Introduction to IPv6

Philip Smith
<philip@apnic.net>
APRICOT 2013
Singapore
19th Feb – 1st March 2013

Presentation Slides

- Will be available on
  - http://thyme.apnic.net/ftp/seminars/APRICOT2013-IPv6-Introduction.pdf
  - And on the APRICOT2013 website
- Feel free to ask questions any time

Agenda

- Background
- Protocols & Standards
- Addressing
- Routing Protocols
- Integration & Transition

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 – RFC1719 and RFC1726, Dec 1994

Early Internet History → 1996

- RFC1883 published in December 1995
  - IPv6 Specification
- 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

Current Situation

- General perception is that "IPv6 has not yet taken hold"
  - IPv4 Address run-out has now made it into "headline news"
    - More discussions and run-out plans proposed
  - Private sector still demanding a business case to "migrate"
    - No easy Return on Investment (RoI) computation
- But reality is very different from perception!
  - Something needs to be done to sustain the Internet growth
  - IPv6 or NAT or both or something else?

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 92 /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?

- 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?

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, E-mail clients, IM, bittorrent,…
- Services
  - DNS, Apache WebServer, E-mail gateways,…
- Content Availability
  - Needs to be on IPv4 and on IPv6

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/mtg-0710/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

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?

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

Agenda

- Background
- Protocols & Standards
- Addressing
- Routing Protocols
- Integration & Transition

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 segmentation
  - Path MTU discovery
- 64 bits aligned
- Authentication and Privacy Capabilities
  - IPsec is mandated
- No more broadcast

IPv4 and IPv6 Header Comparison

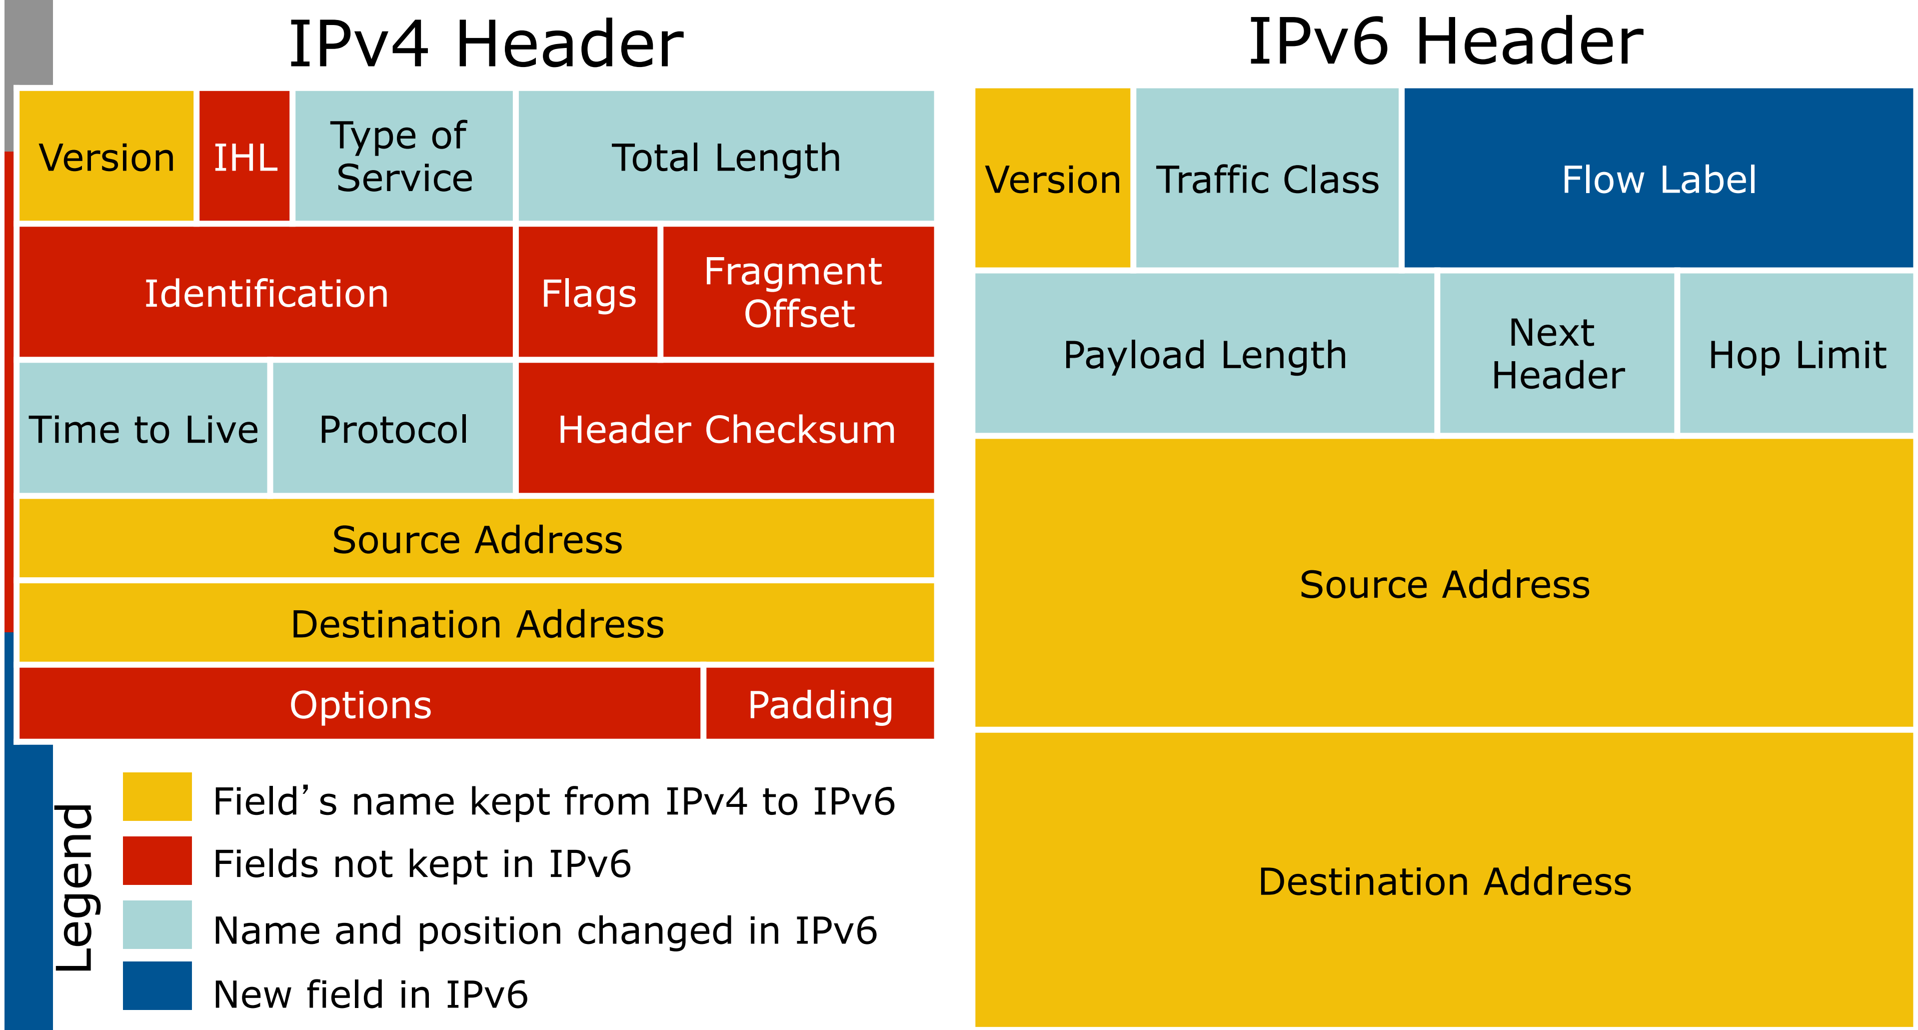

Larger Address Space

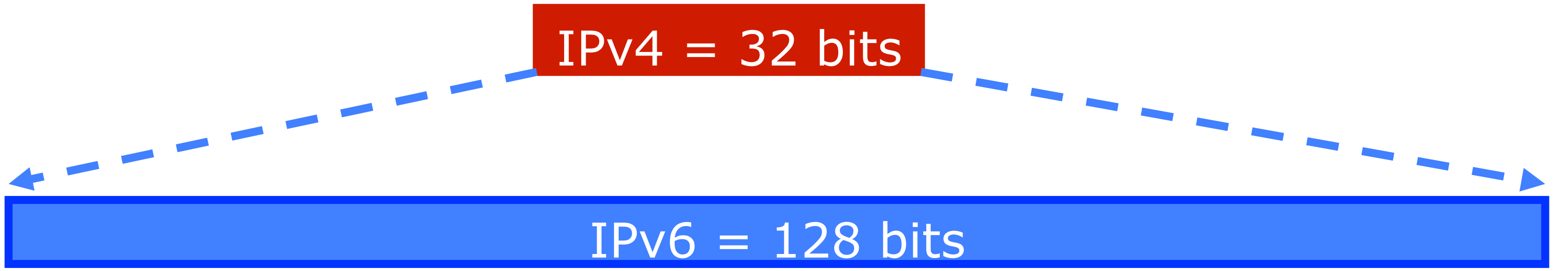

- IPv4
  - 32 bits
  - = 4,294,967,296 possible addressable devices
- IPv6
  - 128 bits: 4 times the size in bits
  - = 3.4 x 10^{38} possible addressable devices
  - = 340,282,366,920,938,463,463,374,607,431,768,211,456
  - ~ 5 x 10^{28} addresses per person on the planet

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

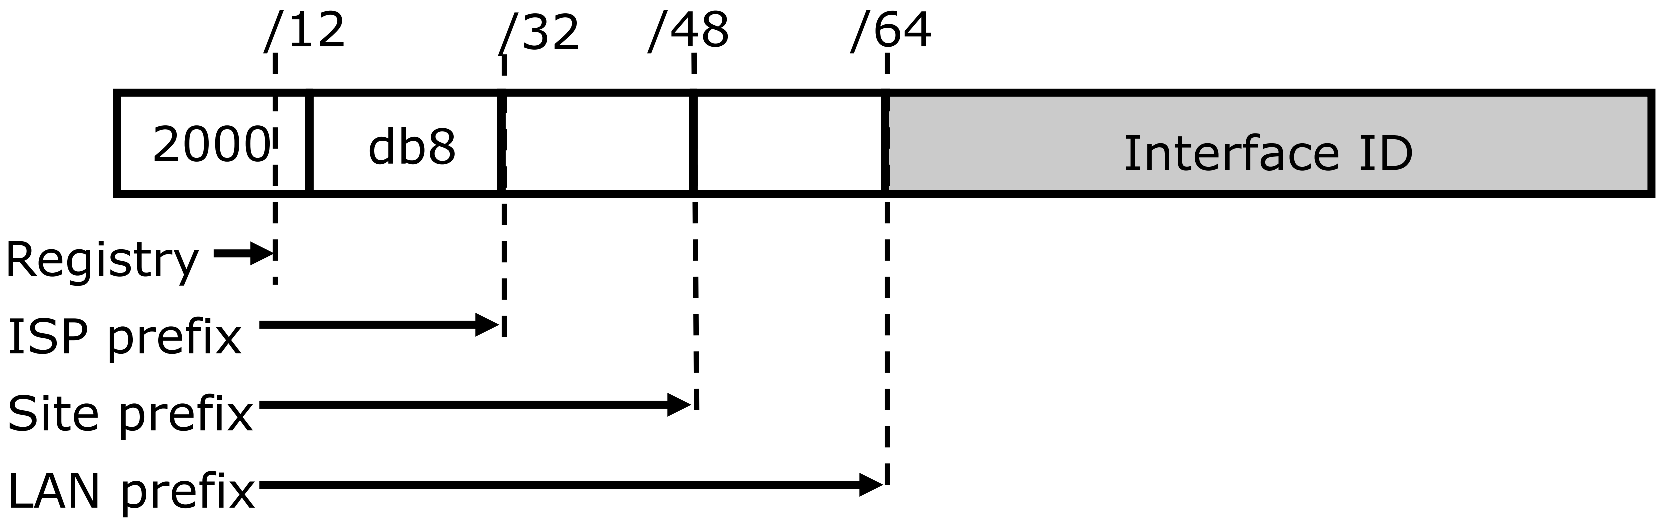

- The allocation process is:
  - The IANA is allocating out of 2000::/3 for initial IPv6 unicast use
  - Each registry gets a /12 prefix from the IANA
  - Registry allocates a /32 prefix (or larger) to an IPv6 ISP
  - Policy is that an ISP allocates a /48 prefix to each end customer

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

Aggregation hopes

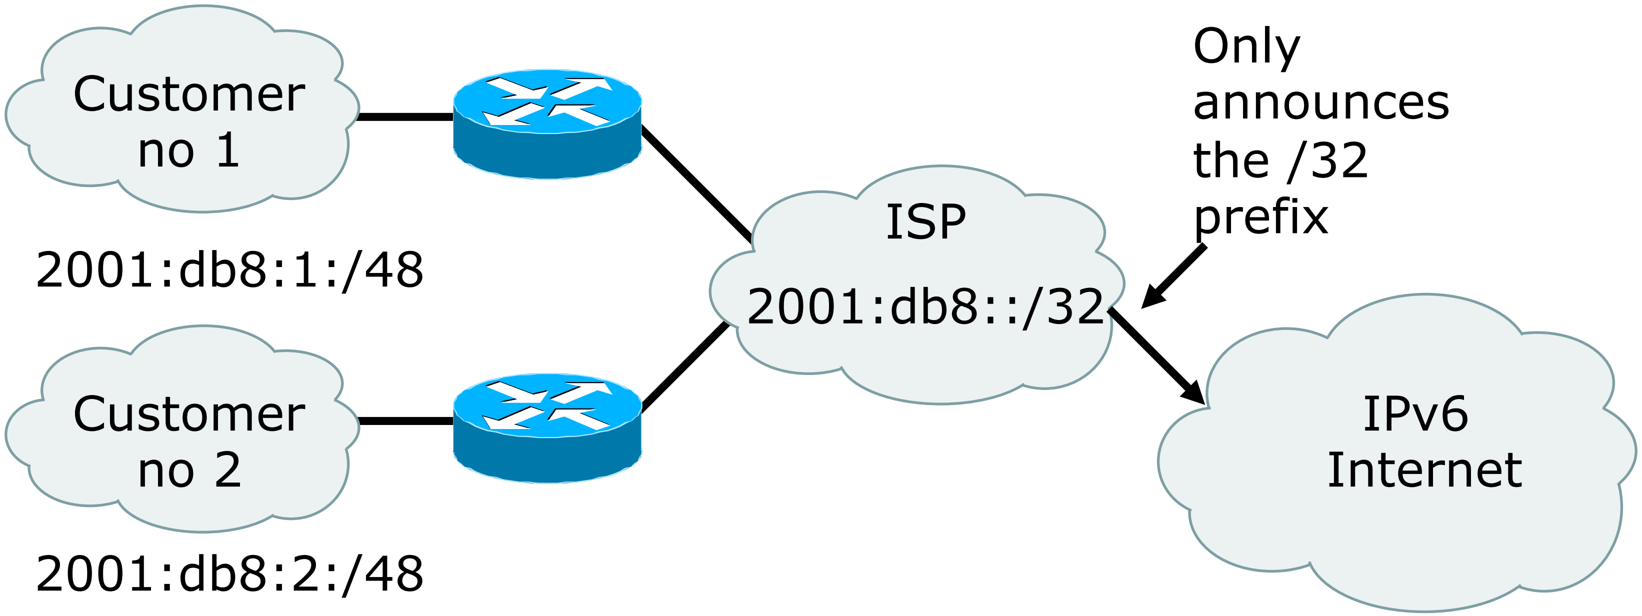

- Larger address space enables aggregation of prefixes announced in the global routing table
- Idea was to allow efficient and scalable routing
- But current Internet multihoming solution breaks this model

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

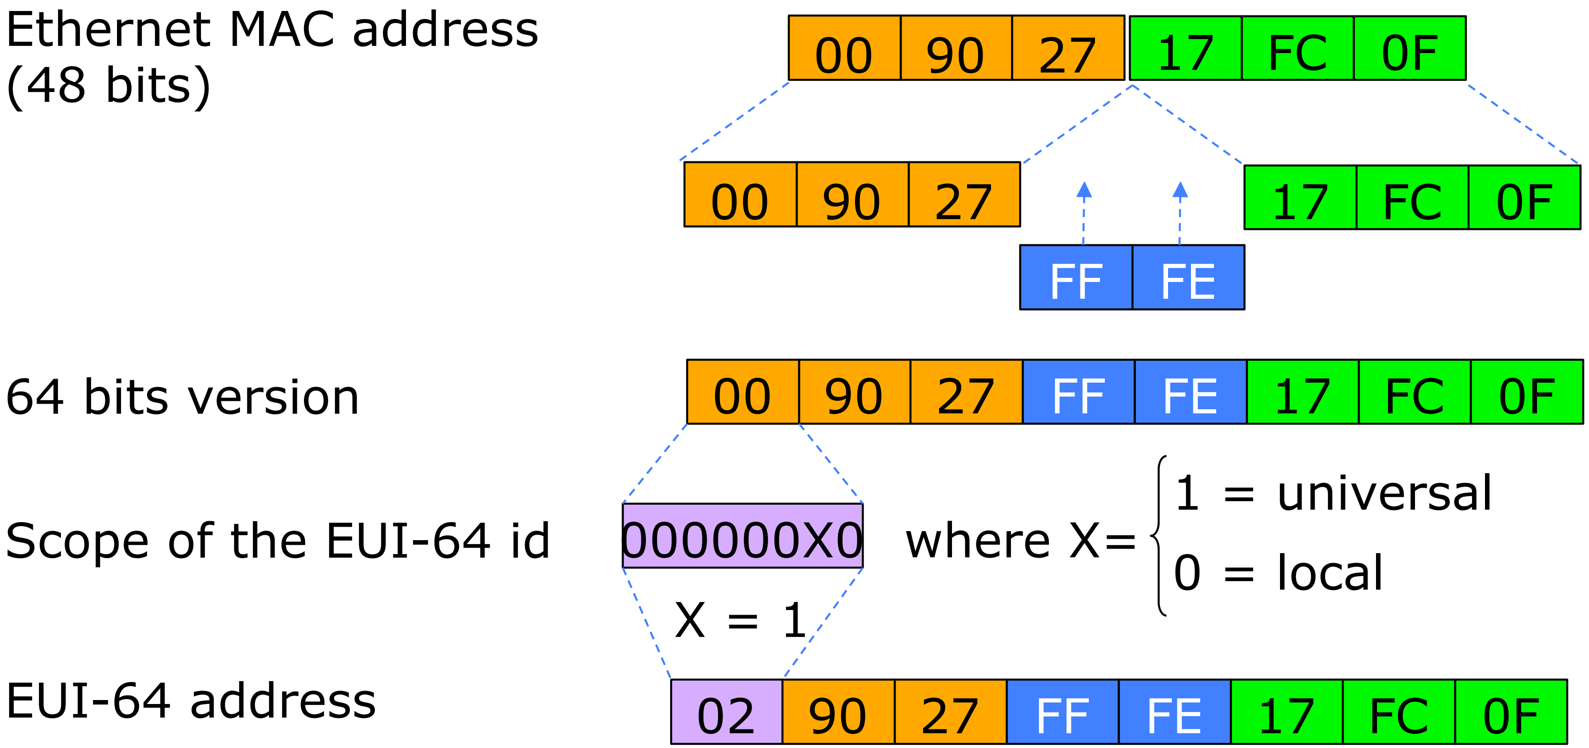

- EUI-64 address is formed by inserting FFFE between the company-id and the manufacturer extension, and setting the "u" bit to indicate scope
  - Global scope: for IEEE 48-bit MAC
  - Local scope: when no IEEE 48-bit MAC is available (eg serials, tunnels)

IPv6 Address Privacy (RFC 4941)

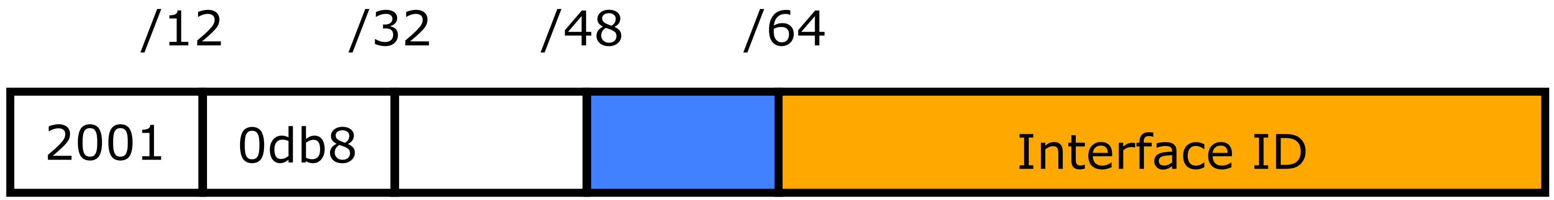

- Temporary addresses for IPv6 host client application, e.g. Web browser
- Intended to inhibit device/user tracking but is also a potential issue
  - More difficult to scan all IP addresses on a subnet
  - But port scan is identical when an address is known
- Random 64 bit interface ID, run DAD before 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

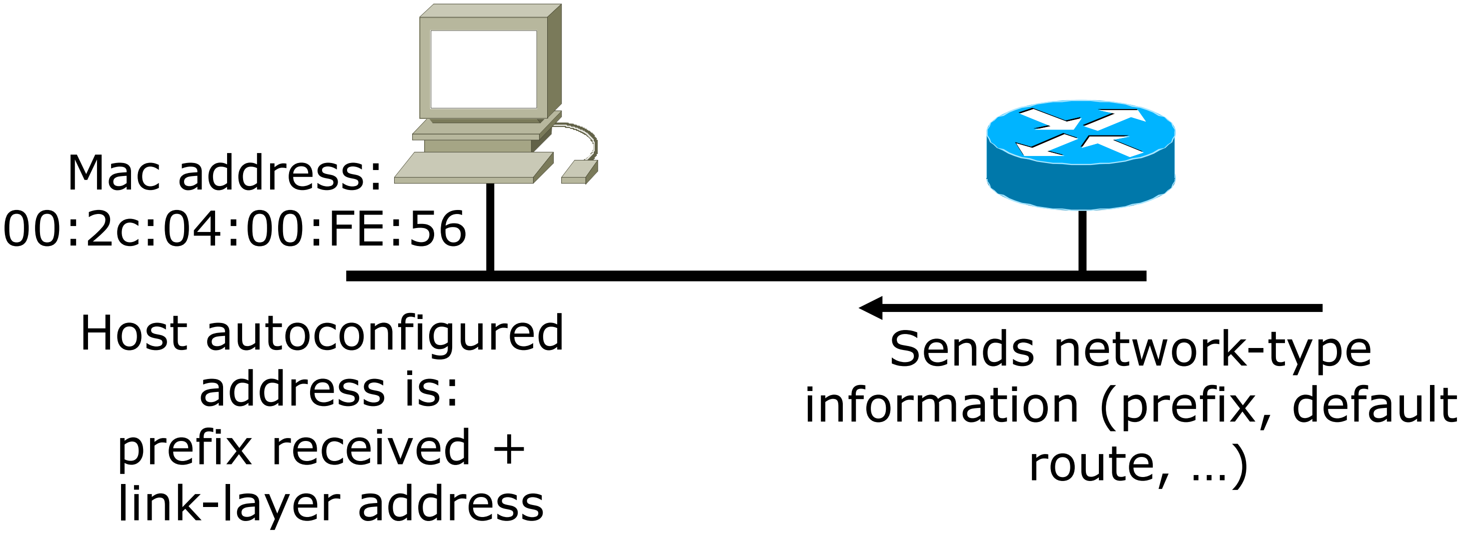

- PC sends router solicitation (RS) message
- Router responds with router advertisement (RA)
  - This includes prefix and default route
  - RFC6106 adds DNS server option
- PC configures its IPv6 address by concatenating prefix received with its EUI-64 address

Renumbering

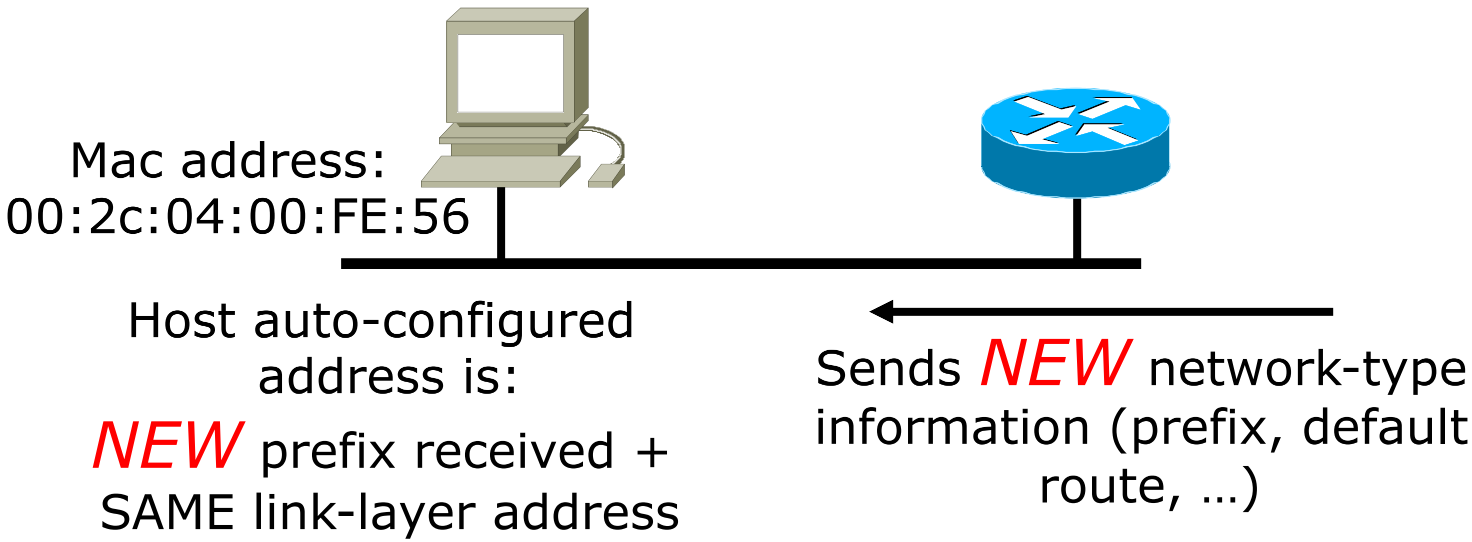

- Router sends router advertisement (RA)
  - This includes the new prefix and default route (and remaining lifetime of the old address)
- PC configures a new IPv6 address by concatenating prefix received with its EUI-64 address
  - Attaches lifetime to old address

Unique-Local

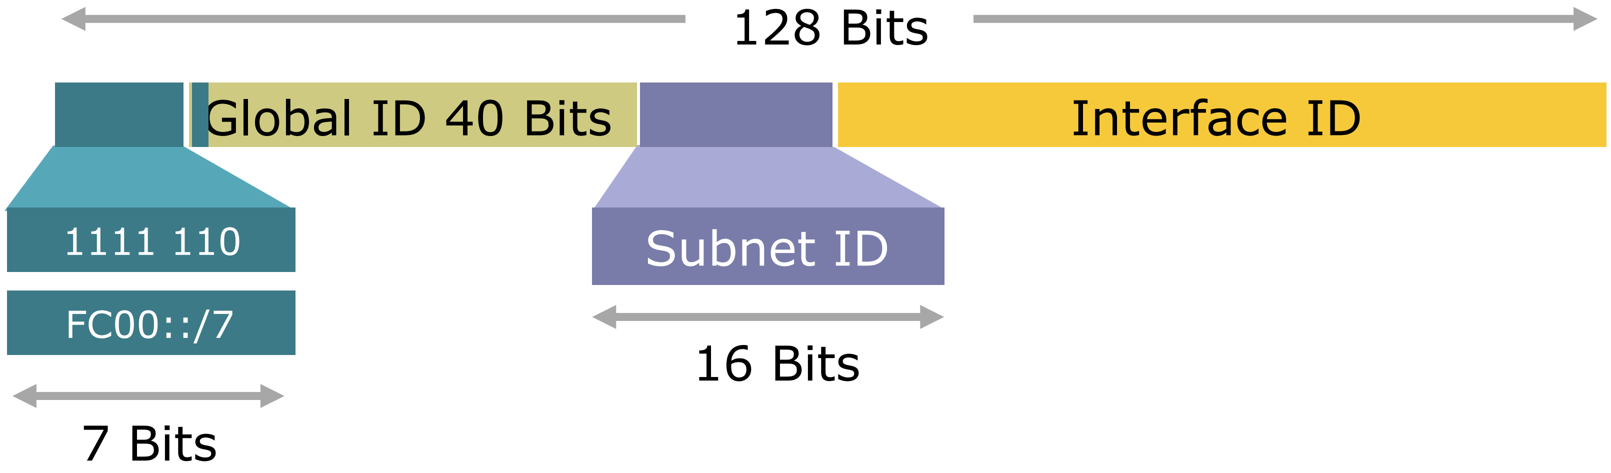

- Unique-Local Addresses Used For:
  - Local communications & inter-site VPNs
  - Local devices such as printers, telephones, etc
  - Site Network Management systems connectivity
- Not routable on the Internet
- Reinvention of the deprecated site-local?

Link-Local

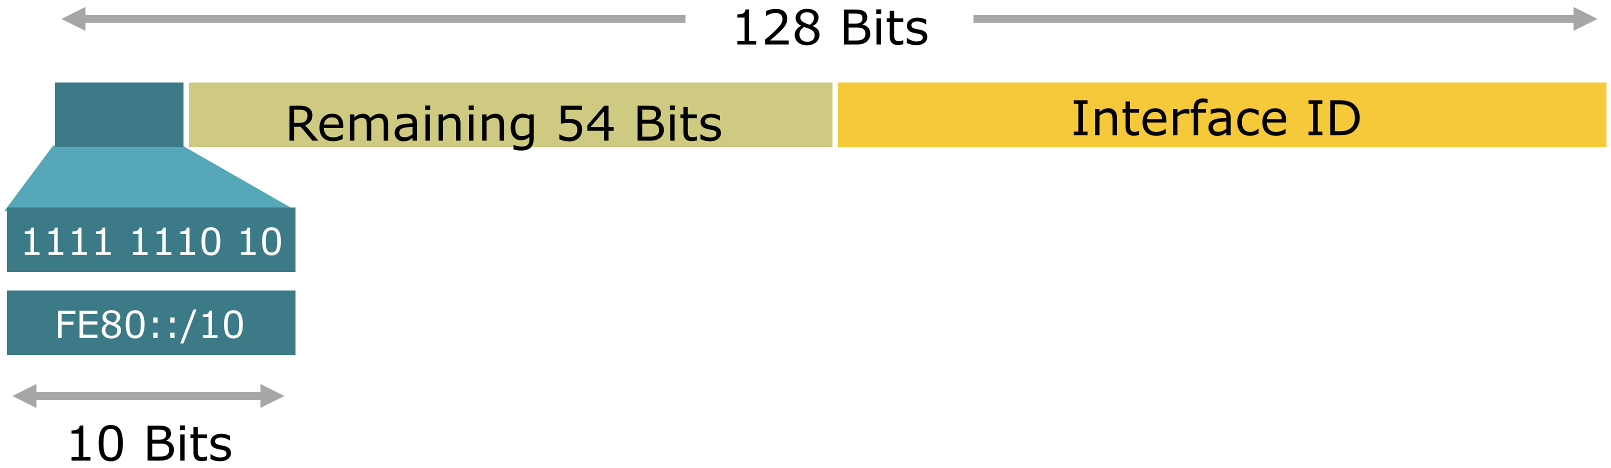

- Link-Local Addresses Used For:
  - Communication between two IPv6 device (like ARP but at Layer 3)
  - Next-Hop calculation in Routing Protocols
- Automatically assigned by Router as soon as IPv6 is enabled
  - Mandatory Address
- Only Link Specific scope
- Remaining 54 bits could be Zero or any manual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

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 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

- Hostname to IP address:

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

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 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)

Agenda

- Background
- Protocols & Standards
- Addressing
- Routing Protocols
- Integration & Transition

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 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

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 – 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 – 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:

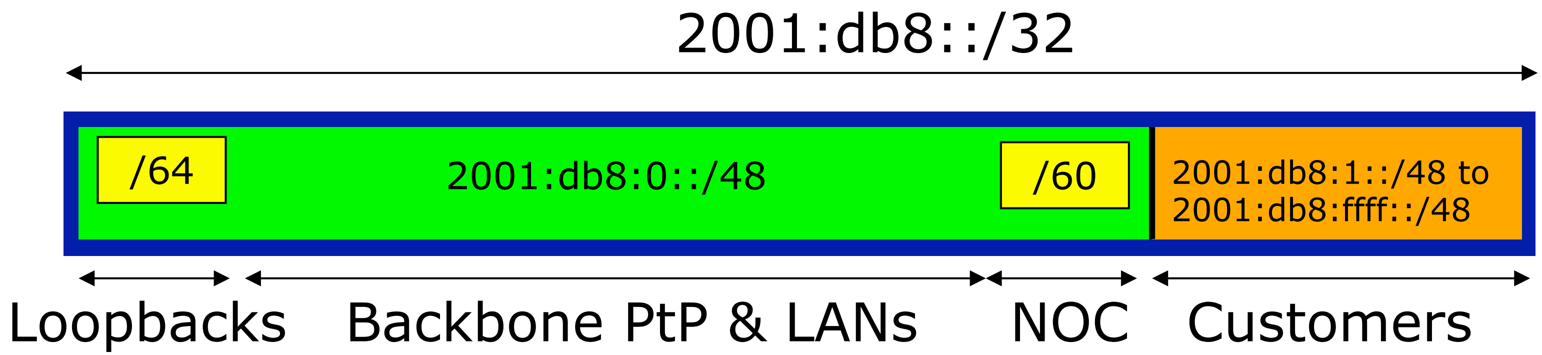

- Alternative:

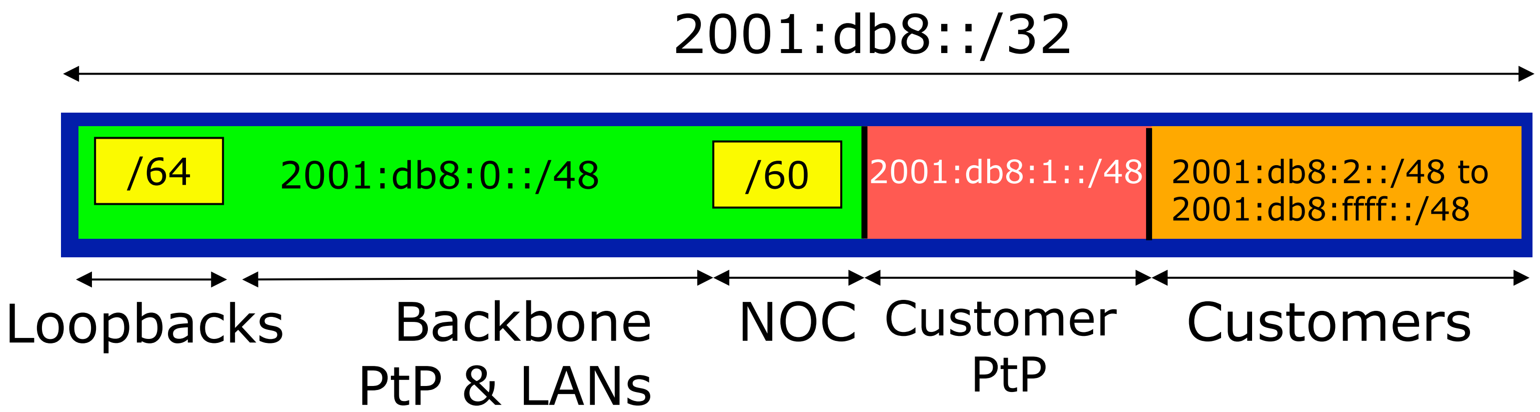

Addressing Plans Planning

- 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 Tools

- Examples of IP address planning tools:
  - NetDot netdot.uoregon.edu (recommended!!)
  - HaCi sourceforge.net/projects/haci
  - 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/

Agenda

- Background
- Protocols & Standards
- Addressing
- Routing Protocols
- Integration & Transition

Static Routing in IPv6

- Unchanged from IPv4
  - Default route is now ::/0
  - On most platforms, the CLI is very similar
- Cisco IOS Static Routing Example:

```
ipv6 route 2001:db8::/64 2001:db8:0:CC::1 110
```

  - Routes packets for network 2001:db8::/64 to a networking device at 2001:db8:0:CC::1 with an administrative distance of 110

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

OSPFv3 overview

- OSPFv3 is OSPF for IPv6 (RFC 5340)
- Based on OSPFv2, with enhancements
- Distributes IPv6 prefixes
- Runs directly over IPv6
- Completely independent of 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

ISIS Standards History

- ISO 10589 specifies the OSI IS-IS routing protocol for CLNS traffic
- RFC 1195 added IPv4 support
  - Also known as Integrated IS-IS (I/IS-IS)
  - I/IS-IS runs on top of the Data Link Layer
- RFC5308 adds IPv6 address family support
- RFC5120 defines Multi-Topology concept
  - Permits IPv4 and IPv6 topologies which are not identical
  - Permits roll out of IPv6 without impacting IPv4 operations

IS-IS for IPv6

- 2 TLVs added to introduce IPv6 routing
  - IPv6 Reachability TLV (0xEC)
  - IPv6 Interface Address TLV (0xE8)
- 4 TLVs added to support multi-topology ISIS
  - Multi Topology
  - Multi Topology Intermediate Systems
  - Multi Topology Reachable IPv4 Prefixes
  - Multi Topology Reachable IPv6 Prefixes
- Multi Topology IDs
  - #0 – standard topology for IPv4/CLNS
  - #2 – topology for IPv6

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)

Agenda

- Background
- Protocols & Standards
- Addressing
- Routing Protocols
- Integration & Transition

IPv4-IPv6 Co-existence/Transition

- A wide range of techniques have been identified and implemented, basically falling into three categories:
  - Dual-stack techniques, to allow IPv4 and IPv6 to co-exist in the same devices and networks
  - Tunneling techniques, to avoid dependencies when upgrading hosts, routers, or regions
  - Translation techniques, to allow IPv6-only devices to communicate with IPv4-only devices
- All of these are being used, in combination

Dual Stack Approach

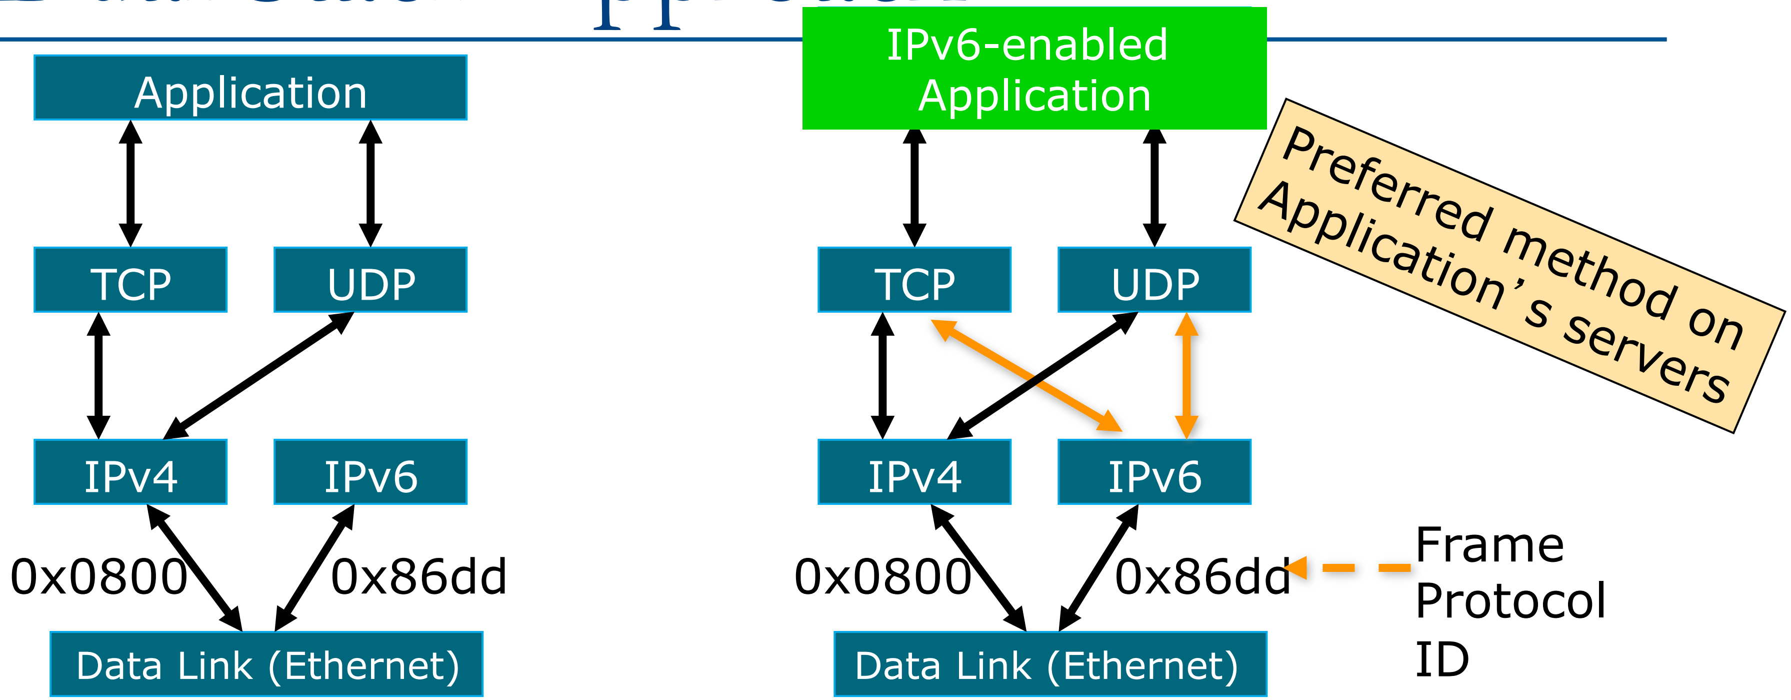

- Dual stack node means:
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on name lookup and application preference

Strategies available for Service Providers

1. Do nothing
   - Wait and see what competitors do
   - Business not growing, so don't care what happens
2. Extend life of IPv4
   - Force customers to NAT
   - Buy IPv4 address space on the marketplace
3. Deploy IPv6
   - Dual-stack infrastructure
   - IPv6 and NATed IPv4 for customers
   - 6rd (Rapid Deploy) with native or NATed IPv4 for customers
   - Or various other combinations of IPv6, IPv4 and NAT

1. Doing Nothing

- Advantages
  - Easiest and most cost effective short term strategy
- Disadvantages
  - Limited to IPv4 address availability (RIRs or marketplace)
  - No access to IPv6
  - Negative public perception of SP as a laggard
  - Strategy will have to be reconsidered once IPv4 address space is no longer available

2. Extending life of IPv4 Network

- Two ways of extending IPv4 network
  - Next step along from "Strategy One: Do nothing"
- Force customers to use NAT
  - Customers moved to RFC1918 address space
  - SP infrastructure moved to RFC6598 address space (or use RFC1918 where feasible)
- Acquire IPv4 address space from another organisation
  - IPv4 subnet trading

3. IPv4/IPv6 coexistence & transition

- Three strategies for IPv6 transition:
  - Dual Stack Network
    - The original strategy
    - Depends on sufficient IPv4 being available
  - 6rd (Rapid Deploy)
    - Improvement on 6to4 for SP customer deployment
  - Large Scale NAT (LSN)
    - SP deploys large NAT boxes to do address and/or protocol translation
    - Typically NAT444, Dual-Stack Lite and NAT64

Conclusions
Potential Scenarios

- Most of the content and applications move to IPv6 only;
- Most of the content and applications are offered for IPv4 and IPv6;
- Most of the users move to IPv6 only
  - Especially mobile operators offering LTE handsets in emerging countries
- No change (the contents/applications stay IPv4 and absence of pro-IPv6 regulation), SP customer expectations devolve to double-NAT;
- No change (the contents/applications stay IPv4) but SP customer expectations do not devolve to double-NAT (or they are ready to pay for peer-to-peer connectivity).
  - Perhaps well established broadband markets like US or Europe

Conclusions
Potential Techniques

Scenario	Potential Techniques
Content and Applications move to IPv6	IPv6 only network; Dual-Stack, 6rd and DS-lite as migration techniques
Content and Applications on IPv4 and IPv6	Dual-Stack (if enough IPv4) or 6rd; SP IPv4-NAT; DS-lite (for greenfield) *
Users are IPv6 only	Stateful/Stateless AFT to get to IPv4 content *
No change (double NAT)	SP IPv4-NAT *
No change (no double NAT)	Do nothing *

*** Transfer Market applicable**

Recommendations

1. Start deploying IPv6 as long term strategy
2. Evaluate current addressing usage to understand if IPv4 to IPv4 NAT is sufficient for transition period
3. Prepare a translation mechanism from the IPv4 Internet to the IPv6 Internet
4. Educate your user base on IPv6 introduction, the use cases and troubleshooting