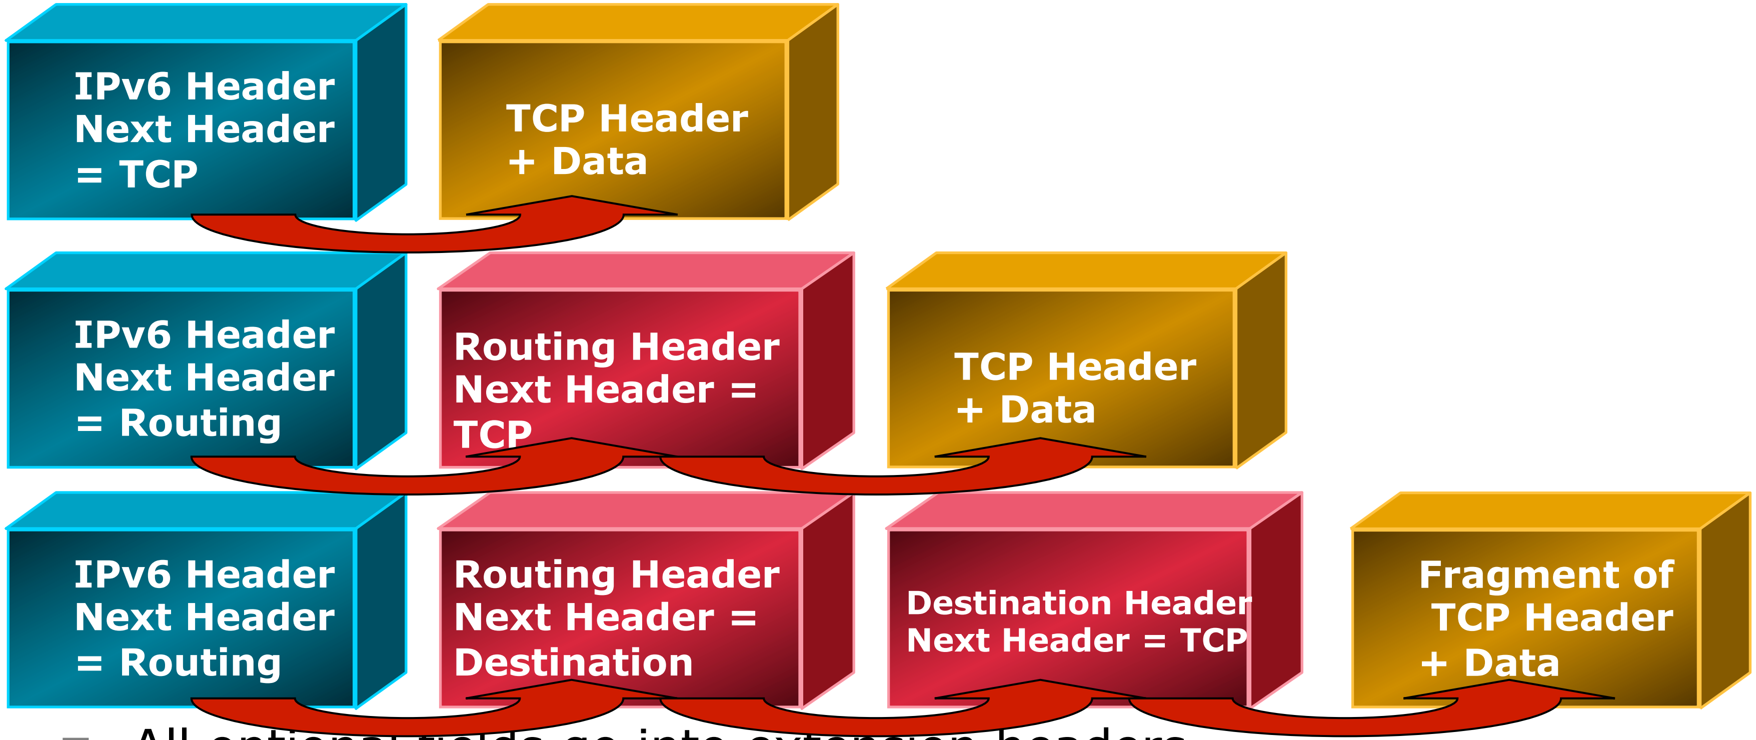

IPv6 Header
Next Header
= TCP

TCP Header
+ Data

IPv6 Header
Next Header
= Routing

Routing Header
Next Header =
TCP

TCP Header
+ Data

IPv6 Header
Next Header
= Routing

Routing Header
Next Header =
Destination

Destination Header
Next Header = TCP

Fragment of
TCP Header
+ Data

- All optional fields go into extension headers
- These are daisy chained behind the main header
  - The last 'extension' header is usually the ICMP, TCP or UDP header
- Makes it simple to add new features in IPv6 protocol without major re-engineering of devices
- Number of extension headers is not fixed / limited

Header Format – Common Headers

- Common values of Next Header field:

0	Hop-by-hop option (extension)
2	ICMP (payload)
6	TCP (payload)
17	UDP (payload)
43	Source routing (extension)
44	Fragmentation (extension)
50	Encrypted security payload (extension, IPSec)
51	Authentication (extension, IPSec)
59	Null (No next header)
60	Destination option (extension)

Header Format – Ordering of Headers

- Order is important because:
  - Hop-by-hop header has to be processed by every intermediate node
  - Routing header needs to be processed by intermediate routers
  - At the destination fragmentation has to be processed before other headers
- This makes header processing easier to implement in hardware

Larger Address Space

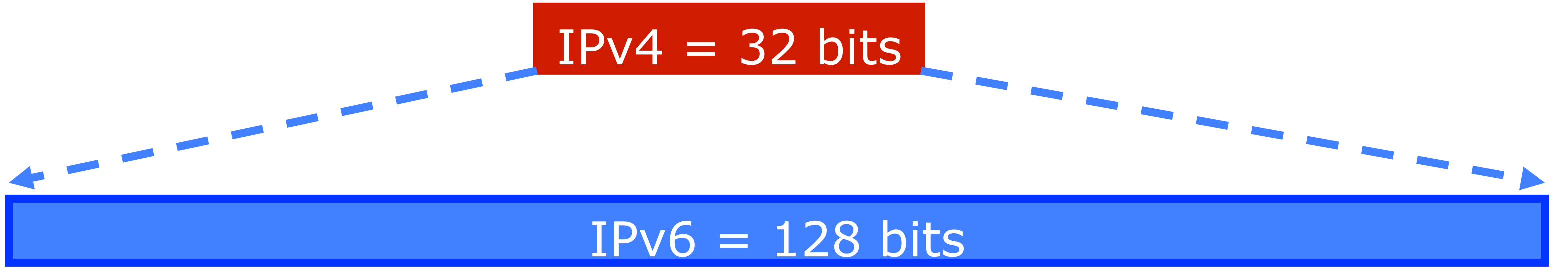

IPv4 = 32 bits

IPv6 = 128 bits

- IPv4
  - 32 bits
  - = 4,294,967,296 possible addressable devices
- IPv6
  - 128 bits: 4 times the size in bits
  - = 3.4 x 10^{38} possible addressable devices
  - = 340,282,366,920,938,463,463,374,607,431,768,211,456
  - ~ 5 x 10^{28} addresses per person on the planet

How was the IPv6 Address Size Chosen?

- Some wanted fixed-length, 64-bit addresses
  - Easily good for 10^{12} sites, 10^{15} nodes, at .0001 allocation efficiency
    - (3 orders of magnitude more than IPv6 requirement)
  - Minimizes growth of per-packet header overhead
  - Efficient for software processing
- Some wanted variable-length, up to 160 bits
  - Compatible with OSI NSAP addressing plans
  - Big enough for auto-configuration using IEEE 802 addresses
  - Could start with addresses shorter than 64 bits & grow later
- Settled on fixed-length, 128-bit addresses

IPv6 Address Representation (1)

- 16 bit fields in case insensitive colon hexadecimal representation
  - 2031:0000:130F:0000:0000:09C0:876A:130B
- Leading zeros in a field are optional:
  - 2031:0:130F:0:0:9C0:876A:130B
- Successive fields of 0 represented as ::, but only once in an address:
  - 2031:0:130F::9C0:876A:130B is ok
  - 2031::130F::9C0:876A:130B is **NOT** ok
  - 0:0:0:0:0:0:0:1 → ::1 (loopback address)
  - 0:0:0:0:0:0:0:0 → :: (unspecified address)

IPv6 Address Representation (2)

- :: representation
  - RFC5952 recommends that the rightmost set of :0: be replaced with :: for consistency
    - 2001:db8:0:2f::5 rather than 2001:db8::2f:0:0:0:5
- IPv4-compatible (not used any more)
  - 0:0:0:0:0:0:192.168.30.1
  - = ::192.168.30.1
  - = ::C0A8:1E01
- In a URL, it is enclosed in brackets (RFC3986)
  - http://[2001:db8:4f3a::206:ae14]:8080/index.html
  - Cumbersome for users, mostly for diagnostic purposes
  - Use fully qualified domain names (FQDN)
  - ⇒ The DNS has to work!!

IPv6 Address Representation (3)

- Prefix Representation
  - Representation of prefix is just like IPv4 CIDR
  - In this representation you attach the prefix length
  - Like IPv4 address:
    - 198.10.0.0/16
  - IPv6 address is represented in the same way:
    - 2001:db8:12::/40

IPv6 Addressing

- IPv6 Addressing rules are covered by multiple RFCs
  - Architecture defined by RFC 4291
- Address Types are :
  - Unicast : One to One (Global, Unique Local, Link local)
  - Anycast : One to Nearest (Allocated from Unicast)
  - Multicast : One to Many
- A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, multicast)
  - No Broadcast Address → Use Multicast

IPv6 Addressing

Type	Binary	Hex
Unspecified	000…0	::/128
Loopback	000…1	::1/128
Global Unicast Address	0010	2000::/3
Link Local Unicast Address	1111 1110 10	FE80::/10
Unique Local Unicast Address	1111 1100 1111 1101	FC00::/7
Multicast Address	1111 1111	FF00::/8

IPv6 Address Allocation

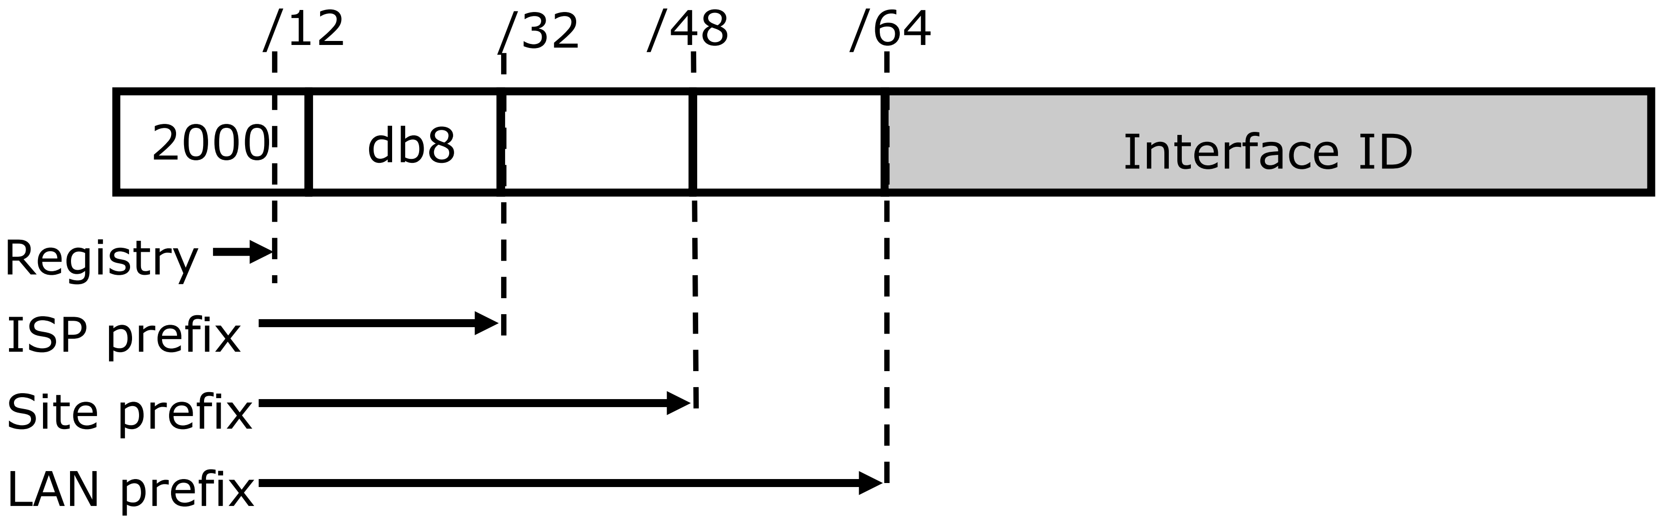

- The allocation process is:
  - The IANA is allocating out of 2000::/3 for initial IPv6 unicast use
  - Each registry gets a /12 prefix from the IANA
  - Registry allocates a /32 prefix (or larger) to an IPv6 ISP
  - Policy is that an ISP allocates a /48 prefix to each end customer

IPv6 Addressing Scope

- 64 bits reserved for the interface ID
  - Possibility of 2^{64} hosts on one network LAN
  - In theory 18,446,744,073,709,551,616 hosts
  - Arrangement to accommodate MAC addresses within the IPv6 address
- 16 bits reserved for the end site
  - Possibility of 2^{16} networks at each end-site
  - 65536 subnets equivalent to a /12 in IPv4 (assuming a /28 or 16 hosts per IPv4 subnet)

IPv6 Addressing Scope

- 16 bits reserved for each service provider
  - Possibility of 2^{16} end-sites per service provider
  - 65536 possible customers: equivalent to each service provider receiving a /8 in IPv4 (assuming a /24 address block per customer)
- 29 bits reserved for all service providers
  - Possibility of 2^{29} service providers
  - i.e. 536,870,912 discrete service provider networks
    - Although some service providers already are justifying more than a /32

How to get an IPv6 Address?

- IPv6 address space is allocated by the 5 RIRs:
  - AfriNIC, APNIC, ARIN, LACNIC, RIPE NCC
  - ISPs get address space from the RIRs
  - Enterprises get their IPv6 address space from their ISP
- 6to4 tunnels 2002::/16
  - Last resort only and now mostly useless
- (6Bone)
  - Was the IPv6 experimental network since the mid 90s
  - Now retired, end of service was 6th June 2006 (RFC3701)

Aggregation hopes

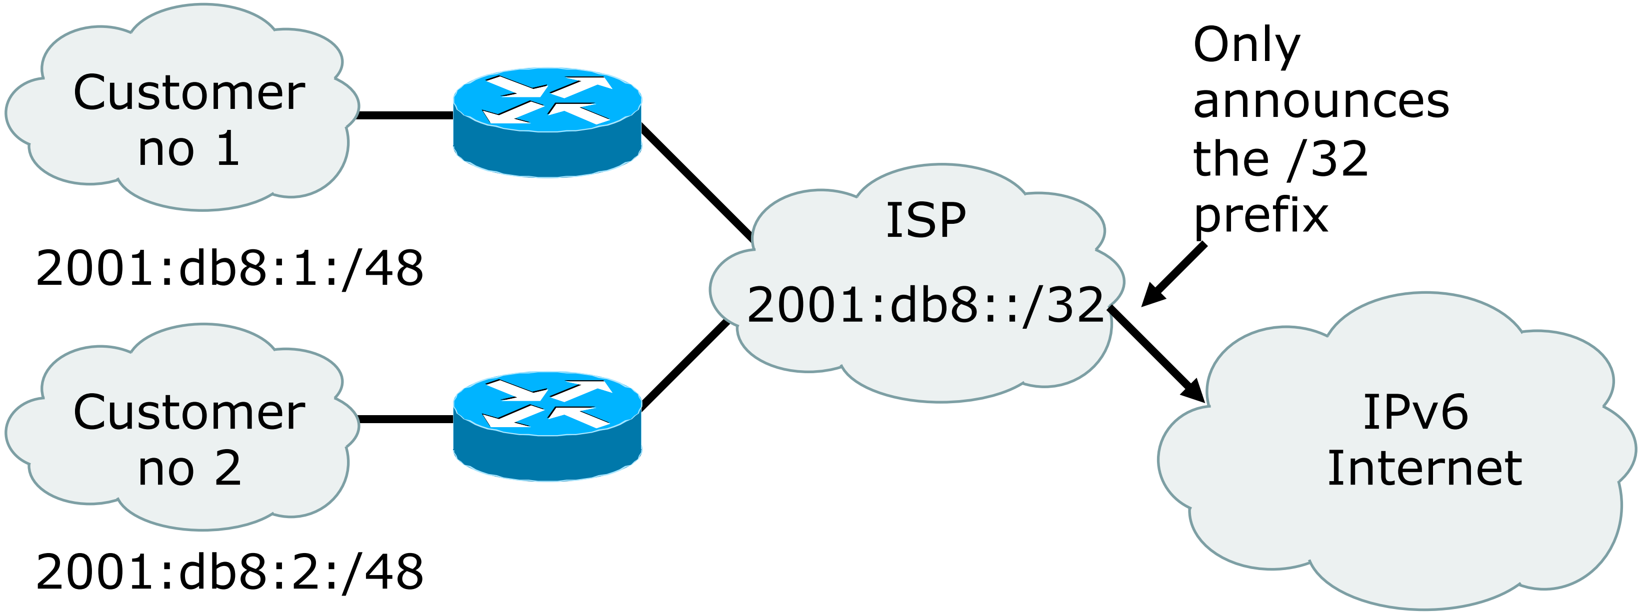

- Larger address space enables aggregation of prefixes announced in the global routing table
- Idea was to allow efficient and scalable routing
- But current Internet multihoming solution breaks this model

Interface IDs

- Lowest order 64-bit field of unicast address may be assigned in several different ways:
  - Auto-configured from a 64-bit EUI-64, or expanded from a 48-bit MAC address (e.g., Ethernet address)
  - Auto-generated pseudo-random number (to address privacy concerns)
  - Assigned via DHCP
  - Manually configured

EUI-64

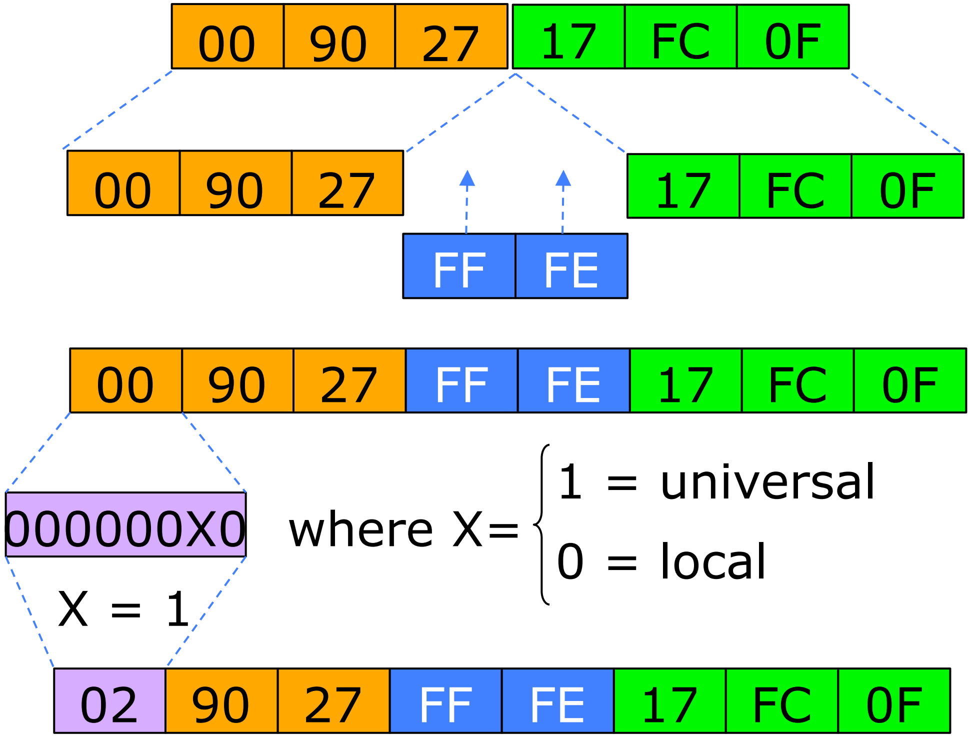

Ethernet MAC address
(48 bits)

00 90 27 17 FC 0F

00 90 27 FF FE 17 FC 0F

64 bits version

00 90 27 FF FE 17 FC 0F

Scope of the EUI-64 id

000000X0 where X= 1 = universal
0 = local

X = 1

EUI-64 address

02 90 27 FF FE 17 FC 0F

- EUI-64 address is formed by inserting FFFE between the company-id and the manufacturer extension, and setting the "u" bit to indicate scope
  - Global scope: for IEEE 48-bit MAC
  - Local scope: when no IEEE 48-bit MAC is available (eg serials, tunnels)

IPv6 Addressing Examples

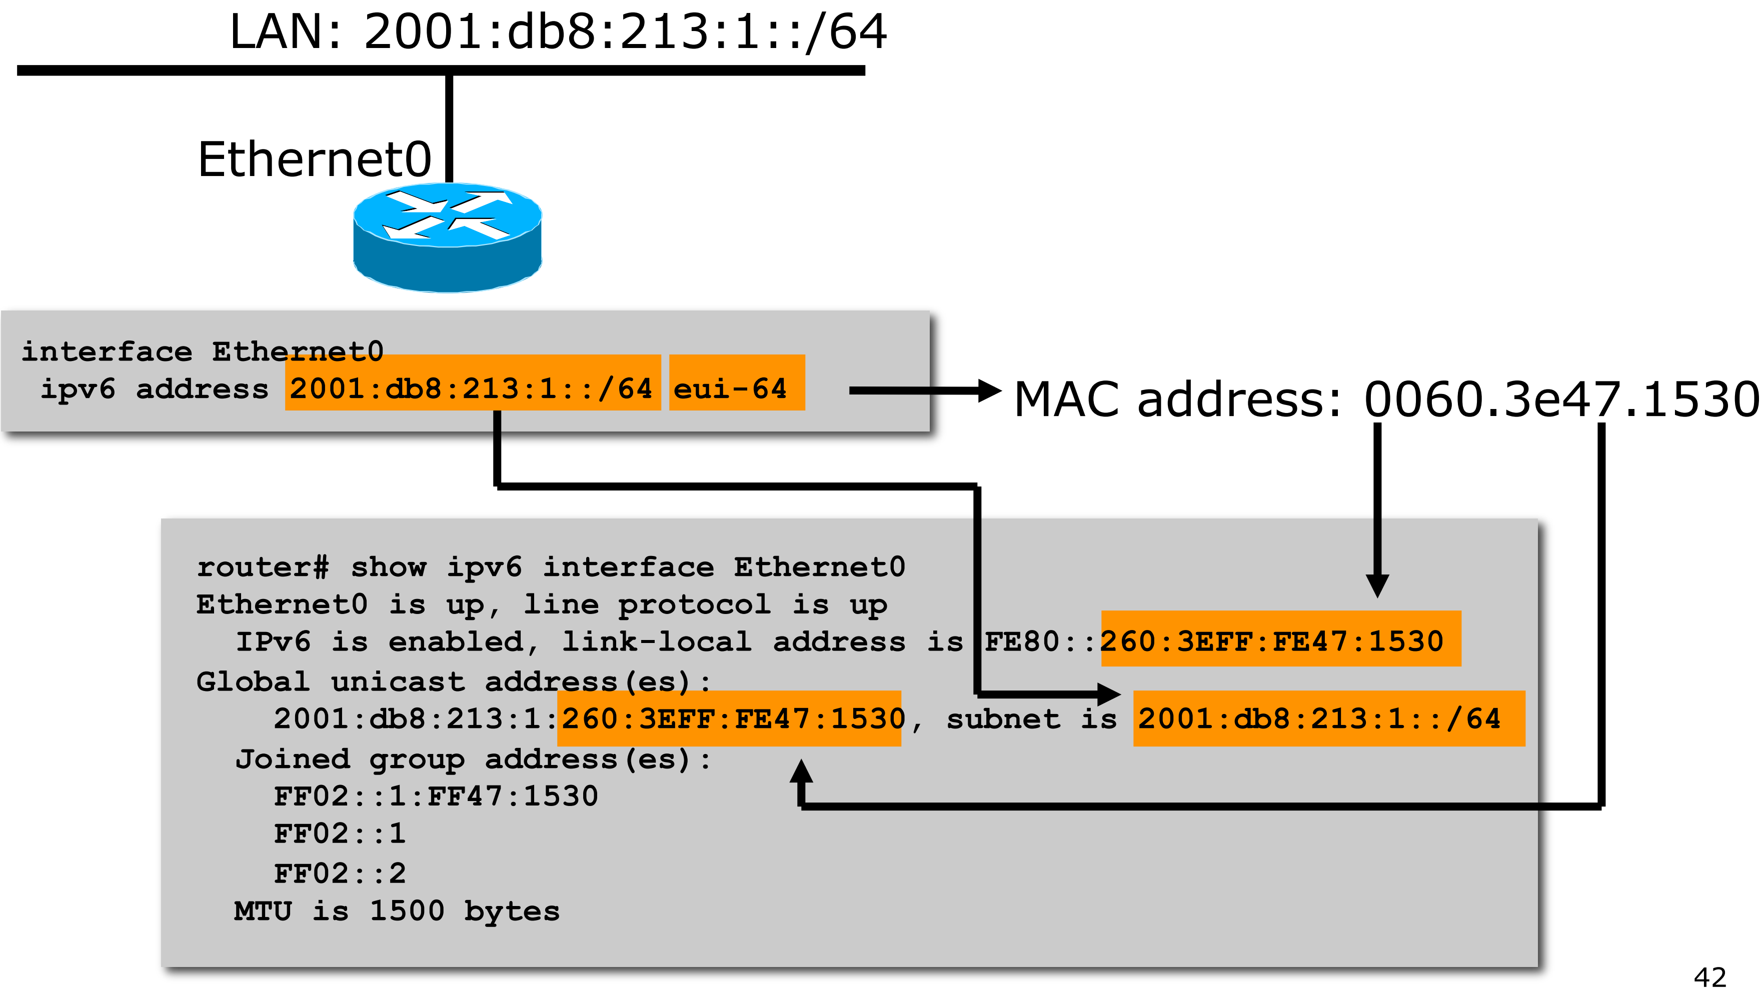

LAN: 2001:db8:213:1::/64

Ethernet0

```
interface Ethernet0
 ipv6 address 2001:db8:213:1::/64 eui-64
```

MAC address: 0060.3e47.1530

```
router# show ipv6 interface Ethernet0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::260:3EFF:FE47:1530
Global unicast address(es):
    2001:db8:213:1:260:3EFF:FE47:1530, subnet is 2001:db8:213:1::/64
  Joined group address(es):
    FF02::1:FF47:1530
    FF02::1
    FF02::2
  MTU is 1500 bytes
```

IPv6 Address Privacy (RFC 4941)

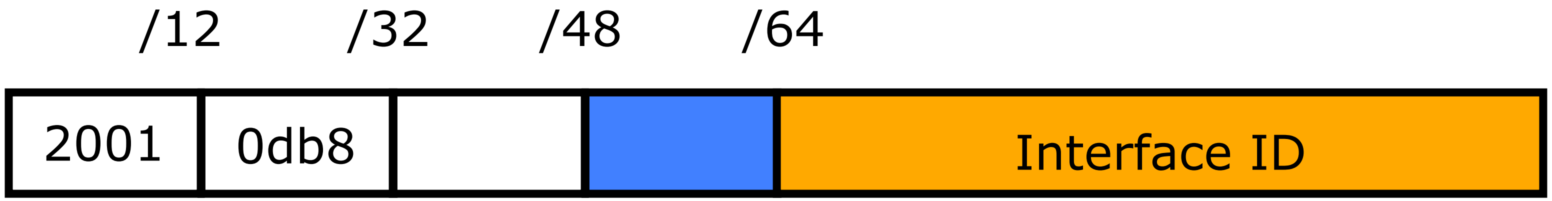

/12 /32 /48 /64

2001 | 0db8 | | | Interface ID

- Temporary addresses for IPv6 host client application, e.g. Web browser
- Intended to inhibit device/user tracking but is also a potential issue
  - More difficult to scan all IP addresses on a subnet
  - But port scan is identical when an address is known
- Random 64 bit interface ID, run DAD before using it
- Rate of change based on local policy
- Implemented on Microsoft Windows XP/Vista/7 and Apple MacOS 10.7 onwards
  - Can be activated on FreeBSD/Linux with a system call

Host IPv6 Addressing Options

- Stateless (RFC4862)
  - SLAAC – Stateless Address AutoConfiguration
  - Booting node sends a “router solicitation” to request “router advertisement” to get information to configure its interface
  - Booting node configures its own Link-Local address
- Stateful
  - DHCPv6 – required by most enterprises
  - Manual – like IPv4 pre-DHCP
    - Useful for servers and router infrastructure
    - Doesn’t scale for typical end user devices

IPv6 Renumbering

- Renumbering Hosts
  - Stateless:
    - Hosts renumbering is done by modifying the RA to announce the old prefix with a short lifetime and the new prefix
  - Stateful:
    - DHCPv6 uses same process as DHCPv4
- Renumbering Routers
  - Router renumbering protocol was developed (RFC 2894) to allow domain-interior routers to learn of prefix introduction / withdrawal
  - **No known implementation!**

Auto-configuration

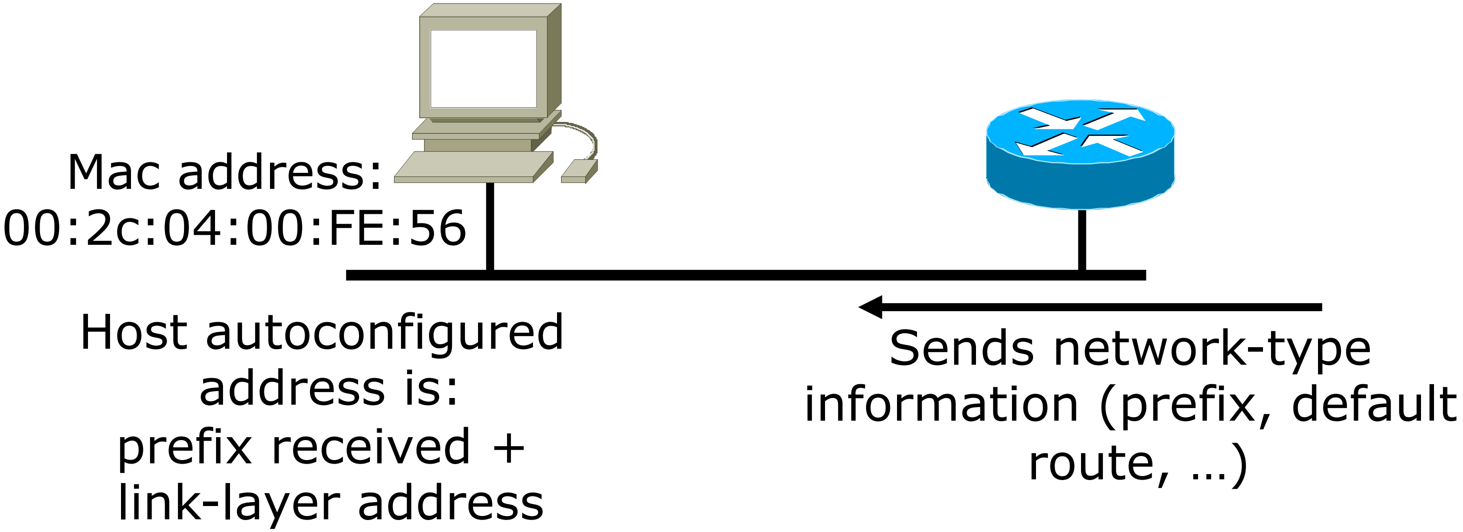

- PC sends router solicitation (RS) message
- Router responds with router advertisement (RA)
  - This includes prefix and default route
  - RFC6106 adds DNS server option
- PC configures its IPv6 address by concatenating prefix received with its EUI-64 address

Renumbering

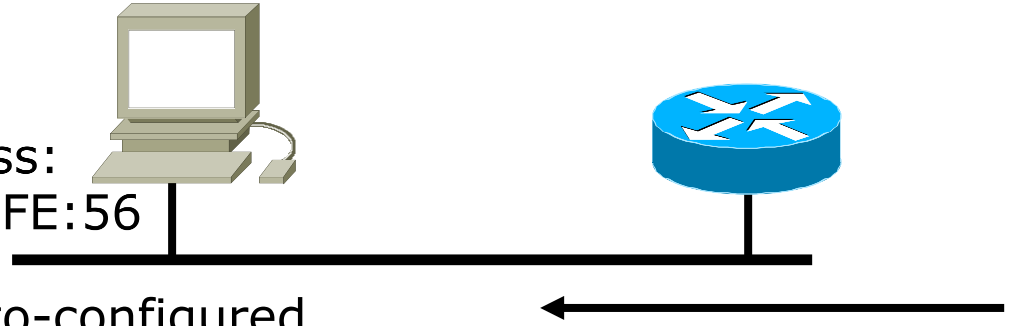

Mac address:
00:2c:04:00:FE:56

Host auto-configured
address is:

NEW prefix received +
SAME link-layer address

Sends *NEW* network-type
information (prefix, default
route, …)

- Router sends router advertisement (RA)
  - This includes the new prefix and default route (and remaining lifetime of the old address)
- PC configures a new IPv6 address by concatenating prefix received with its EUI-64 address
  - Attaches lifetime to old address

Unique-Local

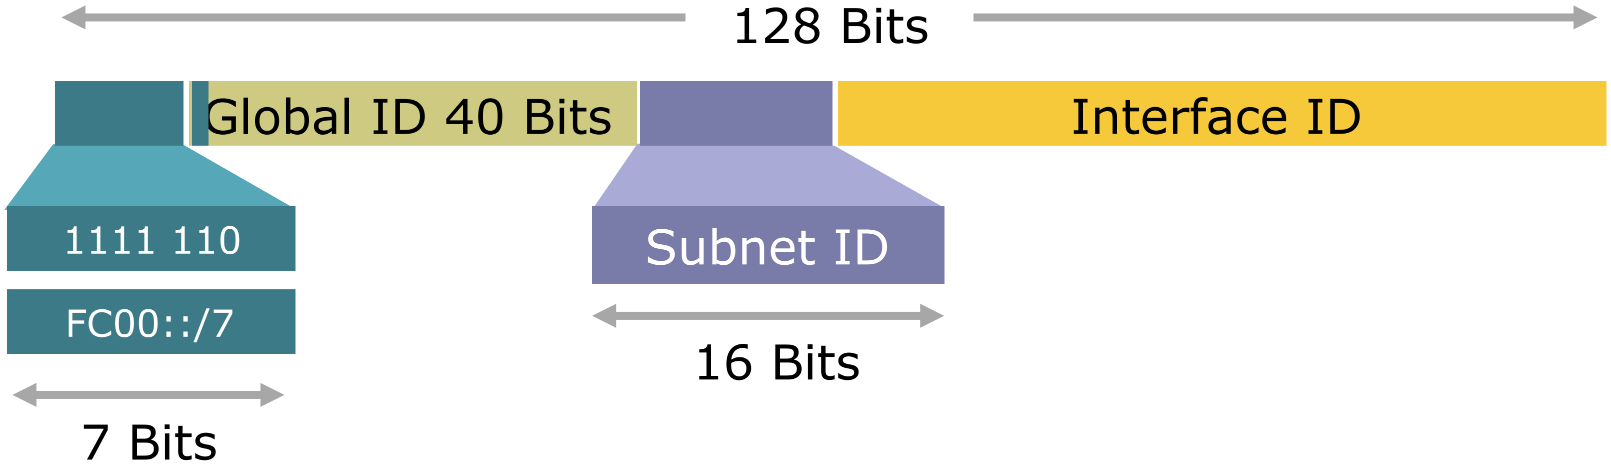

- Unique-Local Addresses Used For:
  - Local communications & inter-site VPNs
  - Local devices such as printers, telephones, etc
  - Site Network Management systems connectivity
- Not routable on the Internet
- Reinvention of the deprecated site-local?

Link-Local

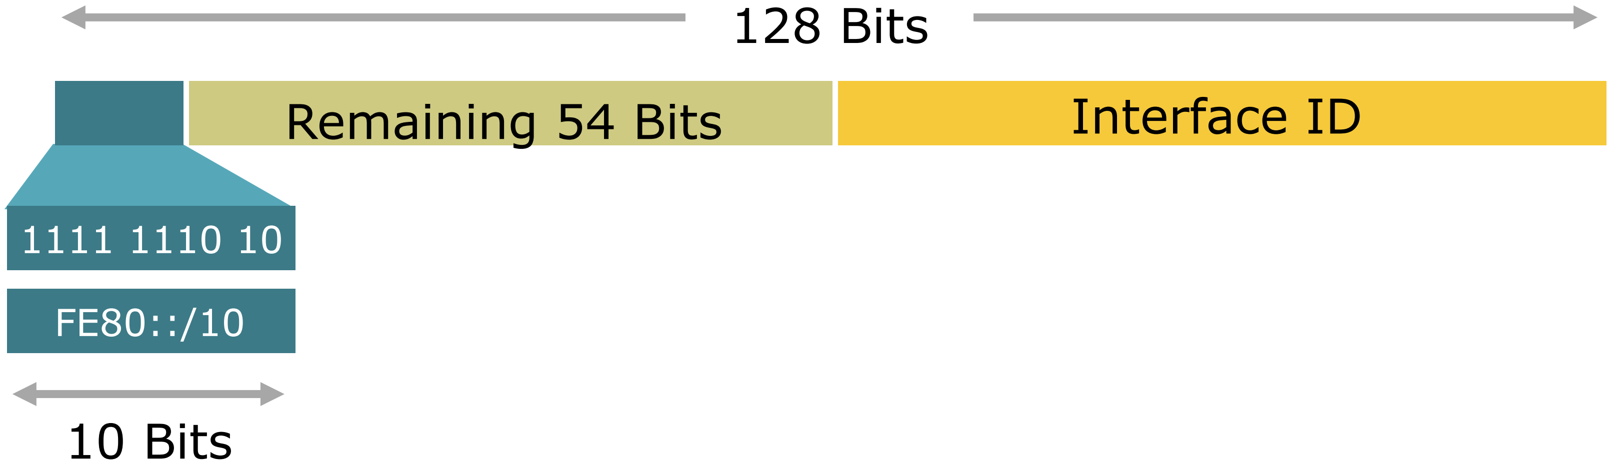

- Link-Local Addresses Used For:
  - Communication between two IPv6 device (like ARP but at Layer 3)
  - Next-Hop calculation in Routing Protocols
- Automatically assigned by Router as soon as IPv6 is enabled
  - Mandatory Address
- Only Link Specific scope
- Remaining 54 bits could be Zero or any manual configured value

Multicast use

- Broadcasts in IPv4
  - Interrupts all devices on the LAN even if the intent of the request was for a subset
  - Can completely swamp the network (“broadcast storm”)
- Broadcasts in IPv6
  - Are not used and replaced by multicast
- Multicast
  - Enables the efficient use of the network
  - Multicast address range is much larger

IPv6 Multicast Address

- IP multicast address has a prefix FF00::/8
- The second octet defines the lifetime and scope of the multicast address.

8-bit	4-bit	4-bit	112-bit
1111 1111	Lifetime	Scope	Group-ID

Lifetime	
0	If Permanent
1	If Temporary

Scope	
1	Node
2	Link
5	Site
8	Organisation
E	Global

IPv6 Multicast Address Examples

- RIPng
  - The multicast address AllRIPRouters is FF02::9
    - Note that 02 means that this is a permanent address and has link scope
- OSPFv3
  - The multicast address AllSPFRouters is FF02::5
  - The multicast address AllDRouters is FF02::6
- EIGRP
  - The multicast address AllEIGRPRouters is FF02::A

Solicited-Node Multicast

- Solicited-Node Multicast is used for Duplicate Address Detection
  - Part of the Neighbour Discovery process
  - Replaces ARP
  - Duplicate IPv6 Addresses are rare, but still have to be tested for
- For each unicast and anycast address configured there is a corresponding solicited-node multicast address
  - This address is only significant for the local link

Solicited-Node Multicast Address

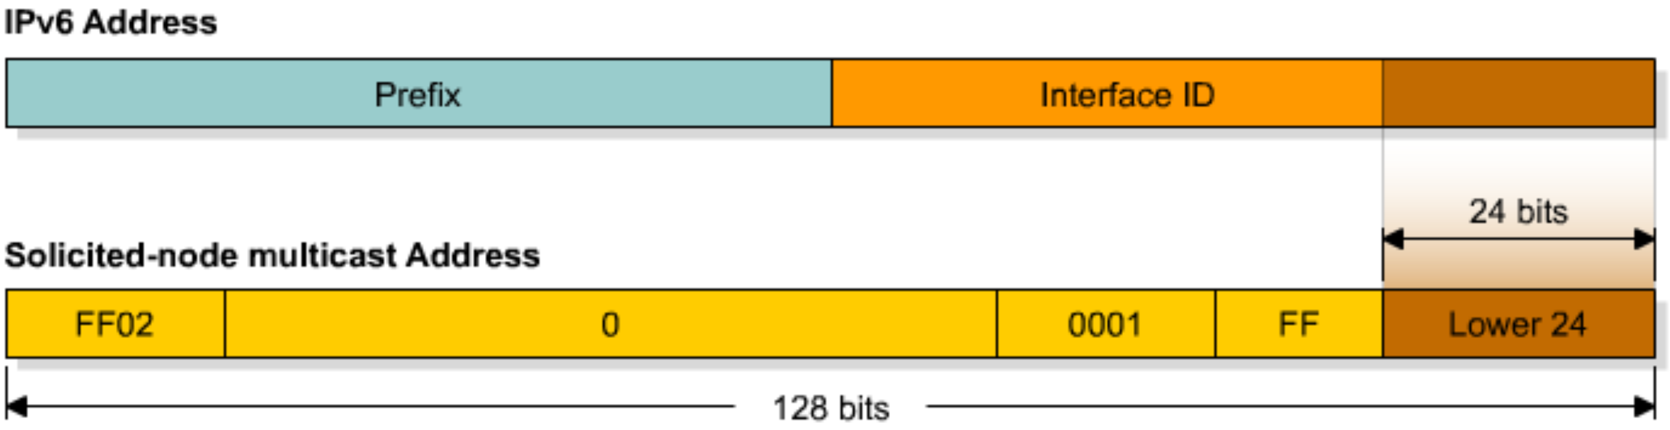

- Solicited-node multicast address consists of FF02:0:0:0:0:1:FF::/104 prefix joined with the lower 24 bits from the unicast or anycast IPv6 address

Solicited-Node Multicast

```
R1#sh ipv6 int e0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF3A:8B18
  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.
R1#
```

Solicited-Node Multicast Address

IPv6 Anycast

- An IPv6 anycast address is an identifier for a set of interfaces (typically belonging to different nodes)
    - A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocol's measure of distance).
    - RFC4291 describes IPv6 Anycast in more detail
- In reality there is no known implementation of IPv6 Anycast as per the RFC
    - Most operators have chosen to use IPv4 style anycast instead

Anycast on the Internet

- A global unicast address is assigned to all nodes which need to respond to a service being offered
    - This address is routed as part of its parent address block
- The responding node is the one which is closest to the requesting node according to the routing protocol
    - Each anycast node looks identical to the other
- Applicable within an ASN, or globally across the Internet
- Typical (IPv4) examples today include:
    - Root DNS and ccTLD/gTLD nameservers
    - SMTP relays and DNS resolvers within ISP autonomous systems

MTU Issues

- Minimum link MTU for IPv6 is 1280 octets (versus 68 octets for IPv4)
  - ⇒ on links with MTU < 1280, link-specific fragmentation and reassembly must be used
- Implementations are expected to perform path MTU discovery to send packets bigger than 1280
- Minimal implementation can omit PMTU discovery as long as all packets kept ≤ 1280 octets
- A Hop-by-Hop Option supports transmission of “jumbograms” with up to 2^{32} octets of payload

IPv6 Neighbour Discovery

- Protocol defines mechanisms for the following problems:
  - Router discovery
  - Prefix discovery
  - Parameter discovery
  - Address autoconfiguration
  - Address resolution
  - Next-hop determination
  - Neighbour unreachability detection
  - Duplicate address detection
  - Redirects

IPv6 Neighbour Discovery

- Defined in RFC 4861
- Protocol built on top of ICMPv6 (RFC 4443)
  - Combination of IPv4 protocols (ARP, ICMP, IGMP,…)
- Fully dynamic, interactive between Hosts & Routers
- Defines 5 ICMPv6 packet types:
  - Router Solicitation
  - Router Advertisement
  - Neighbour Solicitation
  - Neighbour Advertisement
  - Redirect

IPv6 and DNS

- Hostname to IP address:

IPv4	www.abc.test.	A	192.168.30.1
IPv6	www.abc.test	AAAA	2001:db8:c18:1::2

IPv6 and DNS

- IP address to Hostname:

IPv4 1.30.168.192.in-addr.arpa. PTR www.abc.test.

IPv6 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.8.b.d.
0.1.0.0.2.ip6.arpa PTR www.abc.test.

IPv6 Technology Scope

IP Service	*IPv4 Solution*	*IPv6 Solution*
Addressing Range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	Serverless, Reconfiguration, DHCP
Security	IPSec	IPSec Mandated, works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality-of-Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
IP Multicast	IGMP/PIM/Multicast BGP	MLD/PIM/Multicast BGP, Scope Identifier

What does IPv6 do for:

- Security
  - Nothing IPv4 doesn’t do – IPSec runs in both
  - But IPv6 mandates IPSec
- QoS
  - Nothing IPv4 doesn’t do –
    - Differentiated and Integrated Services run in both
    - So far, Flow label has no real use

IPv6 Security

- IPsec standards apply to both IPv4 and IPv6
- All implementations required to support authentication and encryption headers (“IPsec”)
- Authentication separate from encryption for use in situations where encryption is prohibited or prohibitively expensive
- Key distribution protocols are not yet defined (independent of IP v4/v6)
- Support for manual key configuration required

IP Quality of Service Reminder

- Two basic approaches developed by IETF:
  - “Integrated Service” (int-serv)
    - Fine-grain (per-flow), quantitative promises (e.g., x bits per second), uses RSVP signalling
  - “Differentiated Service” (diff-serv)
    - Coarse-grain (per-class), qualitative promises (e.g., higher priority), no explicit signalling
  - Signalled diff-serv (RFC 2998)
    - Uses RSVP for signalling with course-grained qualitative aggregate markings
    - Allows for policy control without requiring per-router state overhead

IPv6 Support for Int-Serv

- 20-bit Flow Label field to identify specific flows needing special QoS
  - Each source chooses its own Flow Label values; routers use Source Addr + Flow Label to identify distinct flows
  - Flow Label value of 0 used when no special QoS requested (the common case today)
- Originally standardised as RFC 3697

IPv6 Flow Label

- Flow label has not been used since IPv6 standardised
  - Suggestions for use in recent years were incompatible with original specification (discussed in RFC6436)
- Specification updated in RFC6437
  - RFC6438 describes the use of the Flow Label for equal cost multi-path and link aggregation in Tunnels

IPv6 Support for Diff-Serv

- 8-bit Traffic Class field to identify specific classes of packets needing special QoS
    - Same as new definition of IPv4 Type-of-Service byte
    - May be initialized by source or by router enroute; may be rewritten by routers enroute
    - Traffic Class value of 0 used when no special QoS requested (the common case today)

IPv6 Standards

- Core IPv6 specifications are IETF Draft Standards → well-tested & stable
  - IPv6 base spec, ICMPv6, Neighbor Discovery, PMTU Discovery,...
- Other important specs are further behind on the standards track, but in good shape
  - Mobile IPv6, header compression,...
  - For up-to-date status: www.ipv6tf.org
- 3GPP UMTS Rel. 5 cellular wireless standards (2002) mandate IPv6; also being considered by 3GPP2

IPv6 Status – Standardisation

- Several key components on standards track…
  - Specification (RFC2460)
  - ICMPv6 (RFC4443)
  - RIP (RFC2080)
  - IGMPv6 (RFC2710)
  - Router Alert (RFC2711)
  - Autoconfiguration (RFC4862)
  - DHCPv6 (RFC3315 & 4361)
  - IPv6 Mobility (RFC3775)
  - GRE Tunnelling (RFC2473)
  - DAD for IPv6 (RFC4429)
  - ISIS for IPv6 (RFC5308)
  - Neighbour Discovery (RFC4861)
  - IPv6 Addresses (RFC4291 & 3587)
  - BGP (RFC2545)
  - OSPF (RFC5340)
  - Jumbograms (RFC2675)
  - Radius (RFC3162)
  - Flow Label (RFC6436/7/8)
  - Mobile IPv6 MIB (RFC4295)
  - Unique Local IPv6 Addresses (RFC4193)
  - Teredo (RFC4380)
  - VRRP (RFC5798)
- IPv6 available over:
  - PPP (RFC5072)
  - FDDI (RFC2467)
  - NBMA (RFC2491)
  - Frame Relay (RFC2590)
  - IEEE1394 (RFC3146)
  - Facebook (RFC5514)
  - Ethernet (RFC2464)
  - Token Ring (RFC2470)
  - ATM (RFC2492)
  - ARCnet (RFC2497)
  - FibreChannel (RFC4338)

Recent IPv6 Hot Topics

- IPv4 depletion debate
  - IANA IPv4 pool ran out on 3rd February 2011
    - http://www.potaroo.net/tools/ipv4/
- IPv6 Transition “assistance”
  - CGN, 6rd, NAT64, IVI, DS-Lite, 6to4, A+P…
- Mobile IPv6
- Multihoming
  - SHIM6 “dead”, Multihoming in IPv6 same as in IPv4
- IPv6 Security
  - Security industry & experts taking much closer look

IPv6 Addressing

How to handle IPv6 addresses
and do address planning

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 (1)

- From your Regional Internet Registry
  - Become a member of your Regional Internet Registry and get your own allocation
    - Membership usually open to all network operators
  - General allocation policies are outlined in RFC2050
    - RIR specific policy details for IPv6 allocations are listed on the individual RIR website
  - Open to all organisations who are operating a network
  - Receive a /32 (or larger if you will have more than 65k /48 assignments)

Getting IPv6 address space (2)

- From your upstream ISP
  - Receive a /48 from upstream ISP's IPv6 address block
  - Receive more than one /48 if you have more than 65k subnets
- If you need to multihome:
  - Apply for a /48 assignment from your RIR
  - Multihoming with provider's /48 will be operationally challenging
    - Provider policies, filters, etc

Using 6to4 for IPv6 address space

- Some entities still use 6to4
  - Not recommended due to operational problems
  - Read http://datatracker.ietf.org/doc/draft-ietf-v6ops-6to4-to-historic for some of the reasoning why
- FYI: 6to4 operation:
  - Take a single public IPv4 /32 address
  - 2002:<ipv4 /32 address>::/48 becomes your IPv6 address block, giving 65k subnets
  - Requires a 6to4 gateway
  - 6to4 is a means of connecting IPv6 islands across the IPv4 Internet

Nibble Boundaries

- IPv6 offers network operators more flexibility with addressing plans
  - Network addressing can now be done on nibble boundaries
    - For ease of operation
  - Rather than making maximum use of a very scarce resource
    - With the resulting operational complexity
- A nibble boundary means subnetting address space based on the address numbering
  - Each number in IPv6 represents 4 bits = 1 nibble
  - Which means that IPv6 addressing can be done on 4-bit boundaries

Nibble Boundaries – example

- Consider the address block 2001:db8:0:10::/61
  - The range of addresses in this block are:

```
2001:0db8:0000:0010:0000:0000:0000:0000
                  to
2001:0db8:0000:0017:ffff:ffff:ffff:ffff
```

  - Note that this subnet only runs from 0010 to 0017.
  - The adjacent block is 2001:db8:0:18::/61

```
2001:0db8:0000:0018:0000:0000:0000:0000
                  to
2001:0db8:0000:001f:ffff:ffff:ffff:ffff
```

  - The address blocks don’t use the entire nibble range

Nibble Boundaries – example

- Now consider the address block 2001:db8:0:10::/60
  - The range of addresses in this block are:

```
2001:0db8:0000:0010:0000:0000:0000:0000
                  to
2001:0db8:0000:001f:ffff:ffff:ffff:ffff
```

  - Note that this subnet uses the entire nibble range, 0 to f
  - Which makes the numbering plan for IPv6 simpler
    - This range can have a particular meaning within the ISP block (for example, infrastructure addressing for a particular PoP)

Addressing Plans – Infrastructure

- All Network Operators should obtain a /32 from their RIR
- Address block for router loop-back interfaces
  - Number all loopbacks out of **one** /64
  - /128 per loopback
- Address block for infrastructure (backbone)
  - /48 allows 65k subnets
  - /48 per region (for the largest multi-national networks)
  - /48 for whole backbone (for the majority of networks)
  - Infrastructure/backbone usually does NOT require regional/geographical addressing
  - Summarise between sites if it makes sense

Addressing Plans – Infrastructure

- What about LANs?
  - /64 per LAN
- What about Point-to-Point links?
  - Protocol design expectation is that /64 is used
  - /127 now recommended/standardised
    - http://www.rfc-editor.org/rfc/rfc6164.txt
    - (reserve /64 for the link, but address it as a /127)
  - Other options:
    - /126s are being used (mimics IPv4 /30)
    - /112s are being used
      - Leaves final 16 bits free for node IDs
    - Some discussion about /80s, /96s and /120s too

Addressing Plans – Infrastructure

- NOC:
  - ISP NOC is “trusted” network and usually considered part of infrastructure /48
    - Contains management and monitoring systems
    - Hosts the network operations staff
    - take the last /60 (allows enough subnets)
- Critical Services:
  - Network Operator’s critical services are part of the “trusted” network and should be considered part of the infrastructure /48
  - For example, Anycast DNS, SMTP, POP3/IMAP, etc
    - Take the second /64
    - (some operators use the first /64 instead)

Addressing Plans – ISP to Customer

- Option One:
  - Use ipv6 unnumbered
  - Which means no global unicast ipv6 address on the point-to-point link
  - Router adopts the specified interface’s IPv6 address
    - Router doesn’t actually need a global unicast IPv6 address to forward packets

```
interface loopback 0
 ipv6 address 2001:db8::1/128
interface serial 1/0
 ipv6 address unnumbered loopback 0
```

Addressing Plans – ISP to Customer

- Option Two:
  - Use the second /48 for point-to-point links
  - Divide this /48 up between PoPs
  - Example:
    - For 10 PoPs, dividing into 16, gives /52 per PoP
    - Each /52 gives 4096 point-to-point links
    - Adjust to suit!
  - Useful if ISP monitors point-to-point link state for customers
    - Link addresses are **untrusted**, so do not want them in the first /48 used for the backbone &c
  - Aggregate per router or per PoP and carry in iBGP (not ISIS/OSPF)

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)
- In typical deployments today:
  - Several ISPs are giving small customers a /56 and single LAN end-sites a /64, e.g.:

/64	if end-site will only ever be a LAN
/56	for small end-sites (e.g. home/office/small business)
/48	for large end-sites

  - This is another very active discussion area
  - Observations:
    - Don’t assume that a mobile endsite needs only a /64
    - Some operators are distributing /60s to their smallest customers!!

Addressing Plans – Customer

- Consumer Broadband Example:
  - DHCPv6 pool is a /48
    - DHCPv6 hands out /60 per customer
    - Which allows for 4096 customers per pool
- Business Broadband Example:
  - DHCPv6 pool is a /48
    - DHCPv6 hands out /56 per customer
    - Which allows for 256 customers per pool
  - If BRAS has more than 256 business customers, increase pool to a /47
    - This allows for 512 customers at /56 per customer
  - Increasing pool to /46 allows for 1024 customers
  - BRAS announces entire pool as one block by iBGP

Addressing Plans – Customer

- Business “leased line”:
  - /48 per customer
  - One stop shop, no need for customer to revisit ISP for more addresses until all 65k subnets are used up
- Hosted services:
  - One physical server per vLAN
  - One /64 per vLAN
  - How many vLANs per PoP?
  - /48 reserved for entire hosted servers across backbone
    - Internal sites will be subnets and carried by iBGP

Addressing Plans – Customer

- Geographical delegations to Customers:
  - Network Operator subdivides /32 address block into geographical chunks
  - E.g. into /36s
    - Region 1: 2001:db8:1xxx::/36
    - Region 2: 2001:db8:2xxx::/36
    - Region 3: 2001:db8:3xxx::/36
    - etc
  - Which gives 4096 /48s per region
  - For Operational and Administrative ease
  - Benefits for traffic engineering if Network Operator multihomes in each region

Addressing Plans – Customer

- Sequential delegations to Customers:
  - After carving off address space for network infrastructure, Network Operator simply assigns address space sequentially
  - Eg:
    - Infrastructure: 2001:db8:0::/48
    - Customer P2P: 2001:db8:1::/48
    - Customer 1: 2001:db8:2::/48
    - Customer 2: 2001:db8:3::/48
    - etc
  - Useful when there is no regional subdivision of network and no regional multihoming needs

Addressing Plans – Routing Considerations

- Carry Broadband pools in iBGP across the backbone
  - Not in OSPF/ISIS
- Multiple Broadband pools on one BRAS should be aggregated if possible
  - Reduce load on iBGP
- Aggregating leased line customer address blocks per router or per PoP is undesirable:
  - Interferes with ISP’s traffic engineering needs
  - Interferes with ISP’s service quality and service guarantees

Addressing Plans – Traffic Engineering

- Smaller providers will be single homed
  - The customer portion of the ISP’s IPv6 address block will usually be assigned sequentially
- Larger providers will be multihomed
  - Two, three or more external links from different providers
  - Traffic engineering becomes important
  - Sequential assignments of customer addresses will negatively impact load balancing

Addressing Plans – Traffic Engineering

- ISP Router loopbacks and backbone point-to-point links make up a small part of total address space
  - And they don't attract traffic, unlike customer address space
- Links from ISP Aggregation edge to customer router needs one /64
  - Small requirements compared with total address space
  - Some ISPs use IPv6 unnumbered
- Planning customer assignments is a very important part of multihoming
  - Traffic engineering involves subdividing aggregate into pieces until load balancing works

Unplanned IP addressing

- ISP fills up customer IP addressing from one end of the range:

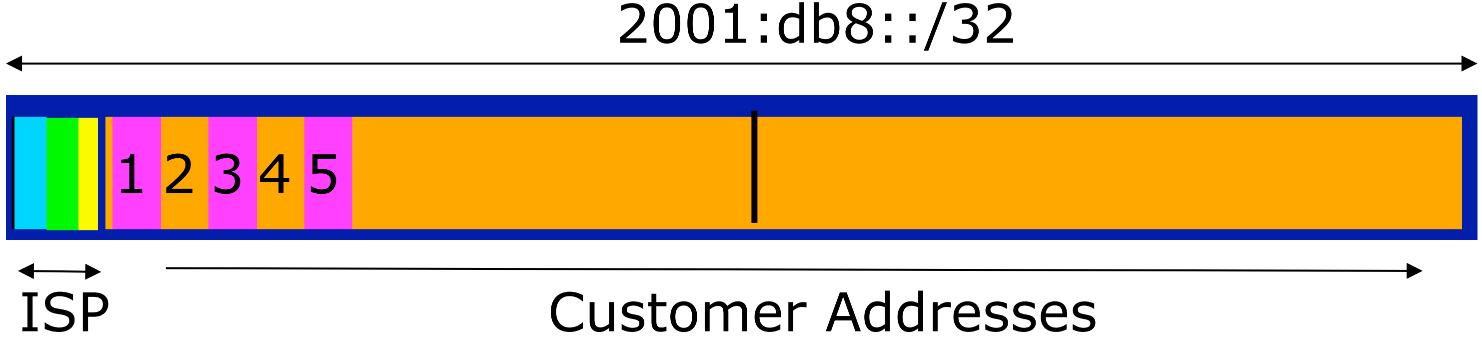

- Customers generate traffic
  - Dividing the range into two pieces will result in one /33 with all the customers and the ISP infrastructure the addresses, and one /33 with nothing
  - No loadbalancing as all traffic will come in the first /33
  - Means further subdivision of the first /33 = harder work

Planned IP addressing

- If ISP fills up customer addressing from both ends of the range:

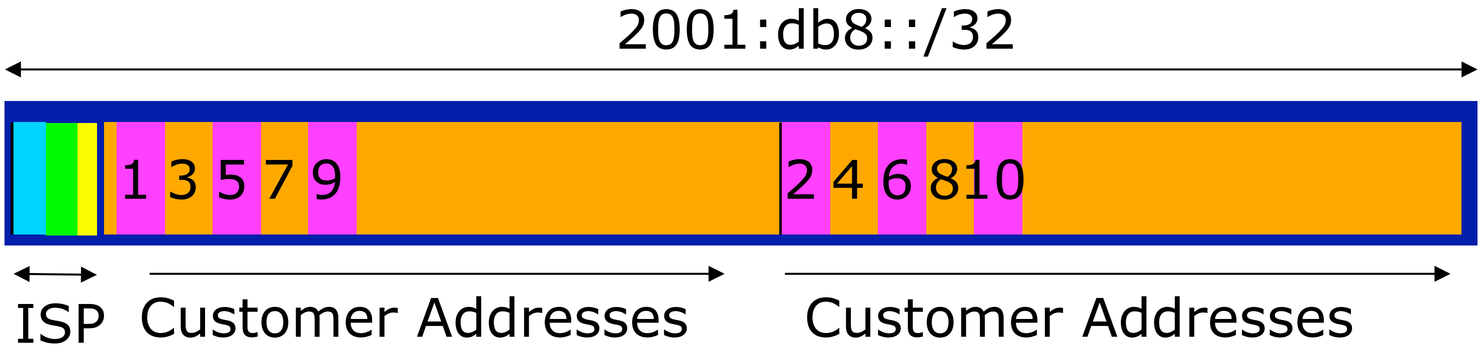

- Scheme then is:
  - First customer from first /33, second customer from second /33, third from first /33, etc
- This works also for residential versus commercial customers:
  - Residential from first /33
  - Commercial from second /33

Planned IP Addressing

- This works fine for multihoming between two upstream links (same or different providers)
- Can also subdivide address space to suit more than two upstreams
  - Follow a similar scheme for populating each portion of the address space
- Consider regional (geographical) distribution of customer delegated address space
- Don’t forget to always announce an aggregate out of each link

Addressing Plans – Advice

- Customer address assignments should not be reserved or assigned on a per PoP basis
  - Follow same principle as for IPv4
  - Subnet aggregate to cater for multihoming needs
  - Consider regional delegation
  - ISP iBGP carries customer nets
  - Aggregation within the iBGP not required and usually not desirable
  - Aggregation in eBGP is very necessary
- Backbone infrastructure assignments:
  - Number out of a single /48
    - Operational simplicity and security
  - Aggregate to minimise size of the IGP

Addressing Plans – Scheme

- Looking at Infrastructure:

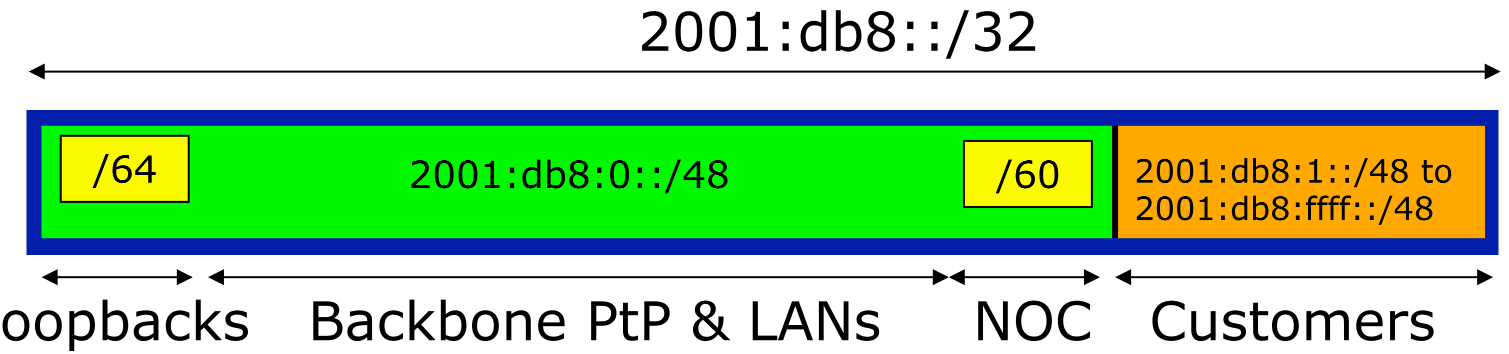

- Alternative:

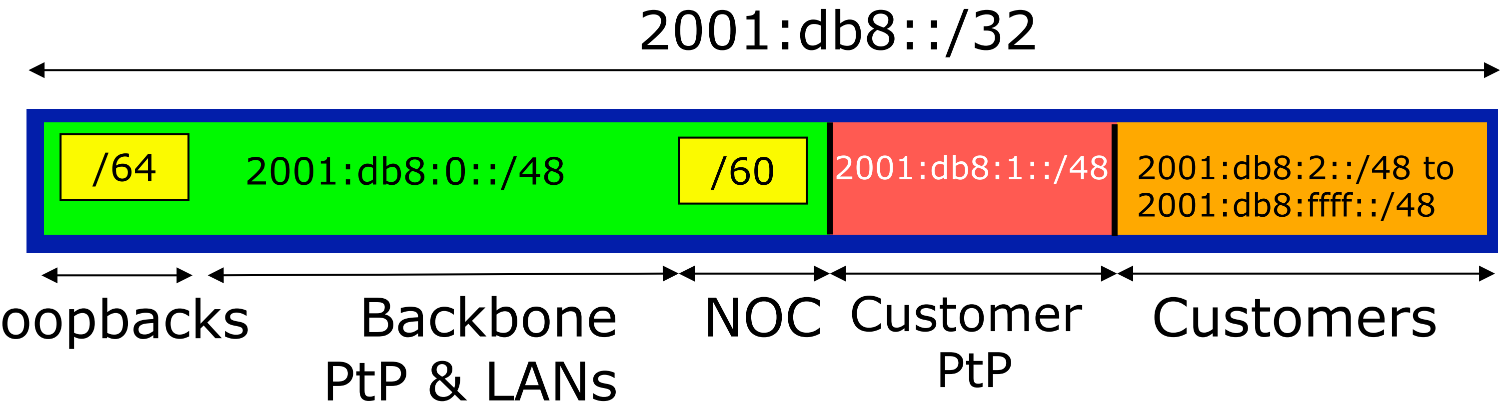

Addressing Plans Planning

- Registries will usually allocate the next block to be contiguous with the first allocation
  - (RIRs use a sparse allocation strategy – industry goal is aggregation)
  - Minimum allocation is /32
  - Very likely that subsequent allocation will make this up to a /31 or larger (/28)
  - 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

Addressing Tools

- Examples of IP address planning tools:
  - NetDot netdot.uoregon.edu (recommended!!)
  - HaCi sourceforge.net/projects/haci
  - Racktables racktables.org
  - IPAT nethead.de/index.php/ipat
  - freeipdb home.globalcrossing.net/~freeipdb/
- Examples of IPv6 subnet calculators:
  - ipv6gen code.google.com/p/ipv6gen/
  - sipcalc www.routemeister.net/projects/sipcalc/

Conclusion

- Presentation has covered:
  - Background of IPv6 – why we are here
  - IPv6 Protocol and Standards status
  - IPv6 Address procurement and address planning

IPv6 Essentials

Philip Smith

SANOG 23

16 January 2014

Thimphu