

Introduction to IPv6
&
Service Provider Migration
SANOG 24
3rd – August 2014



Srinath Beldona
Senior Technical Consultant (Training &
Development)
srinath@apnic.net



Agenda

- ❑ Why do we need IPv6? (Background and History)
- ❑ IPv6 Protocol & IPv6 Standards
- ❑ IPv6 Addressing
- ❑ IPv6 Routing
- ❑ Service Provider IPv4 – IPv6 Coexistence
- ❑ Conclusion



Pre-requisites

- Good understanding of TCP/IP
- Good knowledge of Subnetting
- Intermediate level knowledge of Routing
 - Static Routing
 - OSPFv2 and ISIS for IPv4
 - Multiprotocol BGP

Introduction to IPv6



SANOG 22

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**

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 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?

Do we really need a larger address space?

- Large variety of proposals to “help” with IPv6 deployment
 - NAT444
 - IPv4 NAT in Core and Edge
 - Dual Stack Lite and 464XLAT
 - Running IPv4 over and IPv6 backbone
 - Activity of IETF Softwires and v6ops Working Groups
 - NAT64
 - Translation between IPv6 and IPv4
 - Activity of IETF Behave Working Group
 - 6rd
 - Dynamic IPv6 tunnel from SP to customer
 - Activity of IETF Softwires Working Group

IPv6 Geo-Politics

- Regional and Countries IPv6 Task Force
 - Europe – www.ipv6-taskforce.org/
 - Belgium, France, Spain, Switzerland, UK,...
 - North-America – www.nav6tf.org/
 - Japan IPv6 Promotion Council – www.v6pc.jp/en/index.html
 - China, Korea, India,...
- Relationship
 - Economic partnership between governments
 - China-Japan, Europe-China,...
- Recommendations and project's funding
 - IPv6 2005 roadmap recommendations – Jan. 2002
 - European Commission IPv6 project funding: 6DEPLOY & Euro6IX
- Tax Incentives
 - Japan only – 2002-2003 program

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

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 are we still waiting...?

- That killer application?
 - Internet Gaming or Peer to Peer applications?
- IPv4 to run out?
 - Too late, it has!
- Our competitors?
 - Any network deployed in last 3 years will be IPv6 capable
 - Even if not enabled!
- The end-user?
 - The end-user should not have to choose protocols
 - Remember “Turbo” button on early IBM PC clones?

The On-going Debate (1)

- IPv6 Multihoming
 - Same toolset as IPv4 — long term non-scalable
 - ‘Ultimate Multihoming Solution’ no nearer discovery
 - LISP is making some progress though
- Early rigid IPv6 address allocation model
 - “One size fits all” barrier to deployment:
 - Only ISPs “should” get IPv6 space from RIRs
 - Enterprises “should” get IPv6 space from ISPs only
 - Routing table entries matter, not the nature of business
 - What is an ISP?


The On-going Debate (2)

- Not every IPv4 device is IPv6 capable
 - Do we really need to replicate all IPv4 capability in IPv6 prior to considering deployment?
- “We have enough IPv4”
 - Those with plenty denying those with little/nothing
- Migration versus Co-existence
 - Realistically IPv6 and IPv4 will co-exist for many years
 - Dual-stack operating systems in network equipment makes this trivial

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?

“The times, They are a’ changin’”



Is IPv4 really running out?

- Yes!
 - IANA IPv4 free pool ran out on 3rd February 2011
 - RIR IPv4 free pool will run out soon after
 - www.potaroo.net/tools/ipv4/
 - (depends on RIR soft-landing policies)
- The runout gadgets and widgets are now watching when the RIR pools will run out:
 - inetcore.com/project/ipv4ec/index_en.html
 - ipv6.he.net/statistics/



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)

Issues Today

- Minimal content is available on IPv6
 - Notwithstanding ipv6.google.com
 - World IPv6 Day on 8th June 2011 helped a little
 - World IPv6 Launch on 6th June 2012 helped a little more
- Giving IPv6 to customers might confuse
 - Browsers, e-mail clients, etc are smart
 - But increased tech support if IPv6 version of content is 'down', but IPv4 version works
- Need to “prolong” IPv4 so there is time for all content to be available on IPv6

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

IPv6 Protocol & IPv6 Standards



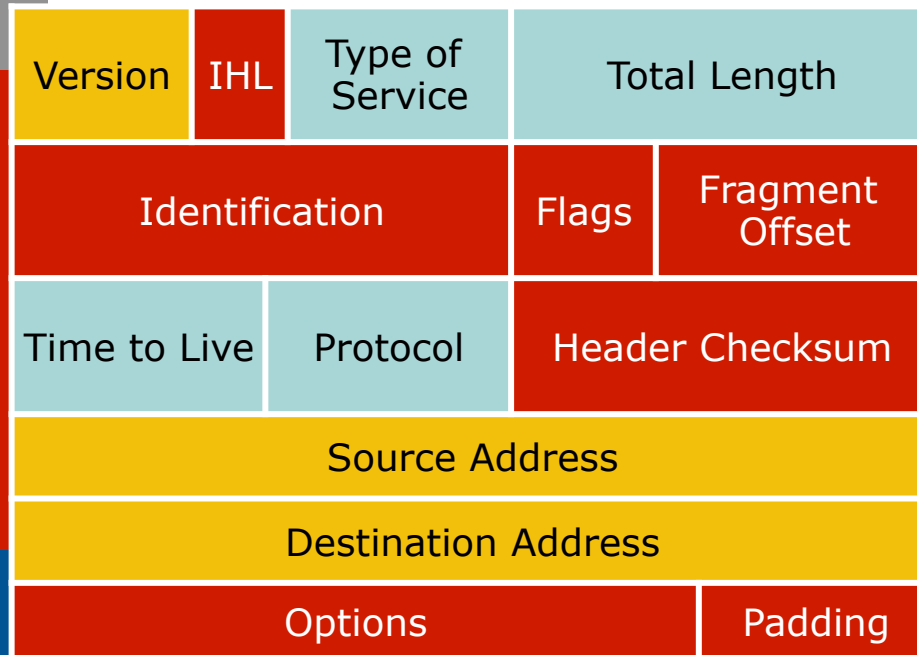
SANOG 24

So what has really changed?

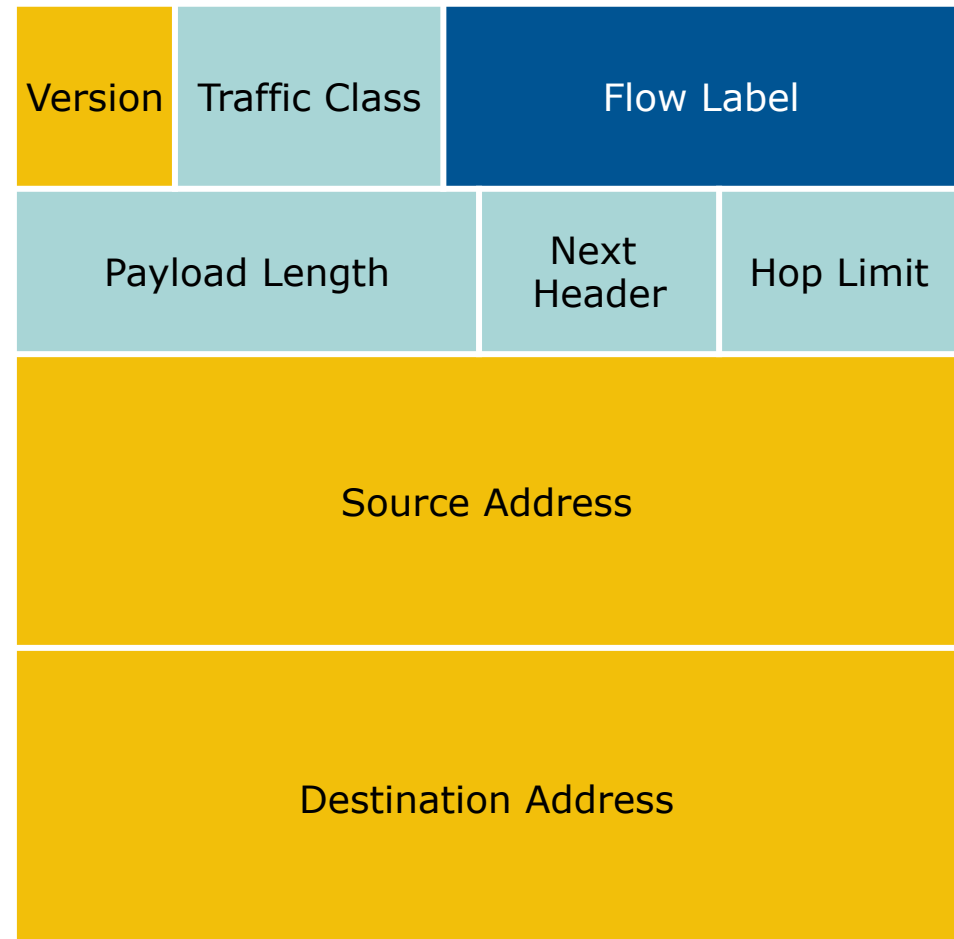
- Expanded address space
 - Address length quadrupled to 16 bytes
- Header Format Simplification
 - Fixed length, optional headers are daisy-chained
 - IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- No checksum at the IP network layer
- No hop-by-hop fragmentation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

IPv4 and IPv6 Header Comparison

IPv4 Header



IPv6 Header



Legend

- Field's name kept from IPv4 to IPv6
- Fields not kept in IPv6
- Name and position changed in IPv6
- New field in IPv6

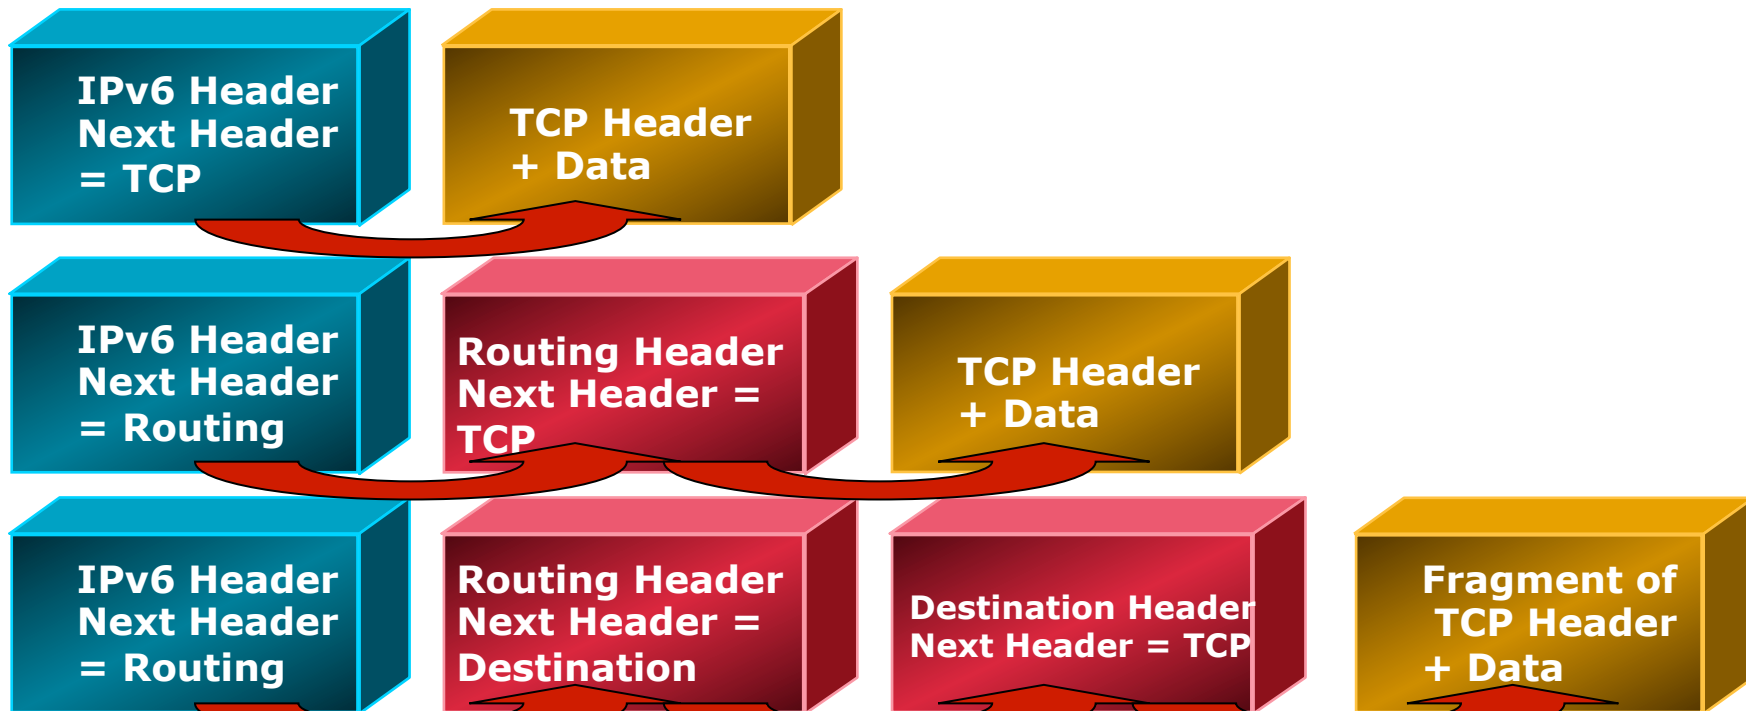
IPv6 Header

- ❑ 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



- ❑ 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
 - = 4,294,967,296 possible addressable devices
- IPv6
 - 128 bits: 4 times the size in bits
 - = 3.4×10^{38} possible addressable devices
 - = 340,282,366,920,938,463,463,374,607,431,768,211,456
 - $\sim 5 \times 10^{28}$ 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](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)



64 bits version



Scope of the EUI-64 id

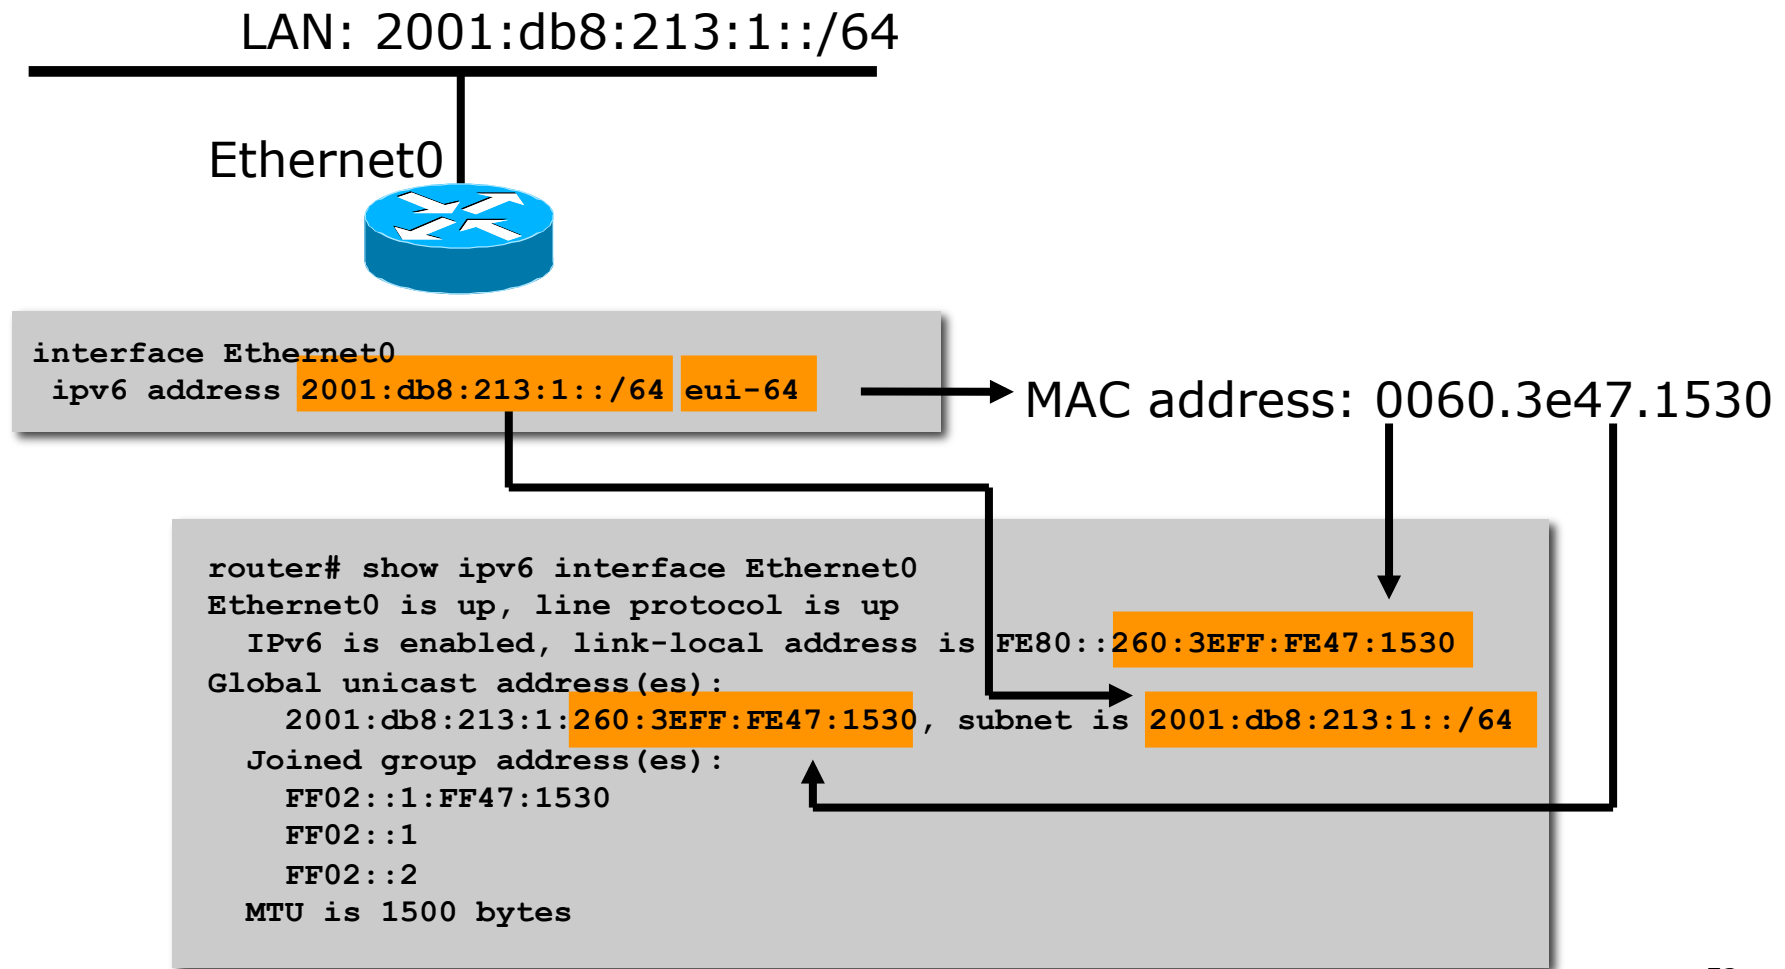


EUI-64 address




- 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, tunnels⁵¹)

IPv6 Addressing Examples



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

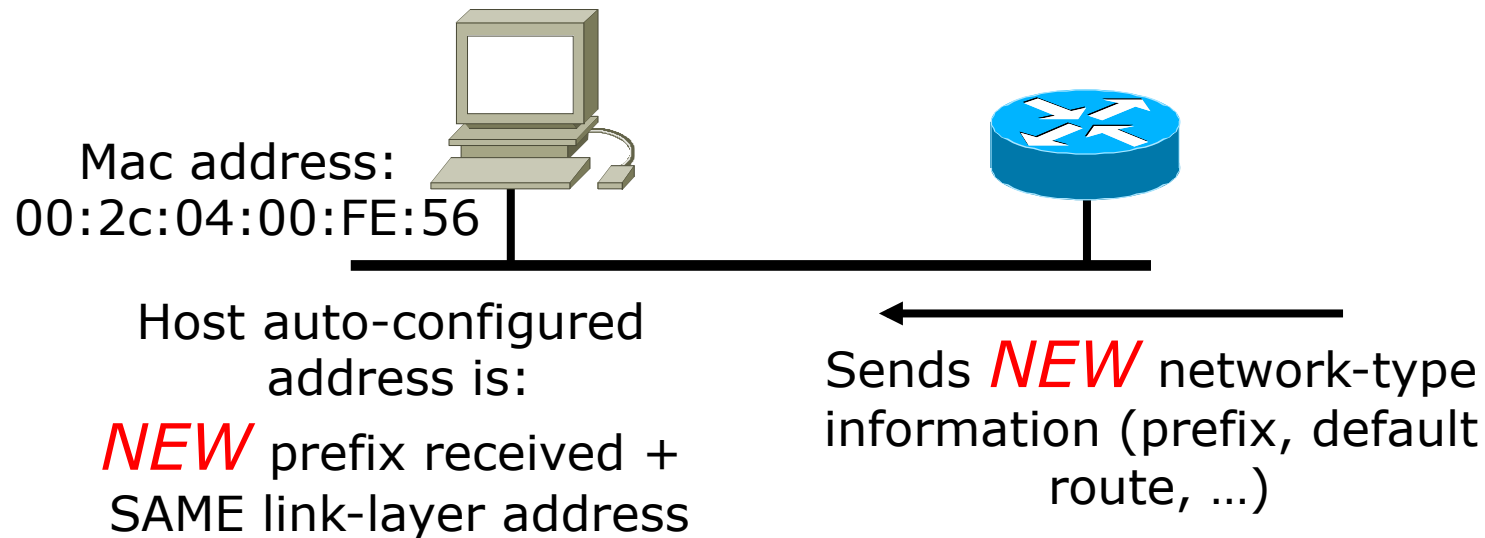
- **No known implementation!**

Auto-configuration



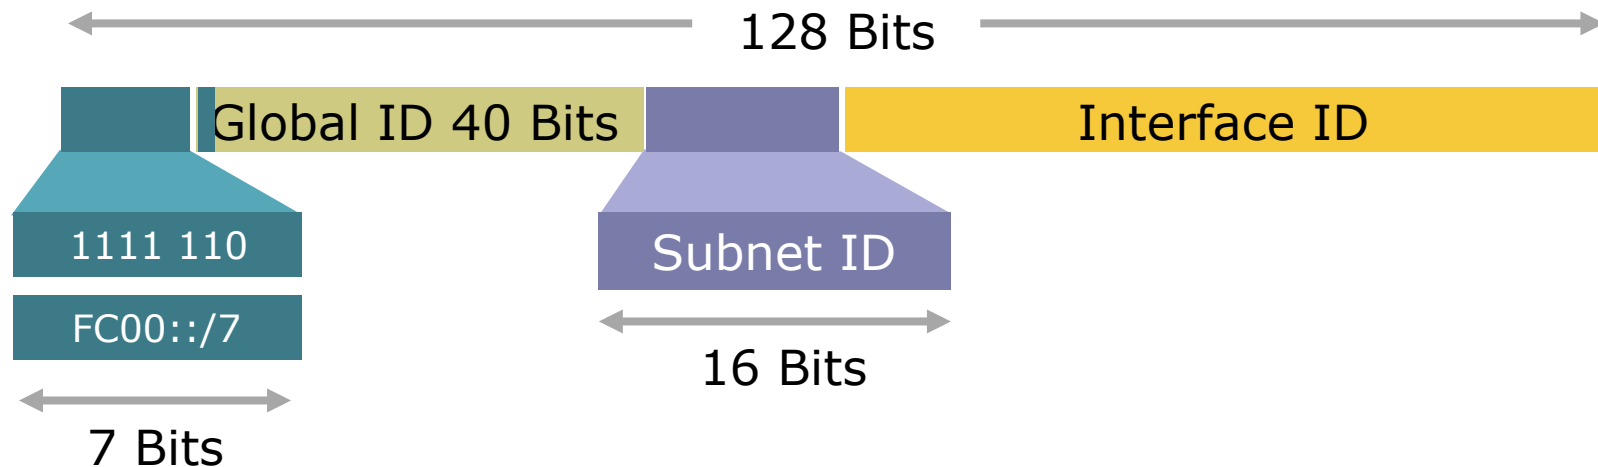
- ❑ PC sends router solicitation (RS) message
- ❑ Router responds with router advertisement (RA)
 - This includes prefix and default route
 - RFC6106 adds DNS server option
- ❑ PC configures its IPv6 address by concatenating prefix received with its EUI-64 address

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**

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 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.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 coarse-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



Conclusion

- Protocol is “ready to go”
- The core components have already seen several years field experience

IPv6 Addressing

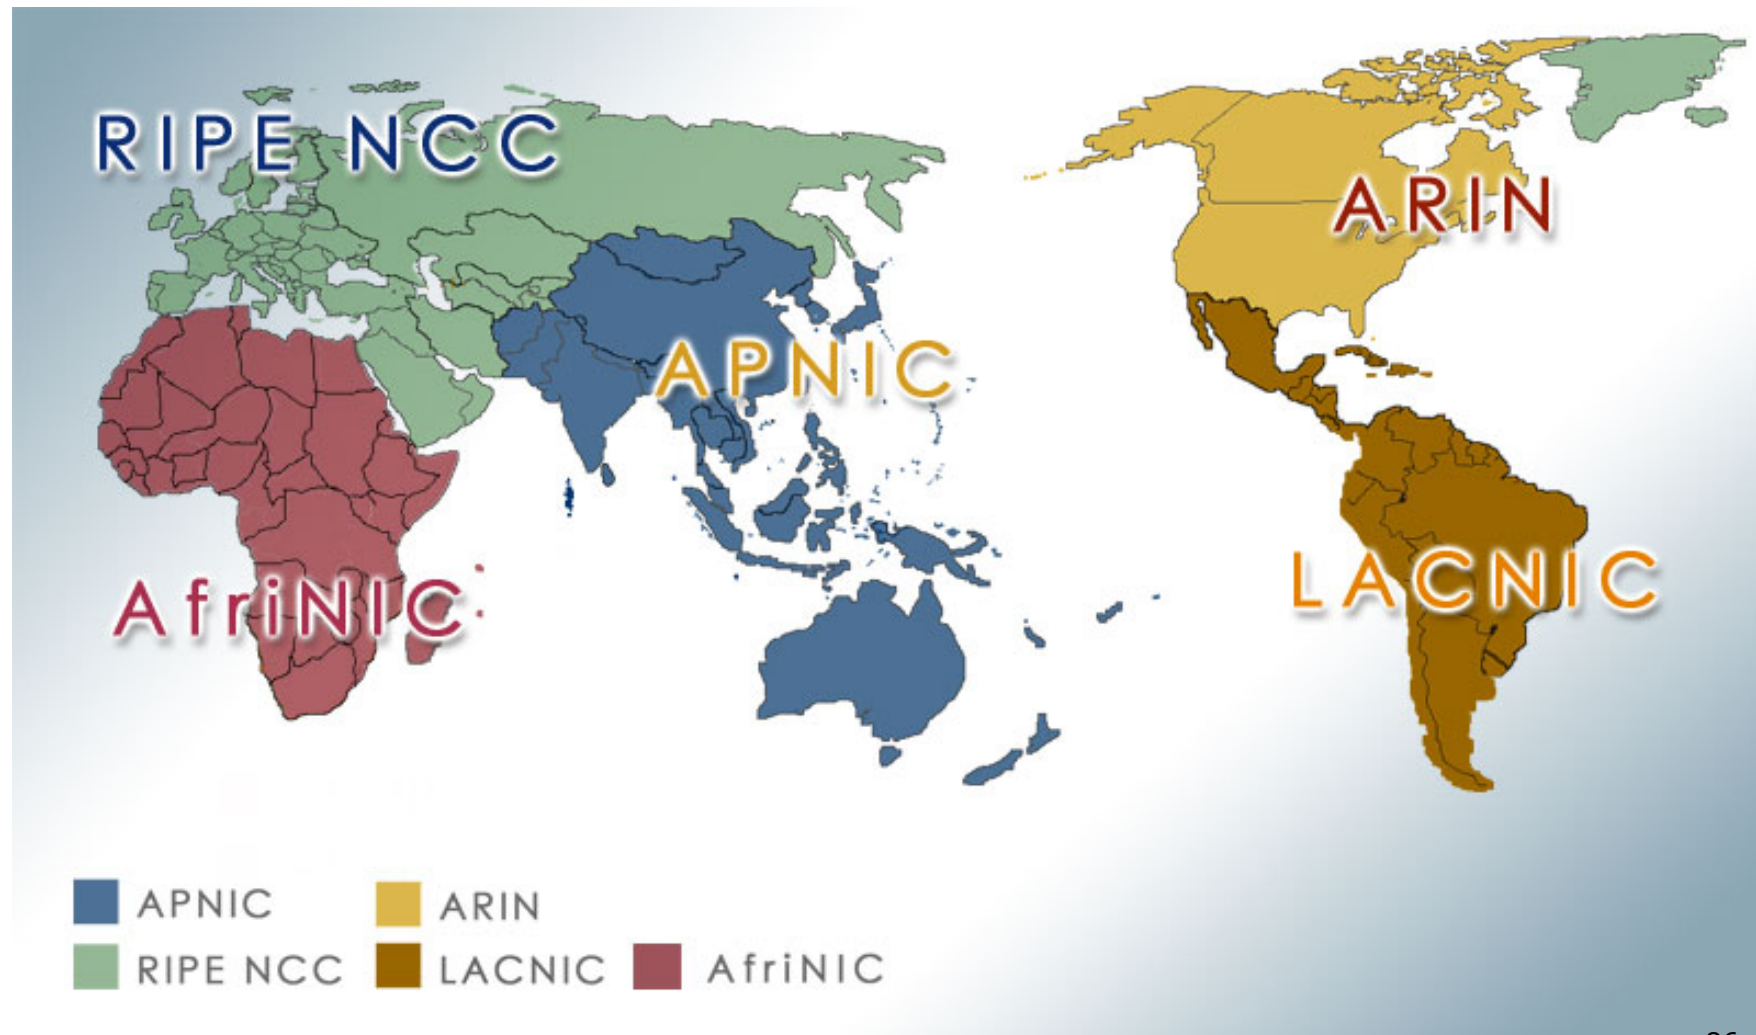


SANOG 22

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
 - 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/

IPv6 Routing Protocols



SANOG 22

IPv6 Configuration on Cisco IOS

- To enable IPv6 the following global commands are required:

- `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
```

- EUI-64 is not helpful on a router and is not recommended

IPv6 Configuration

- ❑ If no global IPv6 address is required on an interface, yet it needs to carry IPv6 traffic:
 - Enable IPv6 on that interface using:
`Router(config-if)# ipv6 enable`
 - Which 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

```
Router1# conf t
Router1(config)# ipv6 unicast-routing
Router1(config)# ipv6 cef
Router1(config)# int fast 0/0
Router1(config-int)# ipv6 enable
Router1(config-int)# ^Z
```

```
Router1#sh ipv6 interface fast 0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8B9:C0FF:FE00:F11D
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:F11D
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
```

IPv6 Configuration – EUI64

```
Router1#sh ipv6 interface fast 0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8B9:C0FF:FE00:F11D
  Global unicast address(es):
    2001:DB8::A8B9:C0FF:FE00:F11D, subnet is 2001:DB8::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:F11D
  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

```
Router1#sh ipv6 int fast 0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8B9:C0FF:FE00:F11D
  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:FF00:F11D
  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.
```



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:ABCD::1 150
```

- Routes packets for network 2001:db8::/64 to a networking device at 2001:DB8:0:ABCD::1 with an administrative distance of 150

Default Routing Example



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

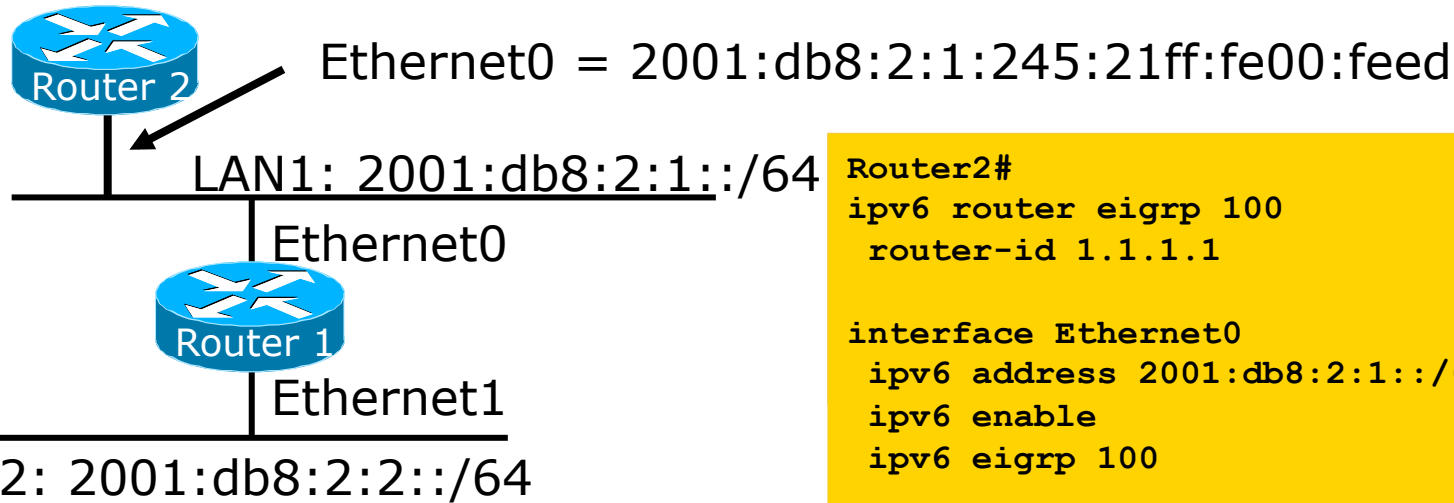
- ❑ Cisco EIGRP has had IPv6 protocol support added
 - Just another protocol module (IP, IPX, AppleTalk) with three new TLVs:
 - ❑ IPv6_REQUEST_TYPE (0X0401)
 - ❑ IPv6_METRIC_TYPE (0X0402)
 - ❑ IPv6_EXTERIOR_TYPE (0X0403)
 - Router-ID is still 32-bit, protocol is still 88
- ❑ Uses similar CLI to existing IPv4 protocol support
- ❑ Easy deployment path for existing IPv4 EIGRP users
- ❑ In Cisco 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 1.1.1.1

interface Ethernet0
ipv6 address 2001:db8:2:1::/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::245:21ff:fe00:feed E0 14 00:01:43 1 4500 0 1
```

```
Router1#show ipv6 eigrp topology all links
IPv6-EIGRP Topology Table for AS(100)/ID(1.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status
P 2001:db8:2:1::/64, 1 successors, FD is 28160, serno 1
via Connected, Ethernet0
via FE80::245:21ff:fe00:feed (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
```


NB: Router-id only required in
absence of any ipv4 configuration



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

Cisco IOS IS-IS dual stack 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::a/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::a/64  
  ip router isis  
  ipv6 router isis  
  
router isis  
  net 42.0001.0000.0000.072c.00  
  metric-style wide
```

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:db8:2:1::f/64  
!  
router bgp 65001  
  bgp router-id 10.10.10.1  
  no bgp default ipv4-unicast  
  neighbor 2001:db8:2:1::1 remote-as 65002  
  address-family ipv6  
    neighbor 2001:db8:2:1::1 activate  
    neighbor 2001:db8:2:1::1 prefix-list bgp65002in in  
    neighbor 2001:db8:2:1::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 separate presentations

SP IPv4 – IPv6 Coexistence



SANOG 22

How can Network Operators face IPv4- Address Exhaustion?




A Review of IPv4-IPv6 Co-Existence
Techniques

Introduction



Why should we care?

“The times, They are a’ changin’”



Is IPv4 really running out?

- Yes!
 - IANA IPv4 free pool ran out on 3rd February 2011
 - RIR IPv4 free pool will run out soon after
 - www.potaroo.net/tools/ipv4/
 - (depends on RIR soft-landing policies)
- The runout gadgets and widgets are now watching when the RIR pools will run out:
 - inetcore.com/project/ipv4ec/index_en.html
 - ipv6.he.net/statistics/



Strategies available for Service Providers

- Do nothing
 - Wait and see what competitors do
 - Business not growing, so don't care what happens
- Extend life of IPv4
 - Force customers to NAT
 - Buy IPv4 address space on the marketplace
- Deploy IPv6
 - Dual-stack infrastructure
 - IPv6 and NATed IPv4 for customers
 - 6rd (Rapid Deploy) with native or NATed IPv4 for customers
 - 464XLAT with native IPv6 and NATed IPv4 for customers
 - Or other combinations of IPv6, IPv4 and NAT

Definition of Terms



Dual-Stack Networks

- Both IPv4 and IPv6 have been fully deployed across all the infrastructure
 - Routing protocols handle IPv4 and IPv6
 - Content, application, and services available on IPv4 and IPv6
- End-users use dual-stack network transparently:
 - If DNS returns IPv6 address for domain name query, IPv6 transport is used
 - If no IPv6 address returned, DNS is queried for IPv4 address, and IPv4 transport is used instead
- It is envisaged that the Internet will operate dual-stack for many years to come

IP in IP Tunnels

- A mechanism whereby an IP packet from one address family is encapsulated in an IP packet from another address family
 - Enables the original packet to be transported over network of another address family
- Allows ISP to provide dual-stack service prior to completing infrastructure deployment
- Tunnelling techniques include:
 - IPinIP, GRE, 6to4, Teredo, ISATAP, 6rd, MPLS

Address Family Translation (AFT)

- Refers to translation of an IP address from one address family into another address family
 - e.g. IPv6 to IPv4 translation
 - Usually called NAT64
 - Or IPv4 to IPv6 translation
 - Usually called NAT46, usually using SIIT

Network Address Translation (NAT)

- ❑ NAT is translation of one IP address into another IP address
- ❑ NAT (Network Address & Port Translation) translates multiple IP addresses into one other IP address
 - TCP/UDP port distinguishes different packet flows
- ❑ NAT-PT (NAT – Protocol Translation) is a particular technology which does protocol translation in addition to address translation
 - NAT-PT is has now been made obsolete by the IETF

Carrier Grade NAT (CGN)

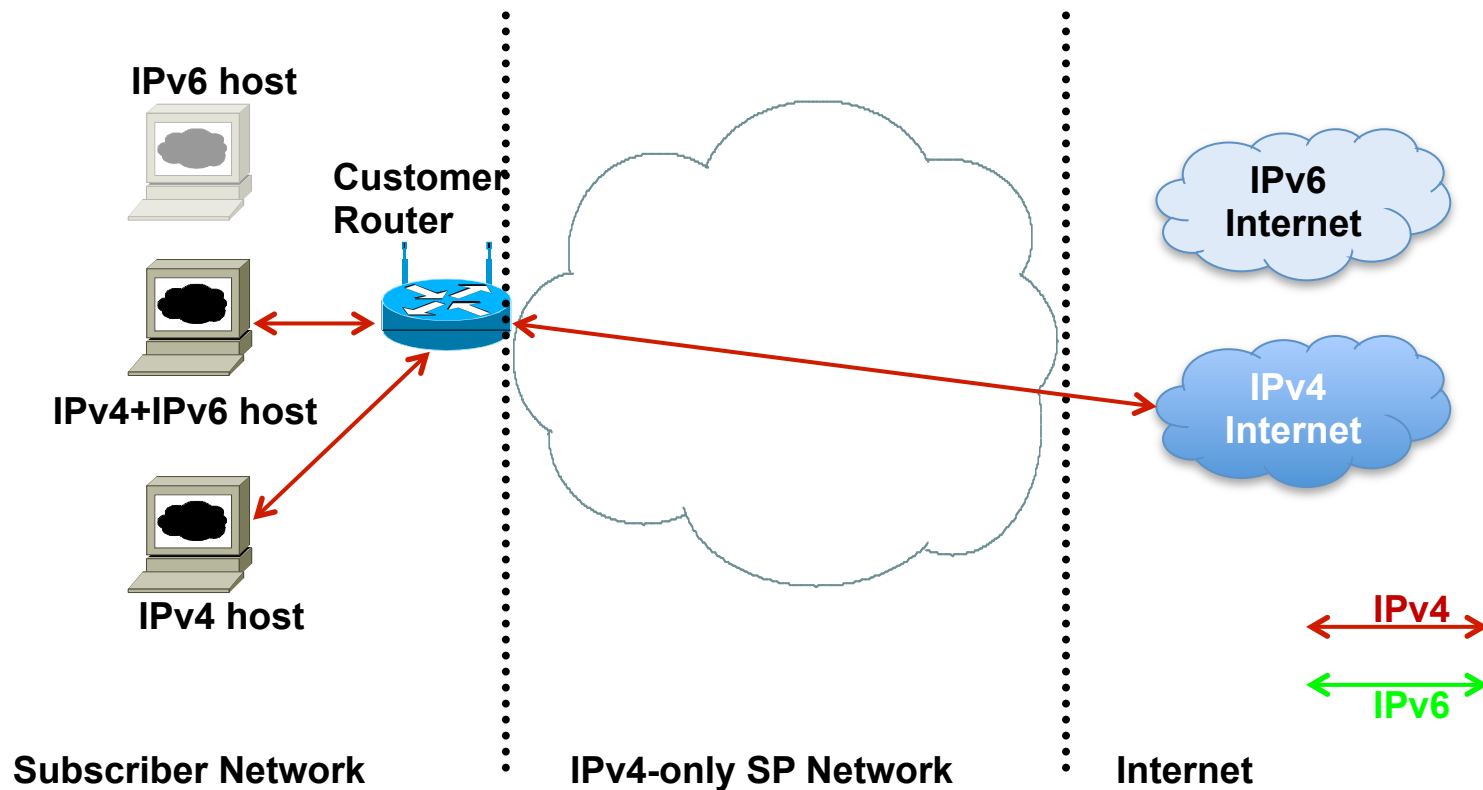
- ❑ ISP version of subscriber NAT
 - Subscriber NAT can handle only hundreds of translations
 - ISP NAT can handle millions of translations
- ❑ Not limited to just translation within one address family, but does address family translation as well
- ❑ Often referred to as Large Scale NAT (LSN)

Strategy One



Do Nothing

IPv4 only Network



- The situation for many SPs today:
 - No IPv6 for consumer
 - IPv4 scaling lasts as long as IPv4 addresses are available

IPv4 only: Issues

□ 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

IPv4 only: Applicability

- For Network Operators who:
 - Have sufficient IPv4 address space for foreseeable future business needs
 - Don't undertake long term planning
 - Are not heeding customer requests regarding IPv6 access
 - Have sufficient funds to purchase IPv4 address space via the marketplace

Strategy Two

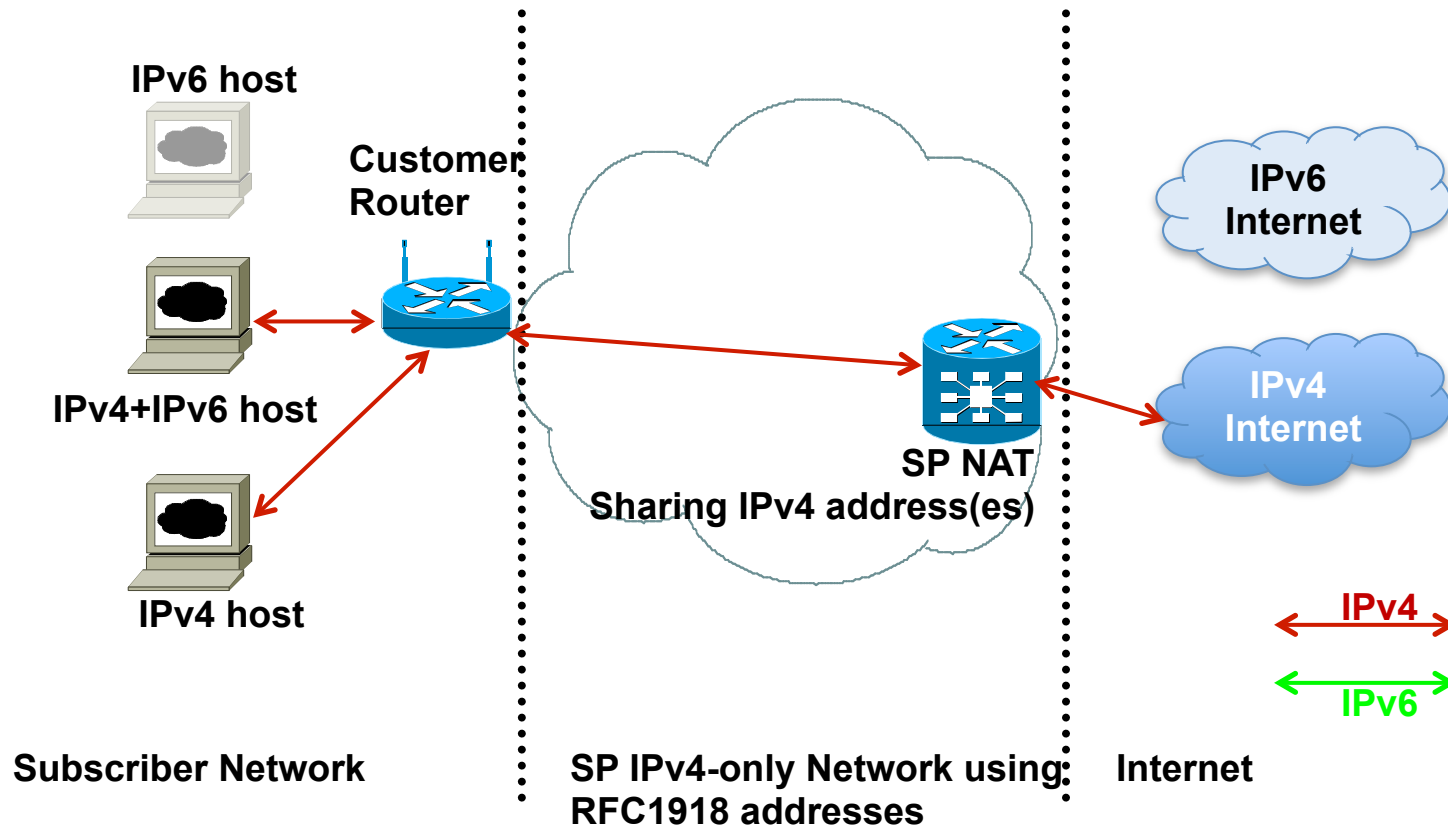


Extend life of IPv4 network

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

SP NAT in IPv4-only network



- Next step on from “doing nothing”:
 - SP introduces NAT in core when IPv4 addresses run out
 - No access to IPv6 Internet for IPv6 enabled hosts

SP NAT in IPv4-only network:

Issues

- Advantages
 - ISPs can reclaim global IPv4 addresses from their customers, replacing with non-routable private addresses and NAT
 - Allows continued IPv4 subscriber growth
- Disadvantages
 - SP needs a large NAT device in the aggregation or core layers
 - Has every well known technical drawback of NAT, including prevention of service deployment by customers
 - Double NAT highly likely (customer NAT as well as SP NAT)
 - Sharing IPv4 addresses could have behavioural, security and liability implications
 - Tracking association of port/address and subscriber, not to mention Lawful Intercept issues, are still under study
 - May postpone IPv6 deployment for a couple of years
 - Prevents subscribers from using IPv6 content, services and applications

SP NAT in IPv4-only network: Applicability

- For Network Operators who:
 - Are content to purchase and operate CGN devices within their core network
 - Are aware of the operational and performance pitfalls of CGN devices
 - Are able to reclaim public addresses from their customers for redeployment in their backbone
 - Are not heeding requests from customers for IPv6 access

IPv4 Subnet Trading

- Today the cost of getting IPv4 address space is low:
 - Service Provider:
 - RIR membership fee
 - Registration service fee (varies according to RIR service region)
 - End-sites usually receive IPv4 address block from SP as part of service
 - Many SPs already charge end-site for privilege of public IPv4 address
- In future when RIRs have no more IPv4 address space to distribute:
 - Cost of IPv4 addresses will be higher (today it's close to 0)
 - SPs may "purchase" IPv4 address space from other organisations

IPv4 Subnet Trading: Issues

□ Advantages

- Valuation of IPv4 addresses may hasten IPv6 adoption by encouraging sellers, perhaps more than offsetting costs to move some or all of their network to v6
- Receivers of transferred IPv4 address space can prolong their IPv4 networks

□ Disadvantages

- Market may not materialise, so organisations hoping to benefit may not
- Depending on region, if RIR doesn't register transfer, there may be no routability
- Risk to integrity of routing system, as RIRs no longer authoritative for address records
- Even more rapid growth of routing system
- Financial pressure on ISPs to dispose of IPv4 addresses they still need

IPv4 Subnet Trading: Applicability

- For Network Operators who:
 - Are have sufficient funds to purchase IPv4 address space on the marketplace
 - Are aware of the operational and performance pitfalls of purchased address space
 - Routability (legacy SP filters)
 - Registration (RIR vs not)
 - Reputation (previous user)
 - Are not heeding requests from customers for IPv6 access

Strategy Three



IPv4/v6 Coexistence/Transition
techniques

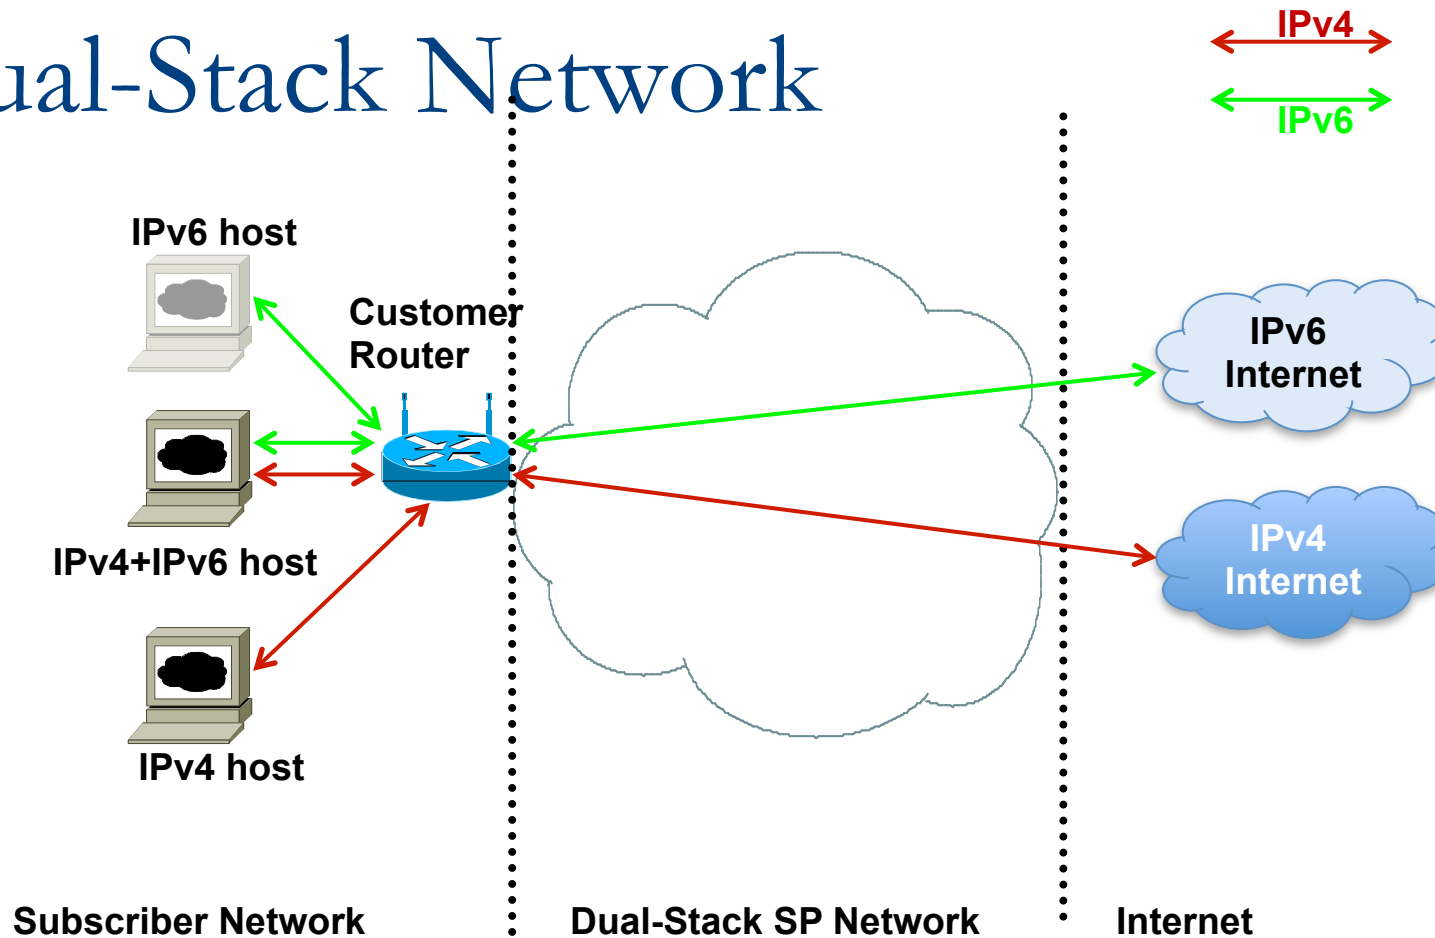
IPv4/IPv6 coexistence & transition

- 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
 - Activity of IETF **Softwires** Working Group
 - Large Scale NAT (LSN)
 - SP deploys large NAT boxes to do address and/or protocol translation

IPv4/IPv6 coexistence & transition

- Large Scale NAT (LSN)
 - NAT444/SP NAT
 - NAT to customer, optionally NAT'ed core.
 - Dual-Stack Lite & 464XLAT
 - IPv4 to IPv4 over IPv6
 - Activity of IETF **Softwires** & **v6ops** Working Groups
 - NAT64
 - Translation between IPv6 and IPv4
 - Activity of IETF **Behave** Working Group

Dual-Stack Network



- The original transition scenario, but dependent on:
 - IPv6 being available all the way to the consumer
 - Sufficient IPv4 address space for the consumer and SP core

Dual-Stack Network: Issues

□ Advantages

- Most cost effective long term model
- Once services are on IPv6, IPv4 can simply be discontinued

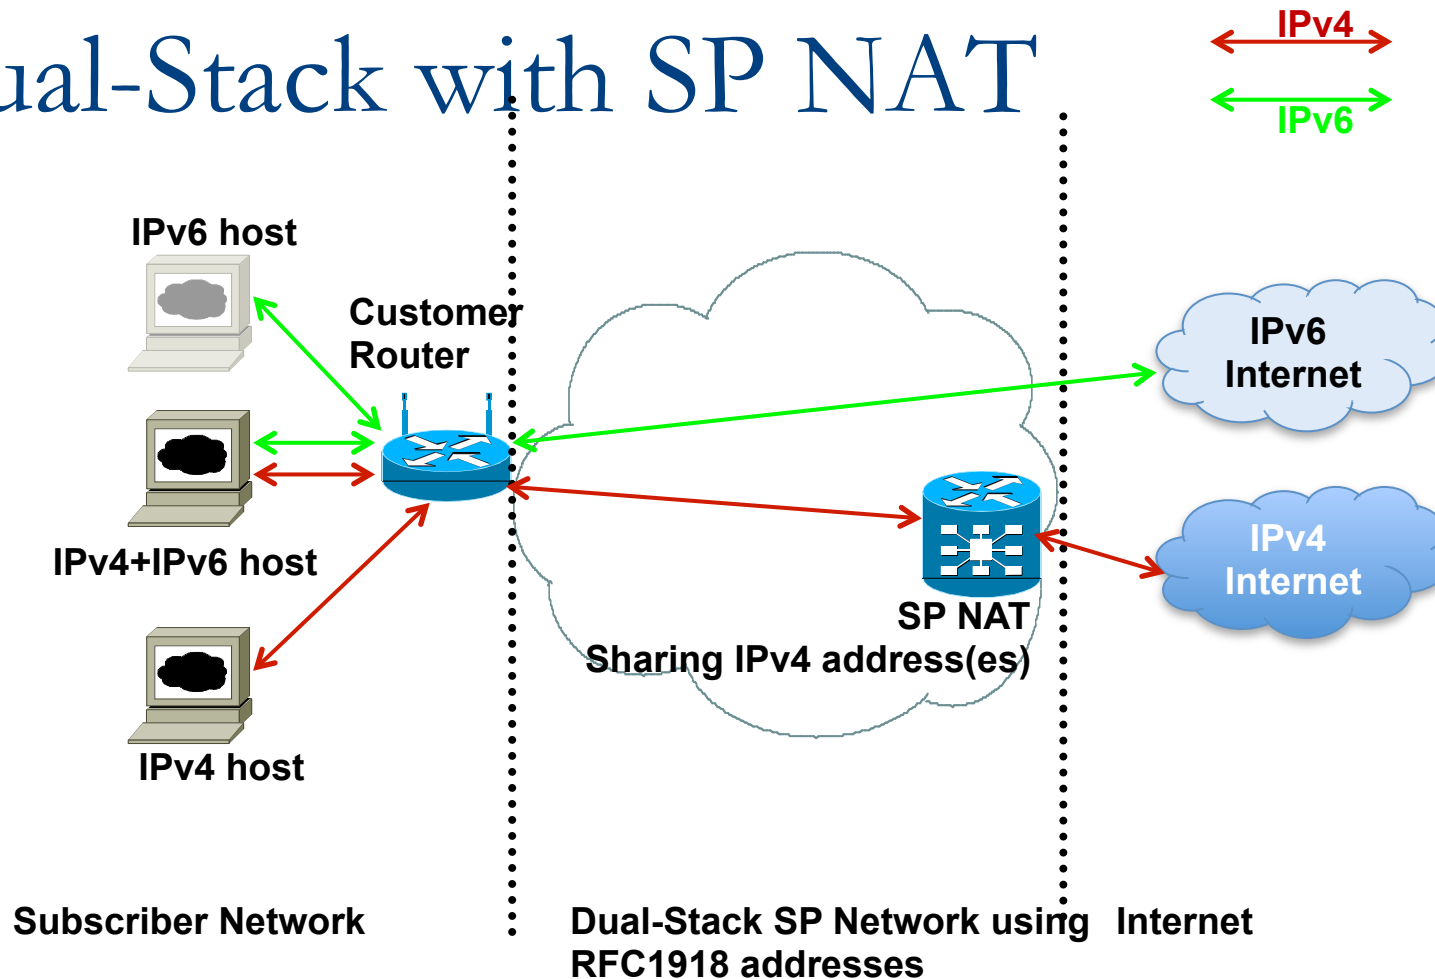
□ Disadvantages

- IPv4 growth limited to available IPv4 address space
- Running dual-stack network requires extra staff training
- IPv6 on existing IPv4 infrastructure might cost extra in terms of hardware changes (RIB and FIB memories)
- IPv6-only end-points cannot access IPv4, but given most IPv6 end-points are dual-stack, require IPv4 address too

Dual-Stack Network: Applicability

- For Network Operators who:
 - Have sufficient IPv4 address space for foreseeable future
 - Also may consider purchasing IPv4 address space on the open market
 - Have no legacy equipment or infrastructure which does not support IPv6
 - Do not wish to deploy CGN (NAT44)
 - **Are willing to support dual-stack CPE**
- Note: this is considered the ideal option
- Example:
 - Typical traditional Internet Service Provider deployment

Dual-Stack with SP NAT



- More likely scenario:
 - IPv6 being available all the way to the consumer
 - SP core and customer has to use IPv4 NAT due to v4 depletion

Dual-Stack with SP NAT: Issues

□ Advantages

- ISPs can reclaim global IPv4 addresses from their customers, replacing with non-routable private addresses and NAT
- Allows continued IPv4 subscriber growth
- SP can offer IPv6 connectivity too
- Does not postpone IPv6 deployment

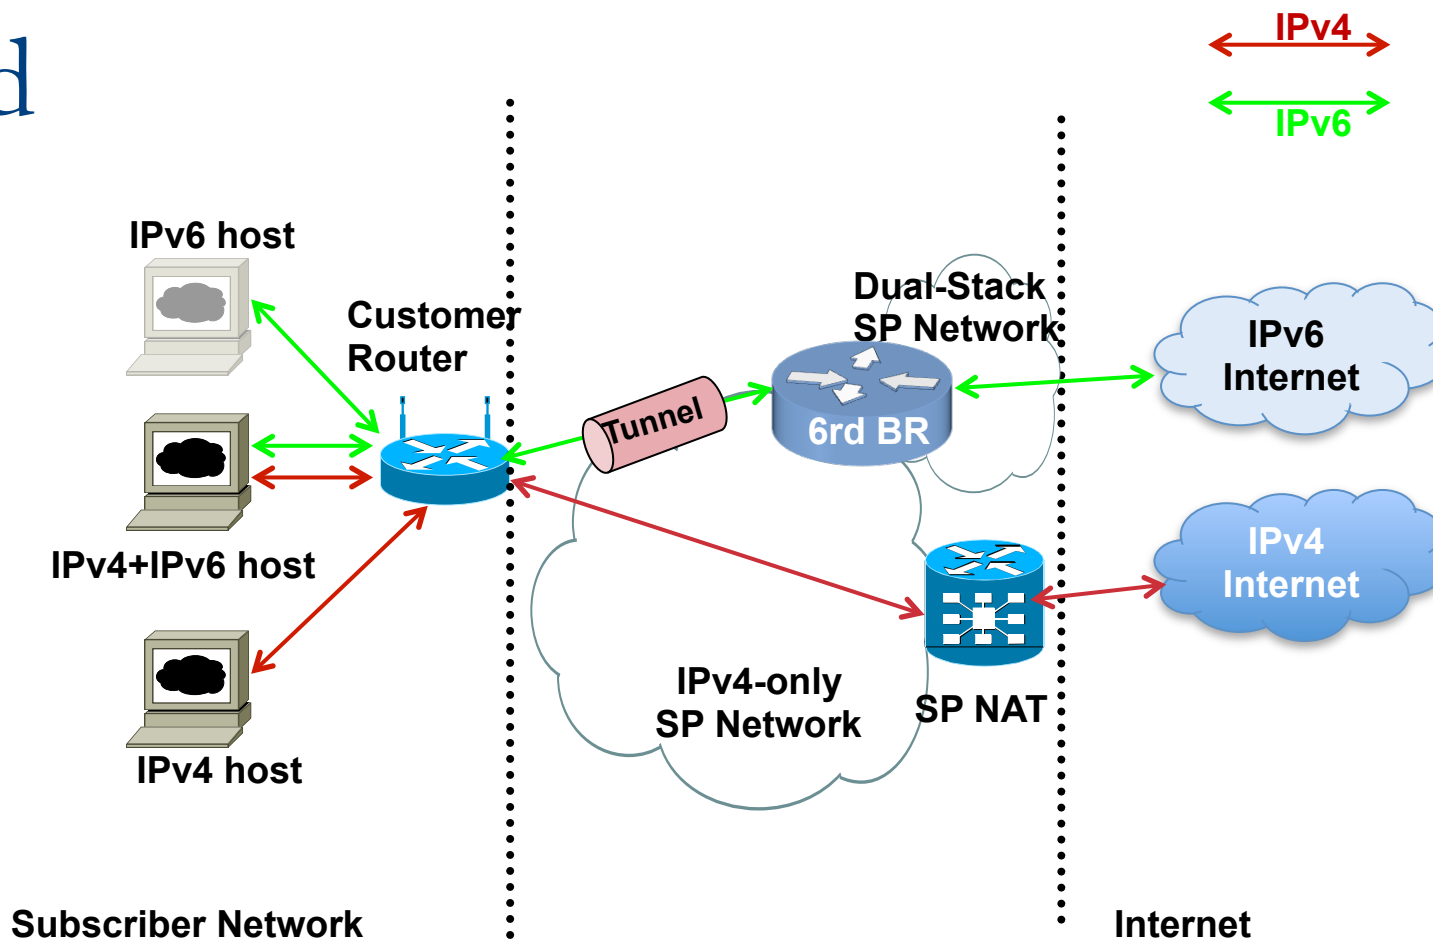
□ Disadvantages

- SP needs a large NAT device in the aggregation or core layers
- Has every well known technical drawback of NAT, including prevention of service deployment by customers
- Double NAT highly likely (customer NAT as well as SP NAT)
- Sharing IPv4 addresses could have behavioural, security and liability implications
- Tracking association of port/address and subscriber, not to mention Lawful Intercept issues, are still under study
- SP incurs additional investment and operational expenditure by deploying an IPv6 infrastructure

Dual-Stack with SP-NAT: Applicability

- For Network Operators who:
 - Have do not sufficient IPv4 address space and are content deploying CGN (NAT44) in the core
 - Are able to reclaim public IPv4 address space from customers for redeployment on their backbone infrastructure
 - Have no legacy equipment or infrastructure which does not support IPv6
 - **Are willing to support dual-stack CPE**
- Note: this is considered the realistic best practice
- Example:
 - Typical traditional Internet Service Provider deployment

6rd



- 6rd (Rapid Deploy) used where ISP infrastructure to customer is not IPv6 capable (eg IPv4-only BRAS)
 - Customer has IPv4 Internet access either natively or via NAT
 - Customer IPv6 address space based on ISP IPv4 block

6rd: Issues

□ Advantages

- The service provider has a relatively quick way of providing IPv6 to their customer without deploying IPv6 across their infrastructure
- Subscribers can readily get access to IPv6
- 6rd relay and CPE are becoming available from vendors
- 6rd operation is completely stateless, does not have the operational drawbacks of 6to4, and does not postpone IPv6 deployment

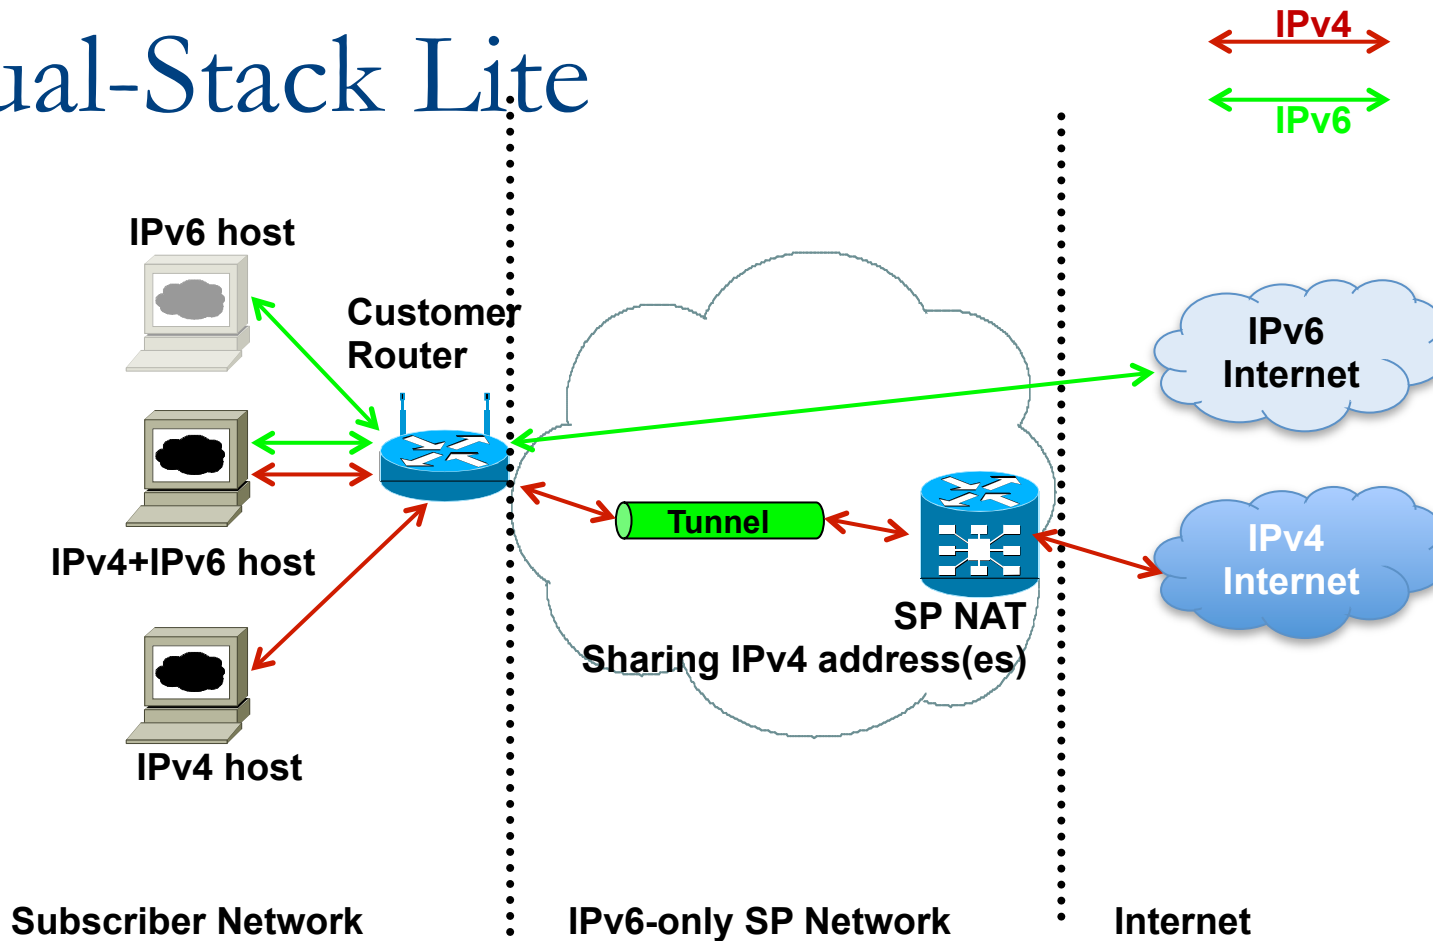
□ Disadvantages

- 6rd is not a long-term solution for transitioning to IPv6 – one further transition step to remove the tunnels
- CPE needs to be upgraded to support 6rd
- The ISP has to deploy one or several 6rd termination devices
- If customer or SP uses NAT for IPv4, all NAT disadvantages are inherited

6rd: Applicability

- For Network Operators who:
 - Have do not sufficient IPv4 address space and are content deploying CGN (NAT44) in the core
 - Are able to reclaim public IPv4 address space from customers for redeployment on their backbone infrastructure
 - Have legacy equipment or infrastructure which does not support IPv6
 - And realise that it will eventually have to be upgraded
 - Are willing to run a 6rd Border Router
 - Are willing to support dual-stack CPE (with 6rd)
- Example:
 - Broadband operators who have legacy DSLAMs or lease a third party's L2 network

Dual-Stack Lite



- Service Provider deploys IPv6-only infrastructure:
 - IPv6 being available all the way to the consumer
 - IPv4 is tunnelled through IPv6 core to Internet via SP NAT device

Dual-Stack Lite: Issues

□ Advantages

- The SP is using IPv6 across their entire infrastructure, avoiding the IPv4 address pool depletion issue totally
- The SP can scale their infrastructure without any IPv4 dependencies
- Consumers can transition from IPv4 to IPv6 without being aware of any differences in the protocols
- IPv6 packets routed natively

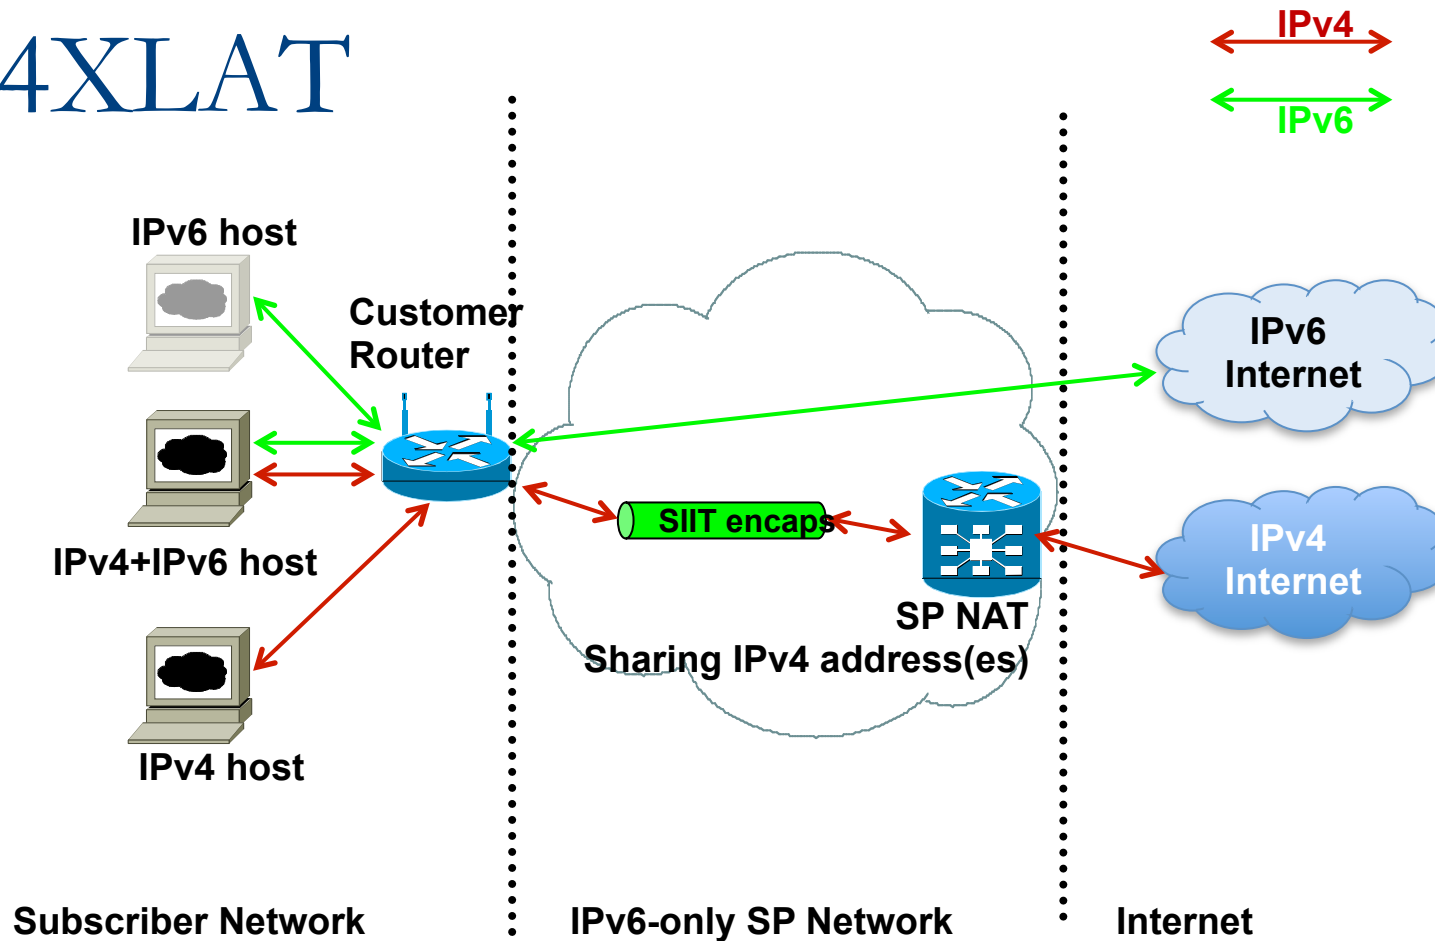
□ Disadvantages

- SP requires NAT device in core supporting DS-Lite
- Subscriber router needs to be IPv6 capable
- Model has all drawbacks of SP NAT model for IPv4 traffic

Dual-Stack Lite: Applicability

- For Network Operators who:
 - Are considering “green-field” deployments
 - Are content running an IPv6-only backbone
 - Are willing to deploy CGN (DS-Lite) in the core
 - **Are willing to support dual-stack CPE (with DS-Lite)**
- Example:
 - Mobile operators rolling out a brand new network, with handsets which have dual-stack radios

464XLAT



- Service Provider deploys IPv6-only infrastructure:
 - IPv6 being available all the way to the consumer
 - IPv4 is transported through IPv6 core to Internet via SIIT on customer router, and NAT64 on SP NAT device

464XLAT: Issues

□ Advantages

- The SP is using IPv6 across their entire infrastructure, avoiding the IPv4 address pool depletion issue totally
- The SP can scale their infrastructure without any IPv4 dependencies
- Consumers can transition from IPv4 to IPv6 without being aware of any differences in the protocols
- Devices not supporting IPv6 can access IPv6-only networks
- IPv6 packets routed natively

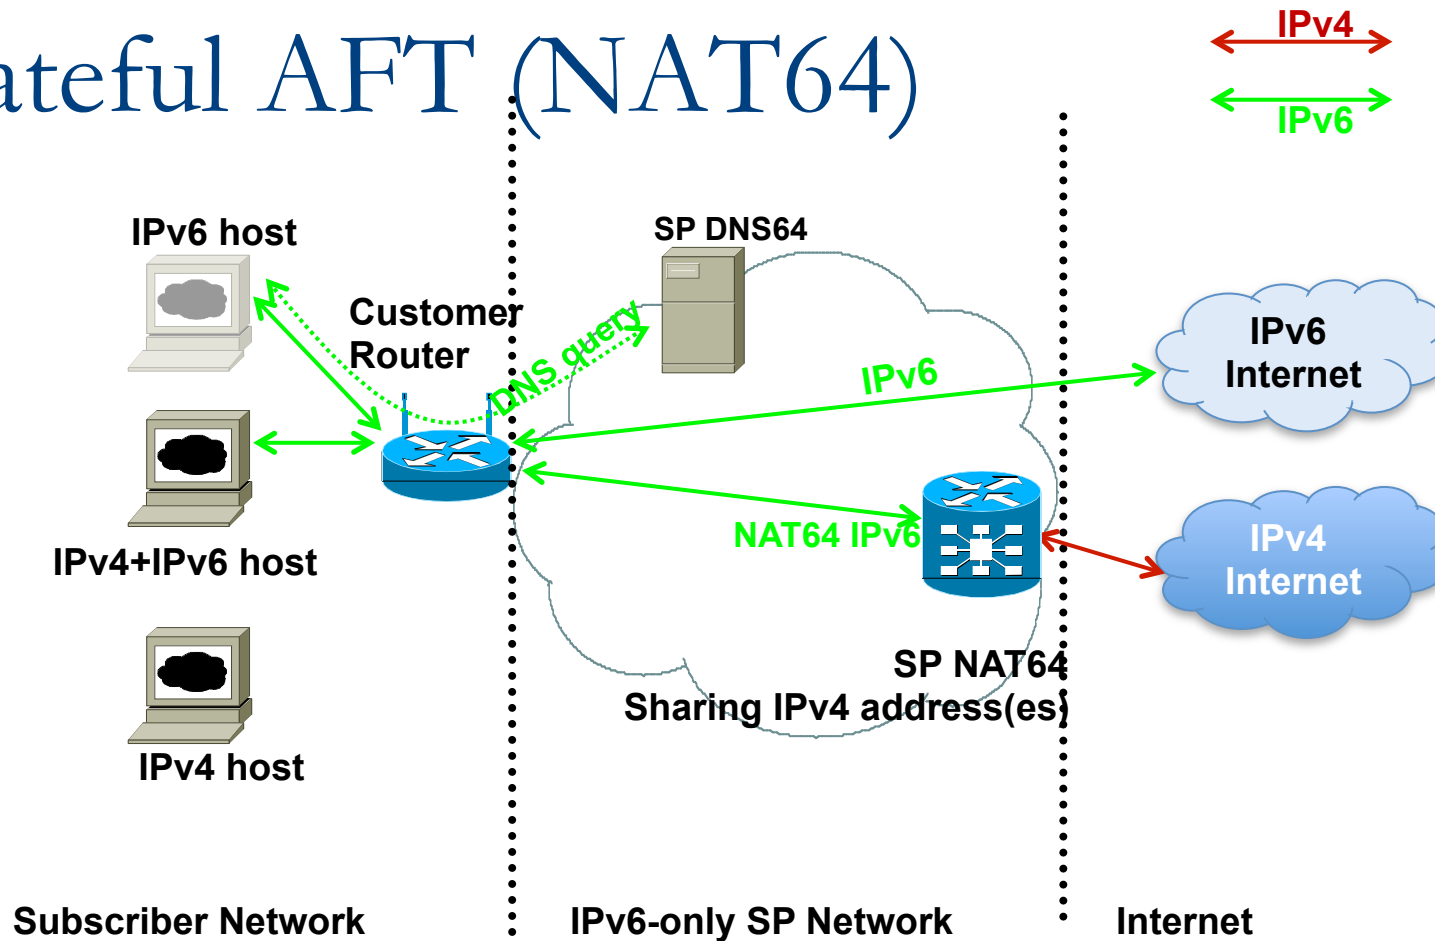
□ Disadvantages

- SP requires NAT device in core (PLAT – NAT64)
- Subscriber router needs to be IPv6 capable and support IPv4/IPv6 header translation (CLAT – SIIT)
- Model has all drawbacks of SP NAT model for IPv4 traffic

464XLAT: Applicability

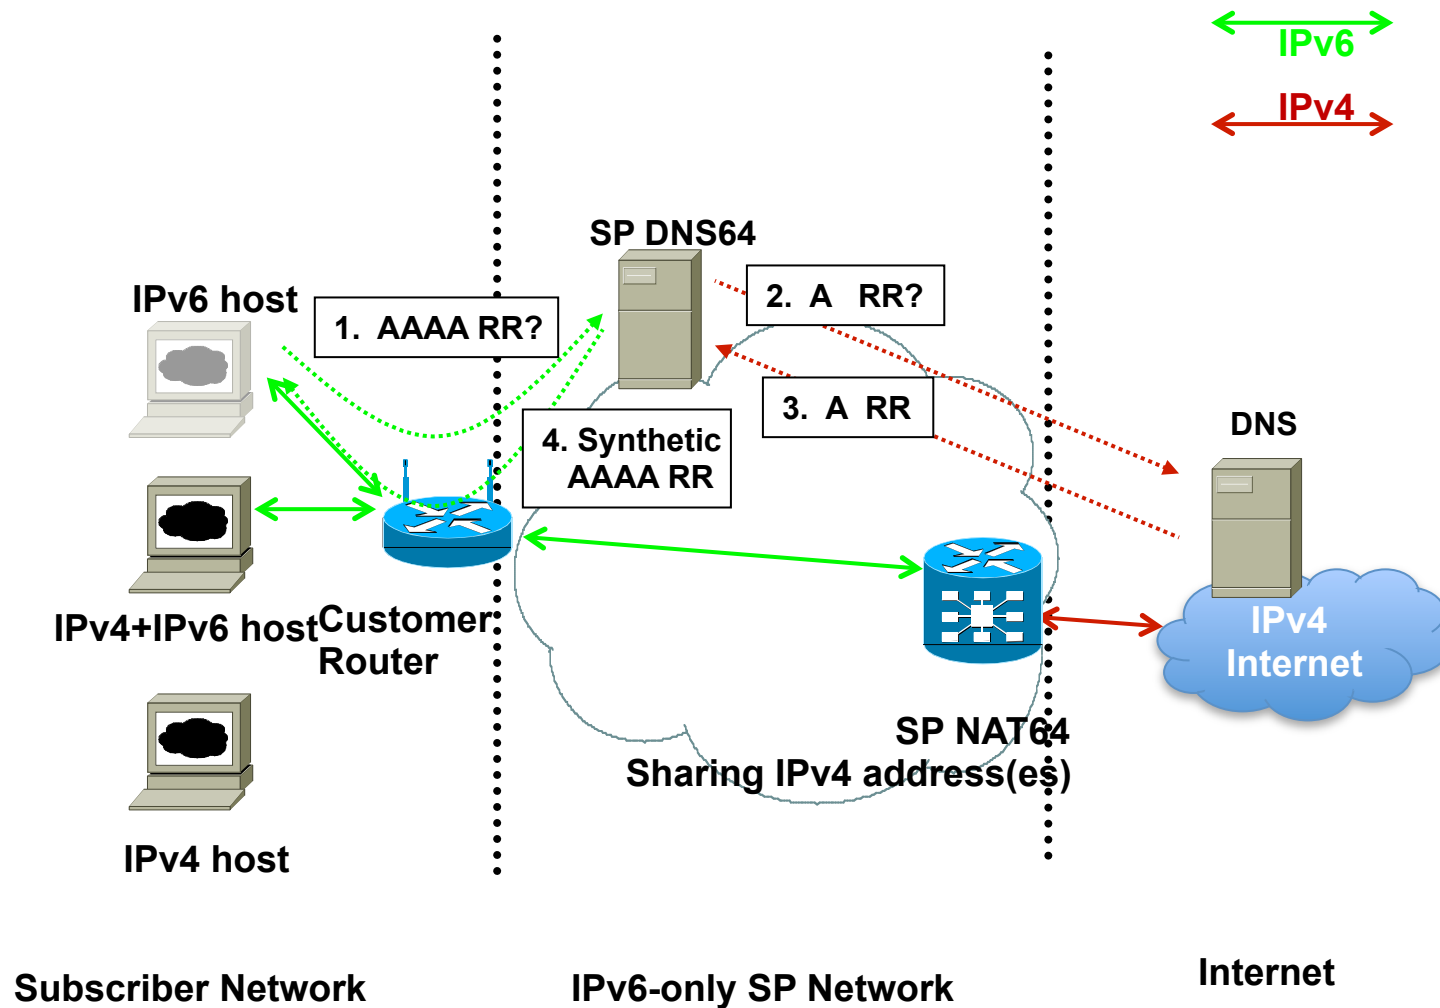
- For Network Operators who:
 - Are considering “green-field” deployments
 - Are content running an IPv6-only backbone
 - Are willing to deploy CGN (PLAT) in the core
 - **Are willing to support dual-stack CPE (with SIIT)**
- Example:
 - Mobile operators rolling out a brand new network, with handsets which have dual-stack radios

Stateful AFT (NAT64)



- Service Provider deploys IPv6-only infrastructure:
 - Only IPv6 is available to the consumer
 - IPv4 Internet available via Address Family Translation on SP NAT device

Stateful AFT (NAT64) Details



Stateful AFT: Issues

□ Advantages

- Allows IPv6 only consumers access to IPv4 based content without giving them IPv4 address resources
- IPv6 services and applications offered natively to consumers
- SP network runs IPv6 only, avoiding IPv4 dependencies

□ Disadvantages

- SP requires NAT device in core
- SP's DNS infrastructure needs to be modified to support NAT64
- Subscriber router needs to be IPv6 capable
- Subscriber devices need to be IPv6 capable (no legacy support)
- Model has all drawbacks of SP NAT model for IPv4 traffic

Stateful AFT: Applicability

- For Network Operators who:
 - Are considering “green-field” deployments
 - Are content running an IPv6-only backbone
 - Are willing to deploy CGN (NAT64) in the core
 - **Are willing to support IPv6-only CPE**
- Example:
 - Mobile operators rolling out a brand new network, with handsets which have single-stack (IPv6-only) radios

Conclusions & Recommendations



Functionalities and Operational Issues

	IPv4-only network	IPv4-only network with IPv4 NAT	Dual-Stack, no IPv4 NAT	Dual-Stack with IPv4 NAT	6rd, no IPv4 NAT	6rd with IPv4 NAT	DS-Lite	464XLAT	Stateful AFT
Prolongs IPv4	No	Yes	No	Yes	No	Yes	Yes	Yes	Yes
Allows Business Growth	No	Yes (scaling issues if content is mostly IPv6)	Limited to IPv4 address availability	Yes (traffic to IPv4-only servers)	Limited to IPv4 address availability	Yes	Yes	Yes	Yes
Requires IPv6 Deployment	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coexists with IPv6 Deployment	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Complexity of Operation	Low	Low	Low	Moderate	Moderate	Moderate	Moderate	Moderate	Moderate
Complexity of Troubleshooting	Low	Moderate	Low	High	Moderate	High	High	High	Moderate
Breaks End-to-End IPv4	No	Yes	No	Yes	No	Yes	Yes	Yes	N/A
NAT Scalability issues to IPv4 services	No	Yes	No	Yes	No	Yes	Yes	Yes	Yes
NAT Scalability issues to IPv6 services	N/A	Yes	No	No	No	No	No	No	No
DNSSEC issues	No	Yes	No	Yes for IPv4 No for IPv6	No	Yes for IPv6 No for IPv4	Yes for IPv4 No for IPv6	Yes for IPv4 No for IPv6	Yes for IPv4 No for IPv6
Lawful Intercept issues	No	Yes	No	Yes for IPv4	No	Yes for IPv4	Yes for IPv4	Yes for IPv4	Yes for IPv4

Functionalities and Operational Issues

- Complexity of operation:
 - Moderate in the case of a single network with two address families
- Complexity of troubleshooting:
 - Running two address families and/or tunnels is assumed to be more complex
- Breaks end-to-end connectivity in IPv4:
 - Subscribers sharing a CGN will have little to no hurdles in their communication
 - Subscribers separated by one or several CGN will experience some application issues

Comparing where changes will occur

	IPv4-only network	IPv4-only network with IPv4 NAT	Dual-Stack, no IPv4 NAT	Dual-Stack with IPv4 NAT	6rd, no IPv4 NAT	6rd with IPv4 NAT	DS-Lite	464XLAT	Stateful AFT
Change CPE	No	No	Only if customer wants IPv6	Only if customer wants IPv6	Yes	Yes	Yes	Yes	Yes
CPE to do AFT to access IPv6	No	No	No	No	No	No	No	No	No
IPv4 NAT in core/edge	No	Yes	No	Yes	No	Yes	Yes	Yes	No
AFT in core/edge to access IPv6	Yes	Yes	No	No	No	No	No	No	Yes

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, 464XLAT or DS-lite as migration techniques
Content and Applications on IPv4 and IPv6	Dual-Stack (if enough IPv4) or 6rd; SP IPv4-NAT; DS-lite or 464XLAT (for greenfield) *
Users are IPv6 only	Stateful 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



Conclusion

- Topics covered
 - Introduction to IPv6, History and Background
 - IPv6 concepts and Protocol Standards
 - IPv6 addressing
 - IPv6 Routing
 - Service Provider IPv4 – IPv6 Coexistence

Thank you!



End of Session



Thank You

