C²S CONSULTING
communications/computing/security

Internet Protocol – Next Generation
Internet Protocol – Version 6

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

2 Billion users on the Internet

IPv4 is has now exhausted

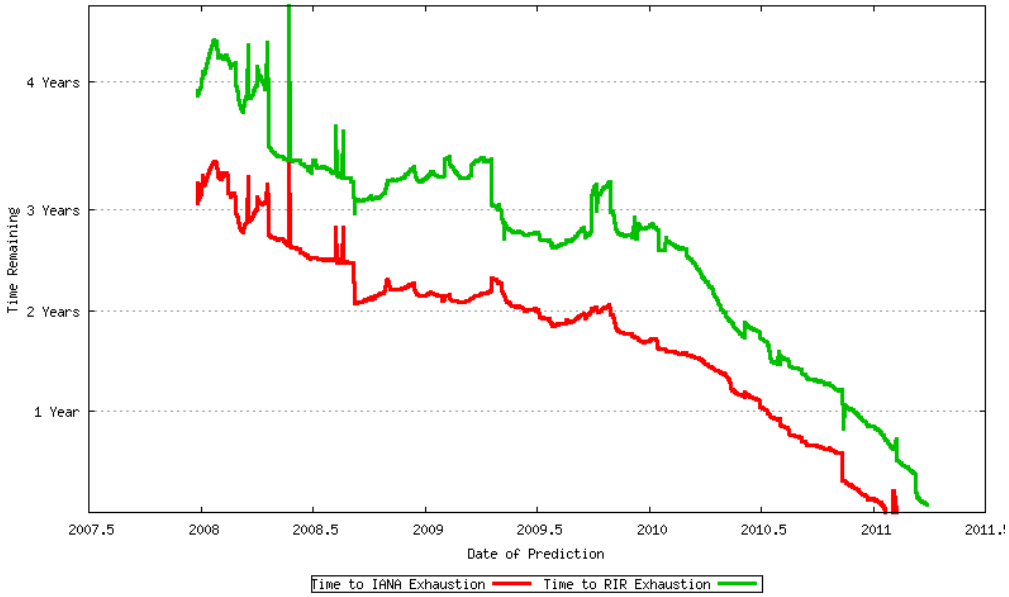

IANA

RIRs

- IANA have no IP addresses
- RIRs are have also run out
- AfriNIC actually ran out on 3 April 2017 !!
- Black market €12/IP and NAT { Bad idea }

Where does she get her IP address ?

Solution ?

IPv4 extension

- RFC 3022 - Traditional IP Network Address Translator
- 2017

Replace IPv4

- Forklift upgrade ??

Solution ?

IPv4 extension

- ~~RFC 3022 - Traditional IP Network Address Translator~~
- 2017 

Replace IPv4 with IPv6

- 128 bit addressing – 3.4×10^{38}
  - IPv4 – 4.3×10^{9}
  - 7.9×10^{28} more addresses
- "**Dual Stack**" for compatibility with IPv4
- Similar routing concepts to IPv4

Why IPv6 ? – Much larger address space

- IPv4 address exhaustion
  - 12 addresses per person in USA
  - 14 people per address outside USA
  - 0.62 addresses per person on the planet.
- IPv6
  - 4.9×10^{28} address per person on the planet.
  - Greater than 1500 addresses per m2 of the planet.
  - Routing table growth and manageability.
  - Current addressing and routing are complicated.
  - Some routes can not be summarised effectively in IPv4.
  - NATs are disadvantageous for the Internet.

Why IPv6 ?

- Multicast
  - Multicast is part of the base protocol suite in IPv6. In IPv4 it is optional.
- Jumbograms
  - Packets payloads greater than 64 KB.
- Faster routing.
- Network Layer Security
  - IPSec.
- Mobility.

IPv6 Security

- Security features are standardised and mandated, i.e. all implementation must offer them.
- Extension of RFC-2401 Security Architecture for the Internet Protocol (IPSec)
  - Authentication and Encryption.
- Invisible to applications as it operates within the IP layer.
- It protects all upper layer protocols.
- It protects both end-to-end and router-to-router communication (secure gateway).

Mobile IPv6

- IPv6 was designed to support mobility.
- Mobility is not an add-on feature
  - IPv6 Neighbour Discovery (ND) and Address Auto-configuration allow hosts to operate in any locations without any special support.
- It is more scalable and the performance is better because less traffic passes through the home link and less redirection and less rerouting.
- No single point of failure.

C²S CONSULTING
communications/computing/security

Address Format

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 Addressing model

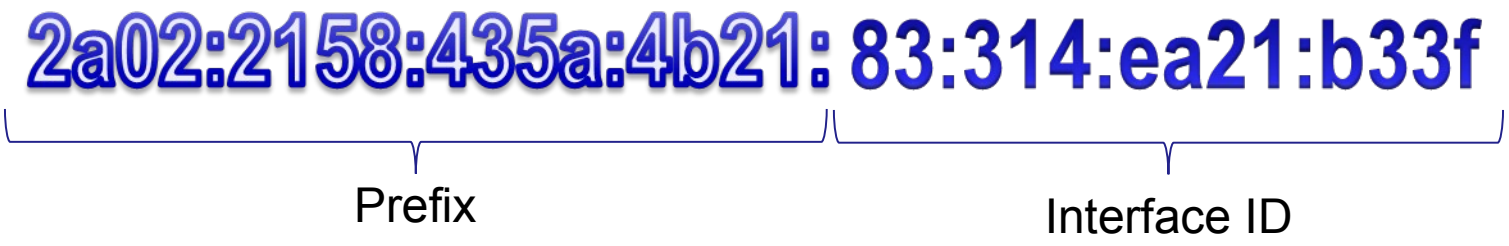

- Separation of “who u are” from “where u are connected to”
  - Prefix: depends on the routing topology.
  - Interface ID: identifies a node (EUI-64 MAC Address).
- New Anycast addresses:
  - Unicast: from one host to another.
  - Multicast: from one to all belonging to a group.
  - Anycast: from one to the nearest belonging to a group.
- IPv4 Broadcast concept disappears
  - Replaced by All hosts multicast.

128 bits
↓
32 nibbles
↓
8 x 4 nibble groups

IPv6 Address model

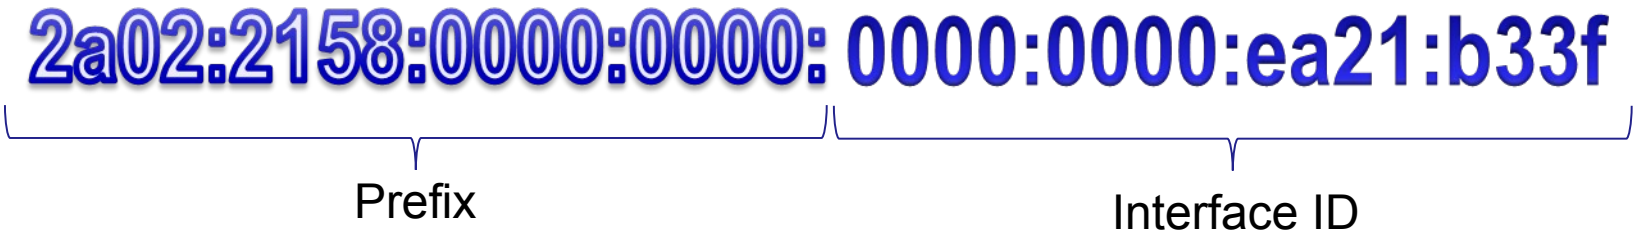

- **The Zero Suppression rule**
  - Leading zeros in a group can be omitted.
- **The Zero Compression rule**
  - Contiguous groups of ‘0‘ can be replaced with ‘::’ as long as there is only one double colon used in an address.

2a02:2158:0:0:0:0:ea21:b33f
2a02:2158:0::0:ea21:b33f
2a02:2158::ea21:b33f

Note: Having more than one double-colon abbreviation in an address is invalid.

IPv6 Network Notation

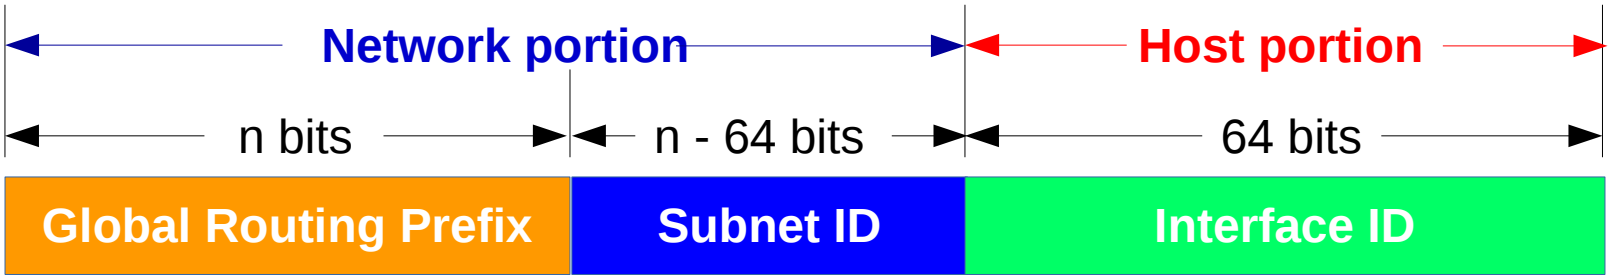

2a02:2158:435a:330::9bc2:45/64

Special IPv6 addresses

- Unspecified address

::

- Default route address

0:0:0:0:0:0:0:0 ⟹ ::0

- Loopback

0:0:0:0:0:0:0:1 ⟹ ::1

- Link-local IPv6 Address (LLA)

fe80::9088:3062:ea21:b33f

IPv6 packet architecture

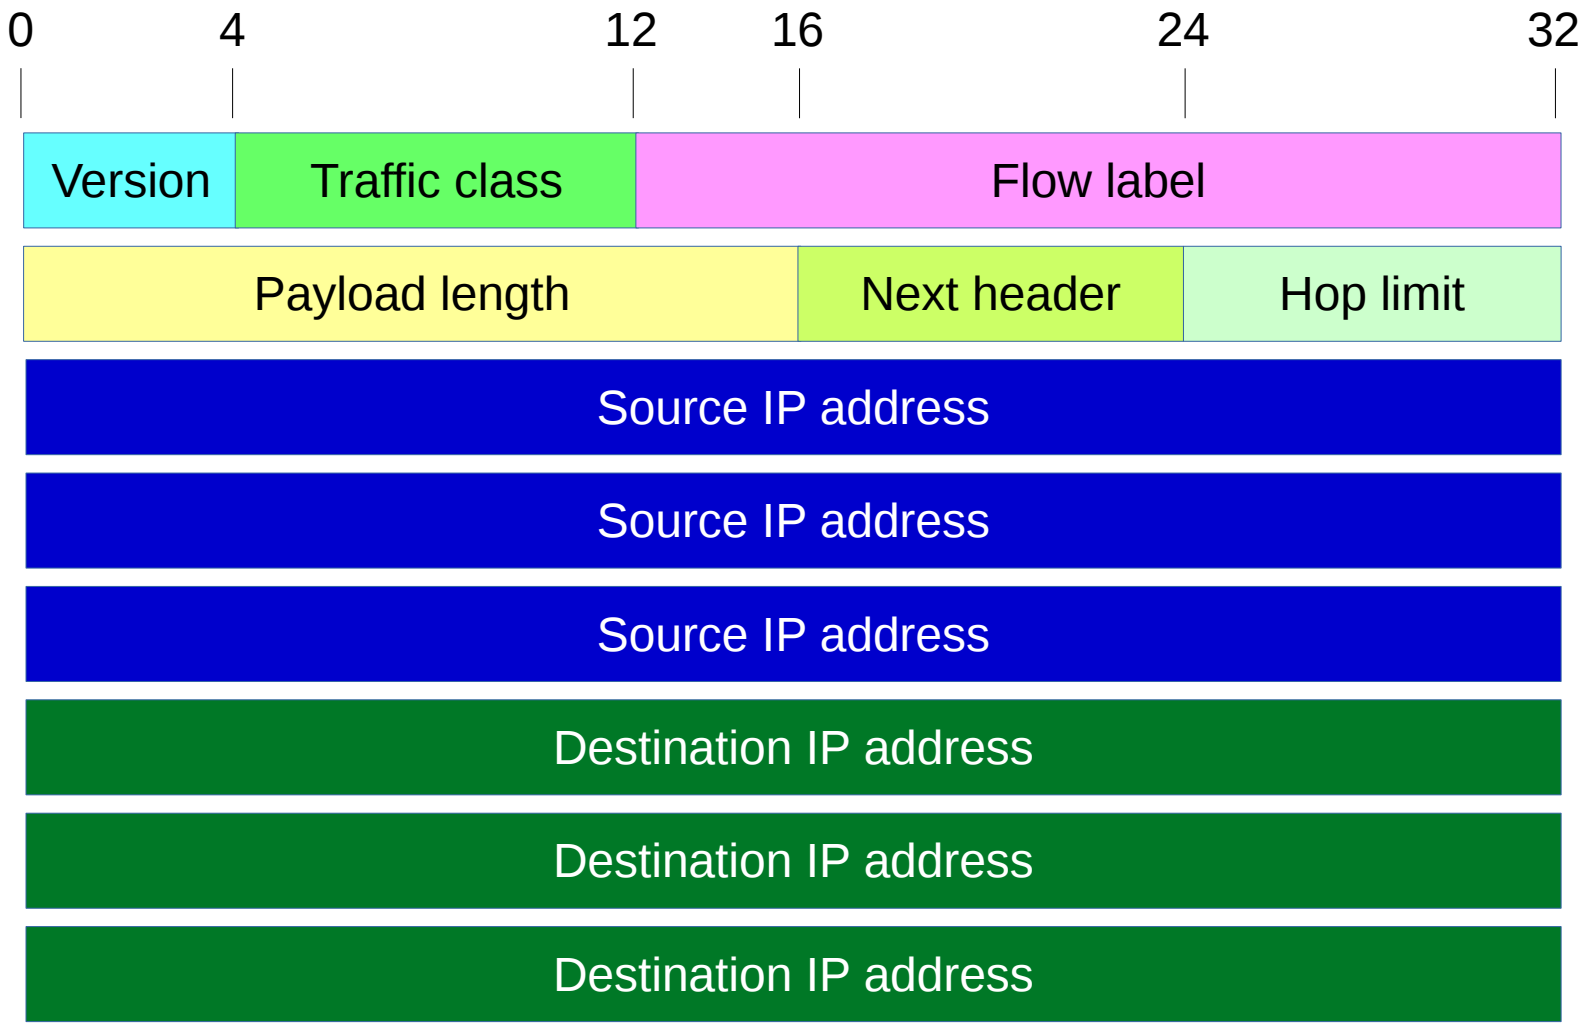

IPv4 Vs IPv6 Fields

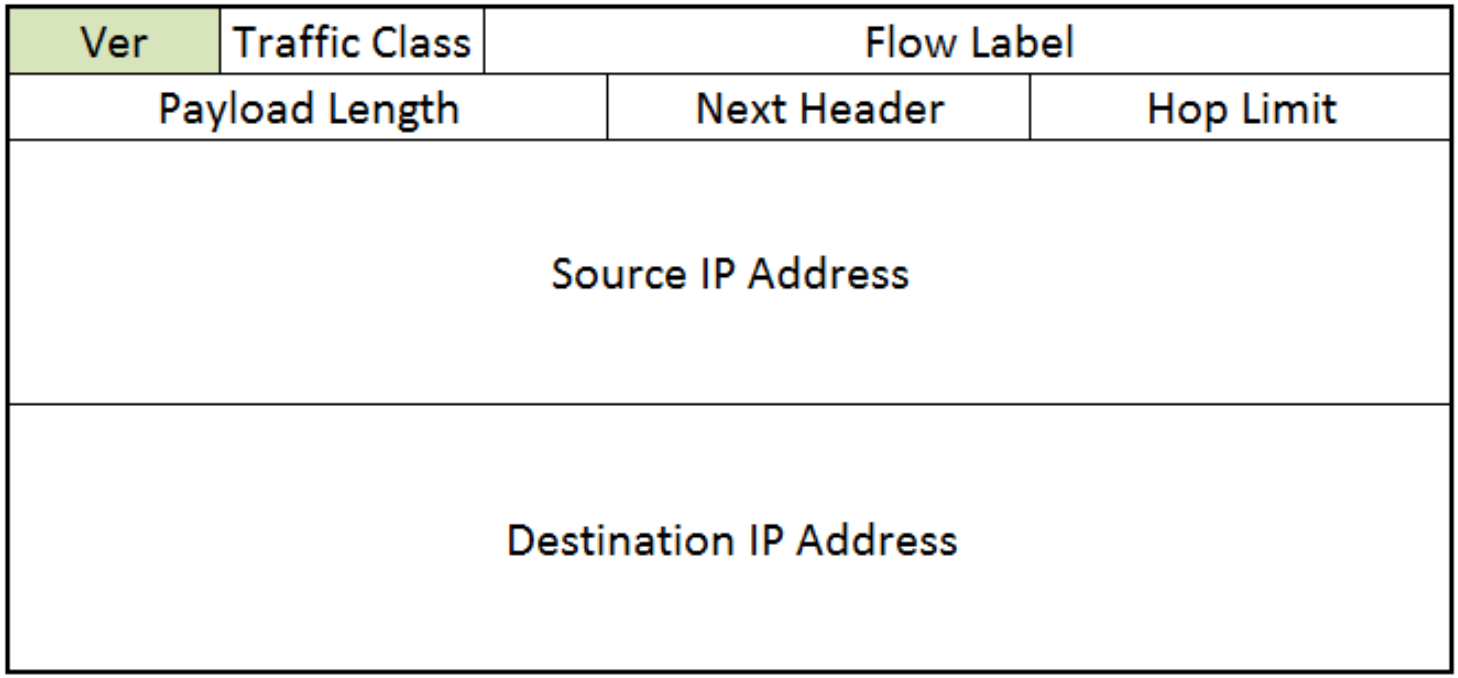

IPv6

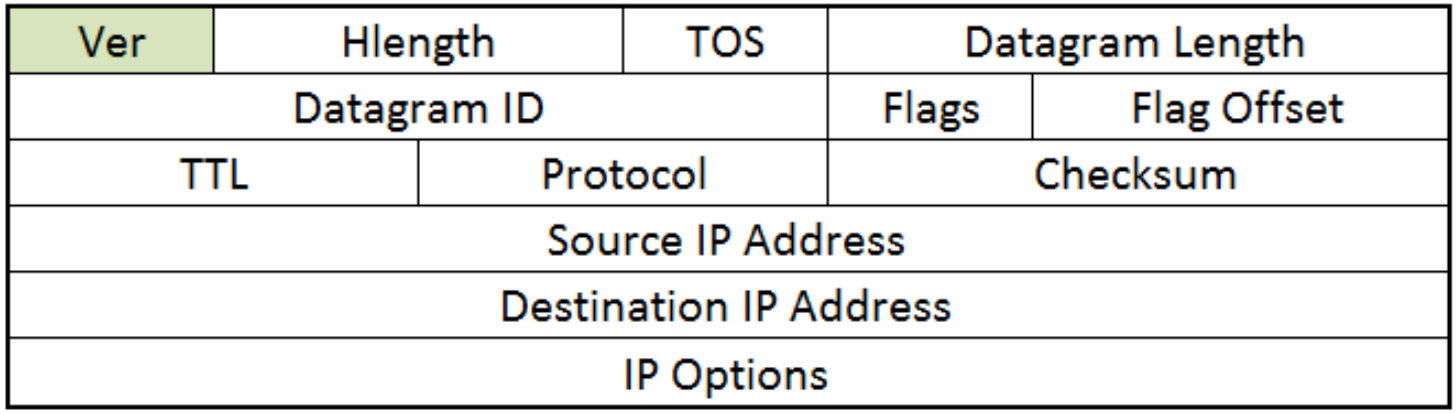

IPv4

“IPv6 Packet

Ethernet II, Src: 00:16:17:ba:0e:74, Dst: 00:12:3f:dc:ab:47
Destination: 00:12:3f:dc:ab:47
Source: 00:16:17:ba:0e:74
Type: **IPv6** (0x86dd)
Internet Protocol Version 6
0110 .... = **Version: 6**
.... 0000 0000 .... = Traffic class: 0x00000000
.... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
Payload length: 40
Next header: **ICMPv6** (0x3a)
Hop limit: 128
Source: 2001::ffff:10
Destination: 2001::ffff:20
Internet Control Message Protocol v6
Type: 128 (Echo request)
Code: 0
Checksum: 0x94bb [correct]
ID: 0x0001
Sequence: 0x000b
Data (32 bytes)

“Similar and Evolving” fields

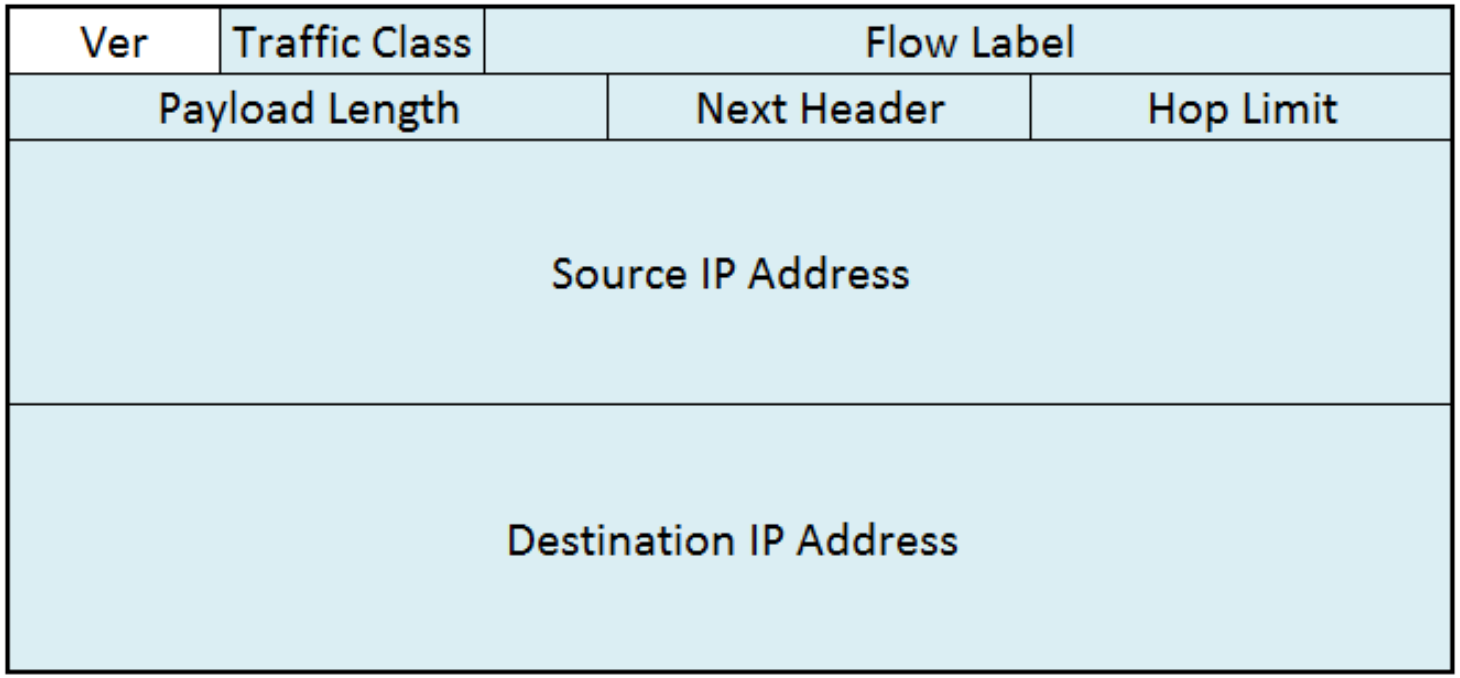

IPv6

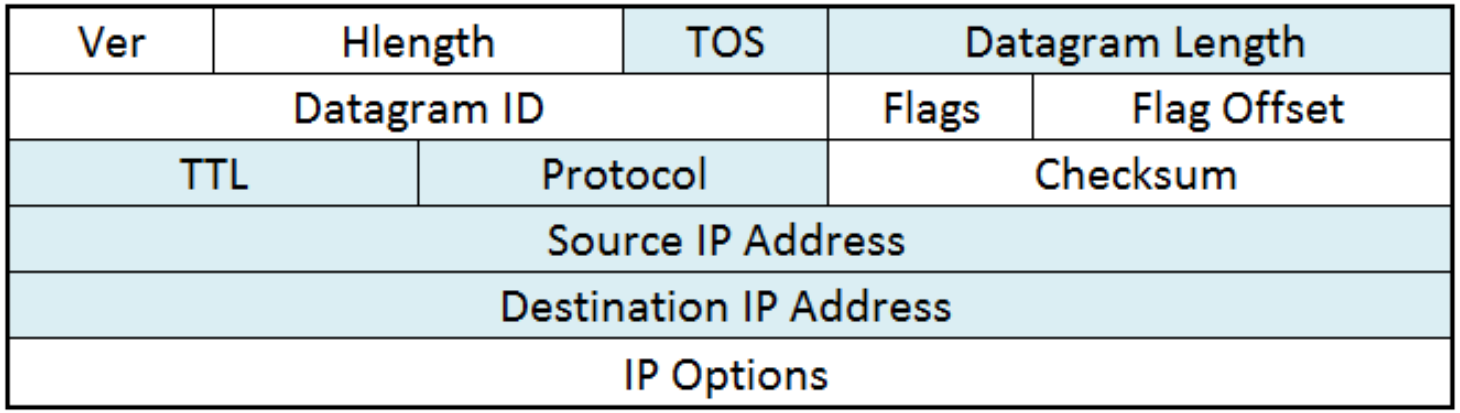

IPv4

“Removed” fields

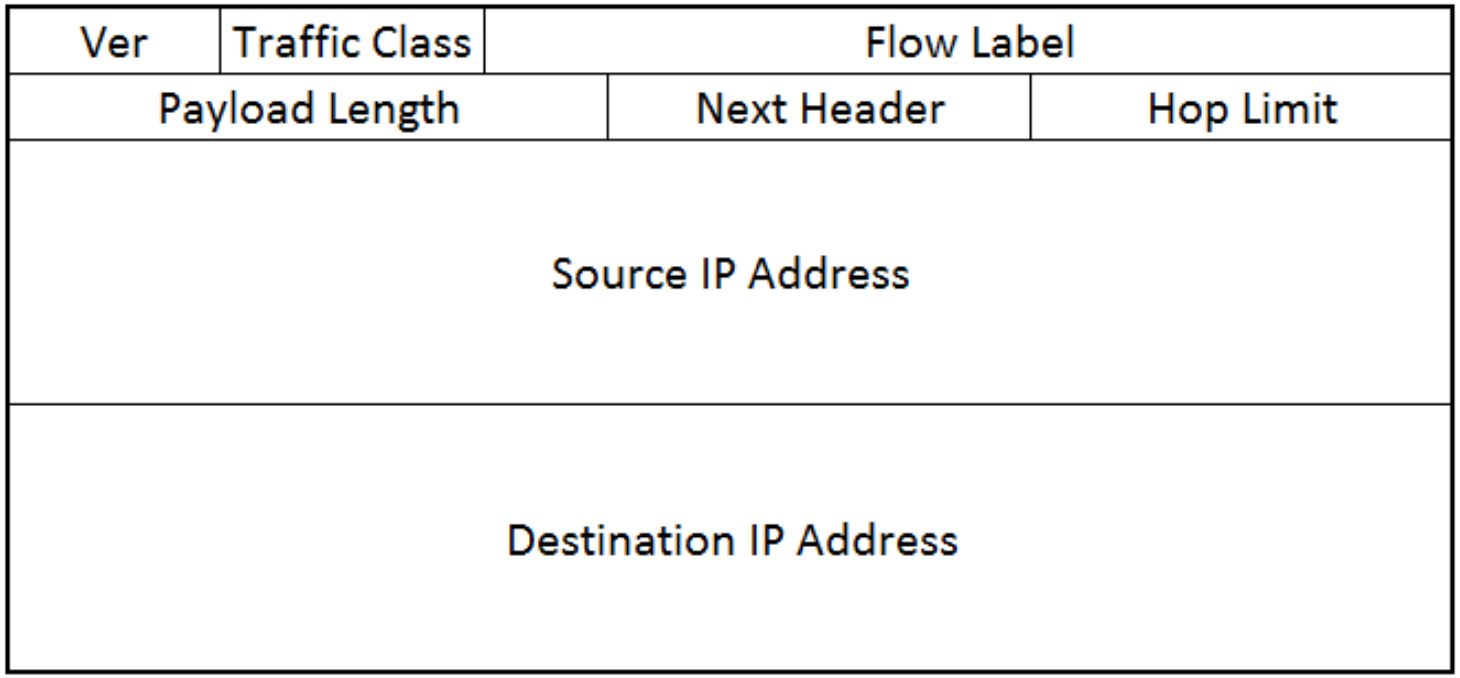

IPv6

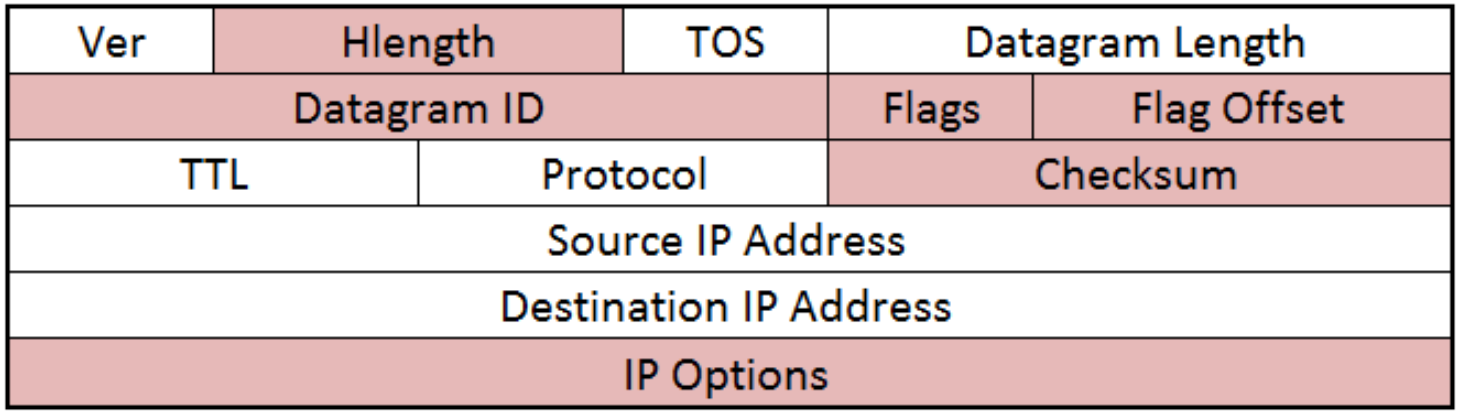

IPv4

Extensions & Options

- Options
  - IPv4 – At end of IP Header.
  - IPv6 – Extensions Headers.
- Extension Headers
  - Only processed as necessary
    - i.e. Only routers process “Hop by Hop options header”.
  - Easier to define new extensions and options.

Extension Headers

- Hop by Hop options header.
- Destination options header.
- Routing header.
- Fragment header.
- Authentication header (AH).
- Encapsulation security payload (ESP) header.

Compare IPv4 and IPv6 headers

IPv4	IPv6
Version	Version
Total length	Total length
Source IP address	Source IP address
Destination IP address	Destination IP address
Type of Service (TOS)	Traffic Class
Time to Live (TTL)	Hop Limit (HL)
Protocol	Next Header
Identification	Extension Headers
Flags	Extension Headers
Fragment Offset	Extension Headers
Options	Extension Headers
Padding	Extension Headers
Internet Header Length (IHL)	
Header Checksum	

Configuring IPv6 Address on UNIX/Linux

- Use the nm-connection-editor

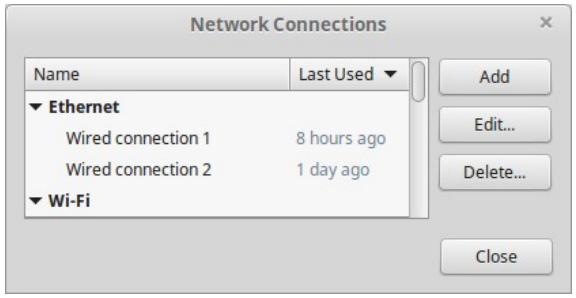


```
$ ping6 2001::ffff:10
PING 2001::ffff:10(2001::ffff:10) 56 data bytes
64 bytes from 2001::ffff:10: icmp_seq=1 ttl=128 time=1.44 ms
64 bytes from 2001::ffff:10: icmp_seq=2 ttl=128 time=0.458 ms
64 bytes from 2001::ffff:10: icmp_seq=3 ttl=128 time=0.477 ms
64 bytes from 2001::ffff:10: icmp_seq=4 ttl=128 time=0.456 ms
64 bytes from 2001::ffff:10: icmp_seq=5 ttl=128 time=0.463 ms

--- 2001::ffff:10 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3998ms
rtt min/avg/max/mdev = 0.456/0.660/1.446/0.393 ms
```

- Use the command shell

```
$ sudo  ip -6 addr add 2001::ffff:0020/112 dev eth0
$ sudo ip -6 route add ::/0 via 2001::ffff:0001

$ sudo ip -6 addr show dev eth0
$ sudo  ip -6 route show
```

Configuring IPv6 Address on UNIX/Linux

- GNU/Linux persistent change using **interfaces** edit

`/etc/network/interfaces`

```
$ sudo vi /etc/network/interfaces
  iface eth0 inet6 static
    address 2001::ffff:0020
    netmask 112
    gateway 2001::ffff:0001
  ~
  :wq!

$ sudo ip link set dev eth0 down
$ sudo ip link set dev eth0 up
```

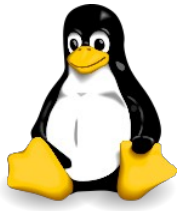

Configuring IPv6 Address on UNIX/Linux

- GNU/Linux persistent change using **netplan** edit

`/etc/netplan/01-network-manager-all.yaml`

```
$ sudo vi /etc/netplan/01-network-manager-all.yaml
network:
  version: 2
  renderer: networkd
  ethernets:
    eth0:
      accept-ra: no
      addresses:
        - 2001::ffff:0020/112
      gateway6: 2001::ffff:0001
~
:wq!

$ sudo netplan apply
```

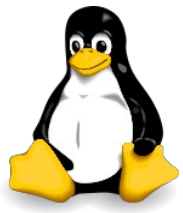

Configuring IPv6 Address on UNIX/Linux

- Add the DNS Server as system-wide DNS server. This is done in **`/etc/systemd/resolved.conf`** file:

```
$ sudo vi /etc/systemd/resolved.conf
[Resolve]
DNS=2001:4860:4860::8888 2001:4860:4860::8844

~
:wq!
```

```
$ sudo systemctl daemon-reload
$ sudo systemctl restart systemd-networkd
$ sudo systemctl restart systemd-resolved
```

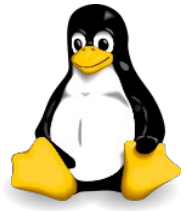

Configuring IPv6 Address on Windows

Start → Settings → Network Connections

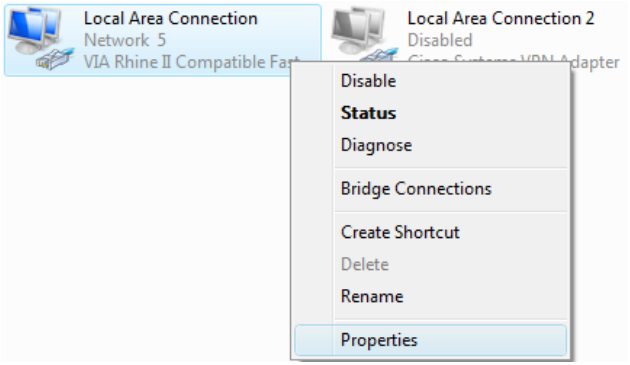

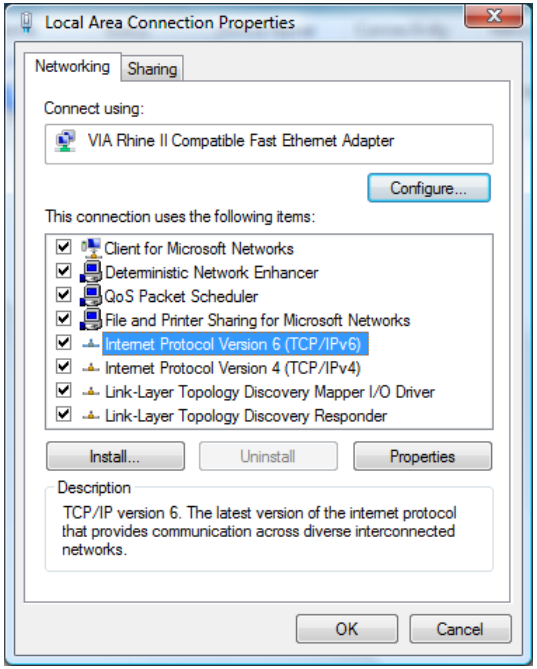

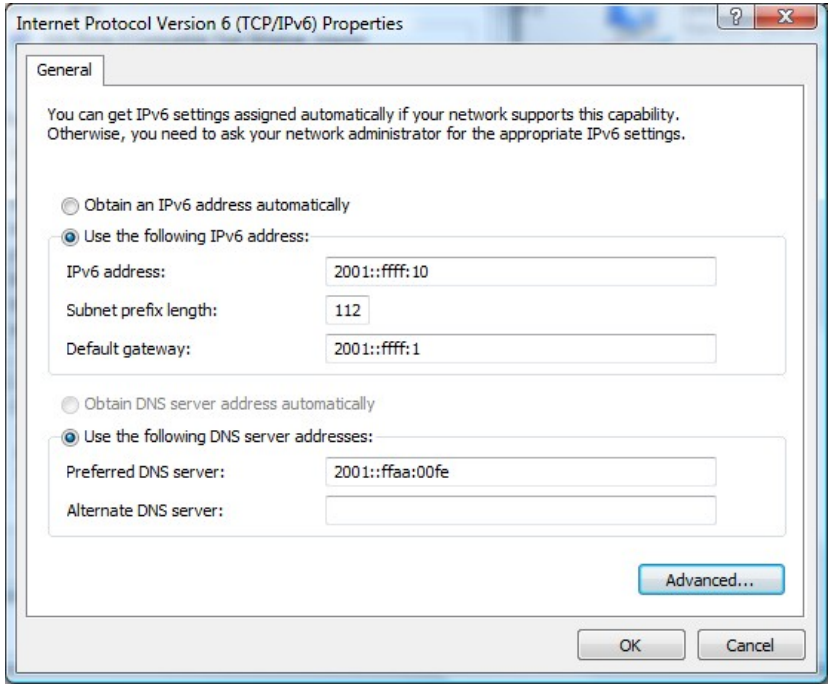

C:\> **ping 2001::ffff:20**

Pinging 2001::ffff:20 from 2001::ffff:10 with 32 bytes of data:

Reply from 2001::ffff:20: time<1ms
Reply from 2001::ffff:20: time<1ms
Reply from 2001::ffff:20: time<1ms
Reply from 2001::ffff:20: time<1ms

Ping statistics for 2001::ffff:20:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms

Prefix Terminology

- Prefix Terms

Prefix Term	Assigned by	Example prefix
Registry Prefix	Assigned to regional registry	2a02::/12
ISP Prefix	Assigned to ISP	2a02:2158::/32
Site Prefix	Assigned to Large Company	2a02:2158:1111/48
Site Prefix	Assigned to Smaller Company	2a02:2158:1111:100::/56
Subnet Prefix	Internal subnet within company	2a02:2158:1111:110::/64
A host address	IT Department in a company	2a02:2158:1111:110::10/128

- IPv6 Relative Network Sizes

/128	1 IPv6 address	A network interface
/64	1 IPv6 subnet	18,446,744,073,709,551,616 IPv6 addresses
/56	256 LAN segments	Popular prefix size for smaller subscriber site
/48	65,536 LAN segments	Popular prefix size for larger subscriber site
/32	65,536 /48 subscriber sites	Minimum IPv6 allocation
/24	16,777,216 subscriber sites	256 times larger than the minimum IPv6 allocation

Prefix assignments

- Provider Aggregatable (PA) Assignments
  - From LIR allocation
  - Register by LIR in RIR IRRDB.
- Provider Independent (PI) Assignments
  - Minimum size is /48
  - Cannot be sub-assigned.

C²S CONSULTING
communications/computing/security

Link Local Addresss (LLA)

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

EUI-48 MAC to EUI-64 to LLA

- EUI-48 : 34ed:8432:5476

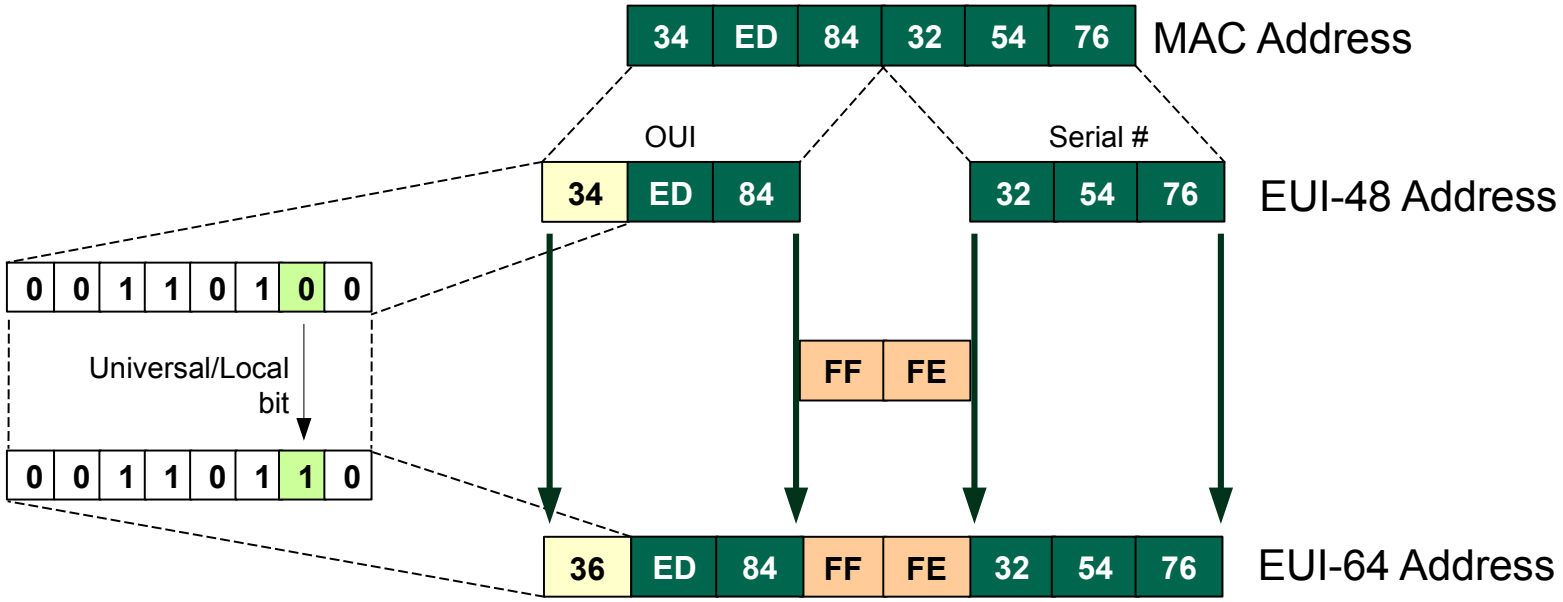

- EUI-64 : 36ed:84ff:fe32:5476
- IPv6 Host ID : fe80::36ed:84ff:fe32:5476

Resolving LLA ambiguity with zone IDs

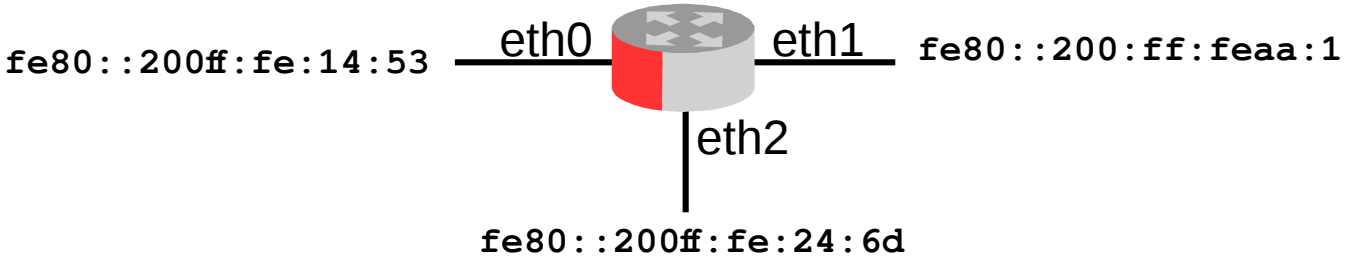


```
n1# ping ipv6 fe80::200:ff:feaa:0%eth1
PING fe80::200:ff:feaa:0%eth1(fe80::200:ff:feaa:0) 56 data bytes
64 bytes from fe80::200:ff:feaa:0: icmp_seq=1 ttl=64 time=0.025 ms
64 bytes from fe80::200:ff:feaa:0: icmp_seq=2 ttl=64 time=0.036 ms
```

C²S CONSULTING
communications/computing/security

IPv6 Address Planning

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IP Address planning

- How many addresses to request from the RIR needs to plan their address space.
- ISP
  - 10 regions
  - 50 Points of Presence (POP)
  - 3,500 clients/POP.

- Calculate the number of bits in the mask for each tier in multiples of 4 (i.e. nibbles).
  - Take the value that gives a result higher that the requirement.
- A similar process is carried out for each item as outlined below.
  - Assuming each client is assigned a /48 the mask
    - POPs can be determined by subtracting 12 from 48 giving a /36
    - Subtract 8 from 36 to give a /28 for Regions
    - Subtract 4 from 28 gives a /24 for the ISP.
  - Therefore in this example the IPS requires a /24 from the RIR.

Item	#	Bits (multiple of 4)	Possible #	Mask
ISP	1	1	2	/24
Regions	10	4	16	/28
POPs	50	8	256	/36
Clients	3,500	12	4,096	**/48**
		24		

C²S CONSULTING
communications/computing/security

Multicast

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

Multicast address

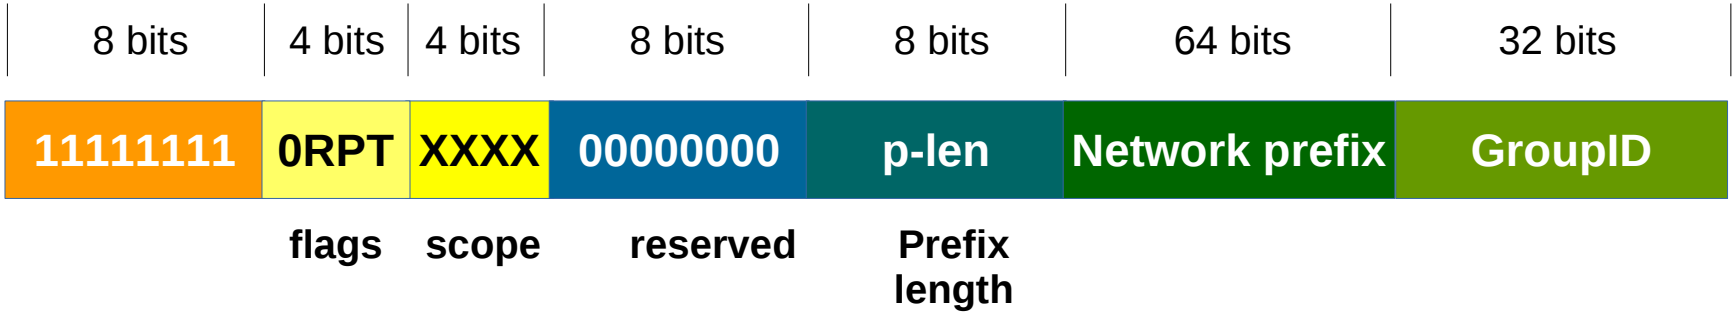

IPv6 Addressing model

- Multicast ff0X::/8

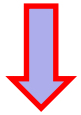

	Multicast Scope
0x1	Interface local
0x2	Link local
0x4	Admin local
0x5	Site local
0x8	Organisation local
0xE	Global
0x0	Reserved
0xF	Reserved

- i.e. ff02::1
  - All nodes on the local link.

Solicited-Node Multicast Address (SNMA)

- EUI-48 : 34ed:8432:5476

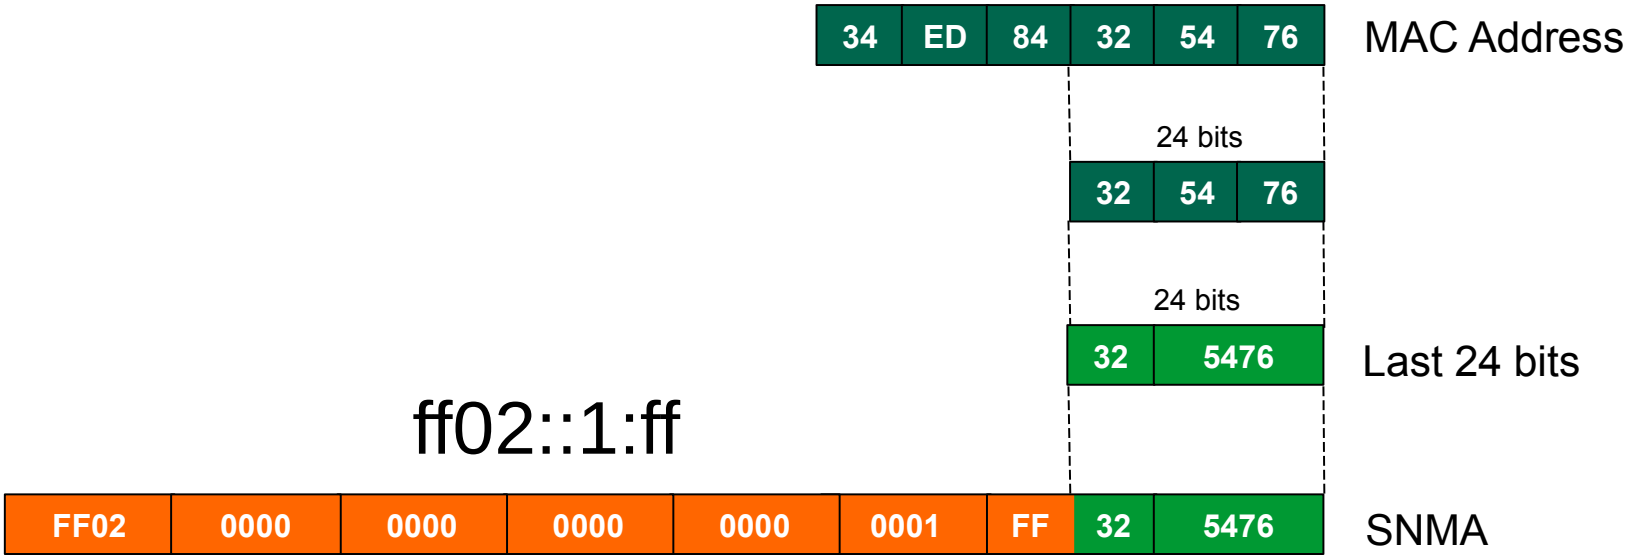

- EUI-48 : 34ed:8432:5476
- SNMA : ff02::1:ff32:5476

IPv6 Assigned Prefixes

- Address Type indicated by Format Prefix (FP).

Type	Format Prefix (FP)	Fraction	Prefix
Loopback			::1/128
IPX	0000 101	1/128	0400::/7
GUA	001	1/8	2000::/3
LLA	1111 1110 10	1/1024	fe80::/10
Local unicast	1111 110	1/128	fc00::/7
Multicast	1111 1111	1/256	ff00::/8
Pre-defined Multicast	1111 1111 0000 0001	1/4096	ff01::/12
IPv4		$1/5.90e^{20}$	::ffff:0:0/96
6to4 Translation	0000 0000 0110 0100 1111 1111 1001 1011	$1/5.90e^{20}$	64:ff9b::/96

- Anycast addresses allocated from unicast prefixes.
  - Assigning a unicast address to more than one interface turns a unicast address into an anycast address.
- Approx 1/8 of the available addressing space has been allocated.

C²S CONSULTING
communications/computing/security

Applications

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 Applications

- DHCPv6
  - Stateless
  - Stateful
- DNSv6
- ICMPv6

C²S CONSULTING
communications/computing/security

Neighbour Discovery (ND)

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 Neighbour Discovery (ND)

- For **Nodes**:
  - Address configuration (SLAAC)
  - Link-layer address resolution
  - Link-layer address change notification
  - Neighbour Unreachability Detection (NUD)

- For **Hosts**:
  - Router discovery
  - Parameter discovery (MTU, prefixes, hop limits).

- For **Routers**:
  - Advertise their presence & parameters
  - Advertise on-link prefixes
  - Determine next hops
  - Redirect hosts to better next hops.

C²S CONSULTING
communications/computing/security

Stateless Autoconfiguration (SLAAC)
Neighbour Discovery (ND)

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 Neighbour Discovery

- Stateless auto-configuration (SLAAC)
  - Creation of Global Unique Address (GUA).
  - Based in ICMPv6.
  - Creation of Link-local Address (LLA)
    - Assumes that each interface can provide a unique identifier.
    - With duplicate address detection.
  - Security to disable hackers plug & play.
- Stateful auto-configuration
  - Use of Stateful DHCPv6.

IPv6 SLAAC process

- Host creates a Solicited-Node Multicast Address (SNMA)
- Host registers a Multicast Listener Report for SNMA to join group
  - from (::) to ff02::16 Multicast Listener Discovery (MLD)
- Host creates a LLA
    - Sends Neighbour Solicitation (NS) (135) from (::) to SNMA with LLA target
    - If Neighbour Advertisement (NA) (136) received auto-configuration stops.
- Host registers a Multicast Listener Report for SNMA address to join group
  - from LLA to ff02::16 MLD
- Host sends Router Solicitation (133) to ff02::2 ‘All routers’ from LLA
- Router sends Router Advertisement (134) to ff02::1 ‘All nodes’ from its LLA with prefix
- Host creates GUA from prefix and MAC
  - Sends NS (135) from (::) to SNMA with GUA target
  - If NA (136) received auto-configuration stops.
- Finish SLAAC.

MLD to join multicast group

Solicited-Node multicast Address (SNMA) **Link-local address**

00:00:00:aa:00:02 – EUI-48

02:00:00:ff:fe:aa:00:02 – EUI-64

ff02::1:ffaa:2 **fe80::200:ff:feaa:2**

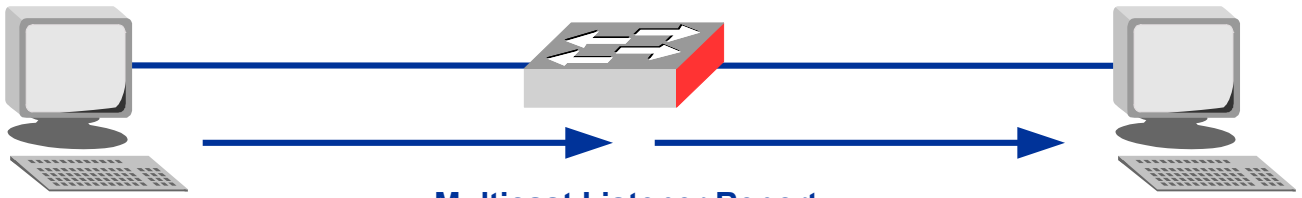

Multicast Listener Report

Ethernet II, Dst: **33:33**:ff:aa:00:02, Src: 00:00:00:aa:00:02; Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: **::**, Dst: **ff02::16**
Payload length: 36
Next header: **IPv6 Hop-by-Hop Option (0)**
Hop limit: 255
Hop-by-Hop Options
Next Header: **ICMPv6 (58)**
Internet Control Message Protocol v6
Type: **Multicast Listener Report Message v2 (143)**
Multicast Address Record Changed to exclude: **fe80::200:ff:feaa:2**

Neighbour Solicitation (135)

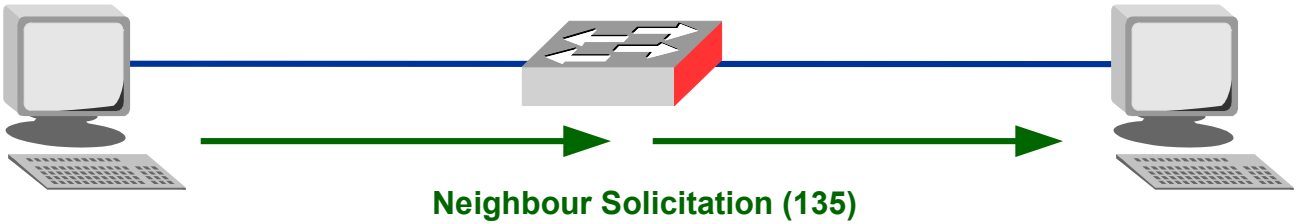


```
Ethernet II, Dst: 33:33:ff:aa:00:02, Src: 00:00:00:aa:00:02; Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: ::, Dst: ff02::1:ffaa:2
    Payload length: 24
    Next header: ICMPv6 (58)
    Hop limit: 255
Internet Control Message Protocol v6
    Type: Neighbor Solicitation (135)
    Target Address: fe80::200:ff:feaa:2
```

- No Neighbour Advertisement (136) is received
- No duplicate for LLA detected

MLD to join multicast group

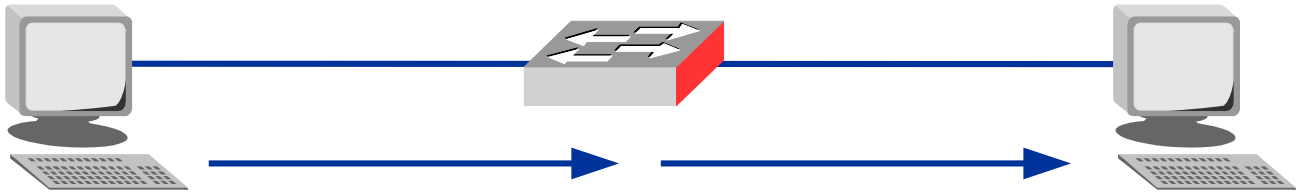

Ethernet II, Dst: 33:33:00:00:00:16, Src: 00:00:00:aa:00:02; Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: fe80::200:ff:feaa:2, Dst: ff02::16
    Payload length: 36
    Next header: IPv6 Hop-by-Hop Option (0)
    Hop limit: 255
    Hop-by-Hop Options
        Next Header: ICMPv6 (58)
Internet Control Message Protocol v6
    Type: Multicast Listener Report Message v2 (143)
    Multicast Address Record Changed to exclude: ff02::1:ffaa:2

Router Solicitation (135)

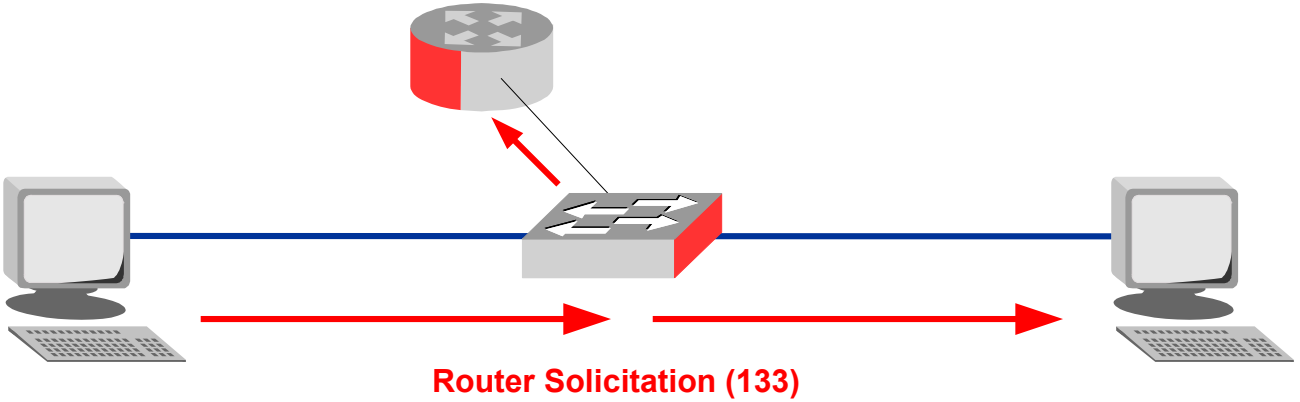

Ethernet II, Dst: 33:33:00:00:00:02, Src: 00:00:00:aa:00:02;
Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: fe80::200:ff:feaa:2, Dst: ff02::2
    Payload length: 16
    Next header: ICMPv6 (58)
    Hop limit: 255
Internet Control Message Protocol v6
    Type: Router Solicitation (133)
    Link-layer address: 00:00:00:aa:00:02

Router Advertisement (134)

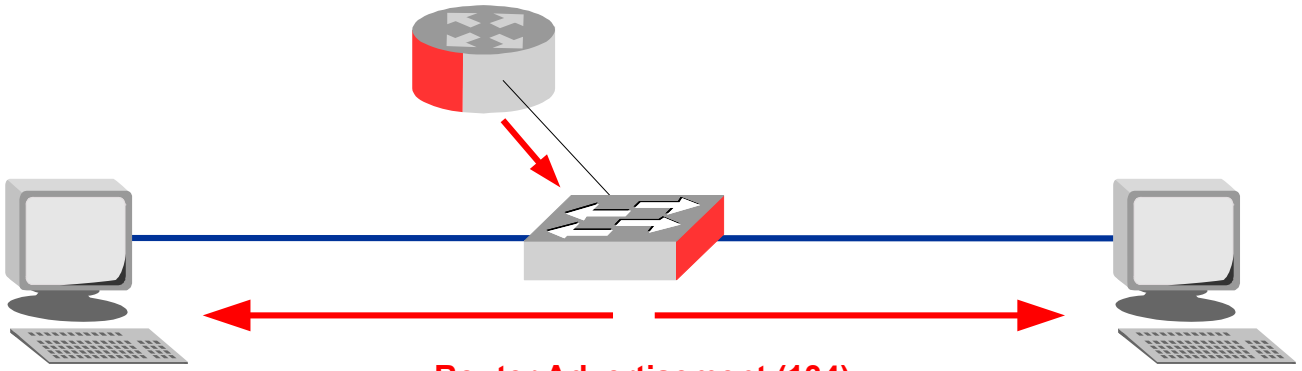


```
Ethernet II, Dst: 33:33:00:00:00:01, Src: 00:00:00:aa:00:03; Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: fe80::200:ff:feaa:3, Dst: ff02::1
    Payload length: 16
    Next header: ICMPv6 (58)
    Hop limit: 255
Internet Control Message Protocol v6
    Type: Router Advertisement (134)
    Flags: 0xc0          .1.. .... = Other configuration: Set
    ICMPv6 Option, Prefix information
        Prefix Length: 64
        Flag: 0xc0       1... .... = On-link flag (L): Set
                         .1.. .... = Autonomous address-configuration flag (A): Set
        Valid Lifetime: 86400, Preferred Lifetime: 86400
        Prefix: 2001:a::
    ICMPv6 Option
        Link-layer address: 00:00:00:aa:00:03
```

RA Flags

- **M - Managed Address Configuration Flag**
- **O - Other Configuration Flag**
- H - Mobile IPv6 Home Agent Flag
- Prf - Router Selection Preferences
- P - Neighbour Discovery Proxy Flag
- R – Reserved

```
ICMPv6 Option (Prefix information)
    Type: Prefix information (3)
    Length: 32
    Prefix length: 64
    Flags: 0xc0
        0... .... = IP Address not DHCPv6
        .1.. .... = Other config on DHCPv6
        ..0. .... = Not router address
        ...0 .... = Not site prefix
    Valid lifetime: 86400
    Preferred lifetime: 86400
    Prefix: 2001:a::
```

- If the M flag is set to 0 & the O flag is set to 1, use DHCPv6 to obtain additional configuration parameters.

RA Flags – Additional Prefix flags

ICMPv6 Option (Prefix information : 2001:a::/64)
- Type: Prefix information (3)
- Length: 4 (32 bytes)
- Prefix Length: **64**
- Flag: **0xc0**
  - **1... .... = On-link flag (L): Set**
  - **.1.. .... = Autonomous address-configuration flag (A): Set**
  - ..0. .... = Router address flag (R): Not set
  - ...0 0000 = Reserved: 0
- Valid Lifetime: 86400
- Preferred Lifetime: 86400
- Reserved
- Prefix: **2001:a::**

- **L Flag -** announces that other devices with the same prefix, use L2 switching not routing.
- **A Flag –** indicates to use SLAAC.

RA Flags – influence on auto-configuration

M	A	Resulting non-Link Local addresses on client
0	**0**	No addresses will be auto-configured
0	**1**	Address generated from prefix in RAs
1	**1**	Address generated from prefix) in RAs or Full address from DHCP server
1	**0**	Full address(es) from DHCP server

- **Note**
  - The hosts must be set to obtain IP address ‘automatically’
  - All hosts always generate and use a Link Local address.

Neighbour Solicitation (135)

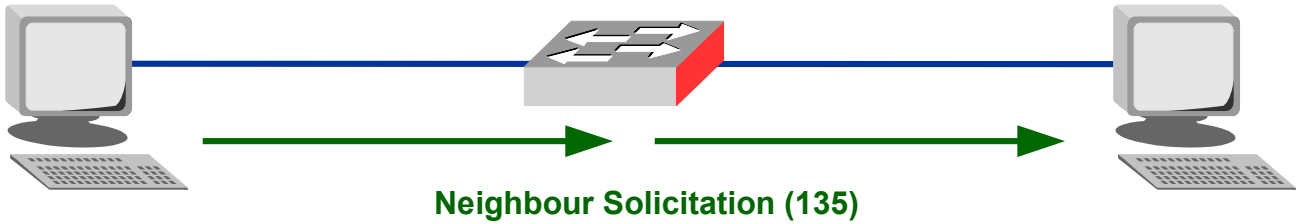


```
Ethernet II, Dst: 33:33:ff:aa:00:02, Src: 00:00:00:aa:00:02; Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: ::, Dst: ff02::1:ffaa:2
    Payload length: 24
    Next header: ICMPv6 (58)
    Hop limit: 255
Internet Control Message Protocol v6
    Type: Neighbor Solicitation (135)
    Target Address: 2001:a::200:ff:feaa:2
```

- No Neighbour Advertisement (136) is received
- No duplicate for GUA detected

MLD to join multicast group

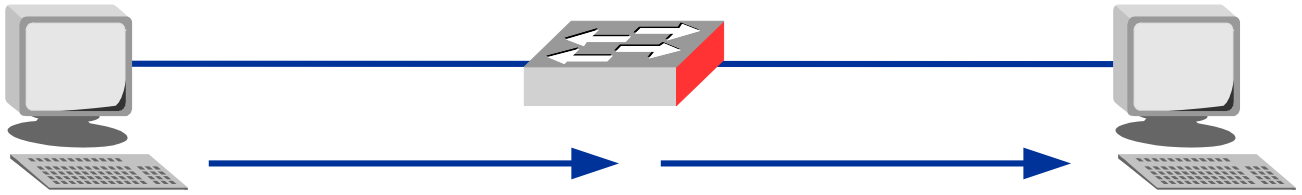

Ethernet II, Dst: 33:33:00:00:00:16, Src: 00:00:00:aa:00:02; Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: fe80::200:ff:feaa:2, Dst: ff02::16
    Payload length: 36
    Next header: IPv6 Hop-by-Hop Option (0)
    Hop limit: 255
    Hop-by-Hop Options
        Next Header: ICMPv6 (58)
Internet Control Message Protocol v6
    Type: Multicast Listener Report Message v2 (143)
    Multicast Address Record Changed to exclude: ff02::1:ffaa:2

Stateful DHCPv6

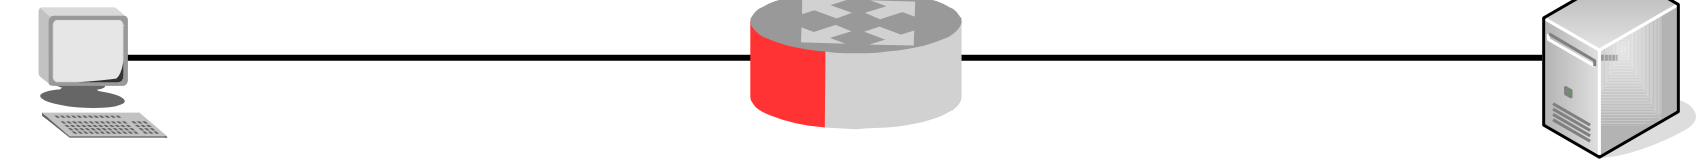

fe80::1:2ff:fe3:cccc | fe80::1:2ff:fe3:4444 | fe80::1:2ff:fe3:5555

Default Gateway
fe80::1:2ff:fe3:4444

RA (M = 1)

Client ID

SOLICIT → ff05::1:3

SOLICIT → ff05::1:3

ADVERTISE

Server ID

Client ID

REQUEST

REPLY

Server ID

IPv6 Prefix: 2010:AA01:10::/64
DNS Server: AAAA:BBBB:10FE:100::15
Domain Name: example.com

DNS via Stateless DHCPv6 in SLAAC

M	0	O	1

- RA message passes the Other flag to inform the client to get additional configuration information from a Stateless DHCPv6 Server.
- Router supplies DNS and domain information.

```
Router(config)# ipv6 dhcp pool dhcp-pool
Router(config-dhcp)# dns-server 2001:db8::d75
Router(config-dhcp)# dns-server 2001:db8::d76
Router(config-dhcp)# domain-name example.com

Router(config)# interface ethernet 0/0
Router(config-if)# ipv6 address 2001:dB8:1234::1/64
Router(config-if)# ipv6 nd other-config-flag
Router(config-if)# ipv6 dhcp server dhcp-pool
```

RADNSS / DNSSL

M 0 **O** 0

- Additional options in RA messages to carry:
  - Address of recursive DNS servers
    - Recursive DNS Server (RDNSS)
  - Search List of DNS suffix domain names
    - DNS Search List Option (DNSSL).

```
Router(config)# interface ethernet 1/0
Router(config-if)# ipv6 nd ra dns server 2001:db8::d75 sequence 0
Router(config-if)# ipv6 nd ra dns server 2001:db8::d76 sequence 1
Router(config-if)# ipv6 nd ra dns search-list example.com sequence 0
```

IPv6 Stateful Autoconfiguration

- Stateful autoconfiguration requires access to a network DHCPv6 server to provide the addressing much like for IPv4.
- Address assignment is centrally managed and clients must obtain configuration information not available through protocols such as address autoconfiguration and neighbour discovery.

NDP Summary

Type	Node	Src addr	Dst addr	Notes
NS/135	Host	IP or ::	IP or SNMA	
NA/136	Host	IP	IP or FF02::1	R - flag
RS/133	Host	LLA or ::	FF02::2	
RA/134	Router	LLA	FF02::1	Flags: (M)anage & (O)ther
RD/137	Router	LLA	IP of node	Next hop

C²S CONSULTING
communications/computing/security

Address Resolution and Redirection

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

Address Resolution

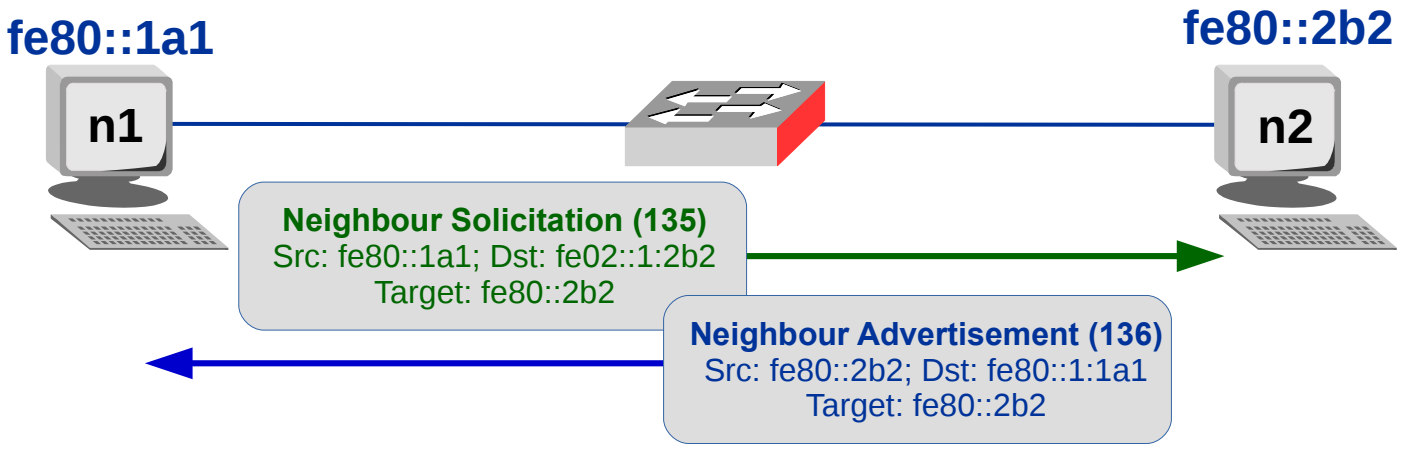

- Host sends a NS (135) message.
- Host retains packets for the requested host in a queue.
- Host received NA (136) from distant host.
- NA (136) is added to the Neighbour cache.

Neighbour Unreachability Detection (NUD)

- Solicit a NA using a unicast probe (NS)

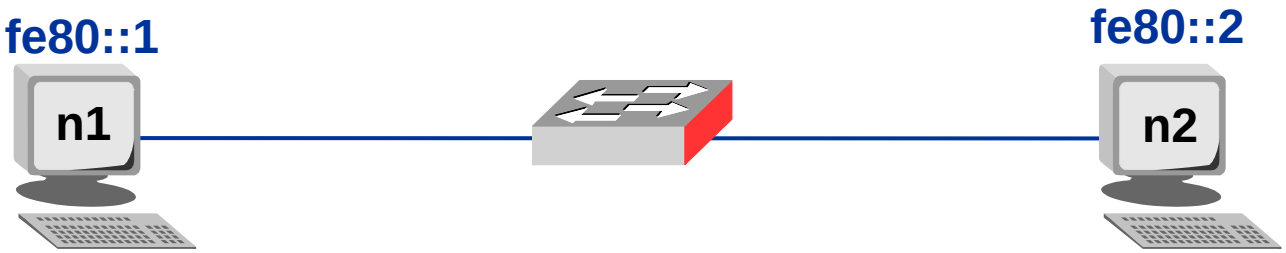

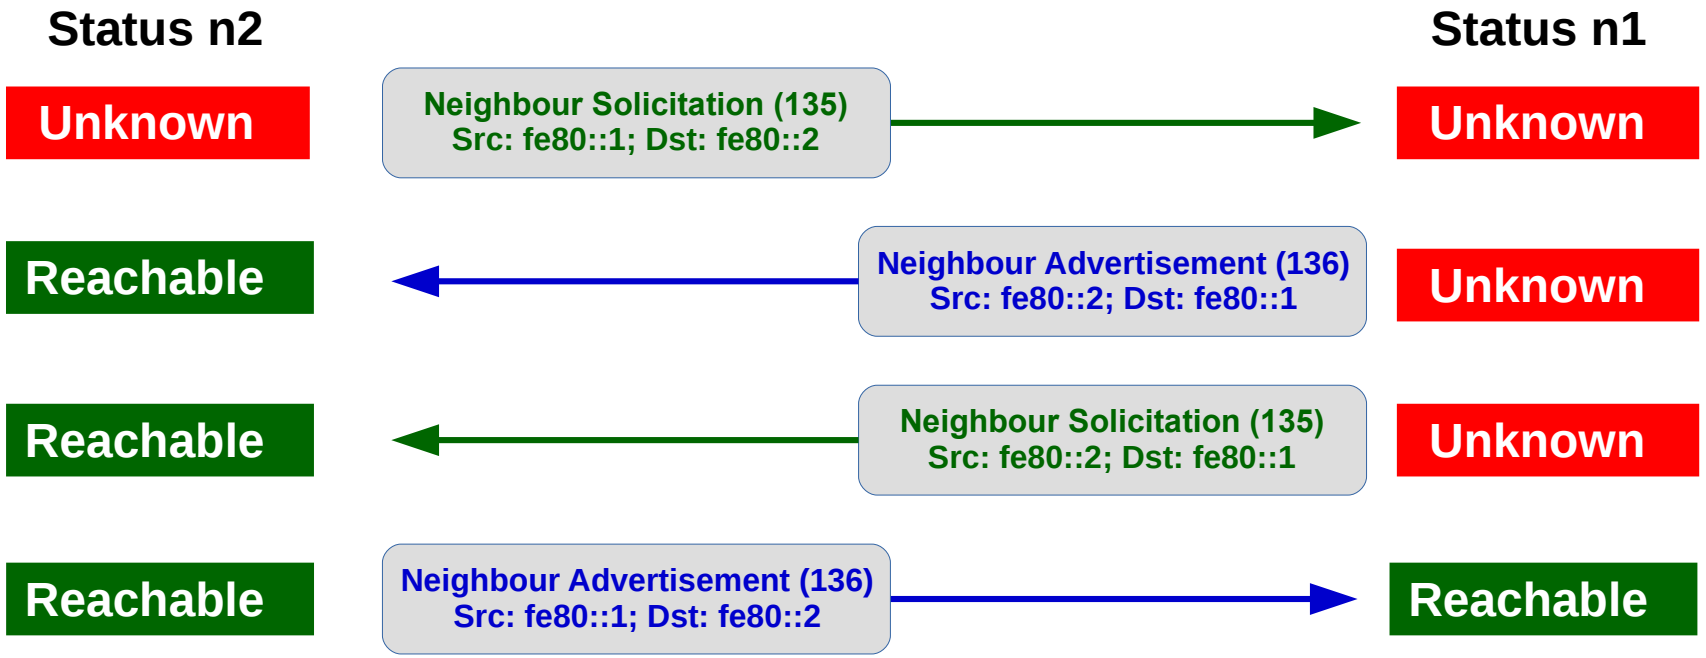

ICMPv6 Redirect

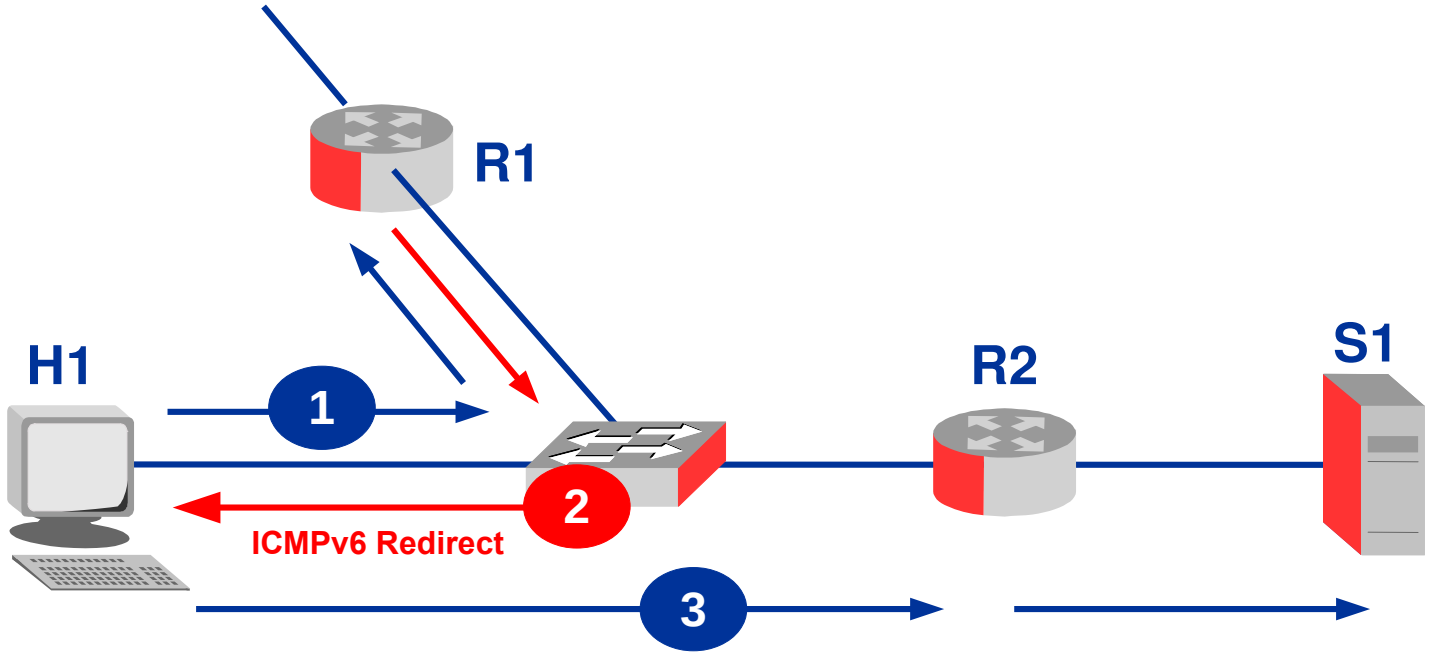

C²S CONSULTING
communications/computing/security

Best practice

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

Configuration best practice

- Don’t fit your network into RIR minimums
  - (/32, /36 & /48)
- Typical prefix lengths
  - Multi-host LAN subnets: /64
  - Inter-router links: /127
  - Loopback addresses: /128
- Plan a hierarchical scheme to optimise for aggregation
- Ensure all prefixes fall on nibble (4 bit) boundaries

Configuration best practice - Inter-router links

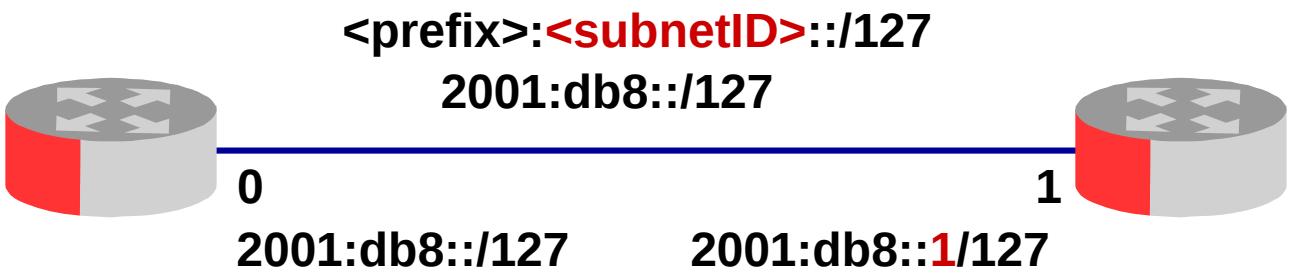

- Improves security by eliminating
  - Forwarding loops (ping pong) on some p2p links
  - Neighbour Exhaustion Attacks
- Addresses with the following 64 bits must NOT be used:
  - 0000:0000:0000:0000
  - ffff:ffff:ffff:ff7f ➠ 0000:0000:0000:ffff

Configuration best practice - Inter-router links

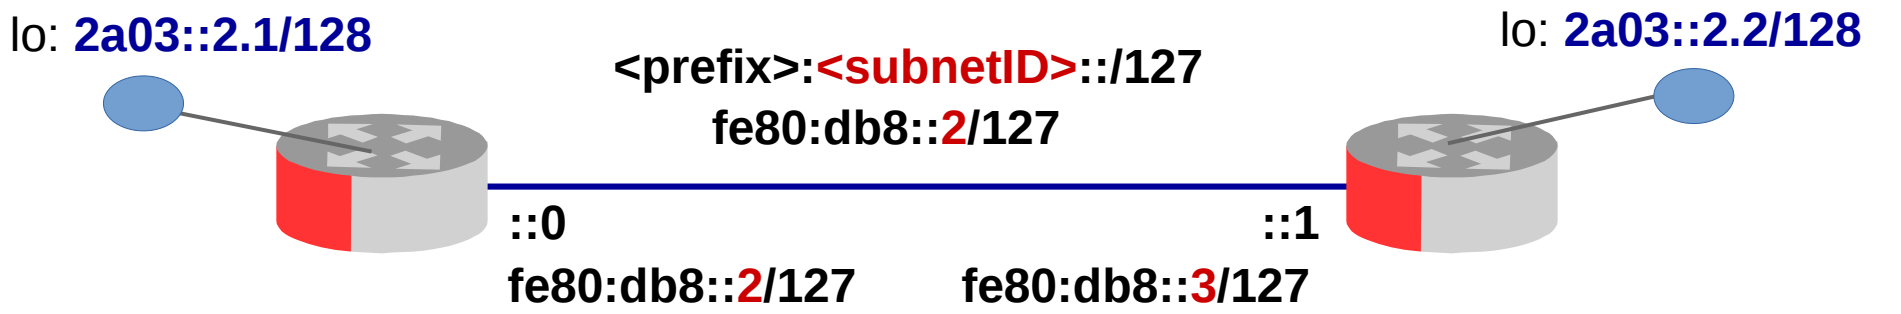

- Improves security by eliminating
  - Forwarding loops (ping pong) on some P2P links
  - Neighbour Exhaustion Attacks
- Addresses with the following 64 bits must NOT be used:
  - 0000:0000:0000:0000
  - ffff:ffff:ffff:ff7f ➠ 0000:0000:0000:ffff

C²S CONSULTING
communications/computing/security

Routing

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 IGP Routing Protocols

- RIPng (RFC 2080)
  - Distance vector algorithm like the IPv4 version.
  - Implementations: GateD, MRTd, Kame, route6d, Zebra, Cisco, Juniper, MikroTik.
- OSPFv3 (RFC 2740)
  - Link State algorithm like the IPv4 version.
  - Recommended IGP of IETF.
  - Changes from OSPFv2
    - Security removed (it uses IPv6's security).
    - Format of addresses for IPv6.
  - Implementations: GateD, MRTd, Kame, route6d, Zebra, Ericsson-Telebit, IBM, Cisco, Juniper, MikroTik.

OSPF

LSA Type	LSA Code	OSPFv3 LSA	OSPFv2 LSA	Flooding Scope
1	0x2001	Router	Router	Area-local
2	0x2002	Network	Network	Area-local
3	0x2003	Inter-Area Prefix	Network Summary	Area-local
4	0x2004	Inter-Area Router	ASBR* Router	Area-local
5	0x4005	AS-external	AS-external	AS
6	0x2006	Group Membership	Group Membership	Not implemented
7	0x2007	Type-7 (NSSA**)	NSSA external	Area-local
8	0x0008	Link		Link-local
9	0x2009	Intra-Area Prefix		Area-local

MP-BGP

- MP-BGP4 (RFC 2545 and RFC 2858)
  - Inter domain routing protocol.
  - Used between ISPs and carriers.
  - RFC 2858 defines BGP4 extensions (IPX, IPv6 etc).
  - RFC 2545 defines how to use IPv6 extensions (Scopes, next hop etc).
  - Implementations: GateD, MTRd, Kame, BGPd, Zebra, Cisco, Juniper, MikroTik.

MP-BGP AFI and SAFI

AFI	Description
0	Reserved
1	IP (IP version 4)
2	IP6 (IP version 6)
3	NSAP
4	HDLC (8-bit multidrop)
6	802 (Ethernet etc..)
16	DNS
18	AS Number

SAFI	Description
1	NLRI for unicast forwarding
2	NLRI for multicast forwarding
3	NLRI for unicast & multicast
4	NLRI with MPLS* labels.

C²S CONSULTING
communications/computing/security

Transition mechanisms

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 Transition techniques

- There are a number of scenarios:
  - Provider doesn’t support IPv6
  - Upstream doesn’t support IPv6
  - IPv6-only network
  - MPLS-based network core
  - IPv6-only services
  - IPv6-only access network.

- Dual stack
- Tunnelling
- Translation

IPv6 transition mechanisms

- RFC 4213 Dual stack
  - Hosts have IPv4 and IPv6 IP Address.
  - OS has hybrid sockets designed to accept both IPv4 and IPv6 packets.
  - When used in IPv4 communications, hybrid stacks use an IPv6 API and represent IPv4 addresses in a special address format, the IPv4-mapped IPv6 address.

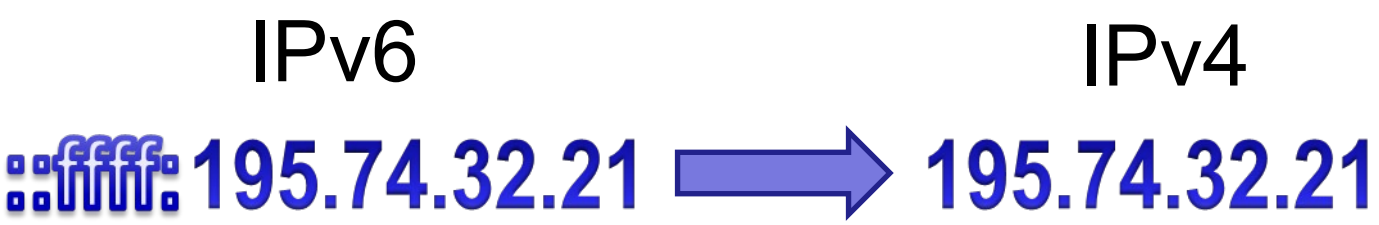

C²S CONSULTING
communications/computing/security

IPIPv6
IPv6 tunnelling over IPv4

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

IPv6 tunnelling mechanisms

- Manual tunnelling
- Automatic 6to4 tunnelling
  - Allows isolated IPv6 sites to easily connect together without having to wait for IPv4 ISPs to deliver native IPv6 services.
  - This is very well suited for extranet and virtual private networks.
    - RFC 3056 Connection of IPv6 Domains via IPv4 Clouds (protocol 41 encapsulation)
    - RFC 5969 IPv6 Rapid Deployment (6rd).

Overlay Tunnels for IPv6

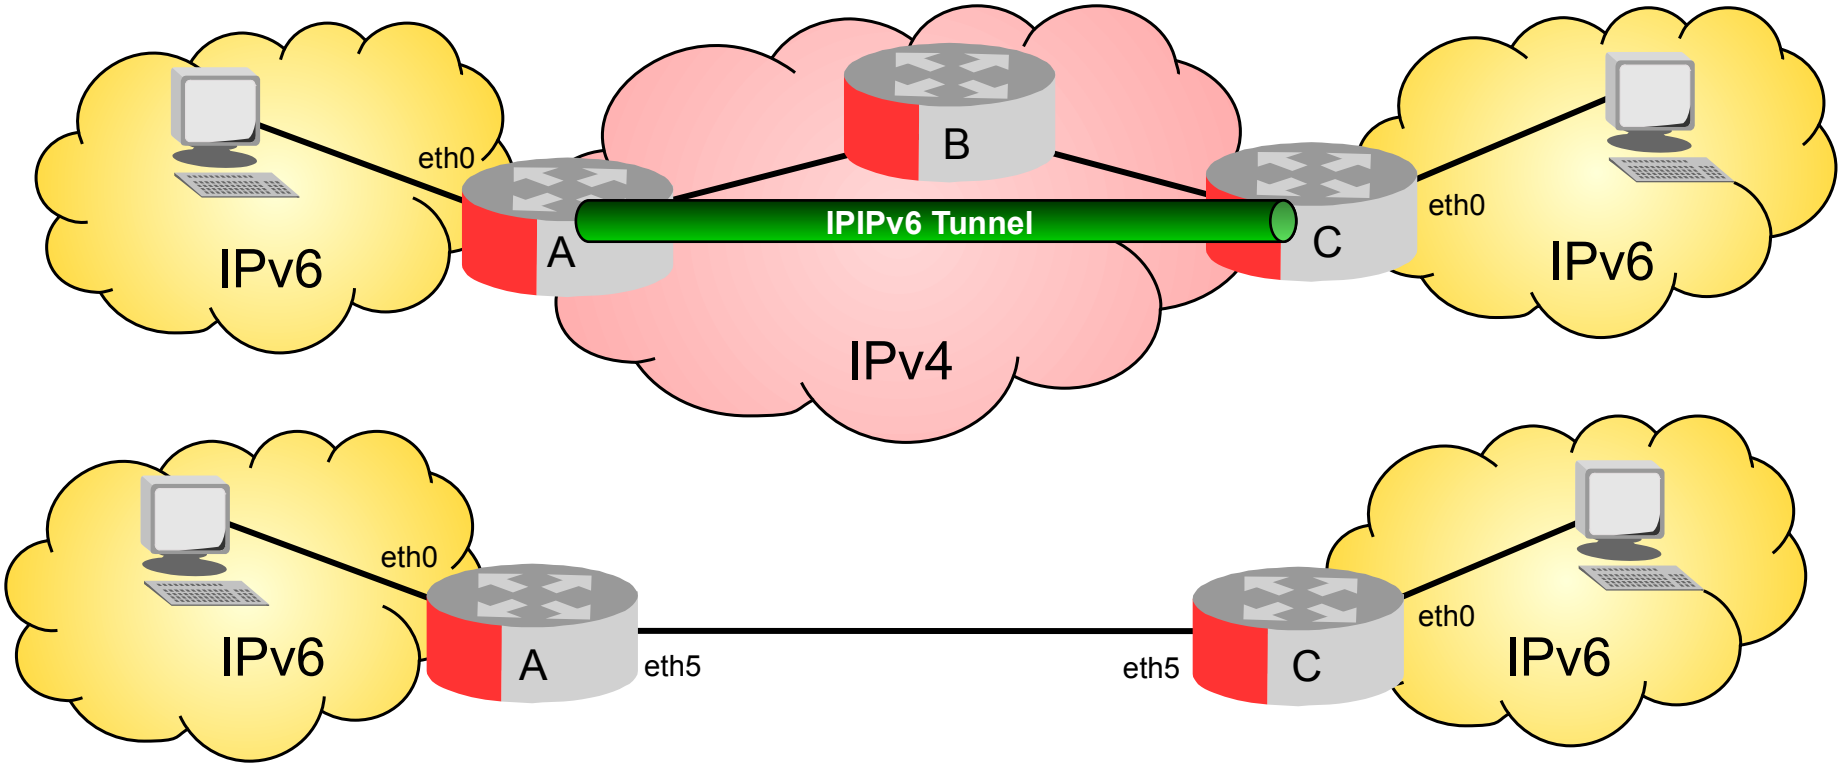

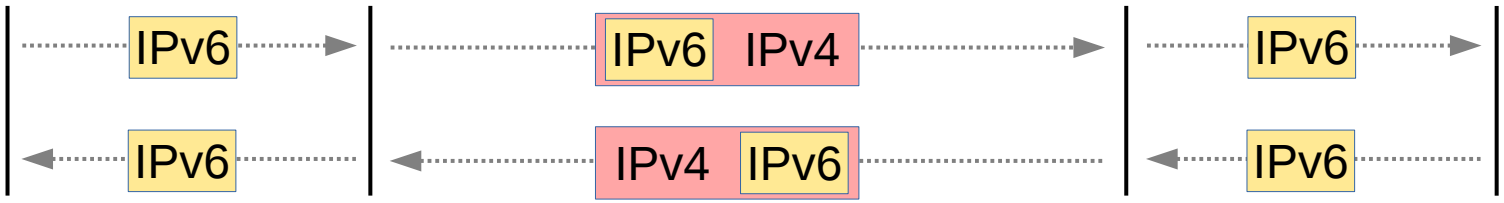

Overlay Tunnels for IPv6

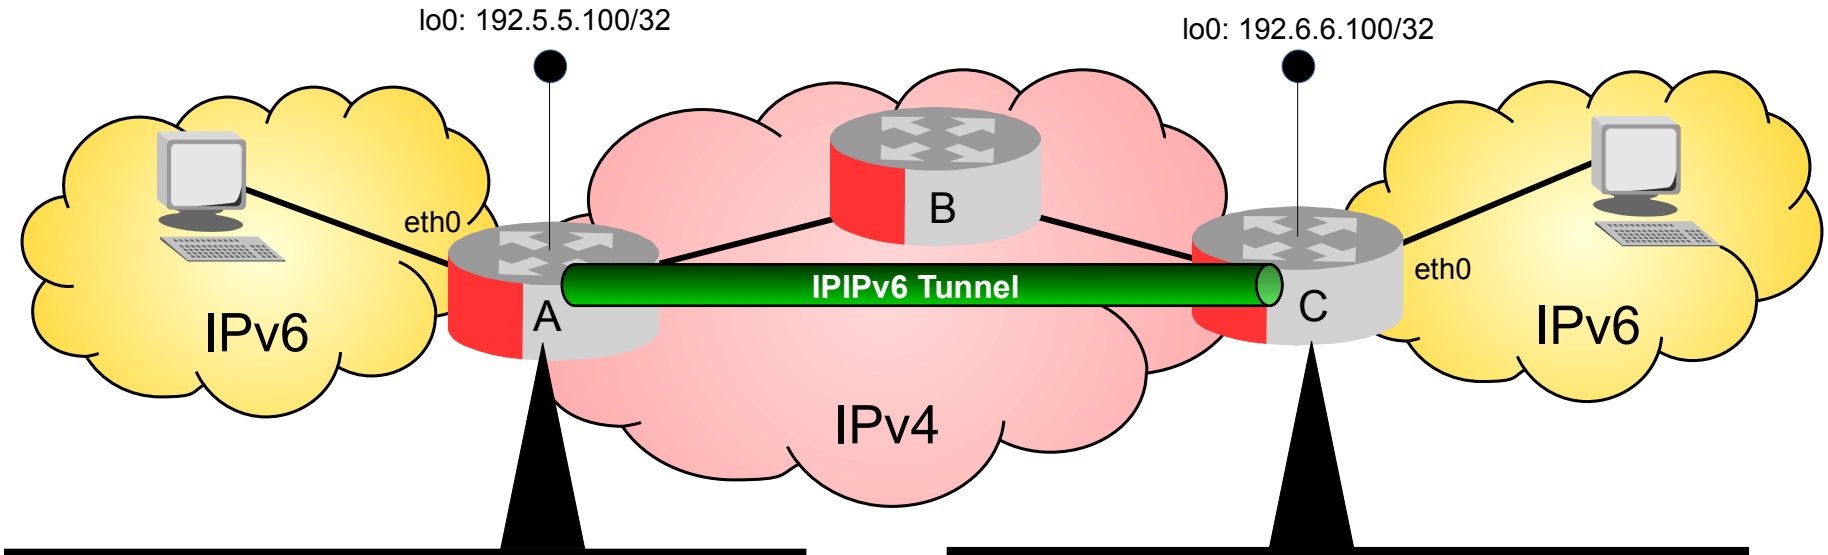


```
rtrA(config)# interface Tunnel0
rtrA(config-if)# no ip address
rtrA(config-if)# ipv6 2001:db8::/127
rtrA(config-if)# tunnel source loopback0
rtrA(config-if)# tunnel destination 192.6.6.100
rtrA(config-if)# tunnel mode ipv6ip
rtrA(config-if)# tunnel path-mtu-discovery
rtrA(config-if)# exit
rtrA(config)# ipv6 route 2001:db8:5555::/64 Tunnel0
```

```
rtrA(config)# interface Tunnel0
rtrA(config-if)# no ip address
rtrA(config-if)# ipv6 2001:db8::1/127
rtrA(config-if)# tunnel source loopback0
rtrA(config-if)# tunnel destination 192.5.5.100
rtrA(config-if)# tunnel mode ipv6ip
rtrA(config-if)# tunnel path-mtu-discovery
rtrA(config-if)# exit
rtrA(config)# ipv6 route 2001:db8:6666::/64 Tunnel0
```

Tunnel Broker

- RFC 3053 Tunnel Broker (TB)
  - Allows isolated users/routers to connect to the IPv6 network via a TB.
    - The router or host establishes an IPv6 encapsulation over IPv4 to the TB.
    - The TB authenticates the connection using RADIUS.
    - The router or host is assigned an IPv6 address and can now route IPv6.

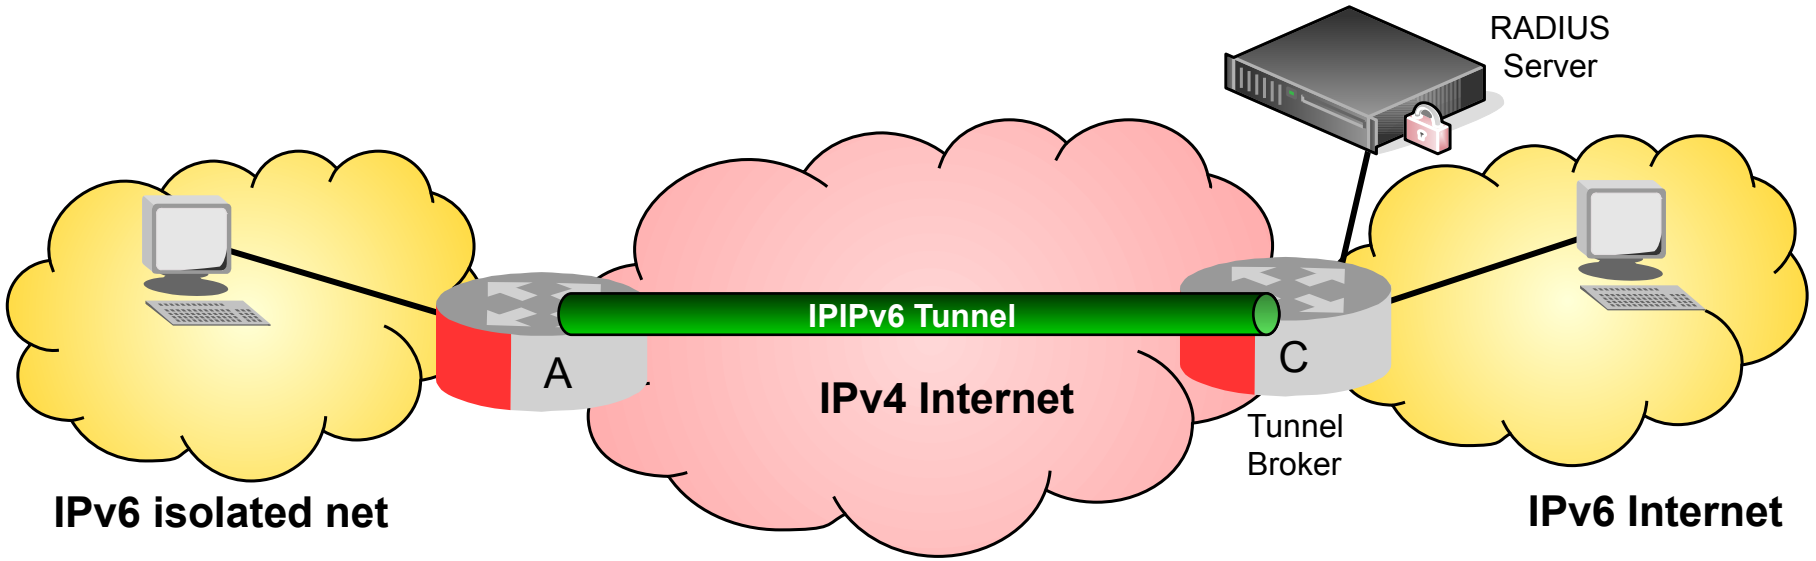

IPIPv6 – IPv6 traffic over IPv4

```
Frame: 138 bytes on wire (1104 bits),
Ethernet II, Src:  00:0c:42:5e:9e:ff, Dst: 00:0c:42:b4:46:ee
Internet Protocol Version 4, Src: 200.200.200.11, Dst: 200.200.200.12
    Version: 4
    Header length: 20 bytes
    Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00: Not-ECT (Not ECN-Capable Transport))
    Total Length: 124
    Identification: 0x0000 (0)
    Flags: 0x02 (Don't Fragment)
    Fragment offset: 0
    Time to live: 64
    Protocol: IPv6 (41)
    Header checksum: 0x18b0 [correct]
    Source: 200.200.200.11
    Destination: 200.200.200.12
Internet Protocol Version 6, Src: 2002:aaaa:1:0:212:3fff:fedc:ab47, Dst: 2002:aaaa:2:0:224:1dff:fe11:30d7
    0110 .... = Version: 6
    .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00000000
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 64
    Next header: ICMPv6 (0x3a)
    Hop limit: 254
    Source: 2002:aaaa:1:0:212:3fff:fedc:ab47
    [Source 6to4 Gateway IPv4: 170.170.0.1]
    [Source 6to4 SLA ID: 0]
    [Source SA MAC: Dell_dc:ab:47]
    Destination: 2002:aaaa:2:0:224:1dff:fe11:30d7
    [Destination 6to4 Gateway IPv4: 170.170.0.2]
    [Destination 6to4 SLA ID: 0]
    [Destination SA MAC: Giga-Byt_11:30:d7]
Internet Control Message Protocol v6
```

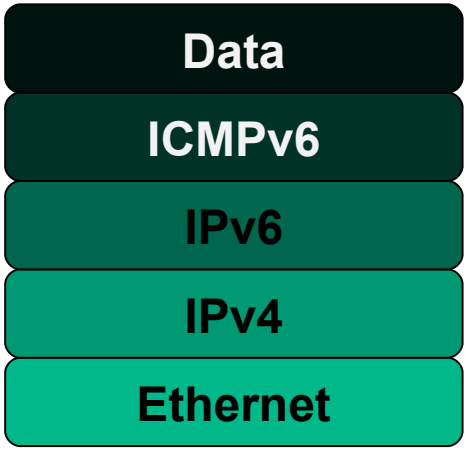

C²S CONSULTING
communications/computing/security

6to4

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

6to4

- Tunnel destination is determined by the IPv4 address of the BR extracted from the IPv6 address that starts with the prefix:
  - 2002::/16 → 2002:<BR-IPv4-addr>::/48
- 16 bits (/48 → /64) for network number
  - 2002:<BR-IPv4-addr><net#>::/64
- Example:
  - 196.55.5.1/24 (c437:0501) → 2002:c437:0501::/64

6to4

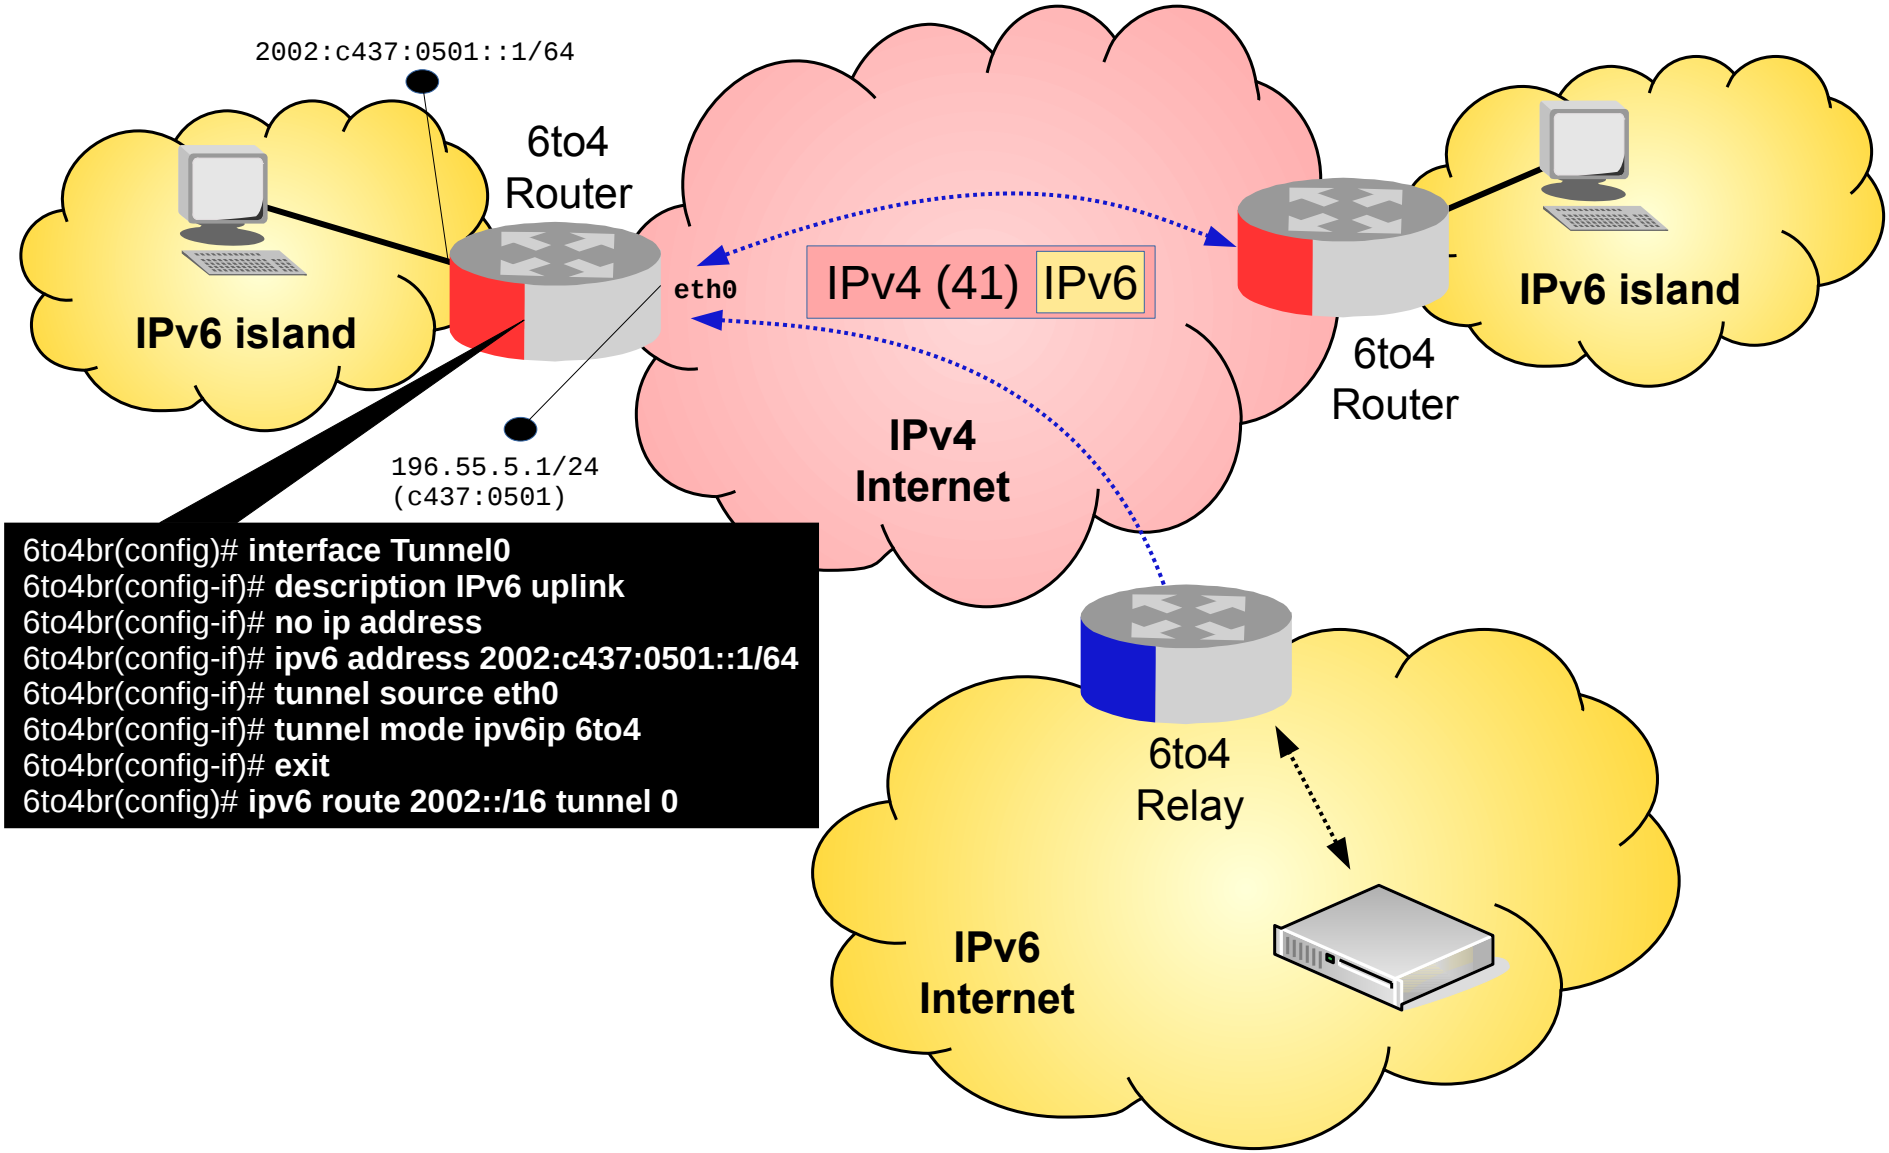


```
6to4br(config)# interface Tunnel0
6to4br(config-if)# description IPv6 uplink
6to4br(config-if)# no ip address
6to4br(config-if)# ipv6 address 2002:c437:0501::1/64
6to4br(config-if)# tunnel source eth0
6to4br(config-if)# tunnel mode ipv6ip 6to4
6to4br(config-if)# exit
6to4br(config)# ipv6 route 2002::/16 tunnel 0
```

C²S CONSULTING
communications/computing/security

6rd

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

6in4

- 6rd is an extension of the 6to4 feature.
- Provider can deliver a unicast IPv6 service over its IPv4 network by using encapsulation of IPv6 in IPv4.
- 6rd does not require addresses to have a 2002::/16 prefix.
- Prefixes come from the providers own IPv6 block.

6rd

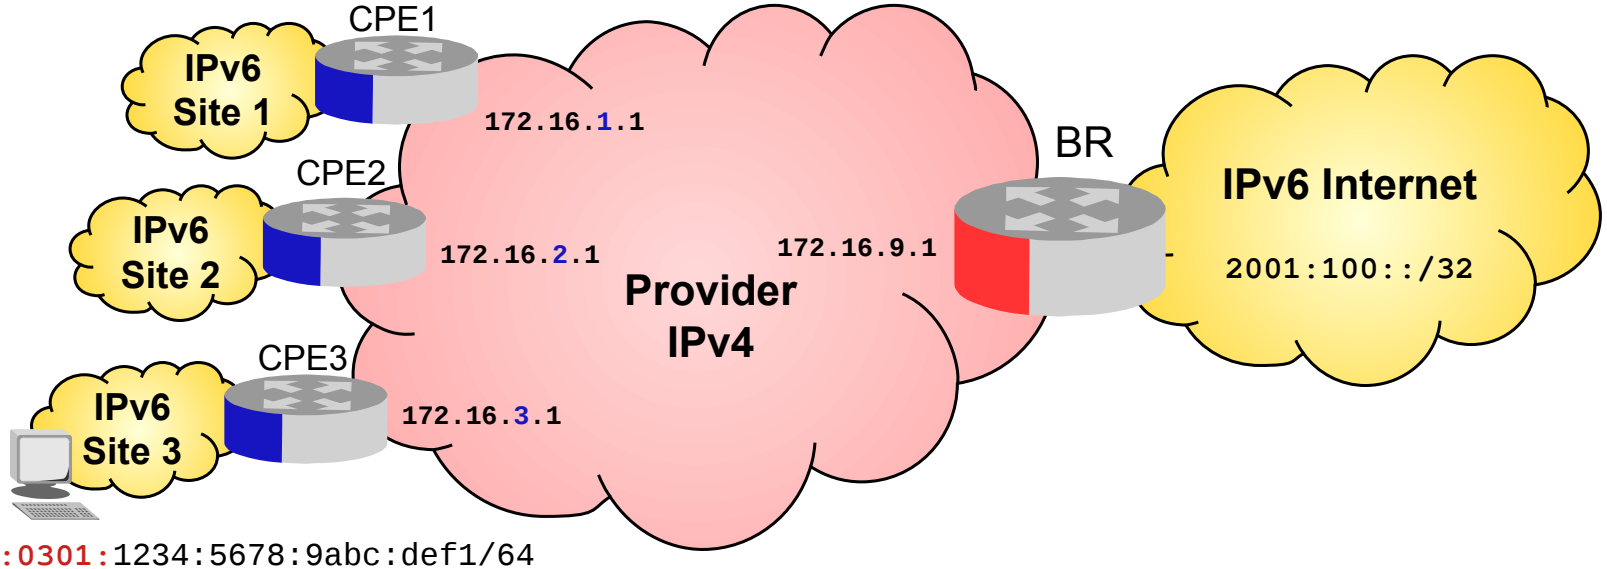

2001:100:as10:0301:1234:5678:9abc:def1/64

Delegated prefix

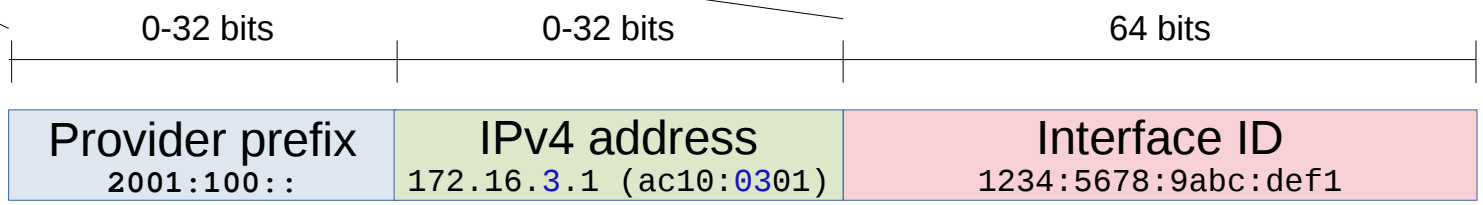

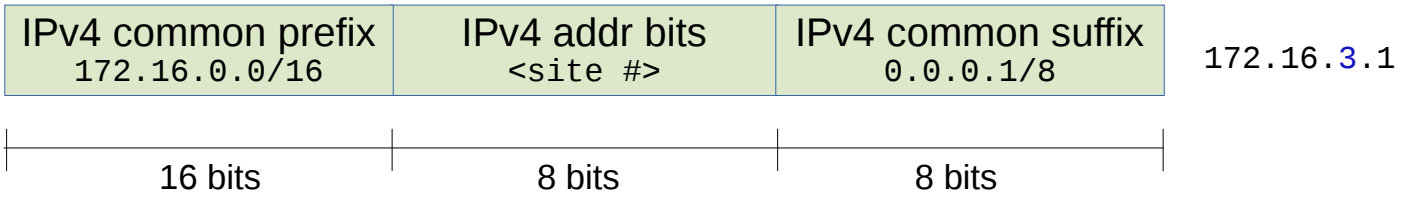

172.16.3.1

C²S CONSULTING
communications/computing/security

Dual Stack (DS)-Lite

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

DS-Lite

- Broadband provider can share IPv4 addresses among customers by combining :
  - IP in IP (IPv4-in-IPv6)
  - NAT.
- Two elements:
  - CPE – Basic Bridging BroadBand (B4)
  - Address Family Transition Router (AFTR)
    - Carrier Grade NAT (CGN)
    - Large Scale NAT (LSN).
- IPv4-in-IPv6 tunnel
  - AFER – 129.0.0.1/29
  - B4 – 129.0.0.2/29.

DS-Lite

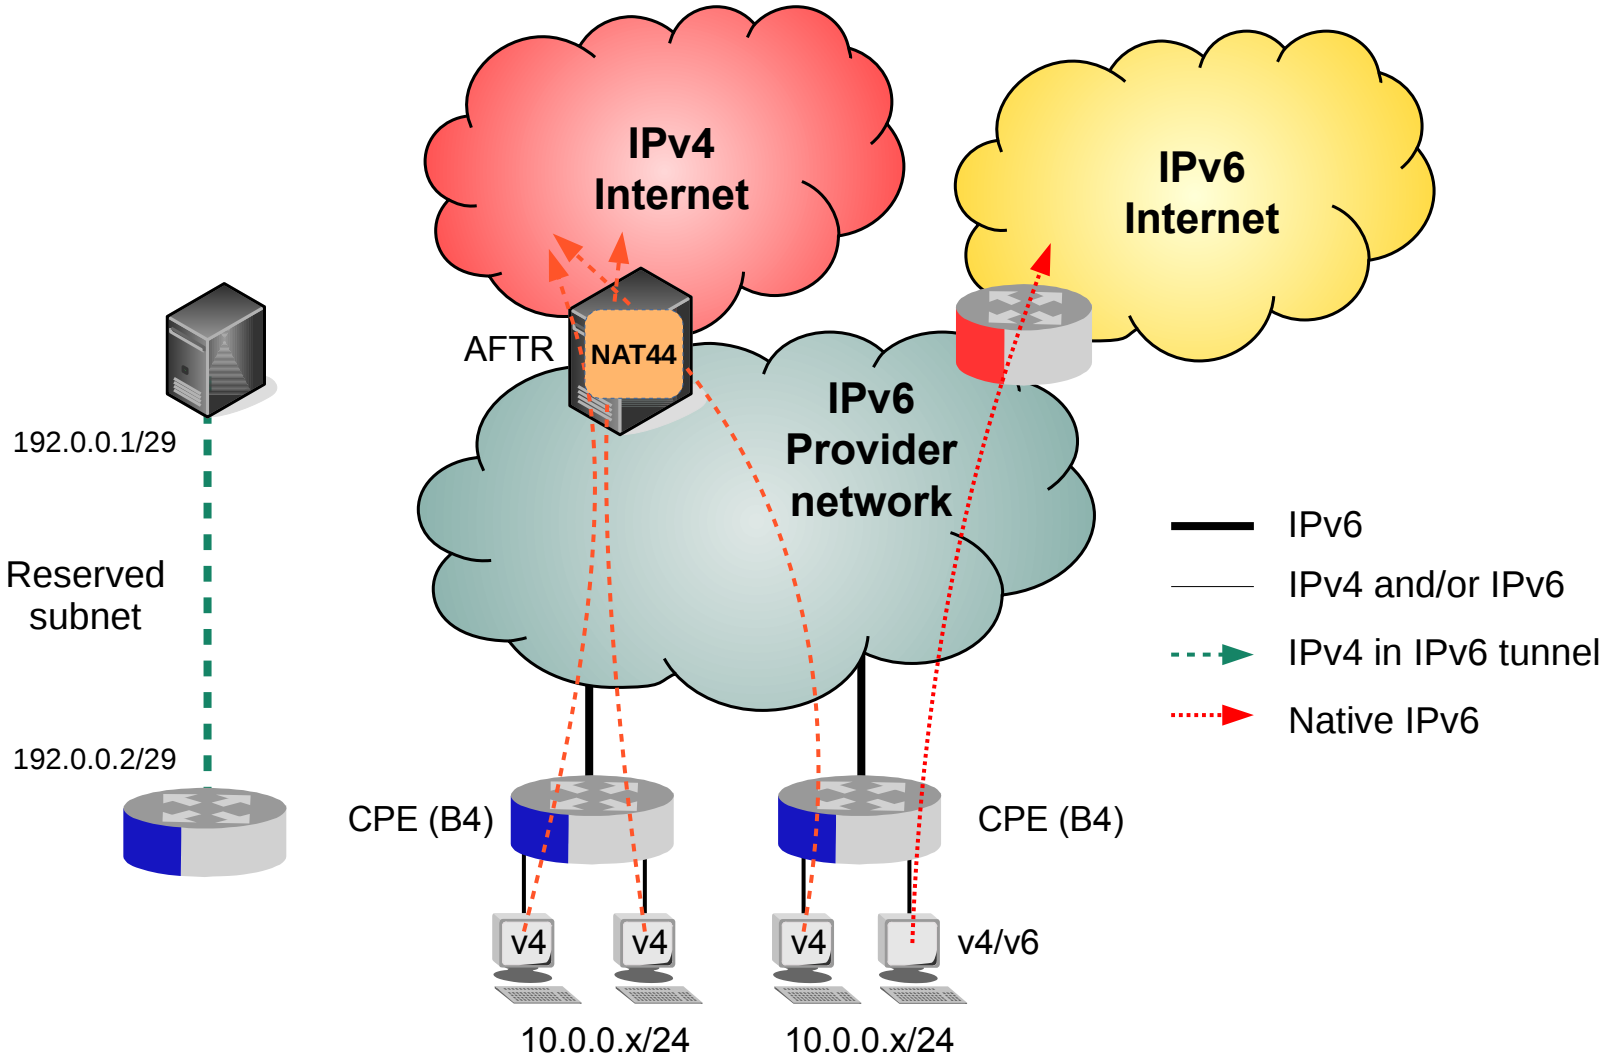
IPv4
Internet
IPv6
Internet
AFTR
NAT44
IPv6
Provider
network
192.0.0.1/29
Reserved
subnet
192.0.0.2/29
CPE (B4)
CPE (B4)
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24
IPv6
IPv4 and/or IPv6
IPv4 in IPv6 tunnel
Native IPv6

C²S CONSULTING
communications/computing/security

lw4o6

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

lw4o6

- lw4o6 extends DS-Lite by moving the NAT functionality from the ISP side to the CPE, eliminating the need to implement CGN.
- There are three main components in the lw4o6 architecture:
  - **lwB4**, performs NAPT and IPv4/IPv6 encapsulation and de-encapsulation
  - **lwAFTR**, performs the IPv4/IPv6 encapsulation and de-encapsulation
  - Provisioning system assigns IPv4 & port to lwB4.

lw4o6

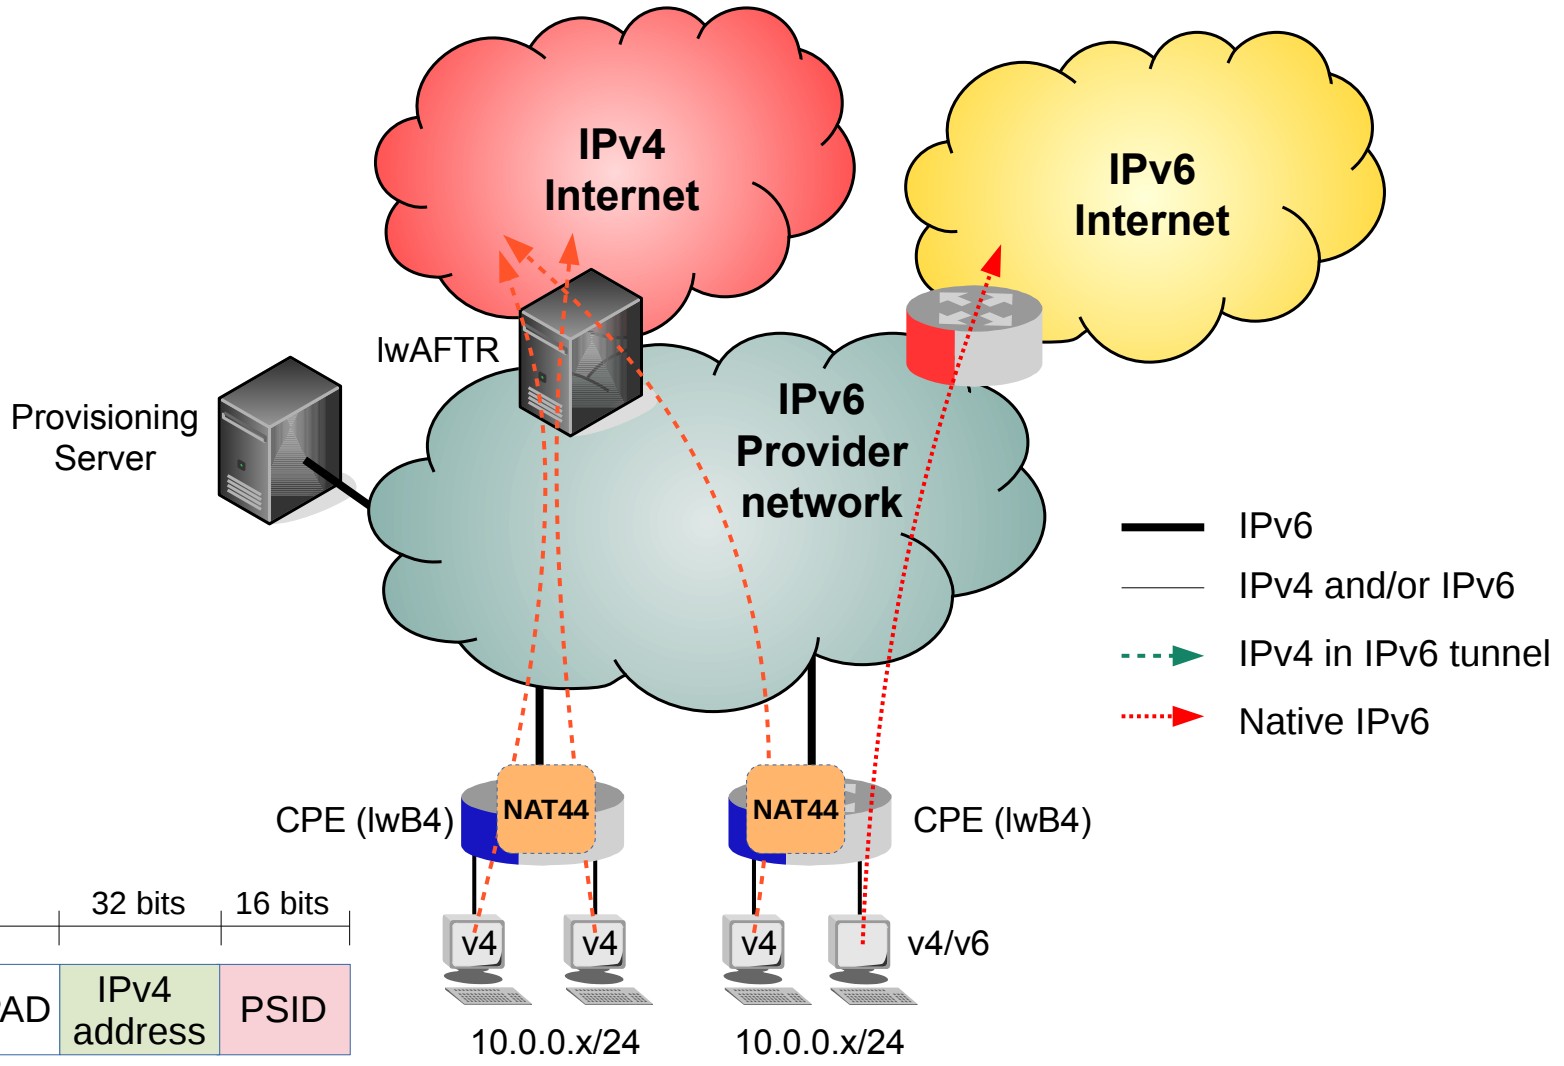
IPv4
Internet
IPv6
Internet
lwAFTR
Provisioning
Server
IPv6
Provider
network
IPv6
IPv4 and/or IPv6
IPv4 in IPv6 tunnel
Native IPv6
CPE (lwB4)
NAT44
NAT44
CPE (lwB4)
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24
64 bits
32 bits
16 bits
Provider
prefix
PAD
IPv4
address
PSID

C²S CONSULTING
communications/computing/security

NAT64/DNS64

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

NAT64 / DNS64

- NAT64 enables v6 only host to communicate with v4 only servers
- Uses the IPv6 prefix 64:ff9b::/96

NAT64 / DNS64

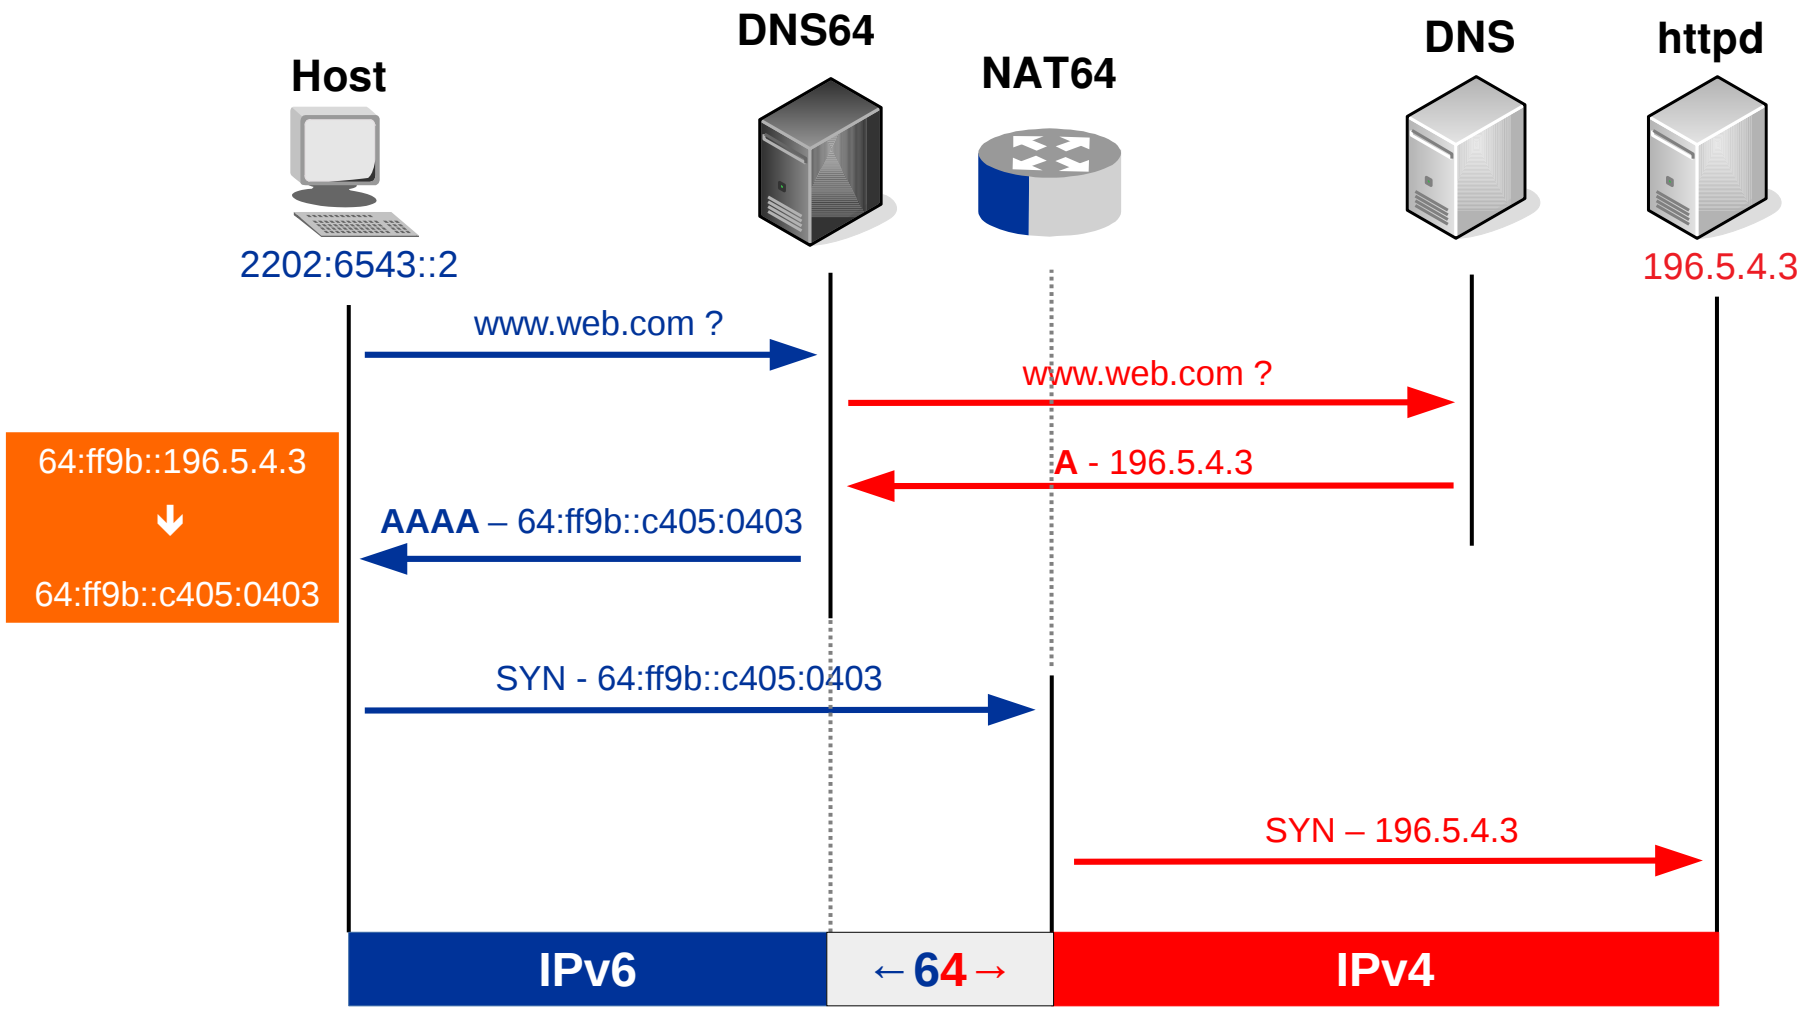
Host
DNS64
NAT64
DNS
httpd
2202:6543::2
196.5.4.3
www.web.com ?
www.web.com ?
A - 196.5.4.3
64:ff9b::196.5.4.3
64:ff9b::c405:0403
AAAA – 64:ff9b::c405:0403
SYN - 64:ff9b::c405:0403
SYN – 196.5.4.3
IPv6
←64→
IPv4

C²S CONSULTING
communications/computing/security

xLAT

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com

xLAT/464LAT

- 464XLAT is a hub and spoke architecture focused on enabling IPv4-only services over IPv6-only networks.
- It is defined in RFC 6877.
- It has two elements:
  - Customer-side transLATor (**CLAT**): translates 1:1 private IPv4 addresses to global IPv6 addresses
  - Provider-side transLATor (**PLAT**): translates N:1 global IPv6 addresses to global IPv4 addresses.

Application and host	Server	Traffic Treatment	Location of Translation
IPv6	IPv6	End-to-End IPv6	None
IPv6	IPv4	Stateful Translation	PLAT
IPv4	IPv4	464XLAT	CLAT/PLAT

XLAT / 464LAT

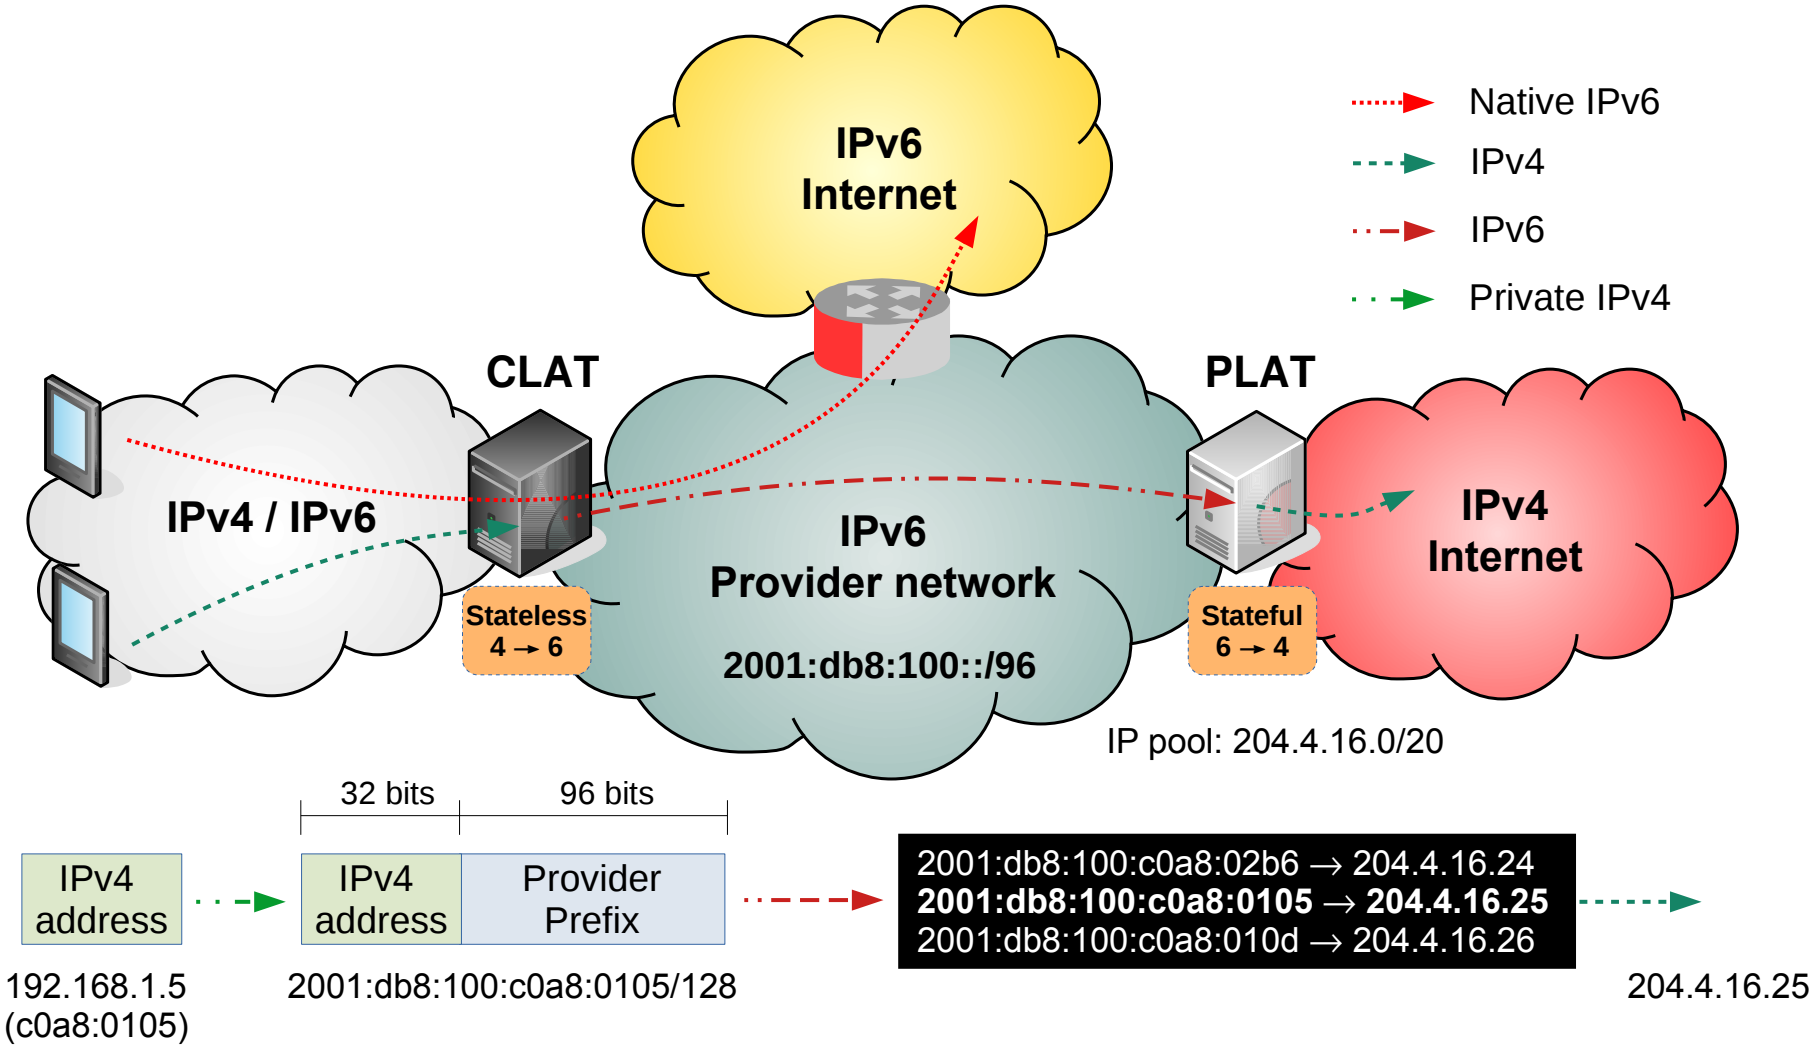
IPv6
Internet
Native IPv6
IPv4
IPv6
Private IPv4
CLAT
PLAT
IPv4 / IPv6
IPv6
Provider network
2001:db8:100::/96
IPv4
Internet
Stateless
4 → 6
Stateful
6 → 4
IP pool: 204.4.16.0/20
32 bits
96 bits
IPv4
address
IPv4
address
Provider
Prefix
2001:db8:100:c0a8:02b6 → 204.4.16.24
2001:db8:100:c0a8:0105 → 204.4.16.25
2001:db8:100:c0a8:010d → 204.4.16.26
192.168.1.5
(c0a8:0105)
2001:db8:100:c0a8:0105/128
204.4.16.25

C²S CONSULTING
communications/computing/security

Thank You

Diarmuid Ó Briain
CEng, FIEI, FIET, CISSP

diarmuid@obriain.com