Internet Protocol Version 6 - IPv6

Getting Familiar with IPv6

22nd May 2017

Deployment of IPv6 in a Campus Network

Thilina Pathirana
(University of Kelaniya)

Overview

- IPv4 addresses were exhausted in IANA on 2011
- RIR started exhausting…
- APNIC in 2011
- RIPE NCC in 2012
- LACNIC in 2014
- ARIN in 2015
- AFRINIC – projected 2018

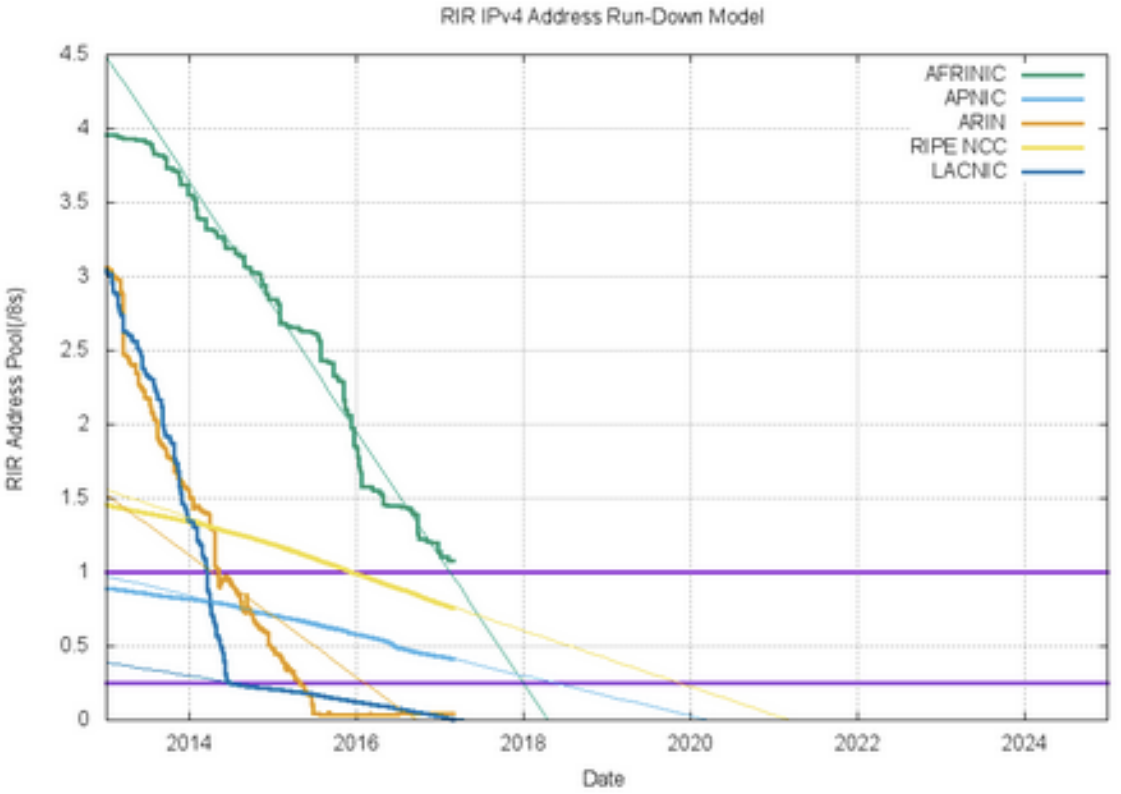

source: ipv4.potaroo.net

Overview

- By 2011 all major Operating systems were supporting IPv6
- In 2012, Internet Society in conjunction with many large companies and organizations, held World IPv6 Launch Day
- In Sri Lanka, IPv6 started way back in 2008 when SLT upgraded their core to support their first IPv6 customer, LEARN.

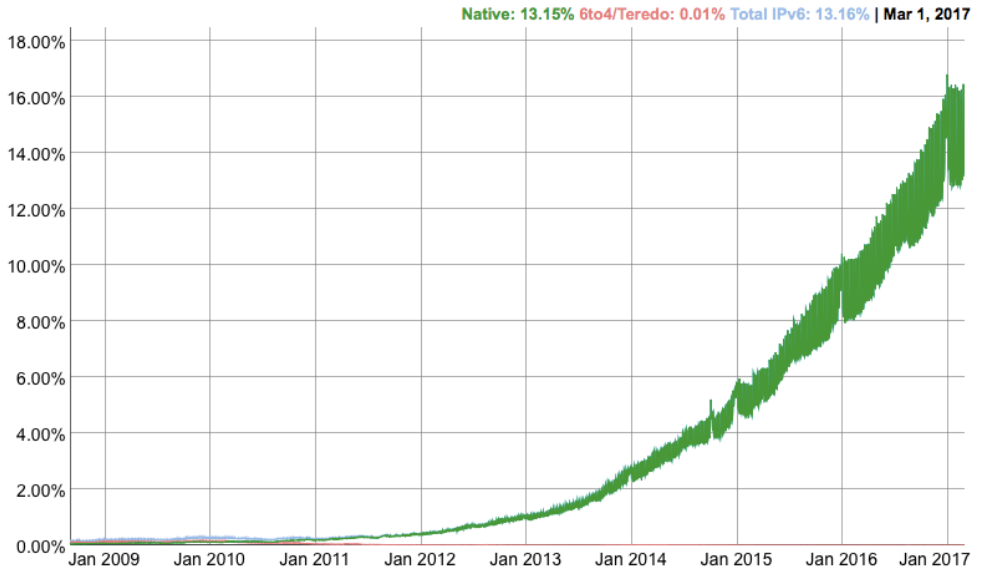

source: https://www.google.com/intl/en/ipv6/statistics.html

Overview

ASN	AS Name	IPv6 Capable	IPv6 Preferred	Samples
AS9329	SLTINT-AS-AP Sri Lanka Telecom Internet	0.01%	0.01%	2,846,150
AS18001	DIALOG-AS Dialog Axiata PLC.	6.06%	5.92%	1,907,332
AS45356	MOBITEL-LK IS Group, No108, W A D Ramanayake Mawatha	0.03%	0.00%	311,405
AS45224	BELLNET-AS-AP Lanka Bells AS	0.00%	0.00%	178,618
AS38229	LEARN-LK Lanka Education Research Network, NREN	26.52%	18.98%	147,437
AS17470	ETISALATLK-AS Etisalat Lanka (Pvt) Ltd.	0.03%	0.00%	123,828
AS132045	AIRTEL-AS-ISP Bharti Airtel Lanka Pvt. Limited	0.02%	0.00%	117,267
AS5087	LANKA-COM Lanka Communication Services	0.02%	0.01%	54,773
AS132447	HUTCHISON-LK 234, Galle Road, Colombo 4	0.01%	0.00%	28,415
AS38573	VIRTUSA-IN-AS Virtusa Global AS	0.33%	0.33%	3,377
AS133051	CBOCP-AS-AP COMMERCIAL BANK OF CEYLON PLC	0.00%	0.00%	2,204
AS132124	ICTA-LK Information and Communication Technology Agency of Sri Lanka	0.00%	0.00%	60
AS16276	OVH OVH SAS	0.00%	0.00%	40
AS17904	SLTASUL-LK Sri Lankan Airlines	0.00%	0.00%	24
AS36351	SOFTLAYER - SoftLayer Technologies Inc.	0.00%	0.00%	5
AS35017	SWIFTWAY-AS Swiftway Sp. z o.o.	0.00%	0.00%	4
AS54334	ROYA - Roya Hosting LLC	0.00%	0.00%	3
AS16276	OVH OVH SAS	0	0	40
AS17904	SLTASUL-LK Sri Lankan Airlines	0	0	24
AS35017	SWIFTWAY-AS Swiftway Sp. z o.o.	0	0	4
AS36351	SOFTLAYER - SoftLayer Technologies Inc.	0	0	5
AS54334	ROYA - Roya Hosting LLC	0	0	3

source: APNIC

Per-Country IPv6 adoption

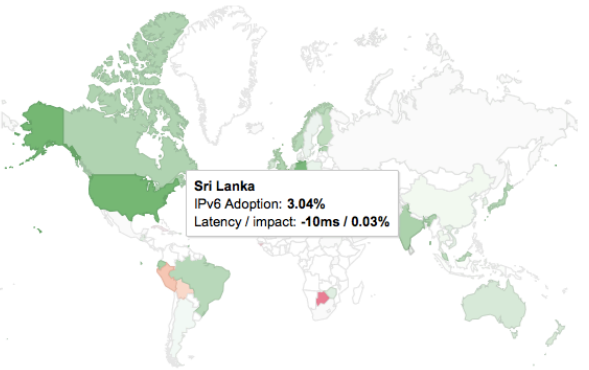

source: https://www.google.com/intl/en/ipv6/statistics.html

IPv6 and LEARN

- IPv6 peering with SLT since 2008-Oct.
- LEARN obtained provider free IPv6 2401:DD00::/32 from APNIC.
- Peering with TEIN3 since 2010.
- All Hosted Services are currently on Dual Stack.
- Member institutes are provided with /48 subnets.
- All institutional access routers except for few are IPv6 enabled
- More info:
  - http://www.learn.ac.lk/ipv6

Why is IPv6?

- Deficiency of IPv4
- Address space exhaustion
- Aggregation-based address hierarchy
  - Efficient backbone routing
- Efficient and Extensible IP datagram
- Stateless Address Auto configuration
- Security (IPsec mandatory)*
- Mobility
- Multicast
- Quality of Service

What's better in IPv6?

- Larger address space
- Better header format
- New options
- Allowance for extension
- Support for resource allocation
- Support for more security
- Support for mobility

IP Evolution

0	IP	March 1977 version	(deprecated)
1	IP	January 1978 version	(deprecated)
2	IP	February 1978 version A	(deprecated)
3	IP	February 1978 version B	(deprecated)
4	IPv4	September 1981 version	(current widespread)
5	ST	Stream Transport	(not a new IP, little use)
6	IPv6	December 1998 version	(formerly SIP, SIPP)
7	CATNIP	IPng evaluation	(formerly TP/IX; deprecated)
8	Pip	IPng evaluation	(deprecated)
9	TUBA	IPng evaluation	(deprecated)
10-15	unassigned		

Source: Cisco.com

128 – bit IPv6 Address

3FFE:08EB:1F2F:0000:0000:0000:00A9:1432

8 groups of 16-bit hexadecimal numbers separated by ":"

For easy handling, leading zeros can be removed

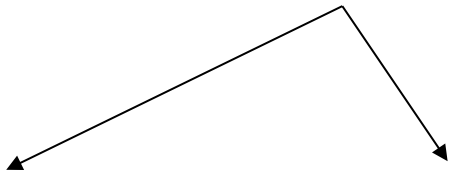

3FFE:8EB:1F2F::A9:1432

:: = all zeros in one or more group of 16-bit hexadecimal numbers

Number of IPv6 Address 2^{128}

340,282,366,920,938,463,463,374,607,431,770,000,000

2^x	10^x	Decimal	IP Quantity	Short Scale	SI Prefix	Equivalent Quantities
2^8	$\approx 10^2$	256	Single IPv4 interface (/24)			
$\approx 2^{10}$	10^3	1,000			kilo	
2^{16}	$\approx 10^5$	65,536	IPv4 Class B (/16)			
$\approx 2^{17}$	10^5	100,000				
$\approx 2^{20}$	10^6	1,000,000		million	mega	
2^{24}	$\approx 10^7$	16,777,216	IPv4 Class A (/8)			
$\approx 2^{30}$	10^9	1,000,000,000		billion	giga	Base pairs in the human genome ($3x10^9$).
2^{32}	$\approx 10^9$	4,294,967,296	Entire IPv4 space			
$\approx 2^{40}$	10^{12}	1,000,000,000,000		trillion	tera	Bacteria on you.
$\approx 2^{50}$	10^{15}	1,000,000,000,000,000		quadrillion	peta	Ants on earth.
$\approx 2^{60}$	10^{18}	1,000,000,000,000,000,000		quintillion	exa	Meters light travels in 100 years.
2^{64}	$\approx 10^{19}$	18,446,744,073,709,551,616	Single IPv6 interface (/64)			
$\approx 2^{70}$	10^{21}	1,000,000,000,000,000,000,000		sextillion	zetta	Grains of sand on earth's beaches.
$\approx 2^{80}$	10^{24}	1,000,000,000,000,000,000,000,000		septillion	yotta	Stars in the universe.
2^{80}	$\approx 10^{24}$	1,208,925,819,614,629,174,706,176	IPv6 Site (/48)			
$\approx 2^{90}$	10^{27}	1,000,000,000,000,000,000,000,000,000		octillion		Atoms in you ($7x10^{27}$).
2^{96}	$\approx 10^{29}$	79,228,162,514,264,337,593,543,950,336	IPv6 ISP/Large enterprise (/32)			
$\approx 2^{100}$	10^{30}	1,000,000,000,000,000,000,000,000,000,000		nonillion		Bacterial cells on earth ($5x10^{30}$).
$\approx 2^{110}$	10^{33}	1,000,000,000,000,000,000,000,000,000,000,000		decillion		Mass of the Sun in grams ($2x10^{33}$).
2^{116}	$\approx 10^{35}$	83,076,749,736,557,242,056,487,941,267,521,536	IPv6, RIR (/12)			
$\approx 2^{120}$	10^{36}	1,000,000,000,000,000,000,000,000,000,000,000,000		undecillion		Ratio of force of electromagnetism to gravity.
2^{125}	$\approx 10^{37}$	42,535,295,865,117,307,932,921,825,928,971,026,432	IPv6 GUA (2000::/3)			
2^{128}	$\approx 10^{38}$	340,282,366,920,938,463,463,374,607,431,768,211,456	Entire IPv6 space			
$\approx 2^{130}$	10^{39}	1,000,000,000,000,000,000,000,000,000,000,000,000,000		duodecillion		Molecules of H_2O in Great Lakes ($53x10^{39}$).

Recall that the character ≈ means *approximately equal to.*

IPv4 Header vs IPv6 Header

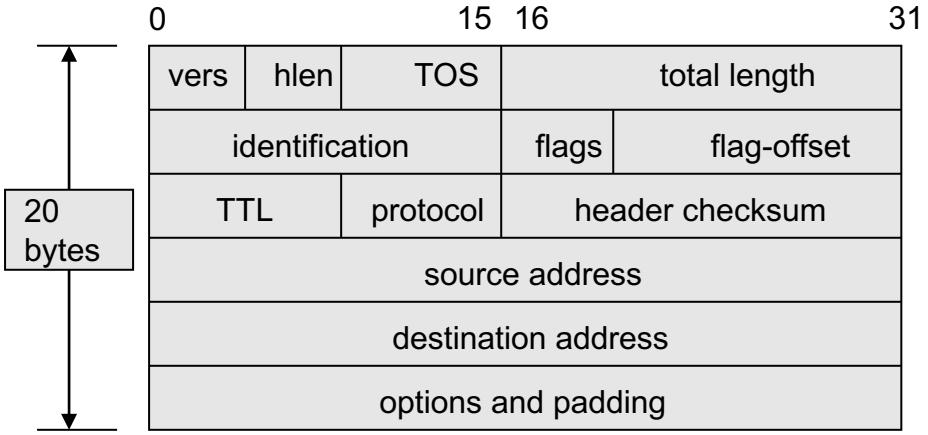

IPv4

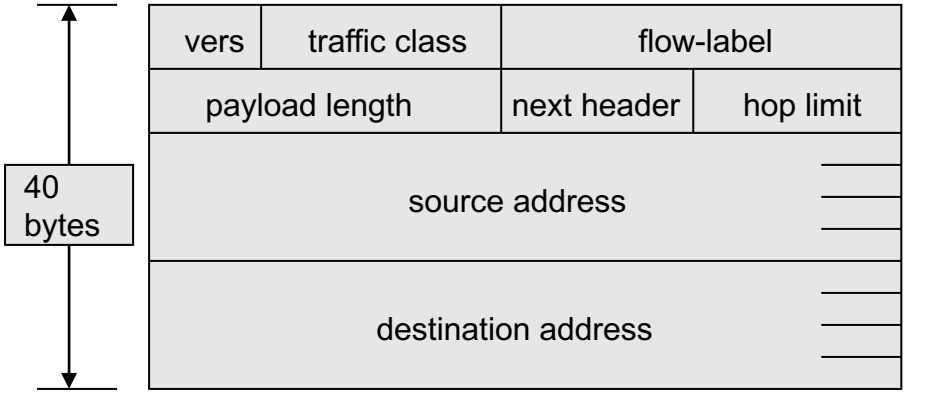

IPv6

Removed (6)

- ID, flags, flag offset
- TOS, hlen
- header checksum

Changed (3)

- total length => payload
- protocol => next header
- TTL => hop limit

Added (2)

- traffic class
- flow label

Expanded

- address 32 to 128 bits

Connection Methods

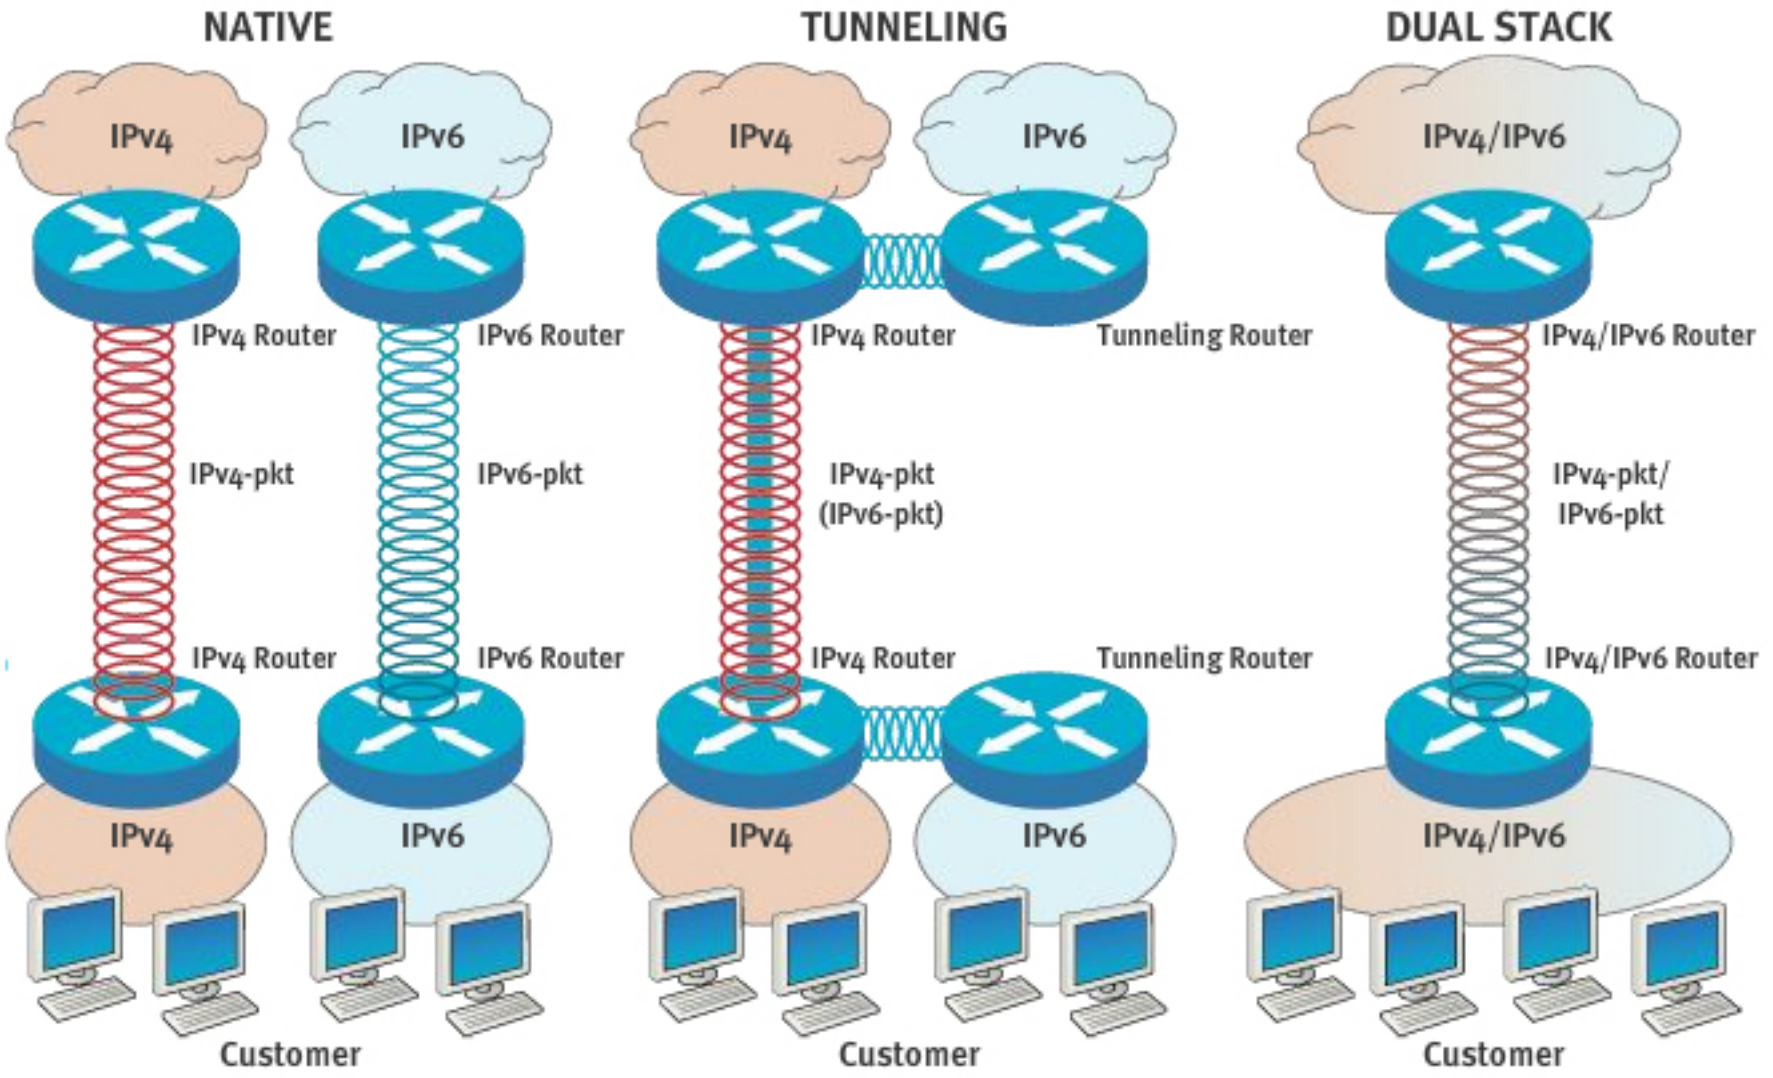

Source: NTT America

IPv6 Address Types

Unicast

- Identifies a single node, and traffic destined to a unicast address is forwarded to a single node.

Multicast

- Identifies a group of nodes, and traffic destined to a multicast address is forwarded to all the nodes in the group.

Anycast

- Identifies a group of nodes, and traffic destined to an anycast address is forwarded to the nearest node in the group.

NO Broadcast Address

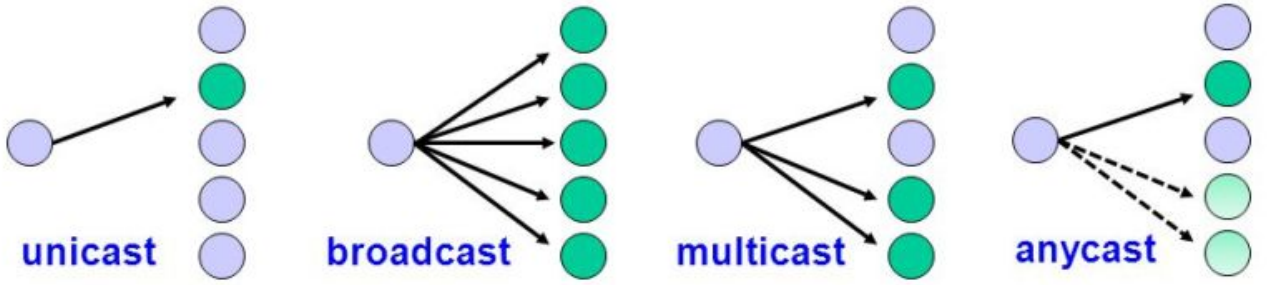

IPv6 Address Allocation

Address	First Bits of the Address	Specifications
0 :: /8	0000 0000	Reserved
100 :: /8	0000 0001	Not Assigned
200 :: /7	0000 0001	ISO Address
400 :: /7)	0000 010	Novell Address (IPX)
600 :: /7	0000 011	Not Assigned
800 :: /5	0000 1	Not Assigned
1000 :: /4	0001	Not Assigned
2000 :: /3	001	Not Assigned
4000 :: /3	010	Addresses of Service Providers
6000 :: /3	011	Not Assigned
8000 :: /3	100	Geographical Addresses of Users
A000 :: /3	101	Not Assigned
C000 ::/3	110	Not Assigned
E000 :: /4	1110	Not Assigned
F000 :: /5	1111 0	Not Assigned
F800 :: /6	1111 10	Not Assigned
FC00 :: /7	1111 110	Not Assigned
FE00 :: /9	1111 1110 0	Not Assigned
FE80 :: /10	1111 1110 10	Local Links Addresses
FEC0 :: /10	1111 1110 11	Local Site Addresses
FF00 :: /8	1111 1111	Multipoint Address

Special Addresses

Address block (CIDR)	Range	Number of addresses	Scope	Purpose
::/128	::	1	Software	Unspecified address
::1/128	::1	1	Host	Loopback address to the local host.
::ffff:0:0/96	::ffff:0.0.0.0 – ::ffff:255.255.255.255	2^{32}	Software	IPv4 mapped addresses
64:ff9b::/96	64:ff9b::0.0.0.0 – 64:ff9b::255.255.255.255	2^{32}	Global Internet[12]	IPv4/IPv6 translation (RFC 6052)
100::/64	100:: – 100::ffff:ffff:ffff:ffff	2^{64}		Discard prefix RFC 6666
2001::/32	2001:: – 2001::ffff:ffff:ffff:ffff:ffff:ffff	2^{96}	Global	Teredo tunneling
2001:10::/28	2001:10:: – 2001:1f:ffff:ffff:ffff:ffff:ffff:ffff	2^{100}	Software	Deprecated (previously ORCHID)
2001:20::/28	2001:20:: – 2001:2f:ffff:ffff:ffff:ffff:ffff:ffff	2^{100}	Software	ORCHIDv2
2001:db8::/32	2001:db8:: – 2001:db8:ffff:ffff:ffff:ffff:ffff:ffff	2^{96}	Documentation	Addresses used in documentation and example source code
2002::/16	2002:: – 2002:ffff:ffff:ffff:ffff:ffff:ffff:ffff	2^{112}	Global Internet	6to4
fc00::/7	fc00:: – fdff:ffff:ffff:ffff:ffff:ffff:ffff:ffff	2^{121}	Private network	Unique local address
fe80::/10	fe80:: – febf:ffff:ffff:ffff:ffff:ffff:ffff:ffff	2^{118}	Link	Link-local address
ff00::/8	ff00:: – ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff	2^{120}	Global Internet	Multicast address

Source: Wikipedia

Special Addresses

Address	Description	Available Scopes
`ff0X::1`	All nodes address, identify the group of all IPv6 nodes	Available in scope 1 (interface-local) and 2 (link-local): • `ff01::1` → All nodes in the interface-local • `ff02::1` → All nodes in the link-local
`ff0X::2`	All routers	Available in scope 1 (interface-local), 2 (link-local) and 5 (site-local): • `ff01::2` → All routers in the interface-local • `ff02::2` → All routers in the link-local • `ff05::2` → All routers in the site-local
`ff02::5`	OSPFIGP	2 (link-local)
`ff02::6`	OSPFIGP Designated Routers	2 (link-local)
`ff02::9`	RIP Routers	2 (link-local)
`ff02::a`	EIGRP Routers	2 (link-local)
`ff02::d`	All PIM Routers	2 (link-local)
`ff02::1a`	All RPL Routers	2 (link-local)
`ff0X::fb`	mDNSv6	Available in all scopes
`ff0X::101`	All Network Time Protocol (NTP) servers	Available in all scopes
`ff02::1:1`	Link Name	2 (link-local)
`ff02::1:2`	All-dhcp-agents	2 (link-local)
`ff02::1:3`	Link-local Multicast Name Resolution	2 (link-local)
`ff05::1:3`	All-dhcp-servers	5 (site-local)
`ff02::1:ff00:0/104`	Solicited-node multicast address. See below	2 (link-local)
`ff02::2:ff00:0/104`	Node Information Queries	2 (link-local)

Source: Wikipedia

IPv6 ISO Address Allocation Format

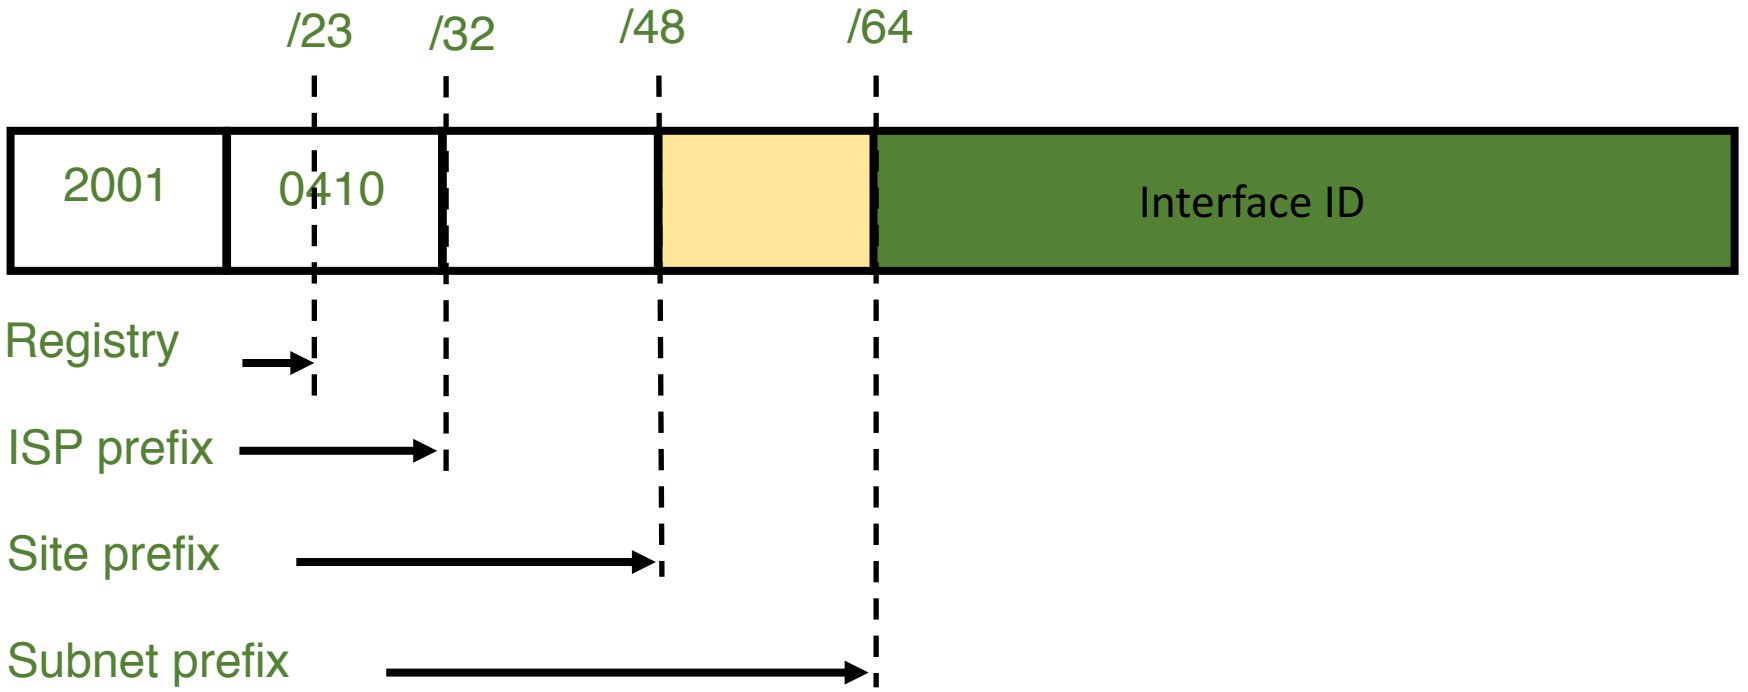

Eg:
LEARN Address Space -- 2401:DD00:0000:0000:0000:0000:0000:0000/32

UoK Address Space -- 2401:DD00:0020:0000:0000:0000:0000:0000/48

Defining Interface ID

- Associate layer 2 address in the Extended Unique Identifier (EUI-64).

  For the interface ID portion of the network, the seventh high-order bit of the EUI-64 format defines a local scope when set to 0 and a global scope (globally unique) when set to 1.

- Auto-generate a random address as defined in RFC 3041.

  This mechanism was developed basically to limit the exposure of a globally reachable address and to increase privacy

- via DHCPv6.

- Manual configuration.

- Cryptographically generated addresses (CGAs) based on RFC 3972

  This is done through a hash that includes a public key. This method of generating an interface ID provides added security and enables address authentication.

Extended Unique Identifier (EUI – 64)

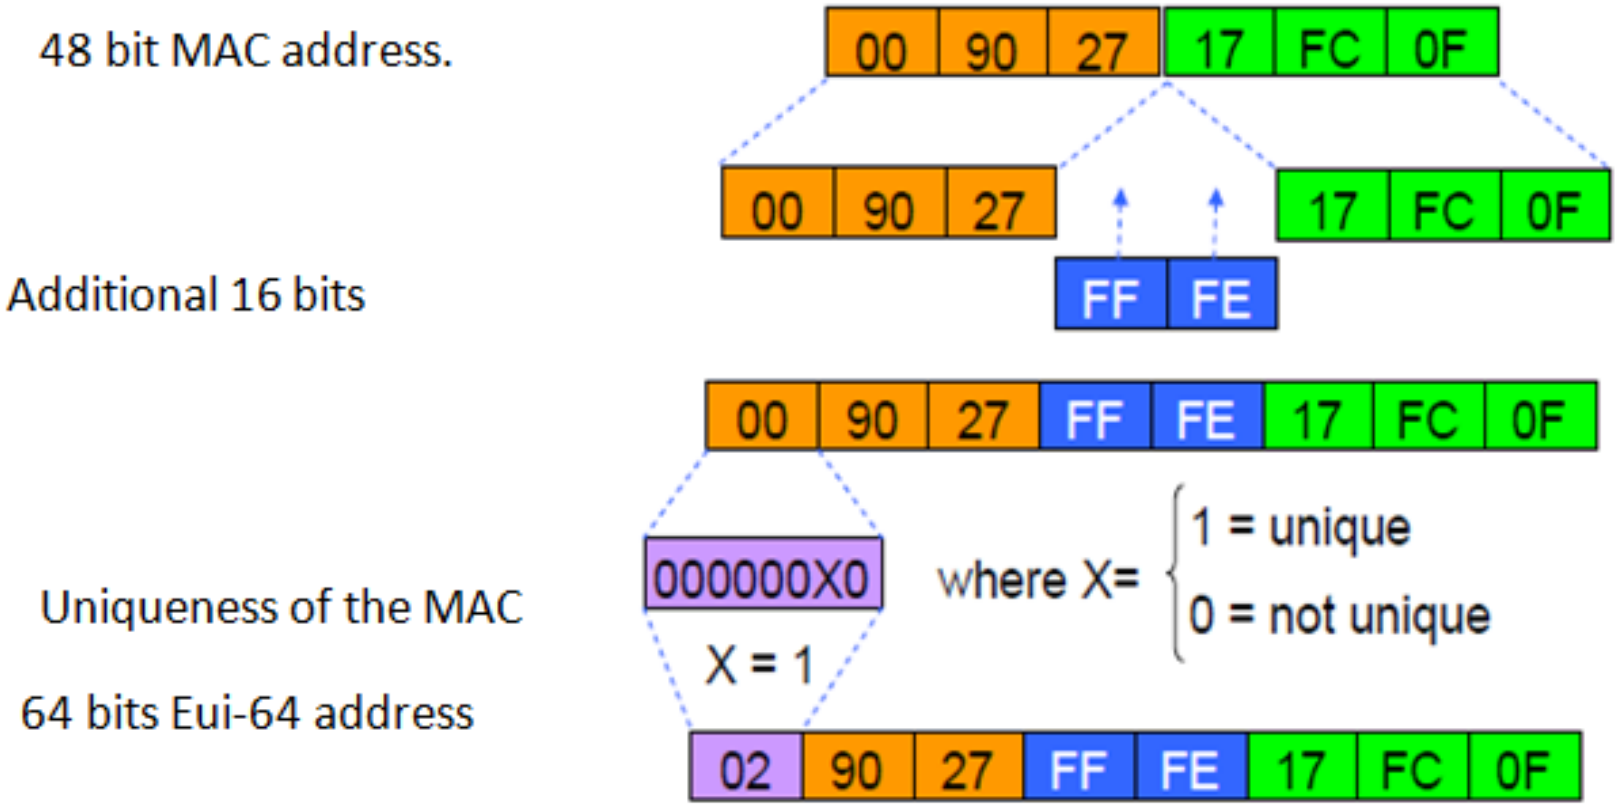

IPv6 Address Types

IPv6 Address Type	Description
Global Unicast	Destined for a single recipient and can be routed on the public Internet
Multicast	Destined for members of a multicast group
Link Local	Valid only on a network segment
Unique Local	Cannot be routed on the public Internet
Loopback	The *localhost* address of a device
Unspecified	Does not specify a source address (all 128 bits in the IPv6 address set to zeros)
Solicited-Node Multicast	A multicast IPv6 address corresponding to a device's IPv6 address(es)

Link-Local Address

- Meaningful only in a single link zone, and may be re-used on other links
- Link-local addresses for use during auto-configuration and when no routers are present
- Required for Neighbor Discovery process, normally automatically configured
- An IPv6 router never forwards link-local traffic beyond the link
- Prefix= FE80::/64

ICMPv6

- An integral part of IPv6 and MUST be fully implement by every IPv6 node (RFC 2463)
- Next Header value= 58
- Report delivery or forwarding errors
- Provide simple echo service for troubleshooting
- Neighbor Discovery (ND): 5 ICMP messages
- Multicast Listener Discovery (MLD): 3 ICMP messages

ICMPv6

-

Message Number	Message Type	Code Field
128	Echo Request	RFC 4443. Used for the ping command
129	Echo Reply	
130	Multicast Listener Query	RFC 2710. Used for multicast group management
131	Multicast Listener Report	
132	Multicast Listener Done	
133	Router Solicitation	RFC 4861. Used for neighbor discovery and autoconfiguration
134	Router Advertisement	
135	Neighbor Solicitation	
136	Neighbor Advertisement	
137	Redirect Message	
200 and 201	Private Experimentation	RFC 4443
255	Reserved for expansion of ICMPv6 informational messages	RFC 4443

Address Auto-Configuration – SLAAC

- SLAAC – Stateless Address Auto Configuration
  - No Need of DHCP Server

- Step1 : A newly connected device will send an ICMPv6 Router Solicitation(133) to the network requesting network prefix from Routers.
- Step 2: Router will reply with a Router Advertisement(134) containing required details
- Step 3: New host will generate its Global Unicast Address using prefix sent and the Interface ID generated using EUI-64 or Random.

Address Auto-Configuration – SLAAC

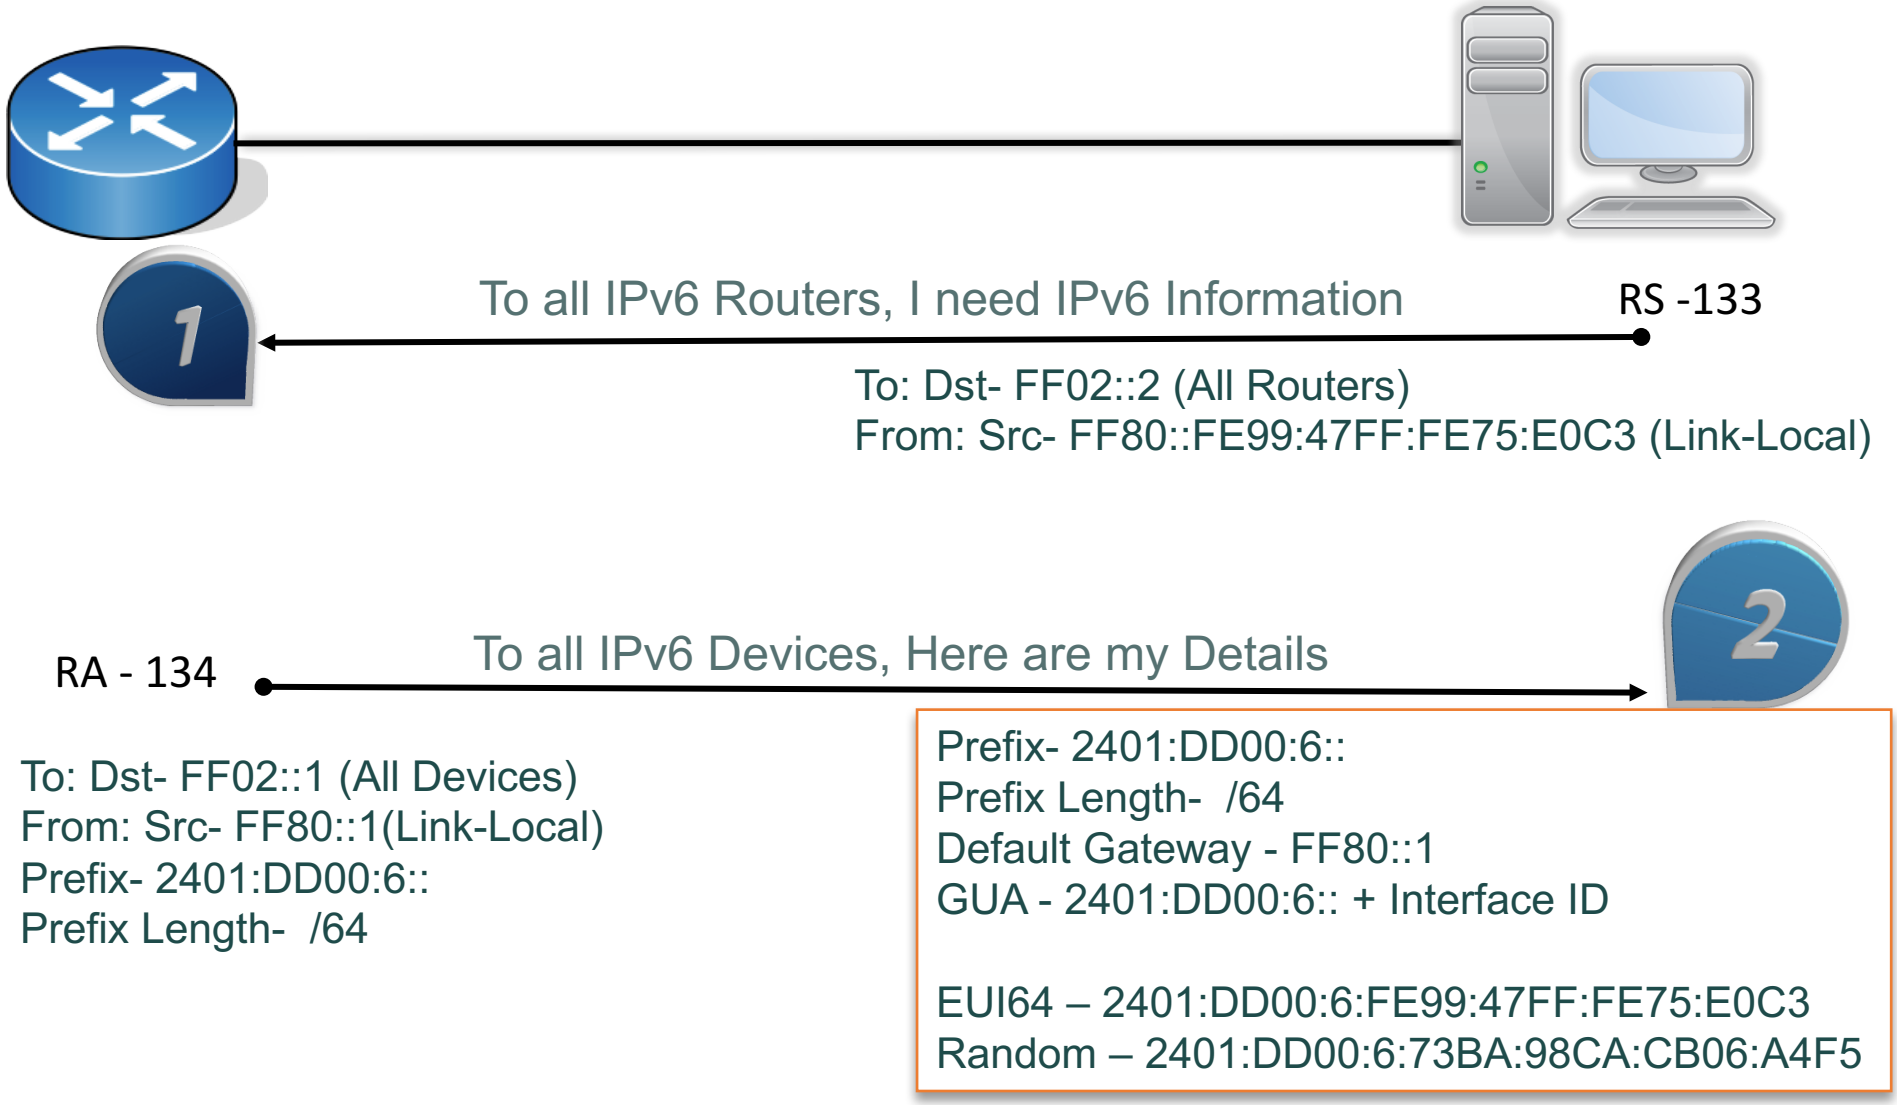

Duplicate Address Detection - DAD

- Since SLACC and Link Local addresses are configured automatically and no inventory like DHCP is recording the usage we need to confirm that the selected address is unique.

- Step1 : Neighbor Solicitation (NS) is sent to own solicited node Multicast address

- Step 2: If the new address is already in use, a Neighbor Advertisement is received through the multicast group FF02::1, if the NS gets timed out then the selected address will be marked as unique and it will be selected as the address

Duplicate Address Detection - DAD

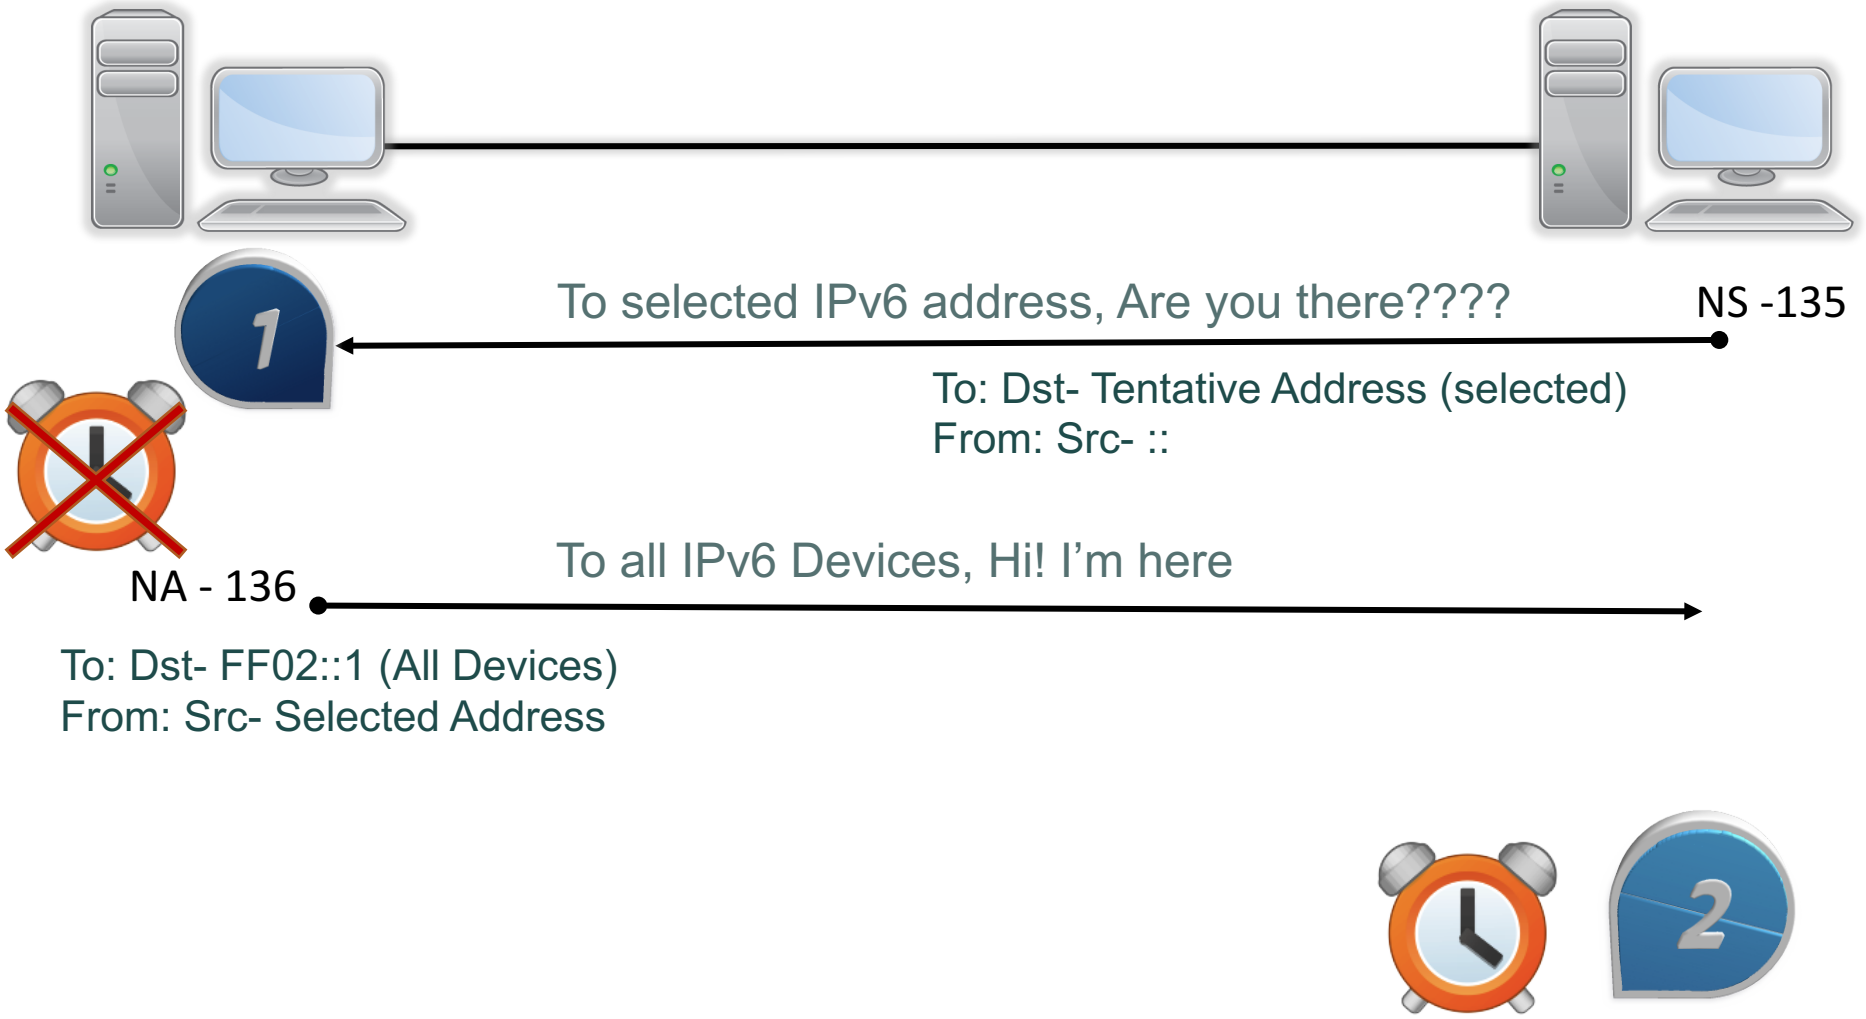

Address Auto Configuration – DHCPv6

- Statefull Configuration
- Problem w/ stateless: Anyone can connect
- Stateless auto-configuration only configures addresses; not “other configuration” information (DNS servers, domain search list, ...)
- Stateless auto-configuration is “one-size fits all” Addresses can not be selectively assigned
- Policies can not be enforced about clients allowed addresses
- So routers ask the new system to go DHCP server (by setting managed configuration bit)
- System multicasts to "All DHCP servers"
- DHCP server assigns an address

Neighbor Discovery Protocol - ND

- Node (Hosts and Routers) use ND to determinate the link-layer addresses for neighbors known to reside on attached links and quick purge cached valued that become invalid
- Hosts also use ND to find neighboring router that willing to forward packets on their behalf
- Nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses
- Replace ARP, ICMP Router Discovery, and ICMP Redirect used in IPv4

Neighbor Discovery Protocol - ND

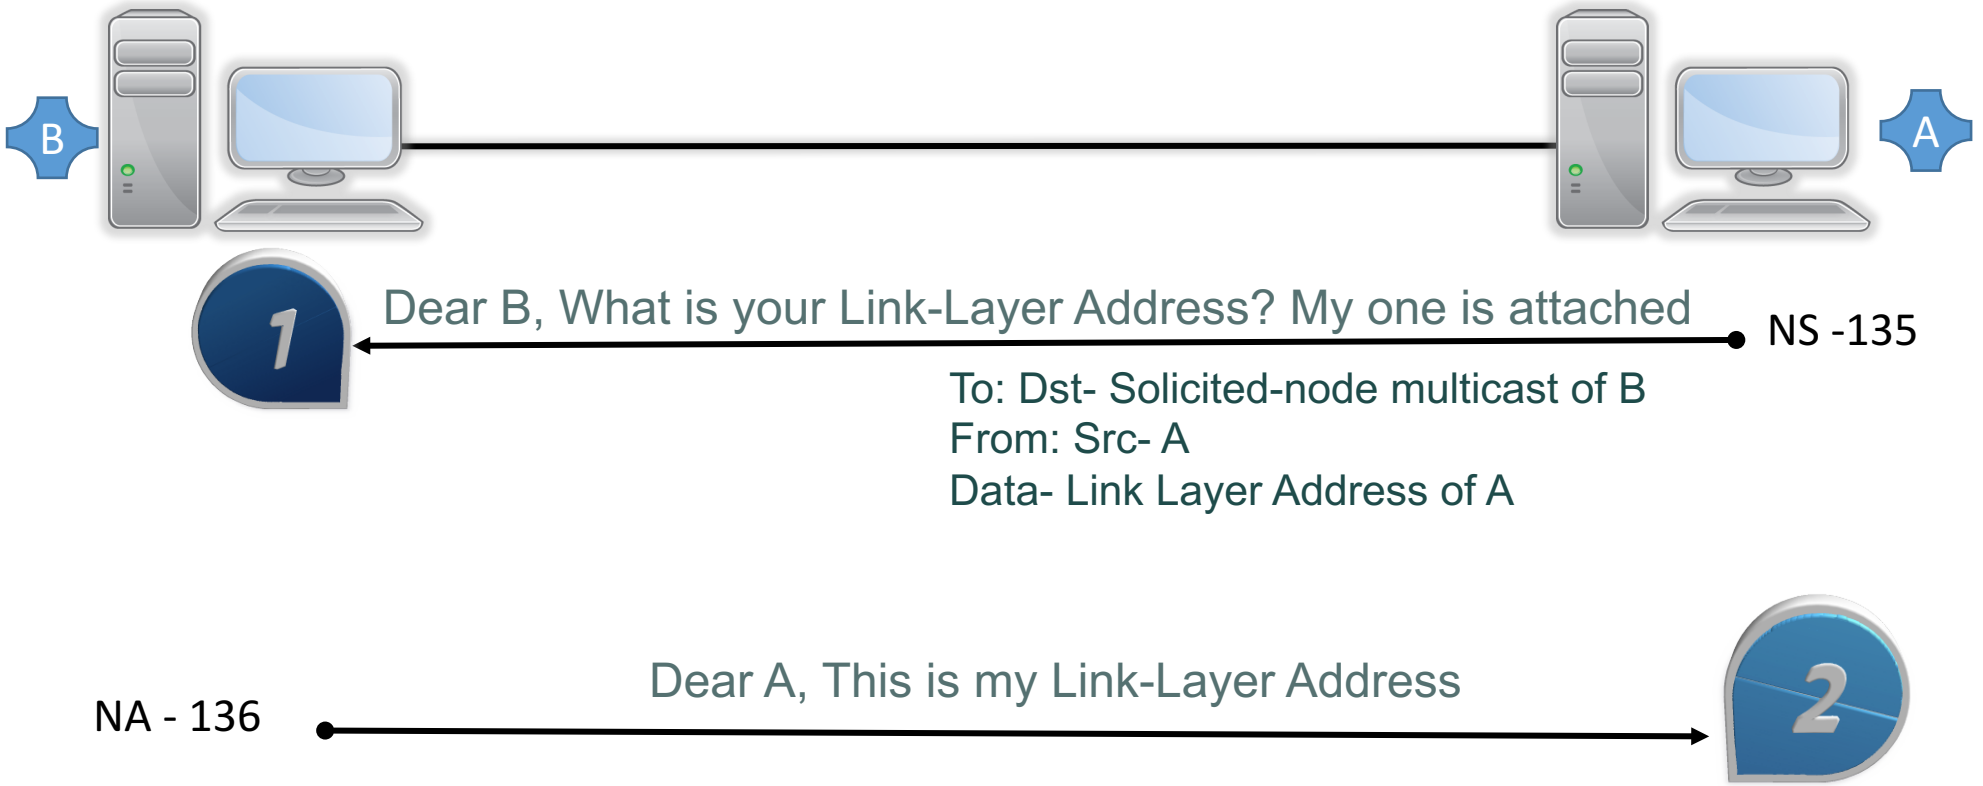

Embedding IPv6 in to Campus Network Design

Executive Steps

- Check whether your devices are IPv6 enabled
  - Today all new OS / HW are IPv6 Capable, please check with your vendors.
- Get your IPv6 range from LEARN (if haven't already)
- Create a suitable subnet plan which matches your existing topology
  - Try to stick with Dual Stack environment
- First configure Core network at your institute with ipv6
- Verify connectivity with Outside
- Configure internally hosted Servers (specially any DNS servers)
- Configure internal vlans
- Confirm End User Connectivity

Campus Network Subnets

- Assume the IP range provided by LEARN as 2401:DD00:SSSS::/48
  (where SSSS can be any pre-defined 0-9,A-F characters)
- Subnetting 2401:DD00:SSSS::/48 into /64s gives 65536 subnets, that is a huge range.
- We will divide /48 into small aggregated units which can be distributed among Faculties, Buildings or Core devices.
- As a standard lets take 2401:DD00:SSSS:XXYY:: where XX represents a Core unit and YY represents a sub unit within that core.
- Then there are 256 /56 when we change XX, and each XX will contain another 256 /64's
- Eg: XX can be Faculties and YY can be departments in it.

Campus Network Subnets

- Example Subnet

Faculty / Aggregation	Department	Vlan	Allocation
Faculty 1 2401:DD00:SSSS:1000::/52	Dept 1	10	2401:DD00:SSSS:1110::/64
	Dept 2 2401:DD00:SSSS:1200::/56	20	2401:DD00:SSSS:1220::/64
		21	2401:DD00:SSSS:1221::/64
Faculty 2 2401:DD00:SSSS:2000::/52	Dept 3	30	2401:DD00:SSSS:2330::/64
	Dept 4	40	2401:DD00:SSSS:2440::/64
Faculty 3 2401:DD00:SSSS:3000::/52	Dept 5	50	2401:DD00:SSSS:3550::/64
	Dept 6	60	2401:DD00:SSSS:3660::/64
Faculty 4 2401:DD00:SSSS:4000::/52	Dept 7	70	2401:DD00:SSSS:4770::/64
	Dept 8	80	2401:DD00:SSSS:4880::/64
Core Network	2401:DD00:SSSS:0000::/64 to 2401:DD00:SSSS:00FF::/64		

Campus Network Design

- Always try to assign easy numbers that can be easily remembered
- Keep :ffff of last 4 nibbles for the gateway/vlan interfaces
- Remember to configure DNS otherwise your network will get slow.
- We still have time to go for Native IPv6 therefore master on Dual Stack

Questions

Thank You

Thilina Pathirana/UoK

Email: thilina@kln.ac.lk