APNIC Training

IPv6 Essentials

Overview

- An overview of IPv6
- APNIC IPv6 Policies
- IPv4 to IPv6 Transition

IPv6 Architecture Overview

INTRODUCTION

What is IPv6?

- IPv6 is a new version of the Internet layer protocol (IP) in the TCP/IP suite of protocols.
- It replaces the current Internet protocol layer commonly referred to as IPv4.
- RFC2460

IPv6: The ‘new’ tech!

- IPv6 is not new!
- RFC 2460 (IPv6) was released in December 1998.
- It has been more than a decade, so where are we with IPv6?

IPv6 current deployment status

- Not widely deployed
- Not many cases of production networks
  – Not many business cases
  – Quite a few research and experimental network
- Routing infrastructure not “native” enough – too many tunnels, but, it works
- Not enough content
  – Must parallel IPv4 to be viable

Issues

- Obviously not many production networks deployed
  - Gap in understanding between front line network engineers and decision makers
    - CEOs and CIOs may not be aware or understand the implications.
    - Concerned about making investments without tangible profit
- Question the commercial value

So why aren't we ready yet?

- It's a simple business reality:
  - Highly competitive environment
    - A company will always spend its available resources on profit-making activities
  - Fundamental nature of IPv6
    - No customers are currently demanding IPv6
    - So, there is currently no pressing business case for deploying IPv6
- However, IPv6 is the only path that enables the Internet to continue to expand
  - Large address space
  - Simpler and cheaper with more efficient networks

The crunch line

- IPv6 is the only technology currently available that is mature enough to provide a viable, long term solution
- It is not perfect – it is not backwards compatible and does not work WITH IPv4
- But it does work ALONGSIDE IPv4 and can integrate via a number of methods providing co-existence and transition

The move to IPv6

- Implementation rather than migration.
- IPv6 can coexist with your current IPv4 infrastructure.
- IPv6 "rollout" can occur during end of life equipment replacement.
- Most recent equipment are IPv6 ready.

Main IPv6 benefits

- Expanded addressing capabilities
- Server-less autoconfiguration (“plug-n-play”) and reconfiguration
- More efficient and robust mobility mechanisms
- Built-in, strong IP-layer encryption and authentication
- Streamlined header format
- Improved support for options / extensions
- Flow labeling capability

EXPANDED ADDRESSING CAPABILITIES

Address Space: IPv4 vs IPv6

- IPv4 Address Space
  - 2^{32}
  - 4294967296
- IPv6 Address Space
  - 2^{128}
  - 340282366920938463463374607431768211456

Size of the IPv6 address space

- This would allow every person on the planet to have their own internet as large as the current Internet
- It is difficult to foresee running out of IPv6 addresses

Incidental benefits of bigger addresses

- Easier address management/delegation
- Room for more levels of hierarchy, for route aggregation
- Ability to do achieve end-to-end connectivity (NAT not needed)

Potential impact of bigger addresses

- Increase of backbone routing table size is a problem
  - Current backbone routing table size > 200K
    - CIDR does not guarantee an efficient and scalable hierarchy
    - There is a lack of uniformity in the current hierarchical
  - IPv6 address architecture is more hierarchical than IPv4
    - Potentially better aggregation
    - But still a concern in IPv6 because of huge address space

STREAMLINED HEADER FORMAT

IPv4/IPv6 Header Changes

IPv4 Header

IPv6 Header

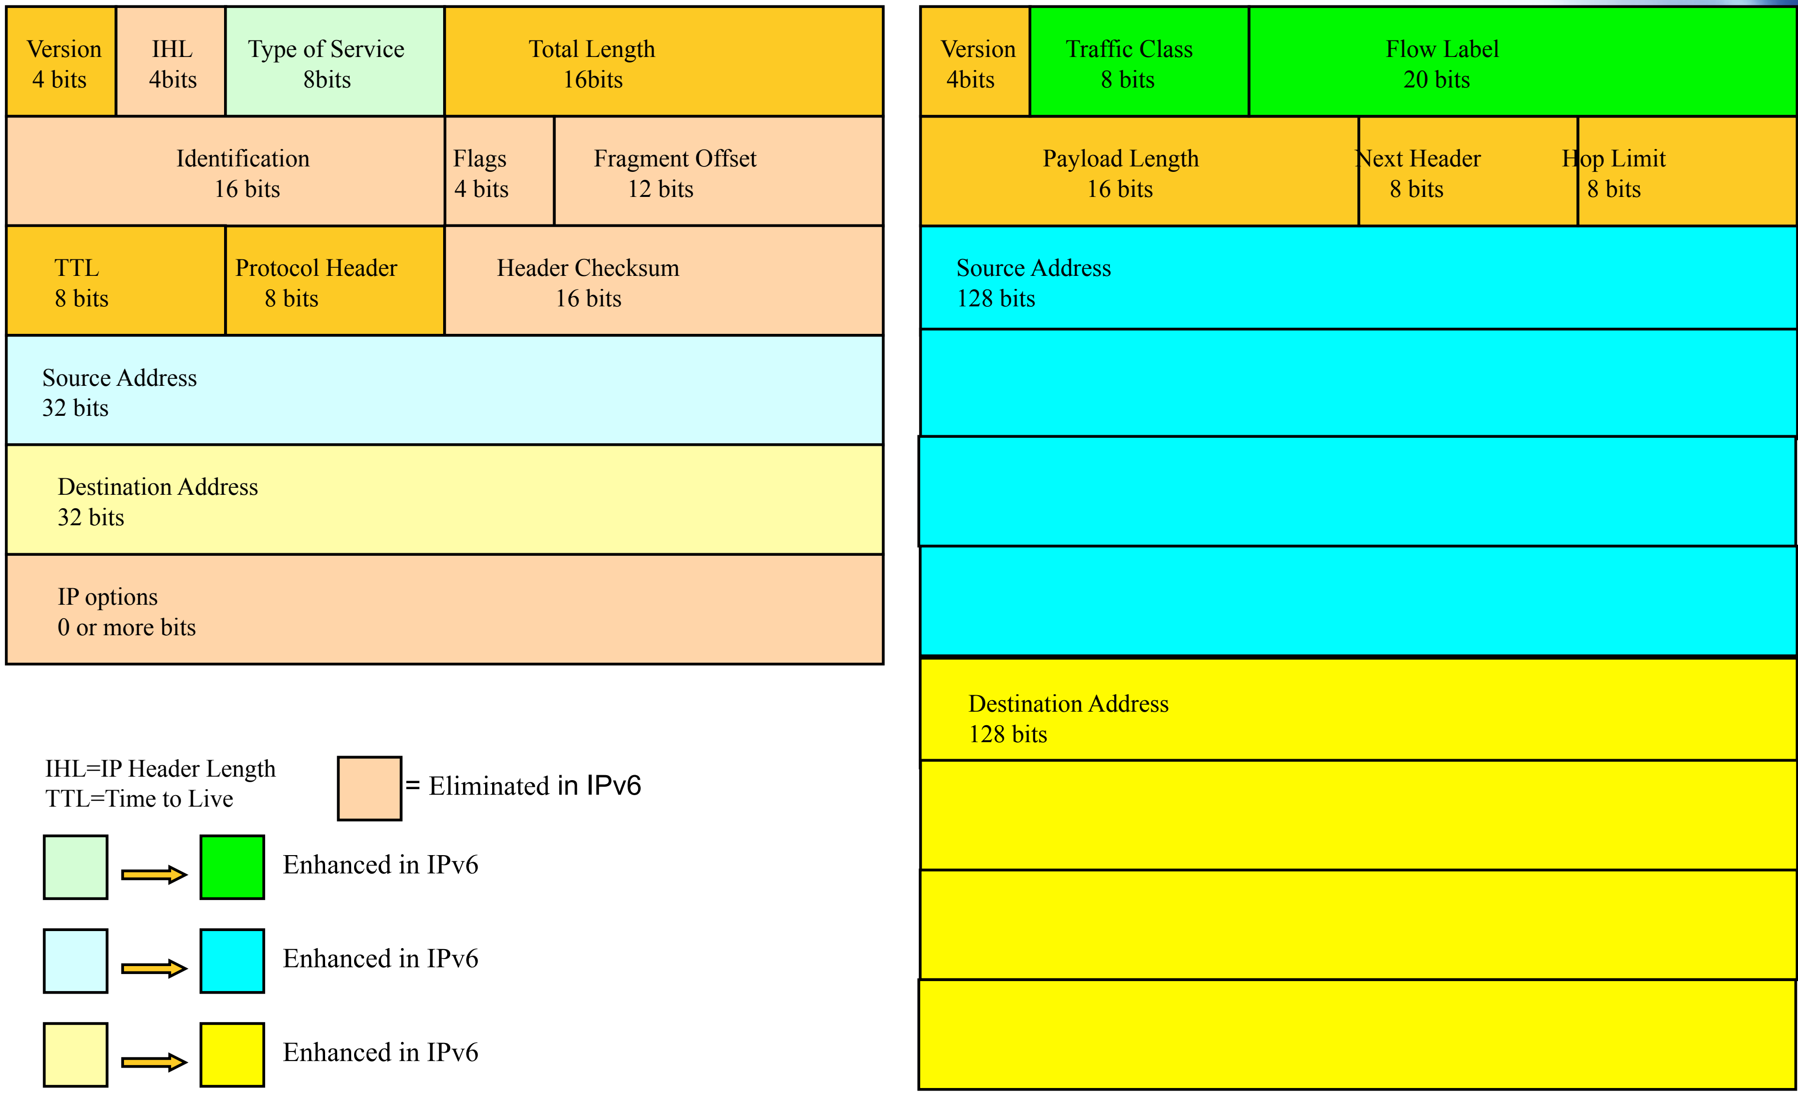

IPv4/IPv6 Header Changes

- Streamlined
  - Fragmentation fields moved out of base header
  - IP options moved out of base header
  - Header Checksum eliminated
  - Header Length field eliminated
  - Length field excludes IPv6 header

IPv4/IPv6 Header Changes

- Revised
  - Time to Live = Hop Limit
  - Protocol = Next Header
  - Precedence & TOS = Traffic Class
  - Addresses increased from 32 bits to 128 bits
- Extended
  - Flow Label field added

IMPROVED SUPPORT FOR OPTIONS / EXTENSIONS

The IPv6 Packet

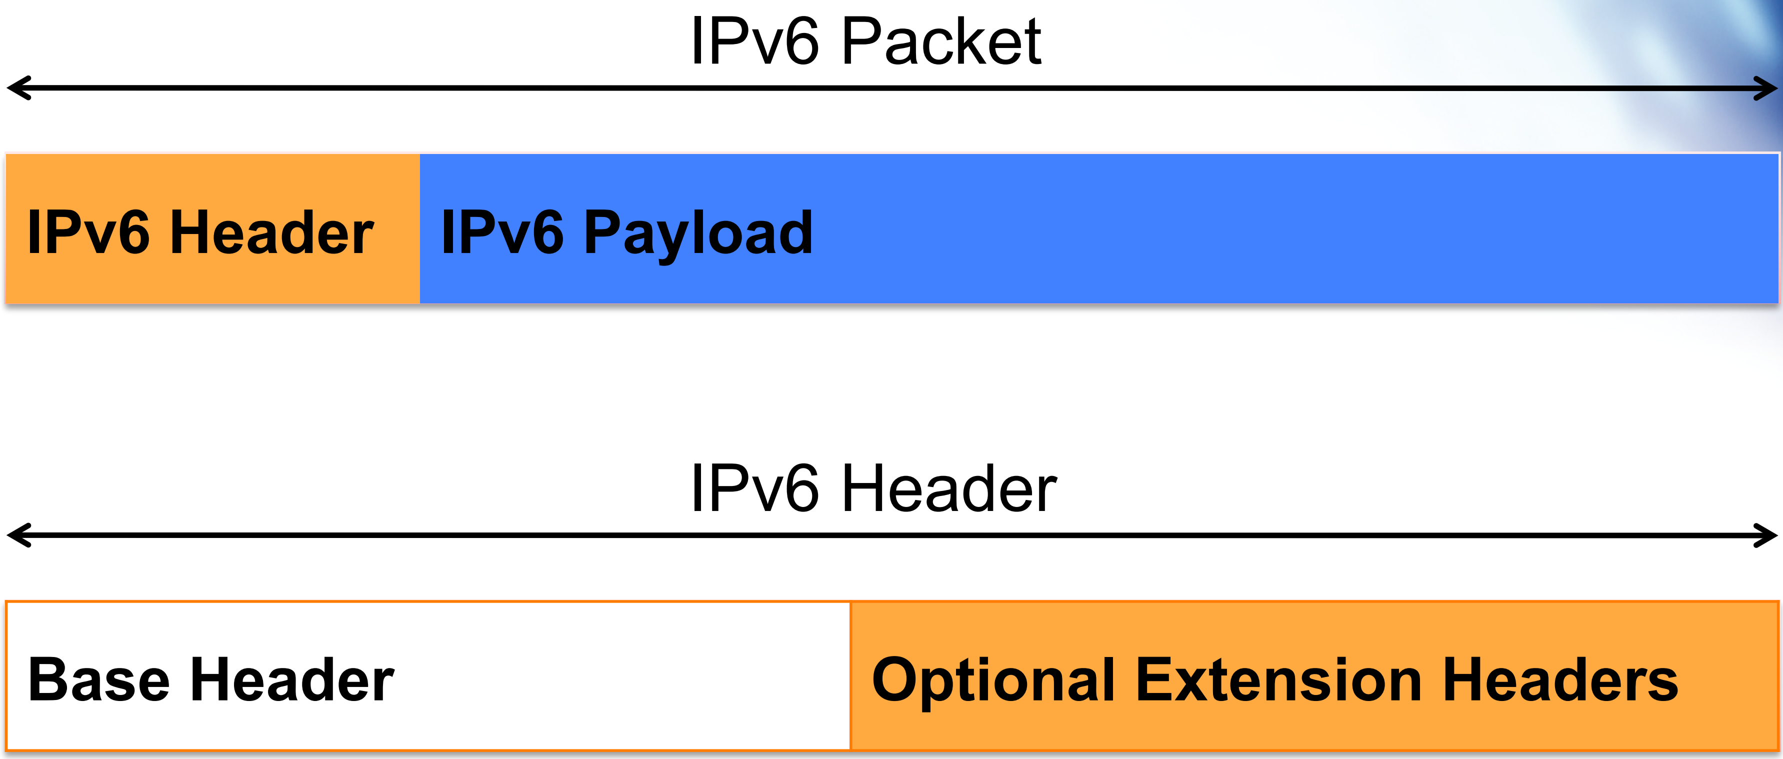

IPv6 Extension Headers

- Carry the additional options and padding features that are part of the base IPv4 header
- Extension headers are optional and placed after the base header
- There can be zero, one, or more Extension Headers between the IPv6 header and the upper-layer protocol header
- Ordering is important

IPv6 Extension Headers

Next Header Value (decimal)	Extension Header Name	Length (bytes)
0	Hop-by-hop Options	Variable
43	Routing	Variable
44	Fragment	8
50	Encapsulating Security Payload (ESP)	Variable
51	Authentication Header	Variable
60	Destination Options	Variable

Extension headers

- Next header field

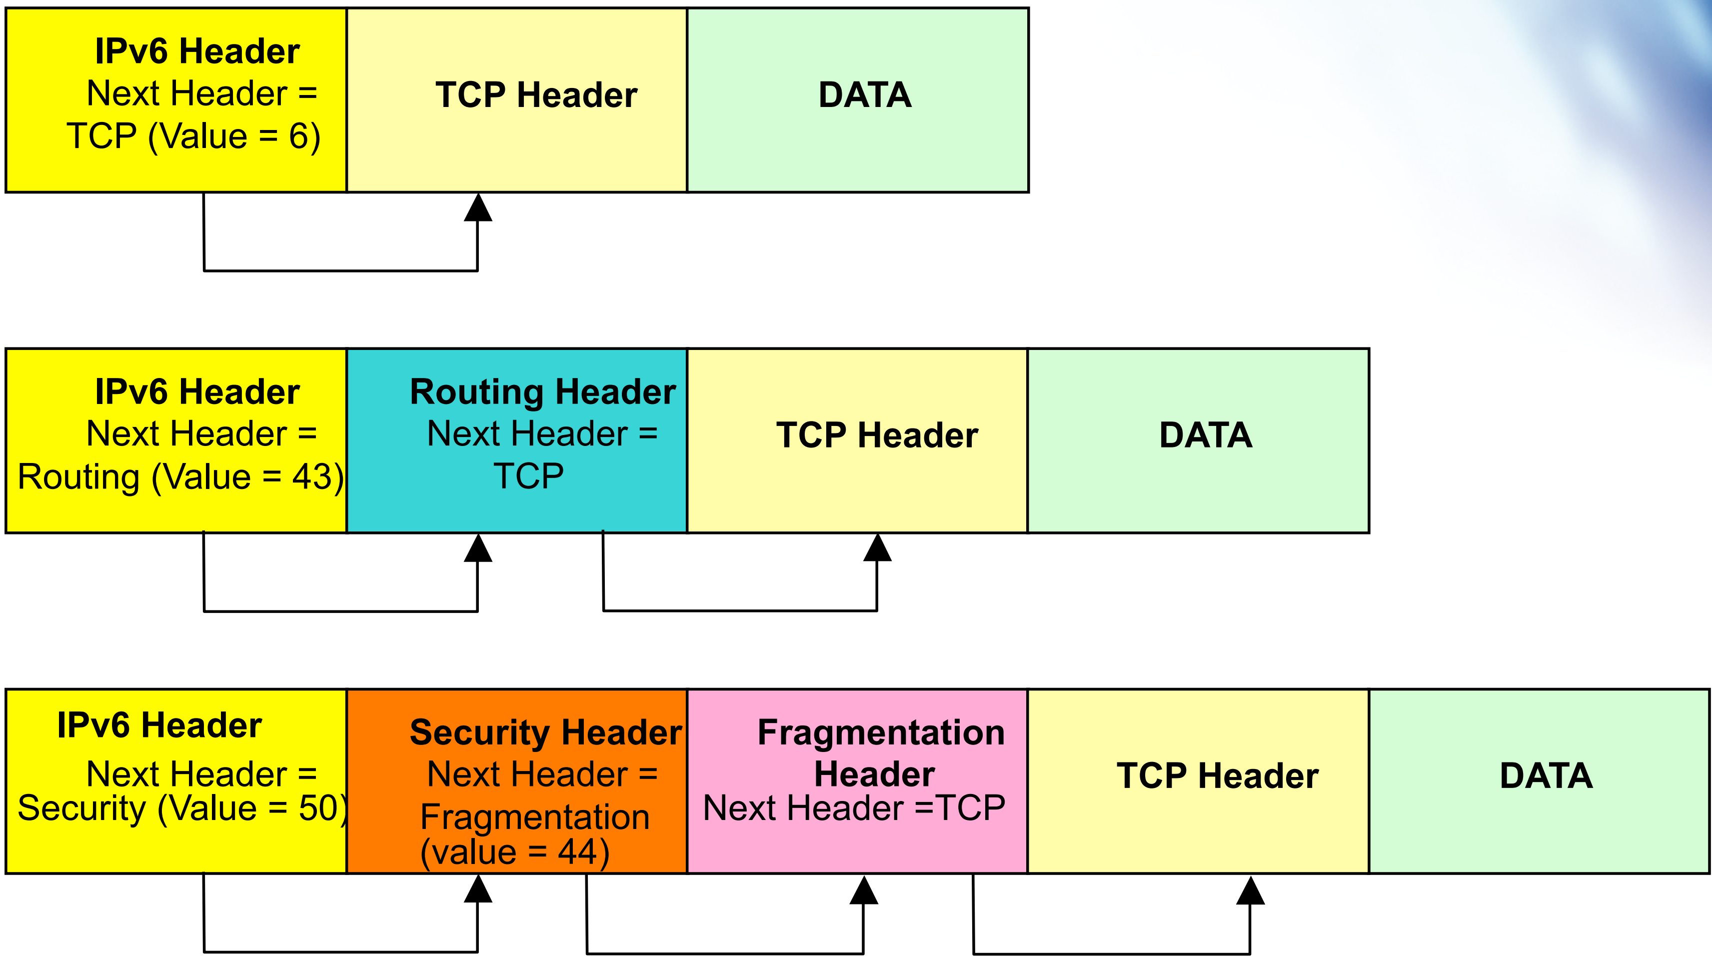

IPV6 ADDRESSING

Numbering Systems

DEC	HEX	BIN
0	0	0000
1	1	0001
2	2	0010
3	3	0011
4	4	0100
5	5	0101
6	6	0110
7	7	0111
8	8	1000
9	9	1001
10	A	1010
11	B	1011
12	C	1100
13	D	1101
14	E	1110
15	F	1111

IPv6 addressing

- 128 bits of address space
- Divided into eight 16 bit fields, each represented as a 4 digit hexadecimal number.
    - X:X:X:X:X:X:X:X (X=16 bit number, ex: A2FE)
- Example:
    - 2001:0DB8:124C:C1A2:BA03:6735:EF1C:683D

IPv6 addressing

- 128 bits of address space
- Divided into eight 16 bit boundaries
- Each represented as a 4 digit hexadecimal number, separated by colons.
  - 2001:0DB8:124C:C1A2:BA03:6735:EF1C:683D
- This representation is called colon-hexadecimal

IPv6 Address Abbreviation

- IPv6 address can be simplified by removing leading zeros.
- Each block must have at least a single digit.
  - Before
    - 2001:0DB8:0023:0000:0000:036E:1250:2B00
  - After
    - 2001:DB8:23:0:0:36E:1250:2B00
- Do not remove trailing zeros!

IPv6 Address Abbreviation

- Address can be further simplified by using "Zero Compression"
- Consecutive fields of 16 bit blocks set to 0 in colon-hexadecimal format can be compressed to ::
  - Before
    - 2001:DB8:23:0:0:36E:1250:2B00
  - After
    - 2001:DB8:23::36E:1250:2B00
  - Can be used only once

IPv6 addressing model

- **IPv6 Address type**
  - Unicast
    - An identifier for a single interface

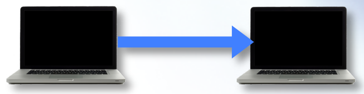

  - Anycast
    - An identifier for a set of interfaces

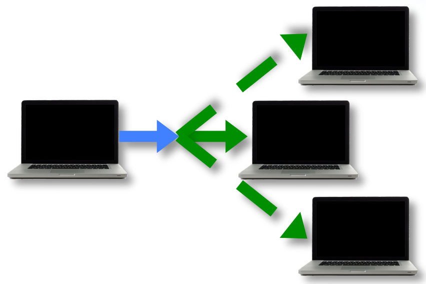

  - Multicast
    - An identifier for a group of nodes

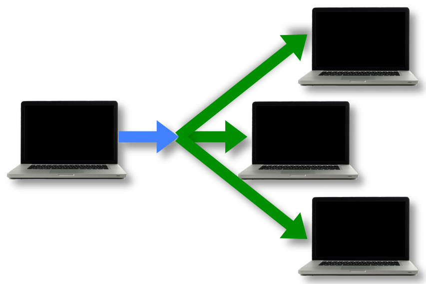

RFC 4291

IPv6 addressing structure

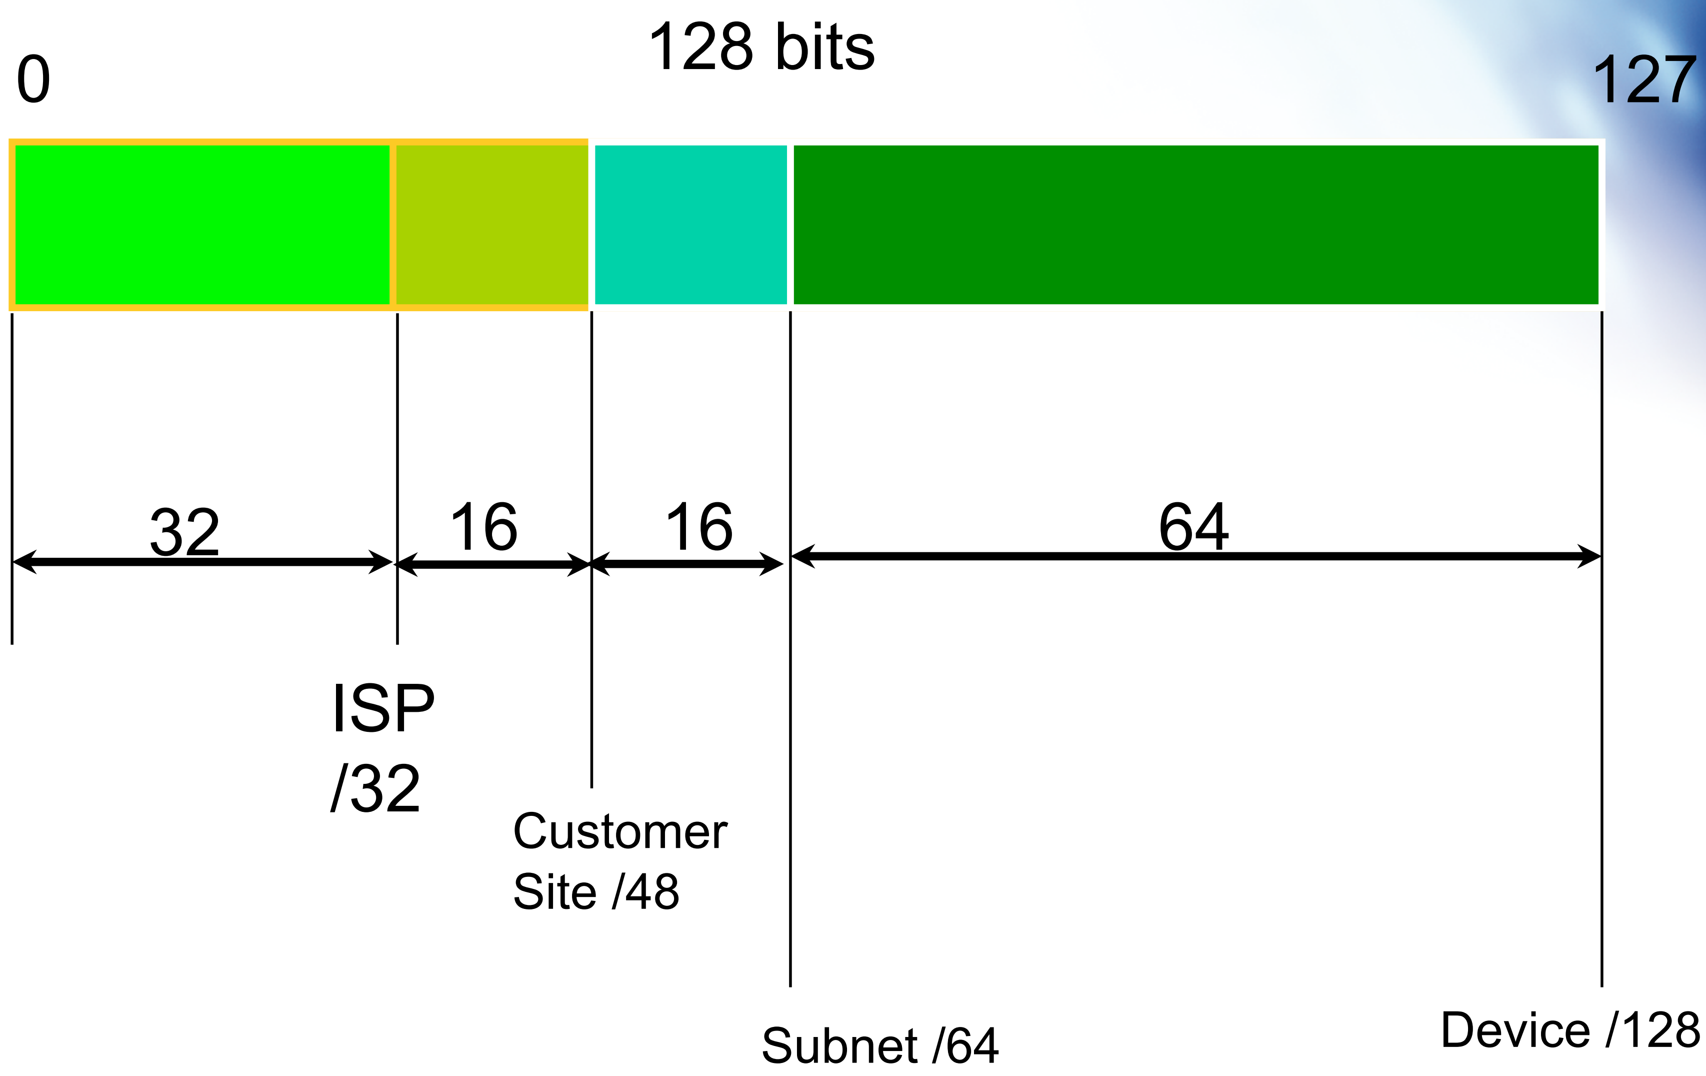

Interface ID

- The lowest-order 64-bit field addresses may be assigned in several different ways:
  - auto-configured from a 48-bit MAC address expanded into a 64-bit EUI-64
  - assigned via DHCP
  - manually configured
  - auto-generated pseudo-random number
  - possibly other methods in the future

EUI-64

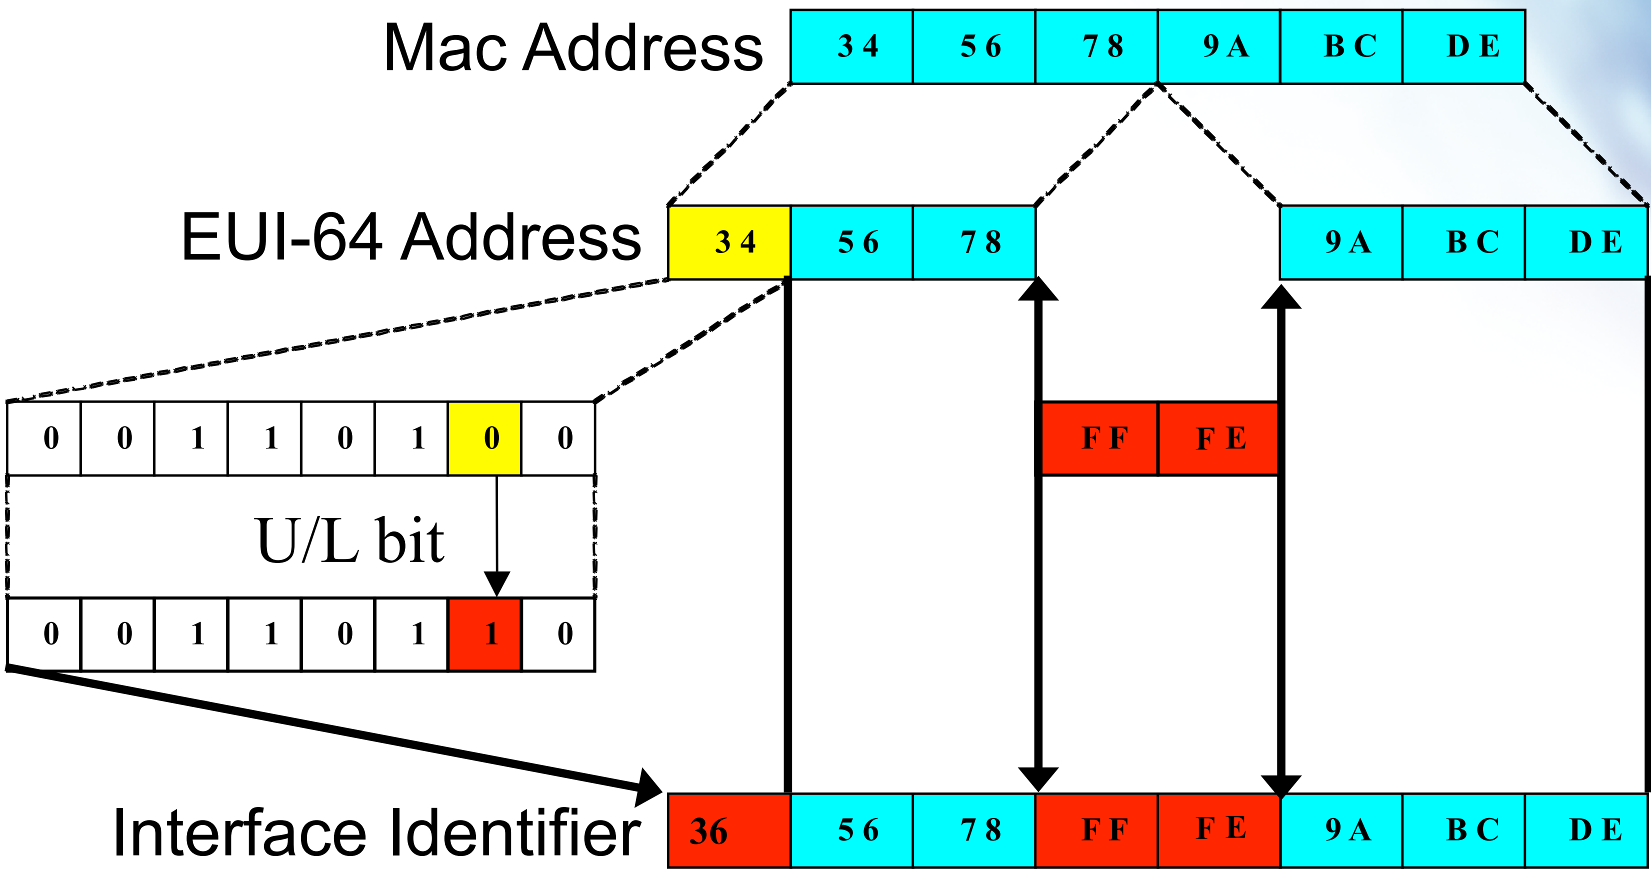
Mac Address
34
56
78
9A
BC
DE
EUI-64 Address
34
56
78
9A
BC
DE
0 0 1 1 0 1 0 0
U/L bit
0 0 1 1 0 1 1 0
FF
FE
Interface Identifier
36
56
78
FF
FE
9A
BC
DE

IPV6 SUBNETTING

IPv6: Think Binary!

- When you do IPv6 subnetting, you need to think in binary values not in hexadecimal value
- 2001:1::/32
  - 2001:0001::/32
  - Hex 2001 = Binary 0010 0000 0000 0001
  - Hex 0001 = Binary 0000 0000 0000 0001
- 2001:2:3::/48
  - 2001:0002:0003::/48
  - Hex 2001 = Binary 0010 0000 0000 0001
  - Hex 0002 = Binary 0000 0000 0000 0010
  - Hex 0003 = Binary 0000 0000 0000 0011

IPv6 address prefix

- /64s in 2001:2:3::/48 are
  - 2001:0002:0003:0001::/64
  - 2001:0002:0003:0002::/64
  - 2001:0002:0003:0003::/64
  - Etc.
  - 16 bits of address space
    - You can have 65536 /64s in one /48 IPv6 address
    - Note:: indicates the remaining 64 bits are all zeros and can then be used to identify hosts::

IPv6 Address Subnetting

- How about /47s in 2001:1::/32?
  - Hex 2001 = Binary 0010 0000 0000 0001 = 16 bits
  - Hex 0001 = Binary 0000 0000 0000 0001 = 32
  - Hex 0000 = Binary 0000 0000 0000 0000 = 47 (32 bits in prefix –"fixed", 15 bits in subnet)
  - So the 15 subnet bits (red) are used to identify the /47s: Subnets numbered using these bits

IPv6 Address Subnetting

- Binary 0000 0000 0000 0000 = Hex 0000
  - The first /47 is 2001:0001:0000::/47
- Binary 0000 0000 0000 0010 = Hex 0002
  - So the second /47 is 2001:0001:0002::/47
- Binary 0000 0000 0000 0100 = Hex 0004
  - So the third /47 is 2001:0001:0004::/47
- Binary 0000 0000 0000 0110 = Hex 0006
  - So the fourth /47 is 2001:0001:0006::/47

Exercise 1: IPv6 Subnetting

- Identify the first four /64 address blocks out of 2001:AA:2000::/48

  1. ______________________________
  2. ______________________________
  3. ______________________________
  4. ______________________________

Exercise 1: IPv6 Subnetting

- Identify the first four /64 address blocks out of 2001:AA:2000:: /64
  1. 2001:AA:2000:0000:: /64
  2. 2001:AA:2000:0001:: /64
  3. 2001:AA:2000:0002:: /64
  4. 2001:AA:2000:0003:: /64

Exercise 2: IPv6 Subnetting

- Identify the first four /36 address blocks out of 2001:ABC::/32
  1. ______________________________
  2. ______________________________
  3. ______________________________
  4. ______________________________

Exercise 2: IPv6 Subnetting

- Identify the first four /36 address blocks out of 2001:ABC::/32
  1. 2001:ABC:0000:: /36
  2. 2001:ABC:1000:: /36
  3. 2001:ABC:2000:: /36
  4. 2001:ABC:3000:: /36

Exercise 3: IPv6 Subnetting

- Identify the first four /37 address blocks out of 2001:AA::/32

1. ______________________________
2. ______________________________
3. ______________________________
4. ______________________________

Exercise 3: IPv6 Subnetting

- Identify the first six /37 address blocks out of 2001:AA::/32
  1. 2001:AA:0000:: /37
  2. 2001:AA:0800:: /37
  3. 2001:AA:1000:: /37
  4. 2001:AA:1800:: /37
  5. 2001:AA:2000:: /37
  6. 2001:AA:2800:: /37

Exercise 4: IPv6 Subnetting

- Identify the first six /54 address blocks out of 2001:AA::/32

  1. ______________________________
  2. ______________________________
  3. ______________________________
  4. ______________________________

Exercise 3: IPv6 Subnetting

- Identify the first six /54 address blocks out of 2001:AA::/32
  1. 2001:00AA:0000:0000::/54
  2. 2001:00AA:0000:0400::/54
  3. 2001:00AA:0000:0800::/54
  4. 2001:00AA:0000:0C00::/54
  5. 2001:00AA:0000:1000::/54
  6. 2001:00AA:0000:1400::/54

AUTOCONFIGURATION

Interface ID

- The lowest-order 64-bit field addresses may be assigned in several different ways:
  - auto-configured from a 48-bit MAC address expanded into a 64-bit EUI-64
  - assigned via DHCP
  - manually configured
  - auto-generated pseudo-random number
  - possibly other methods in the future

EUI-64

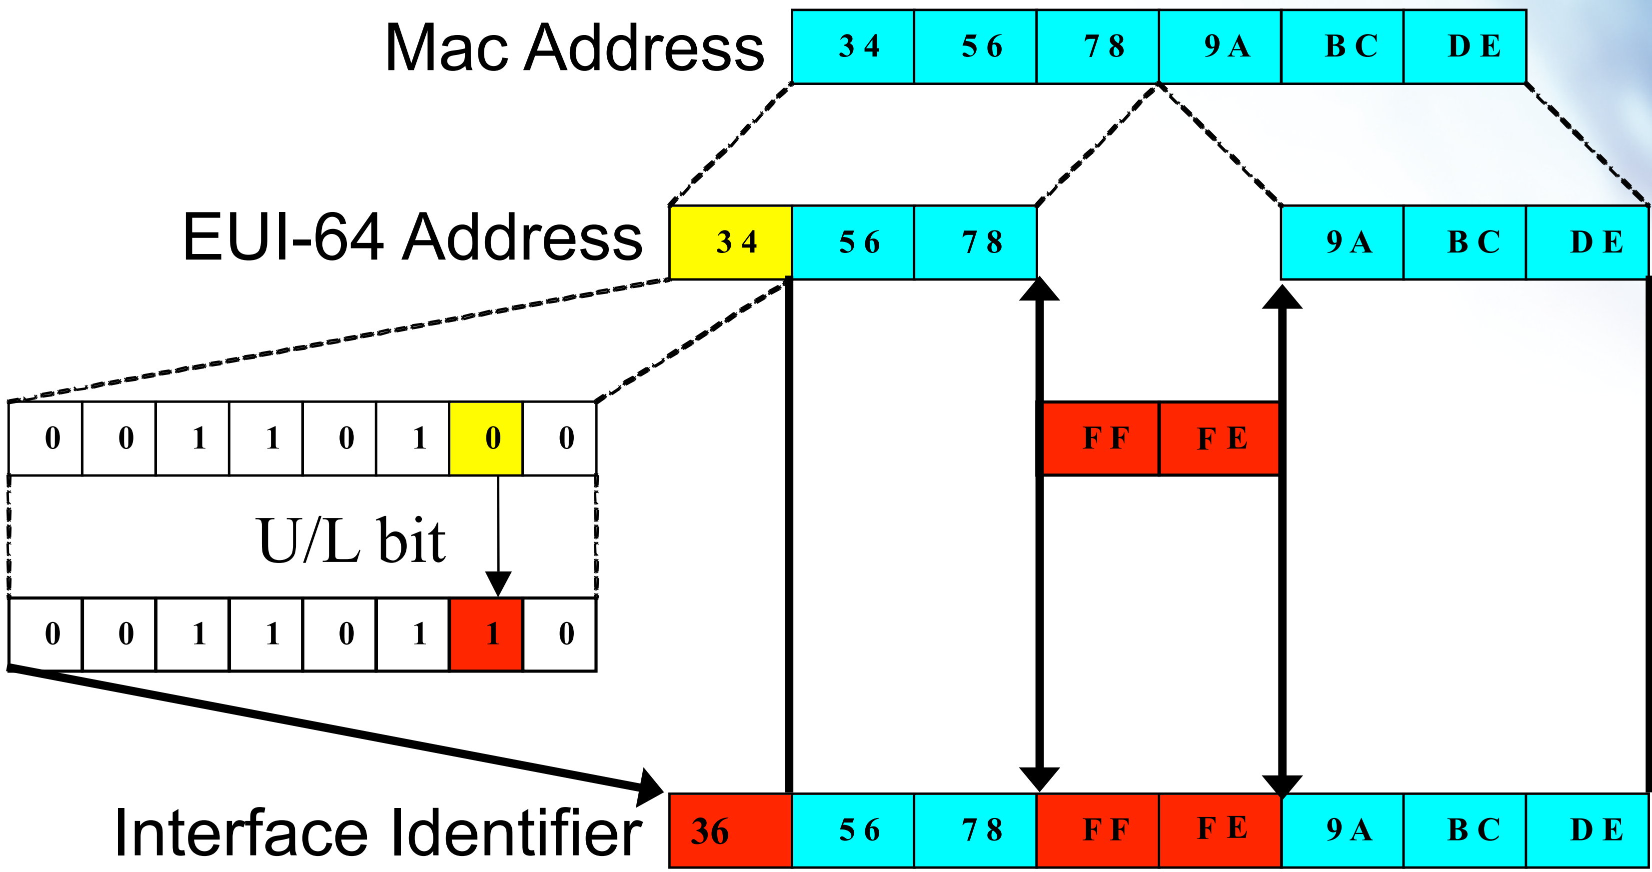

Mac Address
34
56
78
9A
BC
DE
EUI-64 Address
34
56
78
9A
BC
DE
0 0 1 1 0 1 0 0
U/L bit
0 0 1 1 0 1 1 0
FF
FE
Interface Identifier
36
56
78
FF
FE
9A
BC
DE

IPv6 autoconfiguration

- Stateless mechanism
  - For a site not concerned with the exact addresses
  - No manual configuration required
  - Minimal configuration of routers
  - No additional servers
- Stateful mechanism
  - For a site that requires tighter control over exact address assignments
  - Needs a DHCP server
    - DHCPv6

RFC 2462

IPv6 autoconfiguration

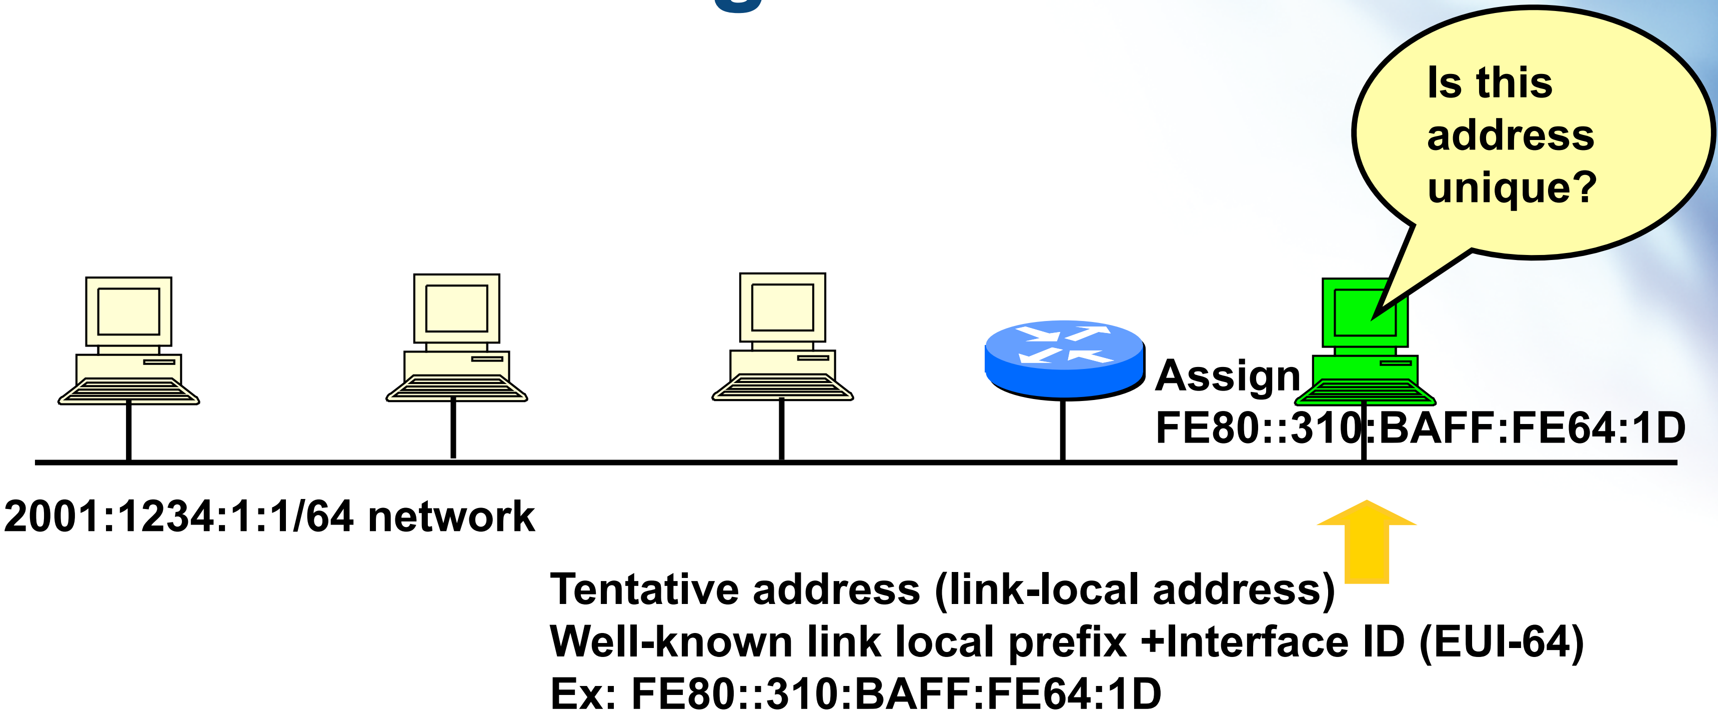

1. A new host is turned on.
2. Tentative address will be assigned to the new host.
3. Duplicate Address Detection (DAD) is performed. First the host transmit a Neighbor Solicitation (NS) message to all-nodes multicast address (FF02::1)
5. If no Neighbor Advertisement (NA) message comes back then the address is unique.
6. FE80::310:BAFF:FE64:1D will be assigned to the new host.

IPv6 autoconfiguration

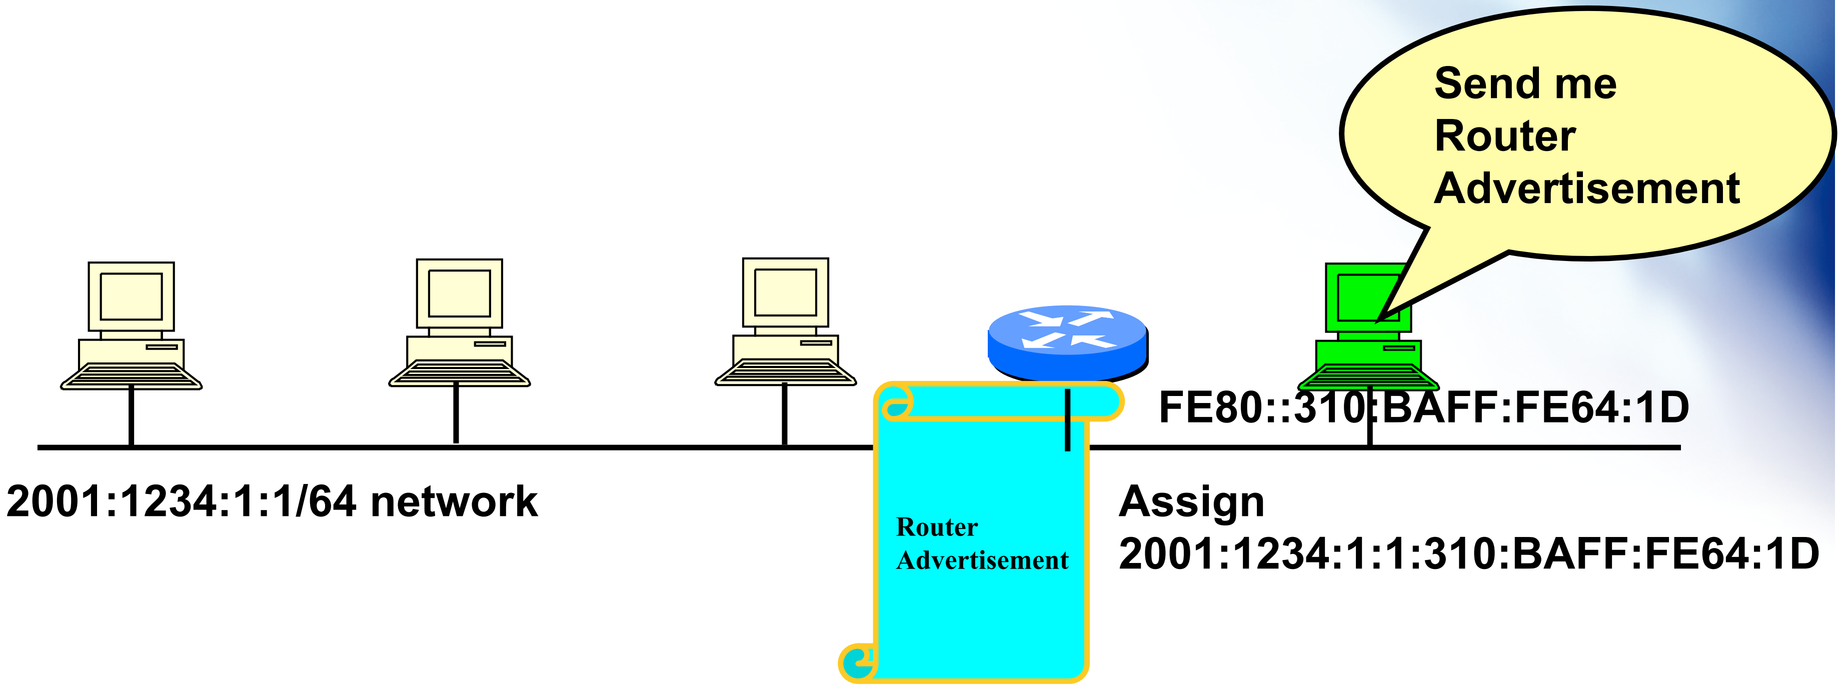

1. The new host will send Router Solicitation (RS) request to the all-routers multicast group (FE02::2).
2. The router will reply Routing Advertisement (RA).
3. The new host will learn the network prefix. E.g, 2001:1234:1:1/64
4. The new host will assigned a new address Network prefix+Interface ID E.g, 2001:1234:1:1:310:BAFF:FE64:1D

IPV6 POLICIES AND PROCEDURES

IPv6 address management hierarchy

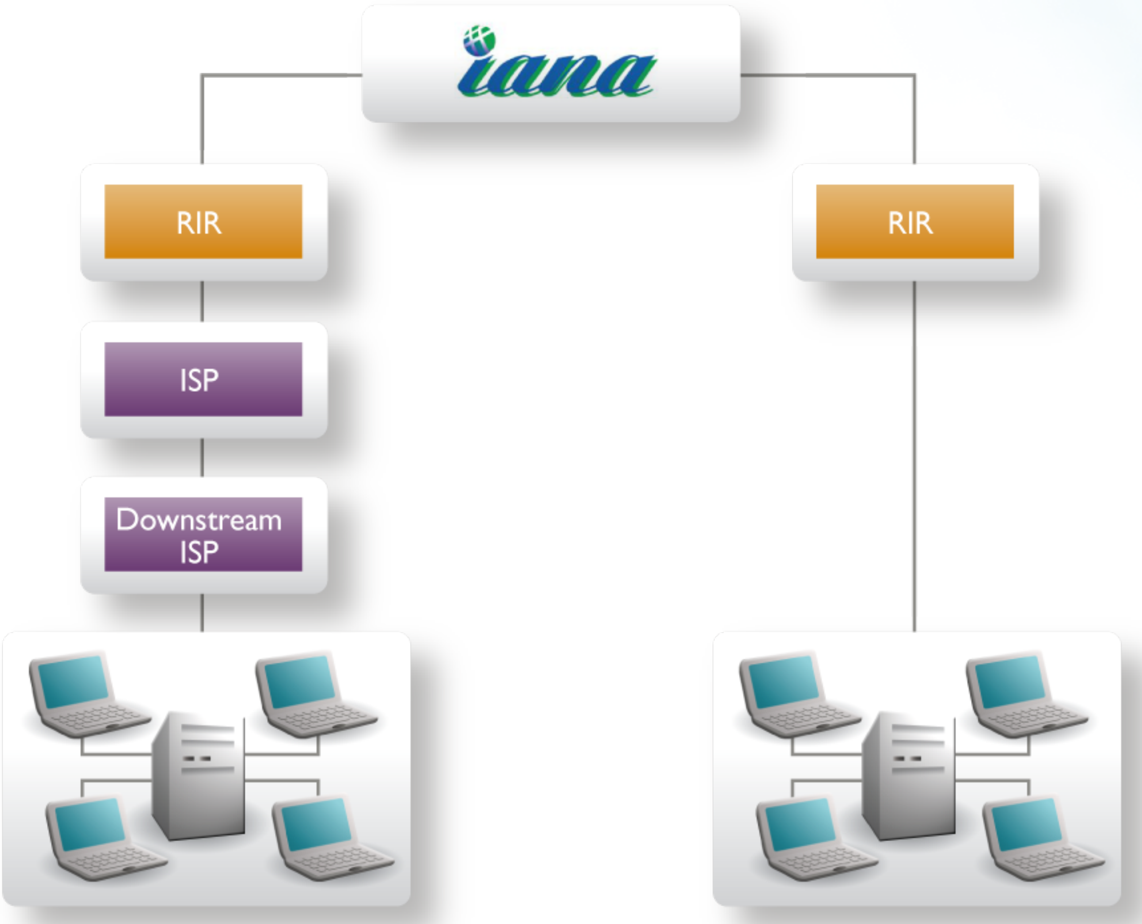

IPv6 address policy goals

- Efficient address usage
  - Avoid wasteful practices
- Aggregation
  - Hierarchical distribution
  - Aggregation of routing information
  - Limiting number of routing entries advertised
- Minimise overhead
  - Associated with obtaining address space
- Registration, Uniqueness, Fairness & consistency

IPv6 initial allocation

- To qualify for an initial allocation of IPv6 address space, an organization must:
  a. Be an ISP
  b. Not be an end site
  c. Plan to provide IPv6 connectivity to organizations to which it will make assignments, by advertising that connectivity through its single aggregated address allocation
  d. Meet one of the two following criteria:
     - Have a plan for making at least 200 assignments to other organizations within two years OR
     - Be an existing ISP with IPv4 allocations from an APNIC or an NIR, which will make IPv6 assignments or sub-allocations to other organizations and announce the allocation in the inter-domain routing system within two years

IPv6 initial allocation

- Private networks (those not connected to the public Internet) may also be eligible for an IPv6 address space allocation provided they meet equivalent criteria to those listed above.
- Initial allocation size is /32
  - Default allocation (“slow start”)

End Sites

- An end site is defined as "an end user (subscriber) who has a business relationship with a service provider”
- End sites do not re-assign any of their IP addresses to other organisations

Typical IPv6 Applicants

- An ISP providing IPv6 connectivity to the global Internet.
- An ISP providing IPv6 services to end sites and restricting connectivity to its own closed network.
- An ISP providing IPv6 services to end sites and restricting connectivity to peering partners.
- A large organisation providing IPv6 connectivity to its group companies or subsidiaries and restricting connectivity to its own network.

IPv6 initial allocation

- Initial allocations larger than /32 may be justified if:
  a. The organization provides comprehensive documentation of planned IPv6 infrastructure which would require a larger allocation; or
  b. The organization provides comprehensive documentation of all of the following:
     - its existing IPv4 infrastructure and customer base,
     - its intention to provide its existing IPv4 services via IPv6, and
     - its intention to move some of its existing IPv4 customers to IPv6 within two years.

End site assignment

- End-users are assigned an end site assignment from their LIR or ISP.
- The exact size of the assignment is a local decision for the LIR or ISP to make:
  - using a minimum value of a /64
  - up to the normal maximum of /48

IPv6 utilisation

- Utilisation determined from end site assignments
  - ISP responsible for registration of all /48 assignments
  - Intermediate allocation hierarchy not considered
- Utilisation of IPv6 address space is measured differently from IPv4
  - Use HD ratio to measure
- Subsequent allocation may be requested when IPv6 utilisation requirement is met

Subsequent allocation

- Must meet HD = 0.94 utilisation requirement of previous allocation (subject to change)
- Other criteria to be met
  - Correct registrations (all /48s registered)
  - Correct assignment practices etc
- Subsequent allocation results in a doubling of the address space allocated to it
  - Resulting in total IPv6 prefix is 1 bit shorter
  - Or sufficient for 2 years requirement

HD Ratio

- The HD ratio threshold is
  - HD=log(/56 units assigned) / log (16,777,216)
  - 0.94 = 6,183,533 x /56 units
- Calculation of the HD ratio
  - Convert the assignment size into equivalent /56 units
    - Each /48 end site = 256 x /56 units
    - Each /52 end site = 16 x /56 units
    - Each /56 end site = 1 x /56 units
    - Each /60 end site = 1/16 x /56 units
    - Each /64 end site = 1/256 x /56 units

IPv6 utilisation (HD = 0.94)

- Percentage utilisation calculation

IPv6 Prefix	Site Address Bits	Total site address in /56s	Threshold (HD ratio 0.94)	Utilisation %
/42	14	16,384	9,153	55.9%
/36	20	1,048,576	456,419	43.5%
/35	21	2,097,152	875,653	41.8 %
/32	24	16,777,216	6,185,533	36.9%
/29	27	134,217,728	43,665,787	32.5 %
/24	32	4,294,967,296	1,134,964,479	26.4 %
/16	40	1,099,511,627,776	208,318,498,661	18.9 %

RFC 3194

> "In a hierarchical address plan, as the size of the allocation increases, the density of assignments will decrease."

IPv6 portable assignment for multihoming

- An organization is eligible to receive a portable assignment from APNIC if it is currently multihomed or plans to be multihomed within three months.
- An organization is considered to be multihomed if its network receives full-time connectivity from more than one ISP and has one or more routing prefixes announced by at least two of its ISPs.

IPv6 portable assignment for multihoming

- The minimum assignment made under these terms is /48.

Internet Exchange Points

- Internet Exchange Points (IXP) are eligible to receive a portable assignment from APNIC to be used exclusively to connect the IXP participant devices to the Exchange Point.
- Criteria
  - Demonstrate ‘open peering policy’
  - 3 or more peers
- Portable assignment size: /48

APNIC IPv6 delegation by economy

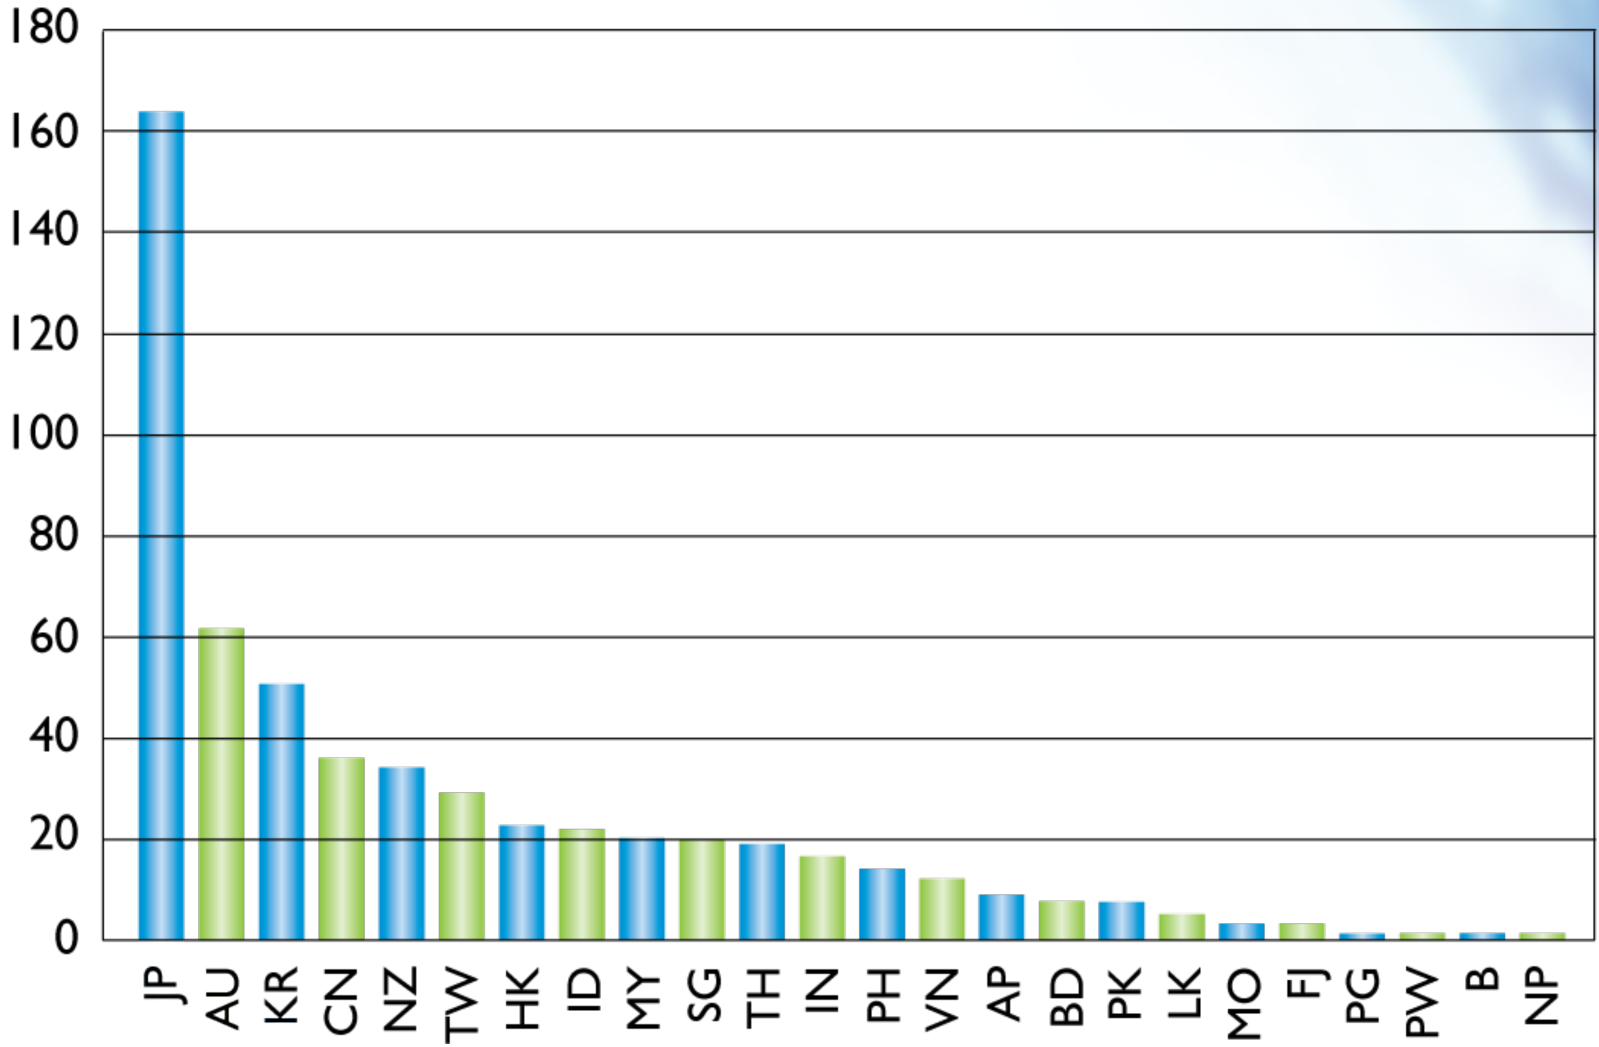

No of delegations (/35, /32)

http://www.apnic.net/stats/o3/ as of 26/03/2009

How do I apply for IPv6 addresses?

Check your eligibility for IPv6 addresses

Read IPv6 policies
http://www.apnic.net/docs/policy/ipv6-address-policy.html

Read IPv6 guideline
http://www.apnic.net/docs/policy/ipv6-guidelines.html

Do you have an APNIC account?

If not, become an APNIC member or open a non-member account

Complete an IPv6 address request form

Submit the form hostmaster@apnic.net

Questions:
email: helpdesk@apnic.net
Helpdesk chat: http://www.apnic.net/helpdesk

APNIC Helpdesk chat

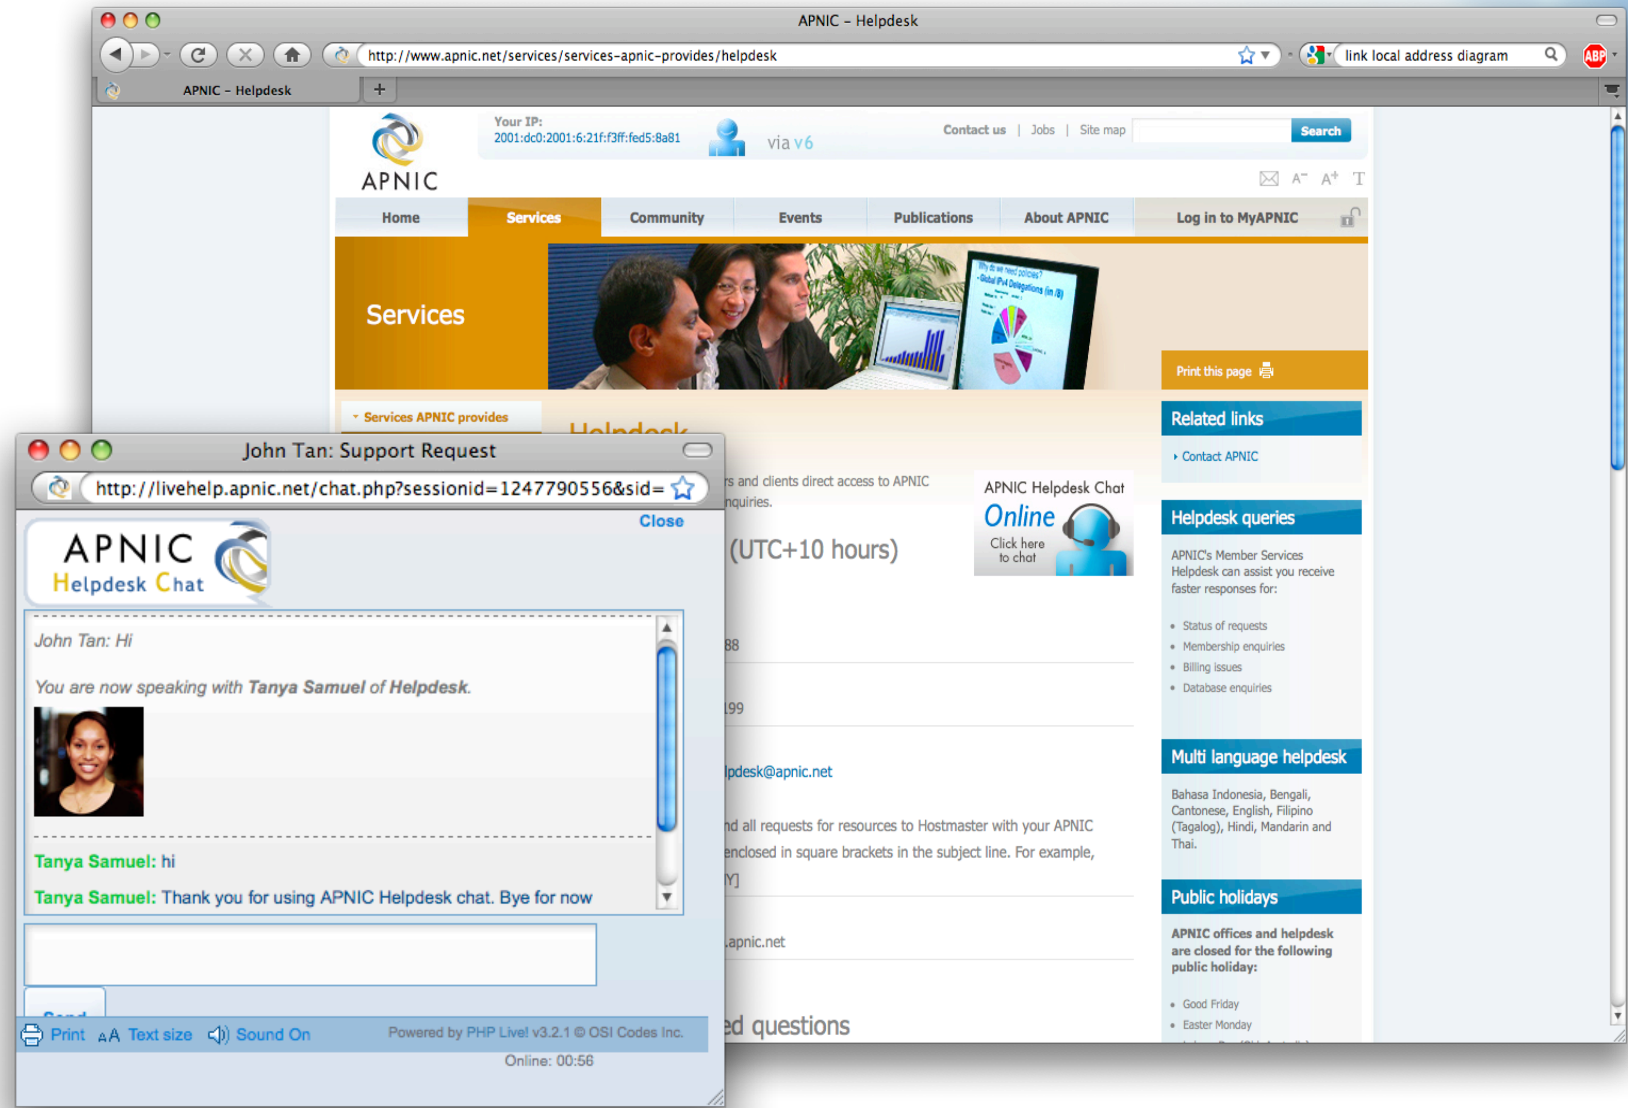

Asia Pacific Network Information Centre

APNIC

Discuss IPv6 Transition – IPv6 Wiki

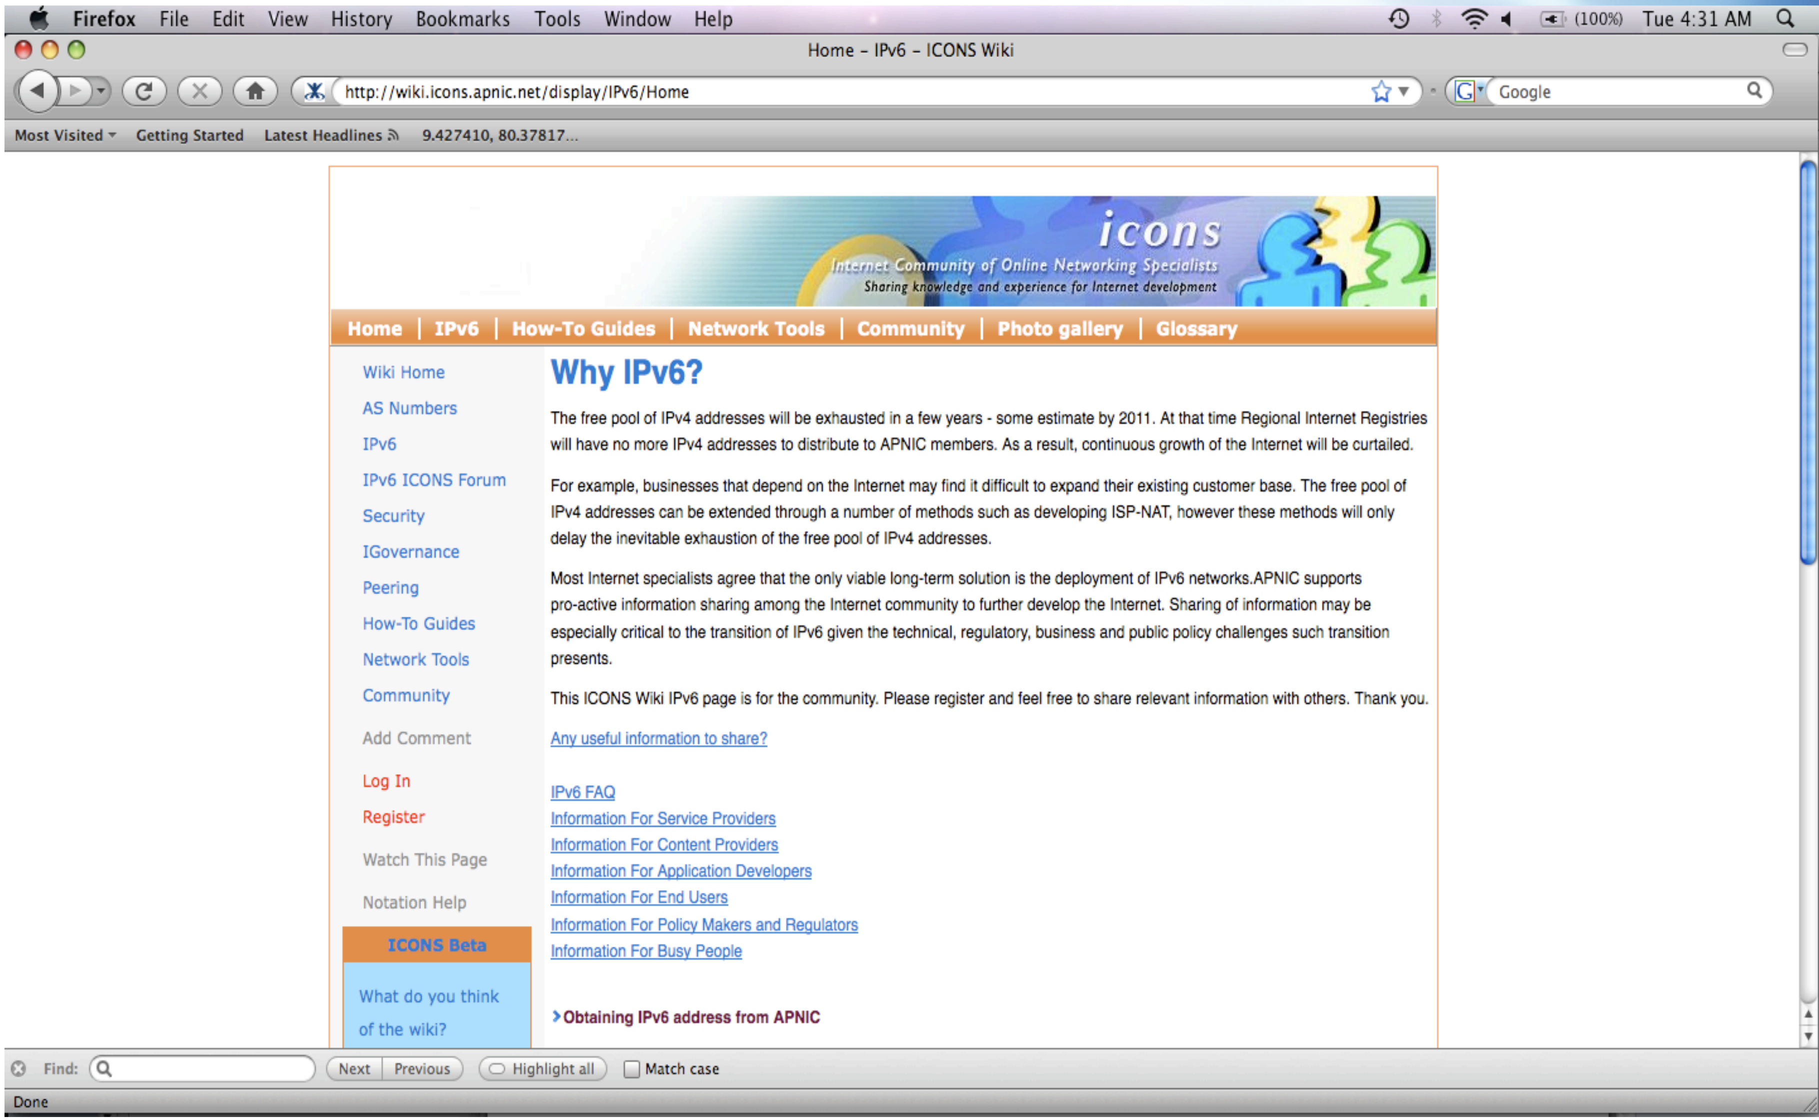

IPv6 policy – have your say!

- Limited experience of policy in action
  - Your feedback very important
  - Policy always subject to change and refinement
- Open discussion list
  - global-v6@lists.apnic.net (all regions)
  - SIG Policy mailing list (APNIC region)
- Documentation
  - FAQ information and more!
    - http://www.apnic.net/services/ipv6_guide.html
  - Guidelines document under development
    - To assist new requestors with policy

IPV4 TO IPV6 TRANSITION

Transition overview

- How to get connectivity from an IPv6 host to the global IPv6 Internet?
  – Via an native connectivity
  – Via IPv6-in-IPv4 tunnelling techniques
- IPv6-only deployments are rare
- Practical reality
  – Sites deploying IPv6 will not transit to IPv6-only, but transit to a state where they support both IPv4 and IPv6 (dual-stack)

http://www.6net.org/book/deployment-guide.pdf p59

IPv4 to IPv6 transition

- Implementation rather than transition
  - No fixed day to convert
- The key to successful IPv6 transition
  - Maintaining compatibility with IPv4 hosts and routers while deploying IPv6
    - Millions of IPv4 nodes already exist
    - Upgrading every IPv4 nodes to IPv6 is not feasible
      - No need to convert all at once
    - Transition process will be gradual

Transition overview

- Transition approaches
  - Dual stack
  - Deploying IPv6 and then implementing IPv6-in-IPv4 tunnelling
  - IPv6 only networking
- Different demands of hosts and networks to be connected to IPv6 networks will determine the best way of transition

Transition overview

- Dual stack
  - Allow IPv4 and IPv6 to coexist in the same devices and networks
- Translation
  - Allow IPv6 only nodes to communicate with IPv4 only nodes
- Tunnelling
  - Allow the transport of IPv6 traffic over the existing IPv4 infrastructure

IPv6 essentials by Silvia Hagen, p255

Dual stack transition

- Dual stack = TCP/IP protocol stack running both IPv4 and IPv6 protocol stacks simultaneously
  - Application can talk to both
- Useful at the early phase of transition

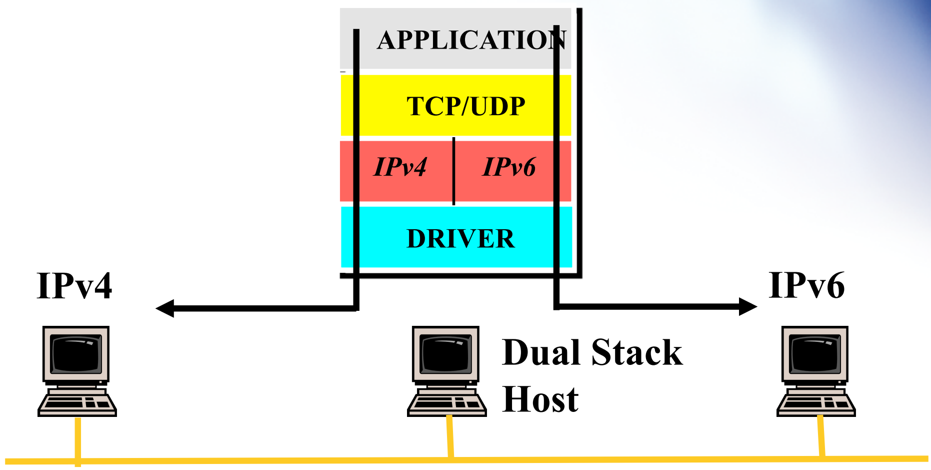

RFC 4213

Dual stack

- A host or a router runs both IPv4 and IPv6 in the protocol TCP/IP stack.
- Each dual stack node is configured with both IPv4 and IPv6 addresses
- Therefore it can both send and receive datagrams belonging to both protocols
- The simplest and the most desirable way for IPv4 and IPv6 to coexist

Dual stack network

- Is an infrastructure in which both IPv4 and Ipv6 forwarding is enabled on routers

Dual stack Challenges

- Compatible software
  - Eg. If you use OSPFv2 for your IPv4 network you need to run OSPFv3 in addition to OPSFv2
- Transparent availability of services
  - Deployment of servers and services
  - Content provision
  - Business processes
  - Traffic monitoring
  - End user deployment

Dual stack and DNS

- DNS is used with both protocol versions to resolve names and IP addresses
  - An dual stack node needs a DNS resolver that is capable of resolving both types of DNS address records
    - DSN A record to resolve IPv4 addresses
    - DNS AAAA record to resolve IPv6 addresses

IPv6 essentials by Silvia Hagen, p256

Translation

- Involves direct conversion of protocols
- May include transformation of both the protocol header and the protocol payload

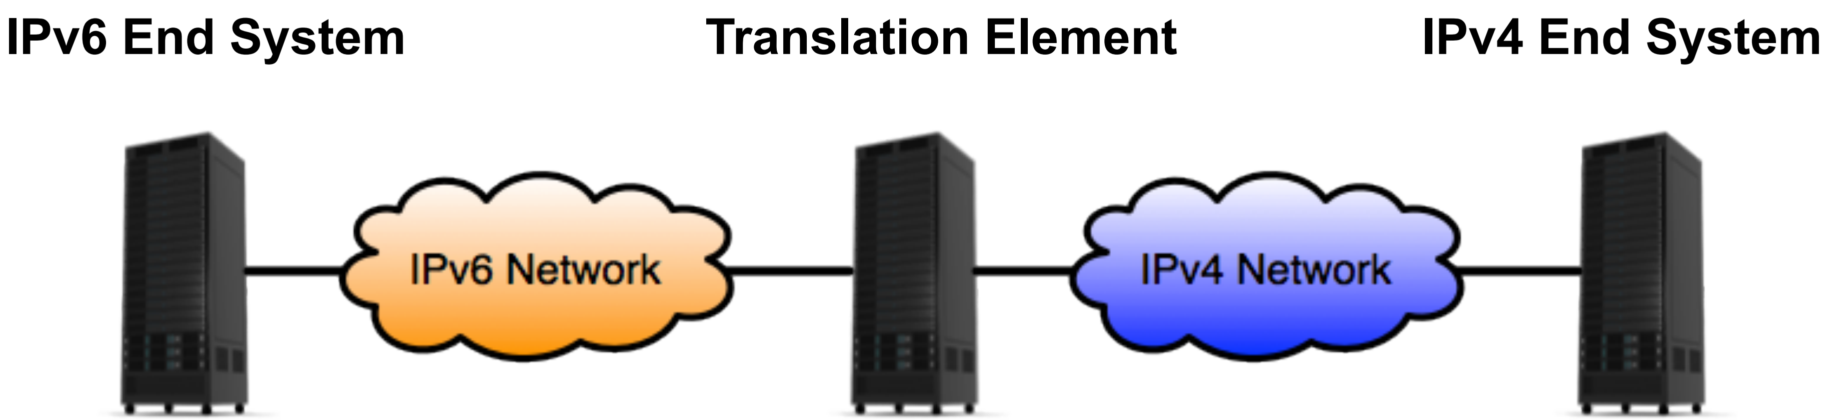

Translation

- Stateless Internet Protocol/Internet Control Messaging Protocol Translation (SIIT)
- NAT-Protocol Translation (Deprecated)
- Transport Relay Translator

Tunnels

- Part of a network is IPv6 enabled
  - Tunnelling techniques are used on top of an existing IPv4 infrastructure and uses IPv4 to route the IPv6 packets between IPv6 networks by transporting these encapsulated in IPv4
  - Tunnelling is used by networks not yet capable of offering native IPv6 functionality
  - It is the main mechanism currently being deployed to create global IPv6 connectivity
- Static, semi-automatic & automatic configured tunnels are available

Tunneling – general concept

- Tunneling can be used by routers and hosts
  - Tunneling is a technique by which one transport protocol is encapsulated as the payload of another.

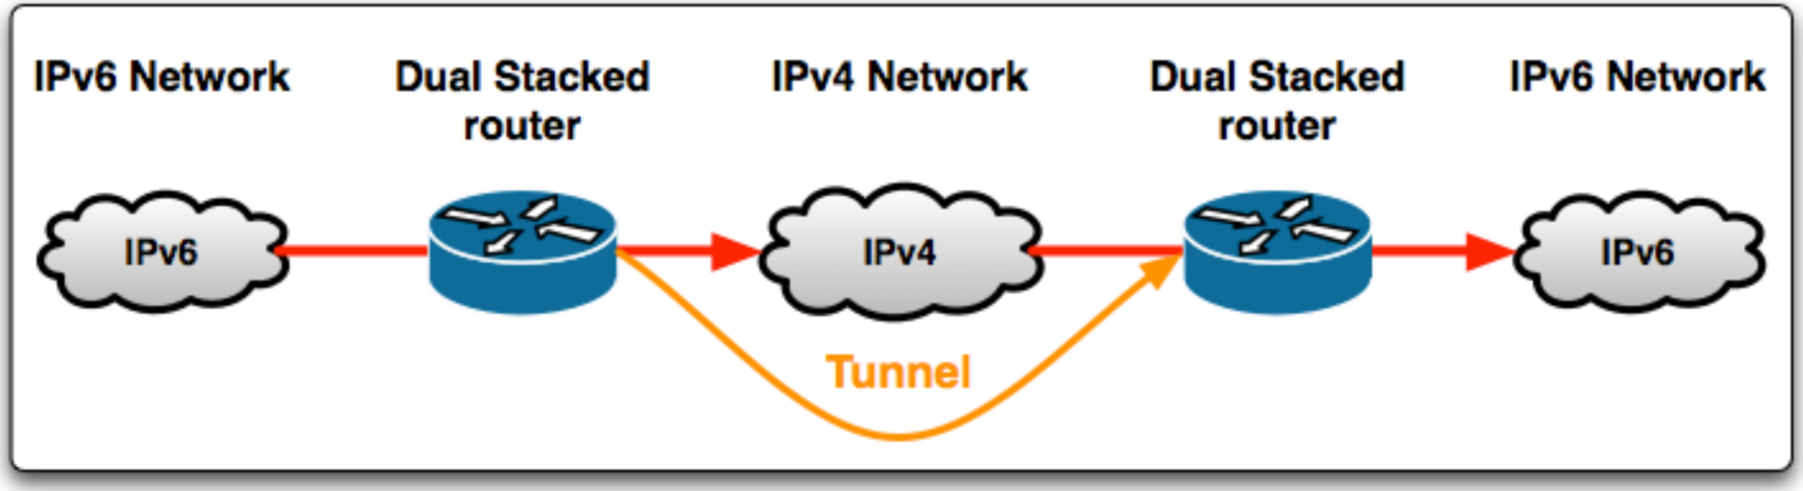

Tunnelling – general concept

- A tunnel can be configured in four different ways:

Router to Router	Spans one hop of the end-to-end path between two hosts Probably the most common method
Host to Router	Spans the first hop of the end-to-end path between two hosts Found in the tunnel broker model
Host to Host	Spans the entire end-to-end path between two hosts
Router to Host	Spans the last hop of the end-to-end path between two hosts

Tunneling – general concept

- Two stepped process
  1. Encapsulation of IPv6 packets to IPv4 packets
  2. Decapsulation of IPv4 packets to IPv6 packets

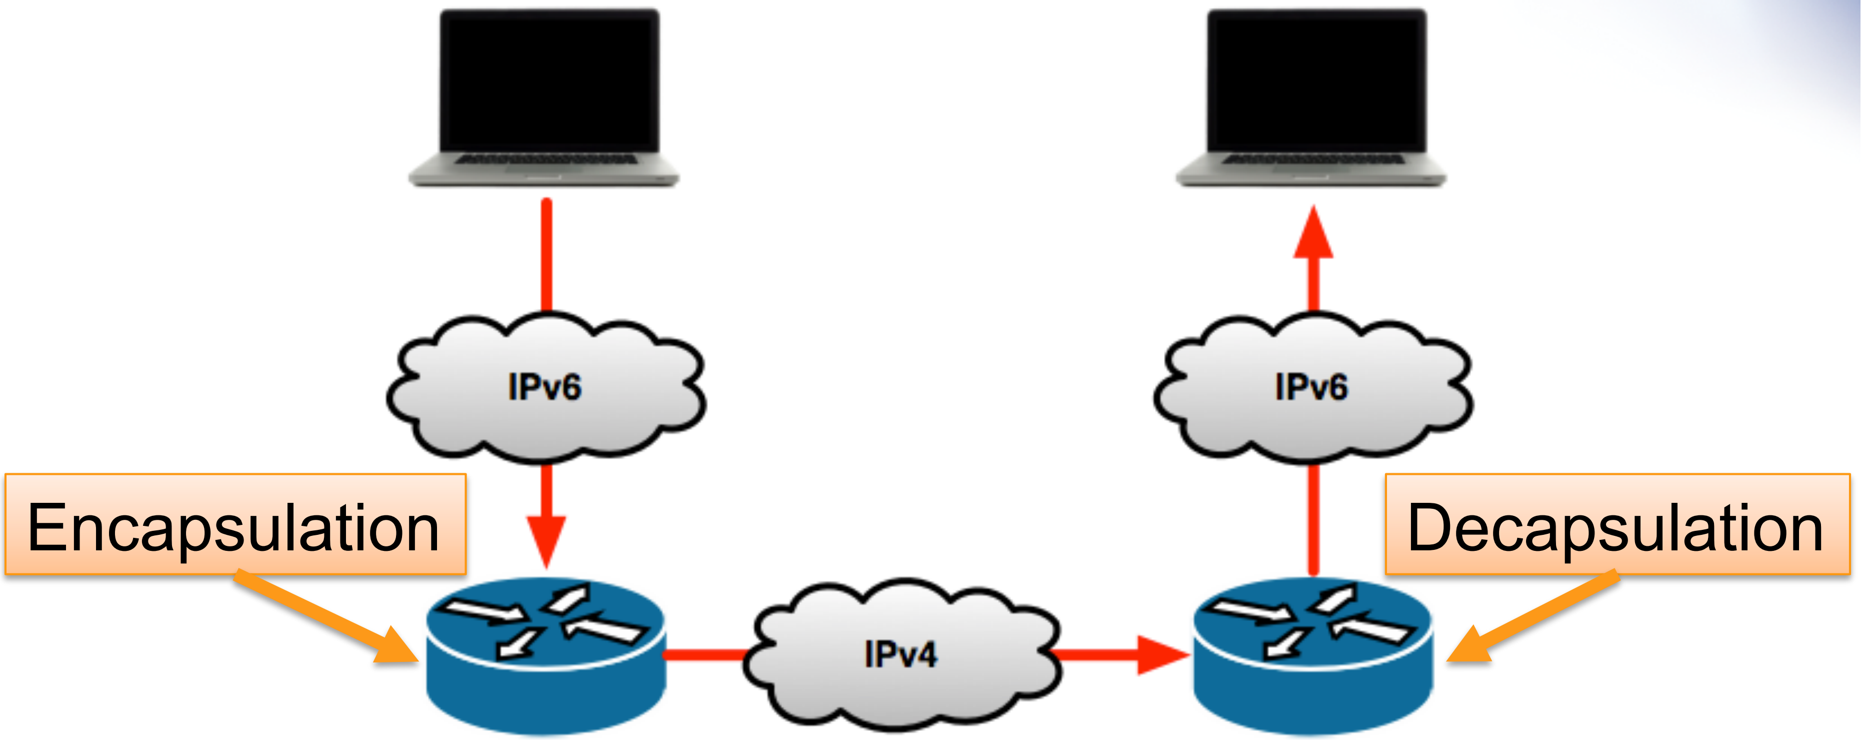

Tunnel encapsulation

- The steps for the encapsulation of the IPv6 packet
  - The entry point of the tunnel decrements the IPv6 hop limit by one
  - Encapsulates the packet in an IPv4 header
  - Transmits the encapsulated packet through the tunnel

IPv6 essentials by Silvia Hagen, p258

Tunnel encapsulation

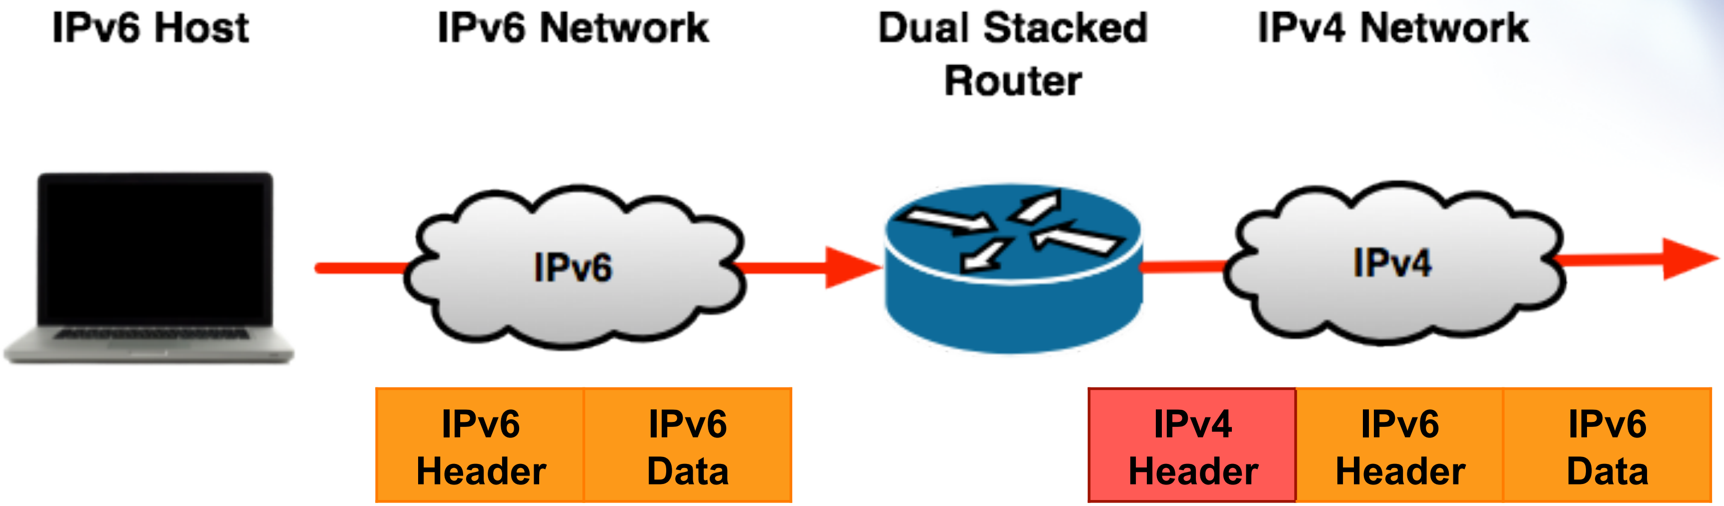

IPv6 essentials by Silvia Hagen, p258

Tunnel decapsulation

- The exit point of tunnel receives the encapsulated packet
  - If necessary, the IPv4 packet is fragmented
- It checks whether the source of the packet (tunnel entry point) is an acceptable source (according to its configuration)
  - If the packet is fragmented, the exit point reassembles it
- The exit point removes the IPv4 header
- Then it forwards the IPv6 packet to its original destination

IPv6 essentials by Silvia Hagen, p258

Tunnel decapsulation

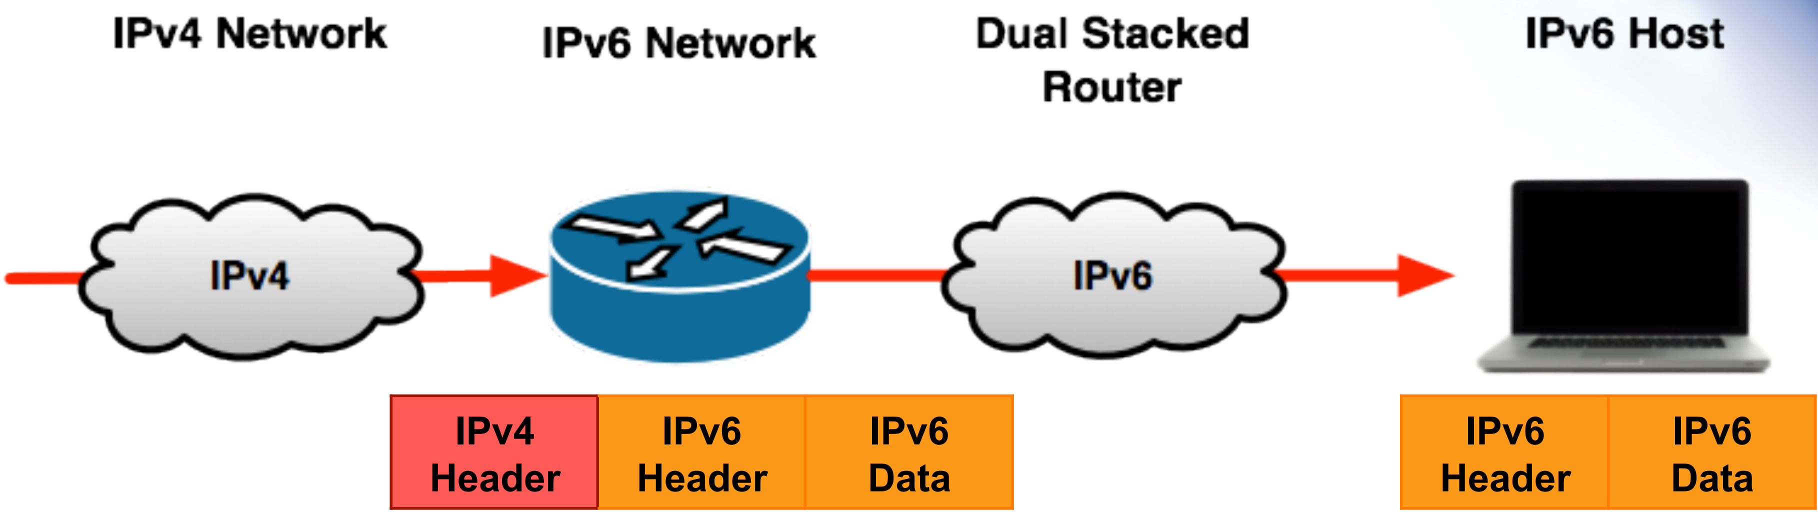

IPv6 essentials by Silvia Hagen, p258

Manual tunneling

RFC 4213

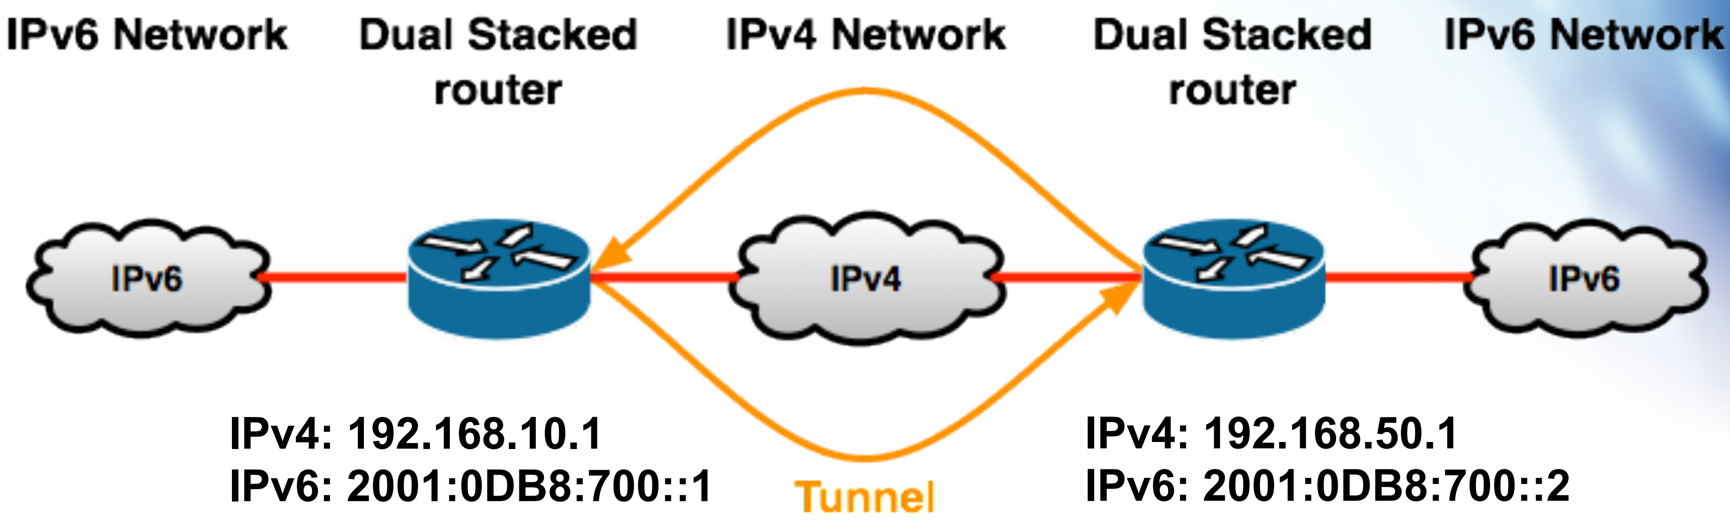

Manually configured tunnels require:

- Dual stack end points
- Explicit configuration with both IPv4 and IPv6 addresses at each end

Concept is borrowed from Cisco, Training material "Ipv6 Seminar" delivered at South Asian IPv6 Summit, Jan 2004

Tunnel broker

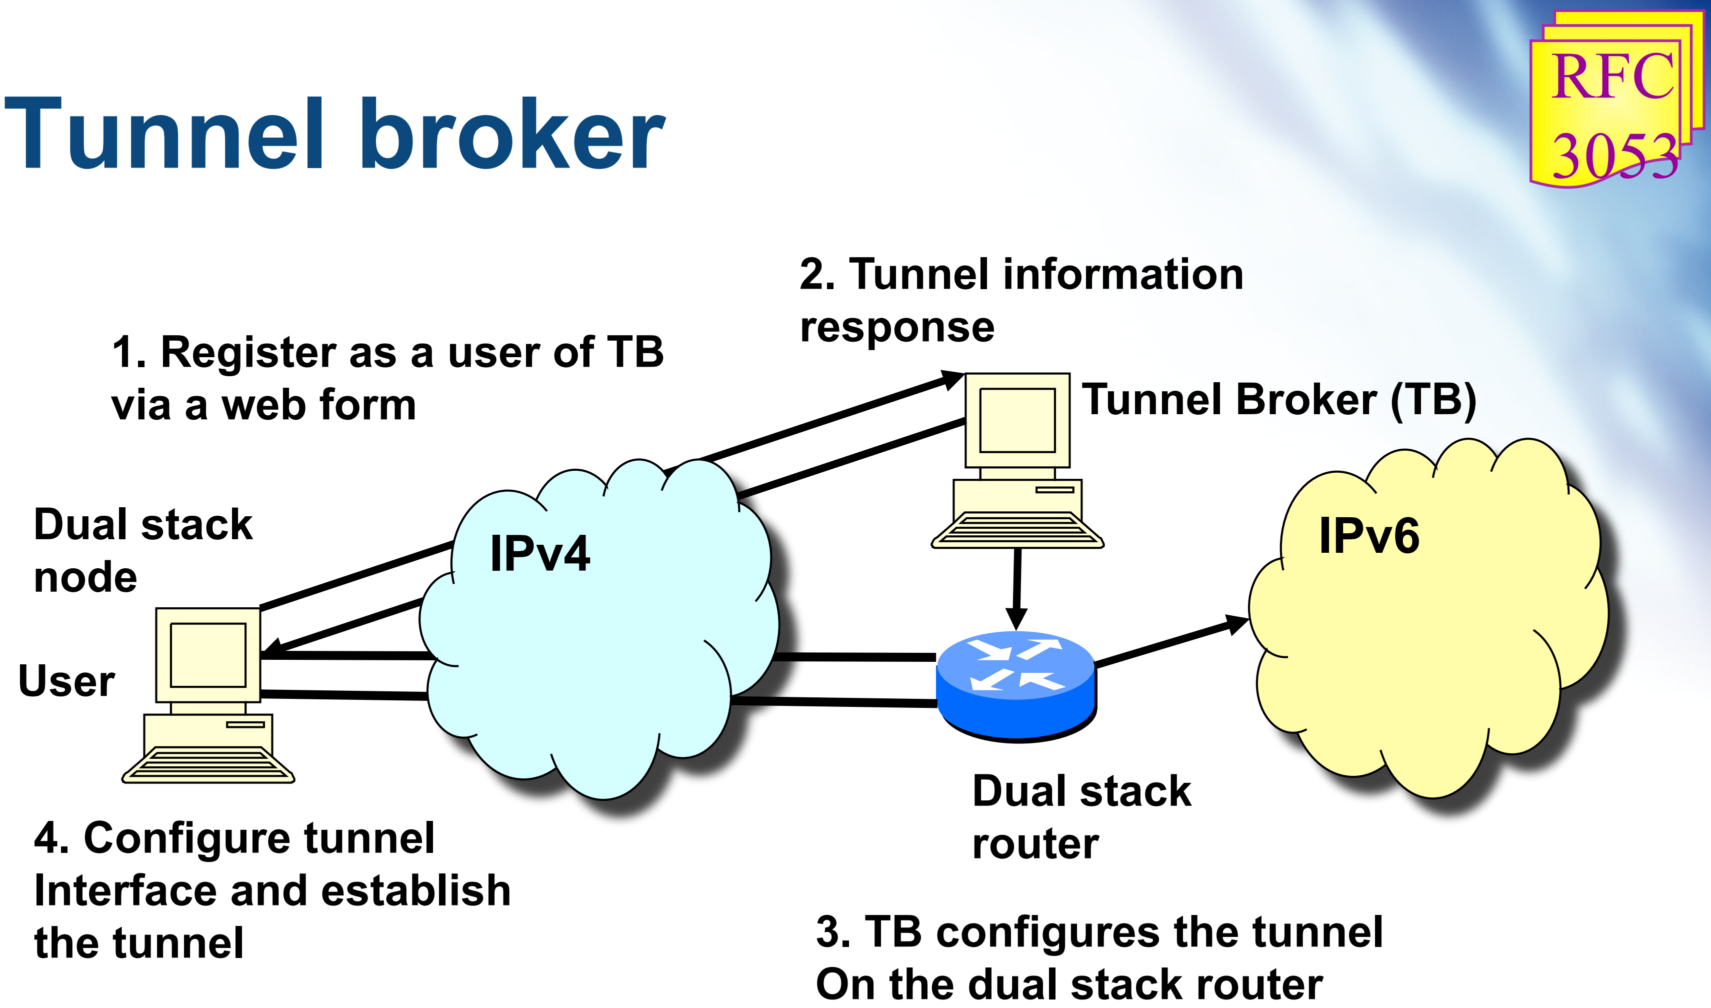

Automatic tunneling – 6to4

RFC 3056 RFC 3068

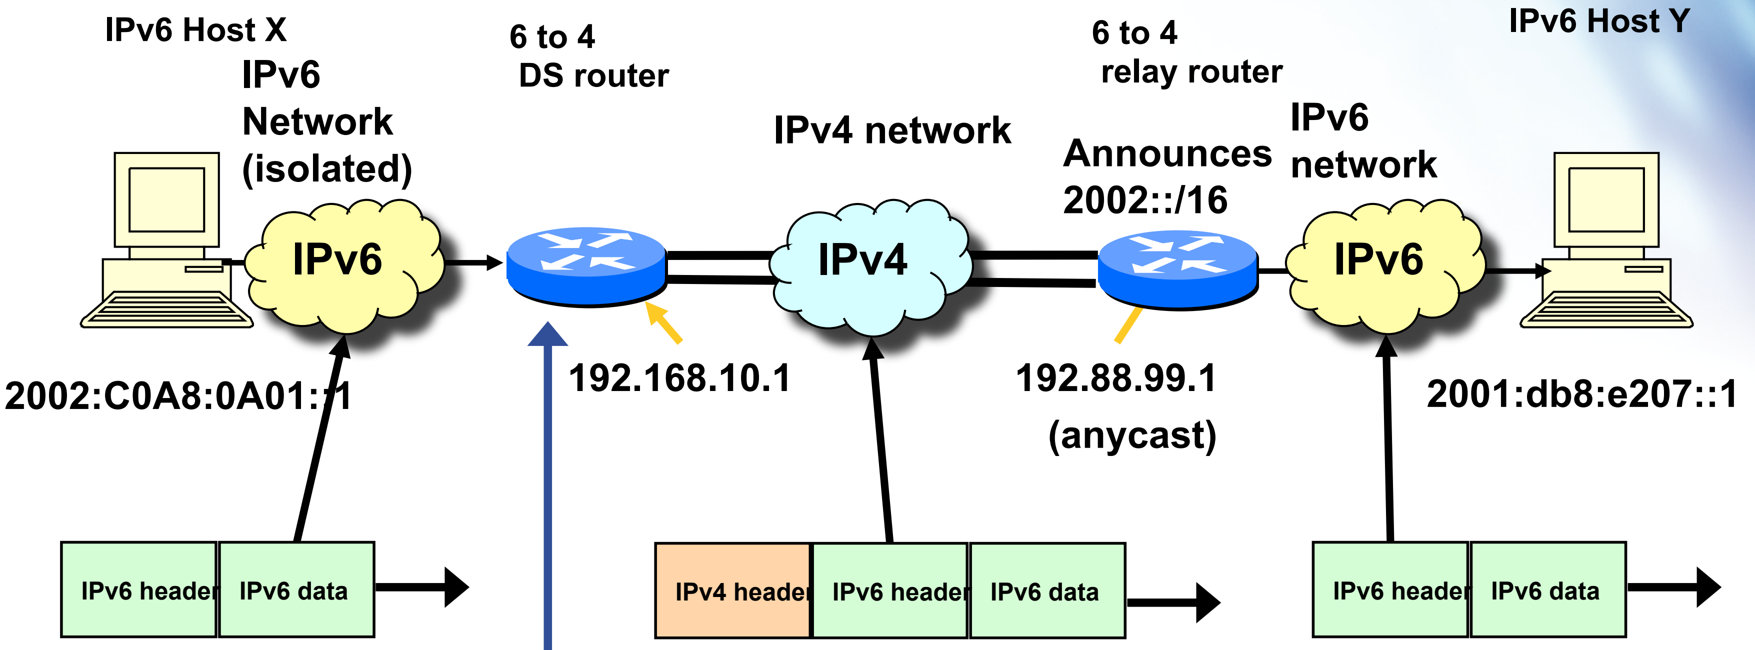

S=2002:C0A8:0A01::1
D=2001:db8:e207::1

S (v4)=192.168.10.1
D (v4)=192.88.99.1
S (v6)=2002:C0A8:0A01::1
D (v6)=2001:db8:e207::1

S=2002:C0A8:0A01::1
D=2001:db8:e207::1

Default IPv6 route is 2002:co58:6301::

A destination route to a 2002::/ prefix is encapsulated in IPv4 and bits 17 – 48 used as the next hop. Ie 192.88.99.1 anycast

6 to 4 Tunnelling

- A IPv6 network prefix is defined for the isolated network using the 2002::/16 prefix and then adding the IPv4 address of the DS router in the next 32 bits (17 – 48).
- The DS router is then assigned an address in this network as is the IPv6 host.

6 to 4 Tunnelling

- The host configures the IPv6 address of the DS router as its default IPv6 route.
- The DS router defines a default IPv6 route using the anycast IPv4 Relay Router address, 192.88.99.1, enbedded in bits 17-48 of an IPv6 address using the 2002::/ prefix.

6 to 4 Tunnelling

- When an IPv6 packet sent by the IPv6 host is routed to the DS router, it in turns uses its default route to the relay router to forward this.
- Because the default route is to a 2002::/ prefix, the packet is encapsulated and the destination IPv4 address is extracted from bits 17 – 48 of the IPv6 default route. That is, it is forwarded to 192.99.88.1

Relay routers

- The Relay Routers are connected to the IPv4 network and announce routes to 192.99.88.0/24
- They are also connected to the IPv6 network and announce routes to 2002::/16
- A packet to the 2002::/16 address of host on a distant network will then be routed from the IPv6 network to a RR.

Relay routers

- Because the destination is a 2002::/16 address the RR will encapsulate this, extract the 32 bit IPv4 address and forward the packet to this address.
- This is the address of the DS router connecting the distant IPv6 network.
- The DS router will receive the packet, decapsulate the packet and forward to the IPV6 address of the host.

Questions?

Thank you!

training@apnic.net