Review Question

Examine the topology below

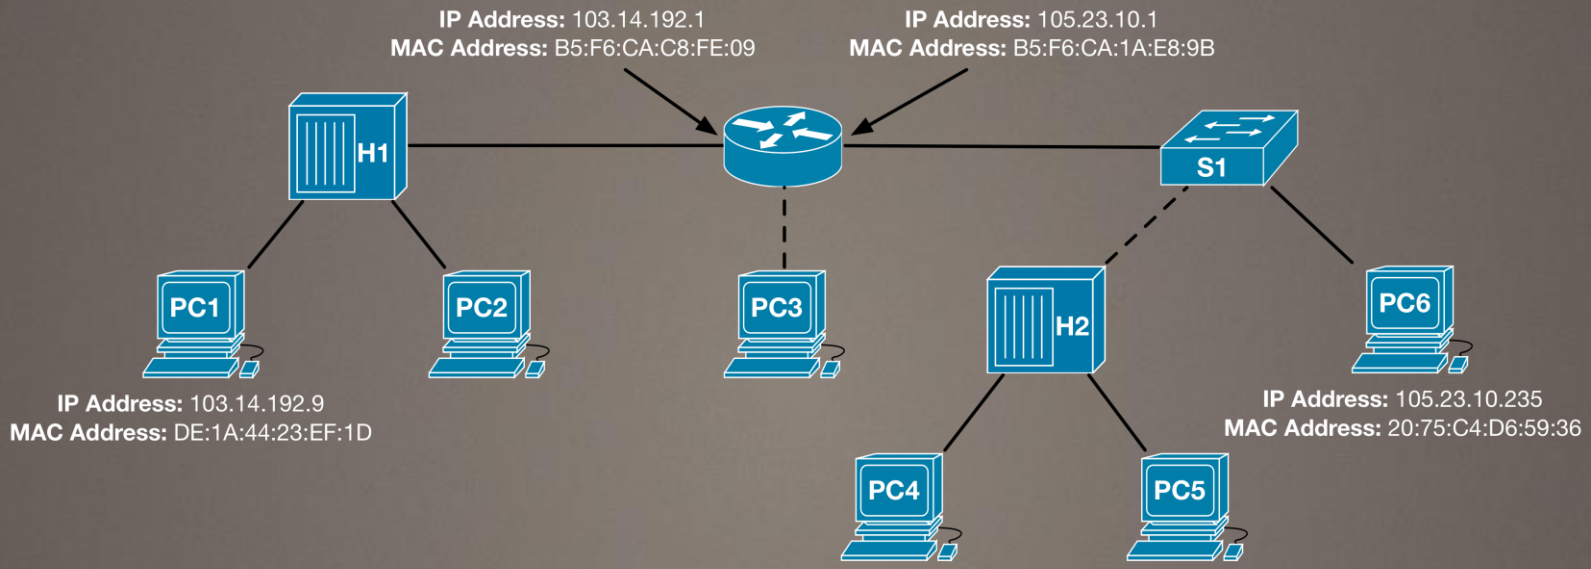

- What destination MAC and IP address will PC1 use to transmit data to PC6?
- What will happen to the packet when it arrives at H1?

Murdoch UNIVERSITY

Review Question

Examine the topology below

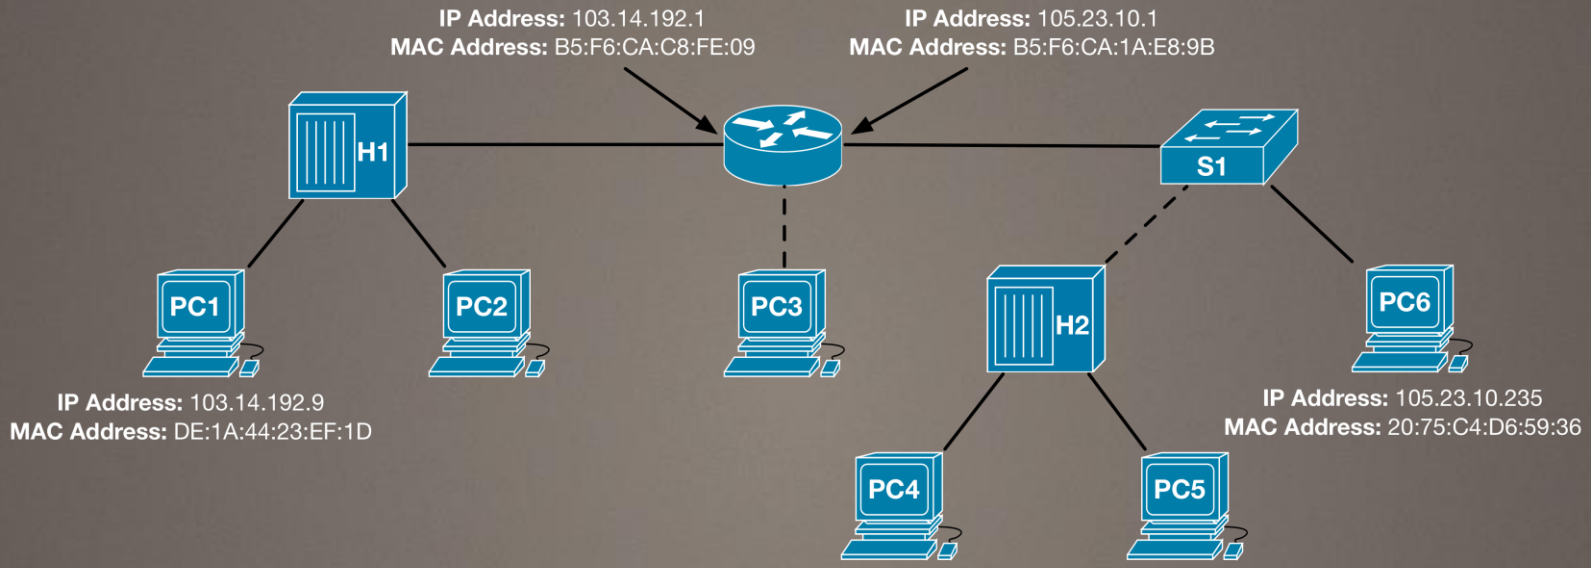

- What destination MAC and IP address will PC1 use to transmit data to PC6?

Destination MAC: B5:F6:CA:C8:FE:09
Destination IP: 105.23.10.235

Murdoch UNIVERSITY

Review Question

Examine the topology below

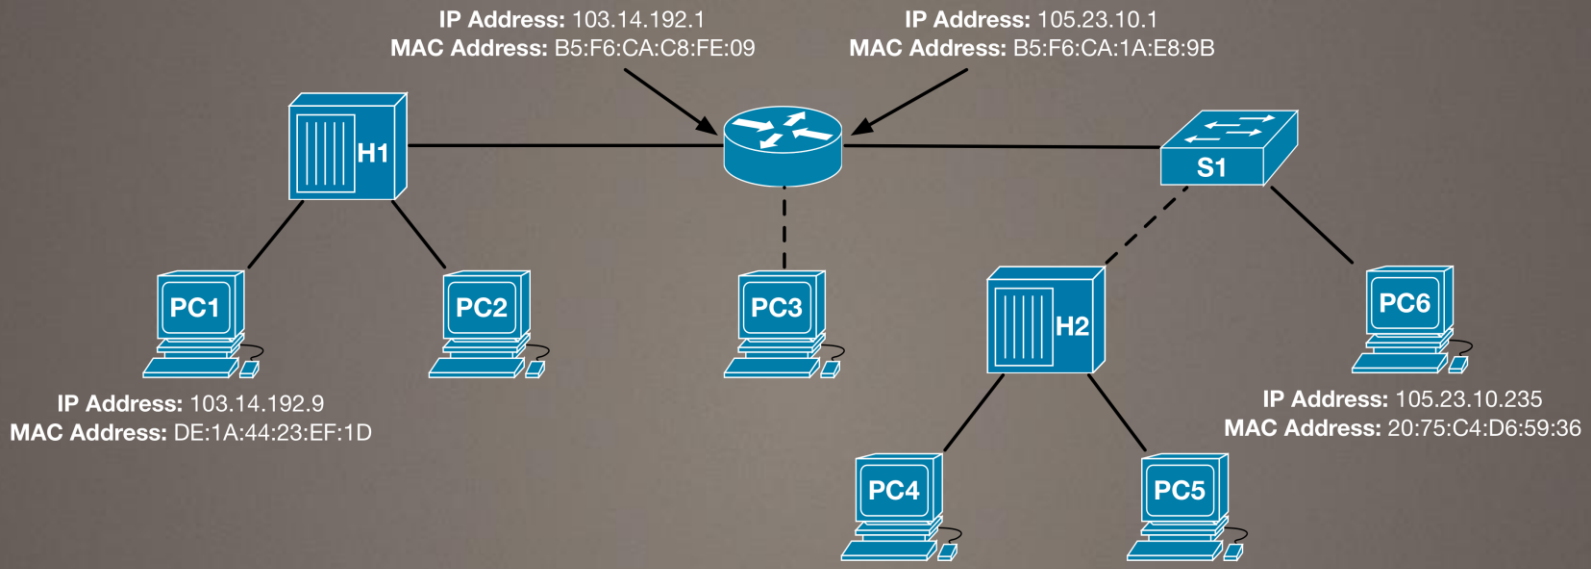

- What will happen to the packet when it arrives at H1?

H1 is an Ethernet hub. Hubs retransmit frames to all ports except the one that received the transmission. Therefore, the packet will be received by PC2 and the Router.

Murdoch UNIVERSITY

Murdoch
UNIVERSITY

IPv6 and Network Address Translation

ICT169

Foundations of Data Communications

Admin

- Mid-Semester Test is this week
    - Internals: During your lab (check which one you're enrolled in)
    - Externals: During announced time slots
- After-hours access to the datacentre

Murdoch UNIVERSITY

Last Week

- We looked at Ethernet, one of the most common Layer 2 technologies in data communications
    - Switch operation
    - Cabling
    - ARP
- VLANs and Inter-VLAN routing

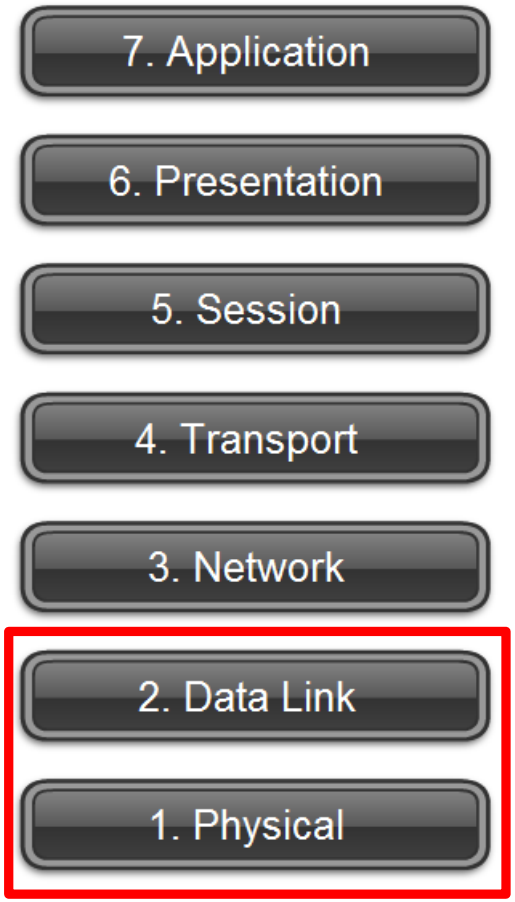

Murdoch UNIVERSITY

Lecture Overview

- Limitations of IPv4; IPv4 address space exhaustion
- The (interim) solution: Network Address Translation (NAT)
- The actual solution: IP version 6 (IPv6)
    - IPv6 operation
    - Addressing and subnetting
    - Address allocation
- Transitioning to IPv6

IPv4 and IPv4 Addresses Revisited

- IPv4 is the most widely used Network layer protocol
- Has been used to run the Internet since inception
- IPv4 addresses are 32 bits long, broken up into 4 octets separated by dots
    - Each octet represents 8 bits (values between 0 and 255)
    - Example: 192.168.10.10

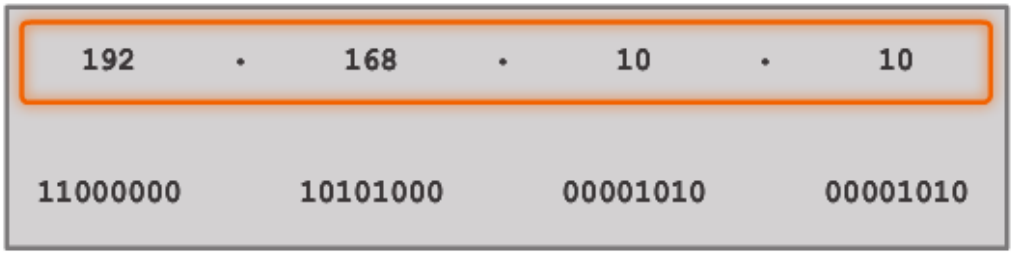

Murdoch UNIVERSITY

IPv4 Addresses and Address Space

- So why 32 bit IPv4 addresses?
  - Processing efficiency (4 bytes)
  - 32 bit addresses were determined to be sufficient when IPv4 was designed
- 32 bit addresses allowed for 2^{32} (4,294,967,296) addresses
  - Should have been (almost) enough for one IPv4 address per person as of 1980
- Requirements for IP addresses have changed since 1980 and early allocation practices were very inefficient

Murdoch UNIVERSITY

Classful IP Addressing Revisited

- Until 1993, IP addressing was classful; the subnet mask used was based solely on the IP address
- Classful addressing is extremely inefficient due to the size of each block of IP addresses
- Now deprecated, but not before many were allocated

Address Class	First Octet Range	First Octet Bits	Subnet Mask	Network / Host Portions	Number of Hosts
A	1—127	**0**0000001—**0**1111111	255.0.0.0	N.H.H.H	16,777,214
B	128—191	**10**000000—**10**111111	255.255.0.0	N.N.H.H	65,534
C	192—223	**110**00000—**110**11111	255.255.255.0	N.N.N.H	254
D	224—239	**1110**0000—**1110**1111	-	N/A (Multicast)	-
E	240—255	**1111**0000—**1111**1111	-	N/A (Experimental)	-

IPv4 Address Allocation

- IP addresses are administered by the **Internet Assigned Numbers Authority** (IANA)
- In the early days of the Internet, assignments were handled by a single individual: Jon Postel
- As the Internet grew, this function was delegated to five **Regional Internet Registries** (RIRs)
- Each RIR services a specific region of the world

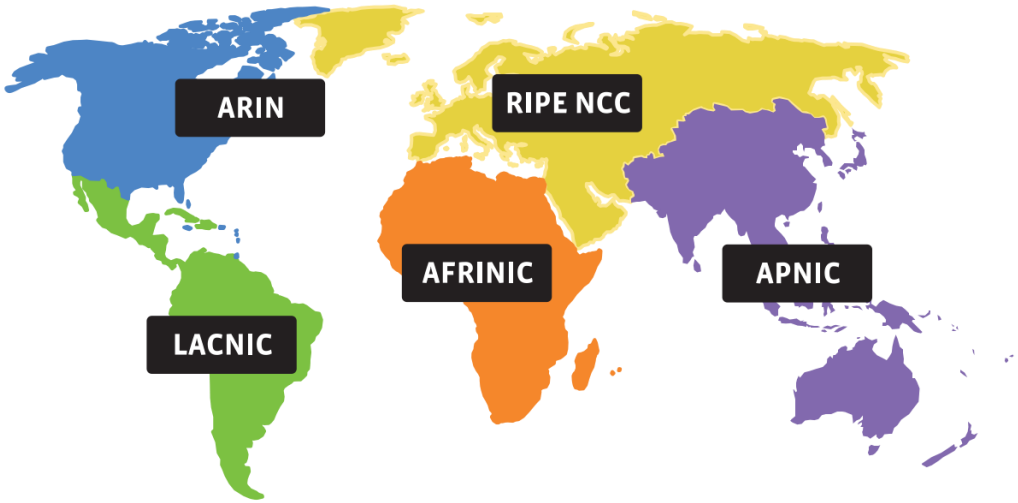

https://www.iana.org/numbers

Murdoch UNIVERSITY

IPv4 Address Allocation – 1993

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Available
Allocated
Unavailable

Murdoch UNIVERSITY

IPv4 Address Allocation – 2000

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Available
Allocated
Unavailable

Murdoch UNIVERSITY

IPv4 Address Allocation – 2007

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Available
Allocated
Unavailable

Murdoch UNIVERSITY

IPv4 Address Allocation – 2010

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Available
Allocated
Unavailable

Murdoch UNIVERSITY

IPv4 Address Exhaustion

- IANA allocated five reserved /8 address blocks to in February 2011
- One block went to each RIR but depleted quickly
    - APNIC depleted in April 2011
    - RIPE initially depleted in September 2012
    - LACNIC depleted in June 2014
    - ARIN depleted September 2015
- AFRINIC the only RIR to still have IPv4 addresses
- Some organisations are returning unused IPv4 addresses, allowing RIRs to reallocate them

Murdoch UNIVERSITY

Combating IPv4 Address Exhaustion

- Classess Inter-Domain Routing (CIDR) was adopted
  - Think VLSM
- Network Address Translation (NAT) became commonplace
- IPv6 was developed to supersede IPv4
- RIRs began limiting the number of IPv4 addresses that organisations could request
- Organisations began purchasing IPv4 addresses
  - Microsoft paid US$7.5m for 666,000 IPv4 addresses in 2011
  - Marketplaces for IPv4 addresses have been established (eg. V4escrow, IPv4 Market Group)

Murdoch UNIVERSITY

Network Address Translation (NAT)

- Historically, each device required a public IP address to communicate over the Internet
- Not sustainable, as the number of Internet-connected devices per household (or even per person) increased
- NAT allows for multiple devices (with private IP addresses) to use a single public IP address

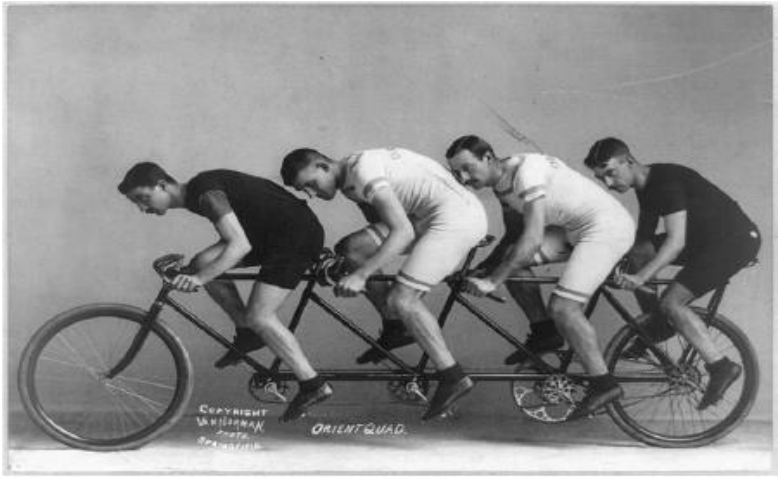

https://commons.wikimedia.org/wiki/File:Bike_for_four.jpg

Murdoch UNIVERSITY

NAT Types

- **One-to-one (basic) NAT:** maps private IP address to a single public IP address
  - Need one public IP address per host that should be able to access the Internet simultaneously
  - Mappings are released when no longer used, allowing another host to have its address translated
- **One-to-many NAT:** maps private IP address and port number combination to a public IP address and port combination
  - Allows multiple hosts to share a single public IP address
  - Sometimes referred to as Port Address Translation (PAT), Network Address Port Translation (NAPT), or NAT Overload
- One-to-many NAT is more commonly used for Internet connections, so we'll focus our discussion on it

NAT Operation

- NAT creates mappings between private IP address and port number combination (eg. 10.1.1.10:32453) to a public IP address and port combination
- Allows multiple devices with private IP addresses to share a single Internet connection
- Breaks the end-to-end nature of the Network layer

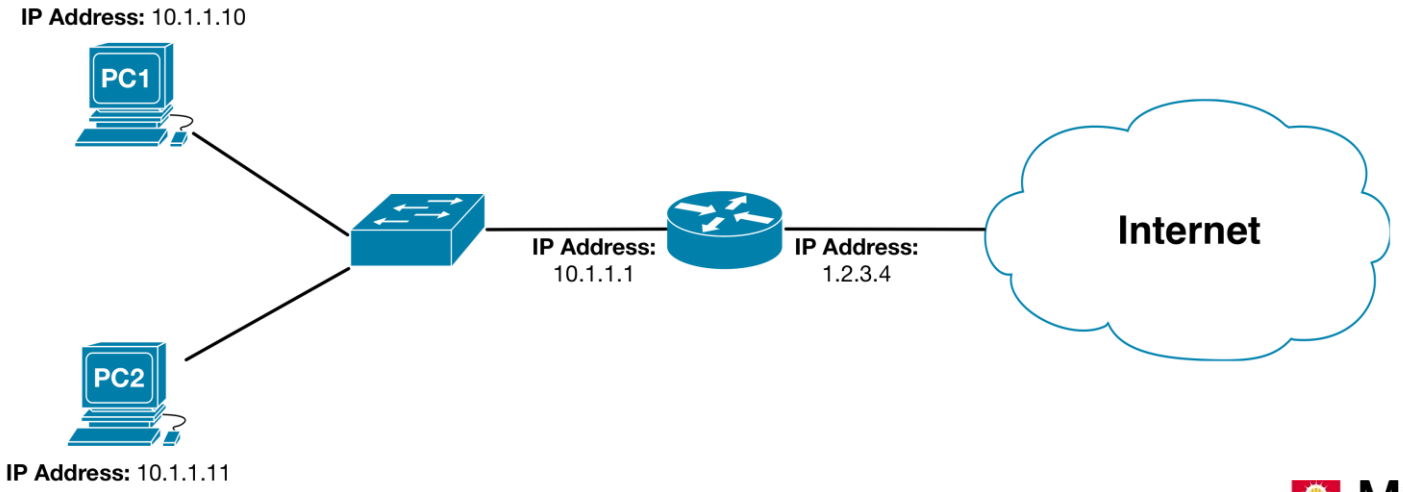

Murdoch UNIVERSITY

NAT Operation (cont.)

- Mappings are created when a device tries to communicate outside of the local network
    - Internal and external port numbers don't have to match
- Stored in a NAT translation table

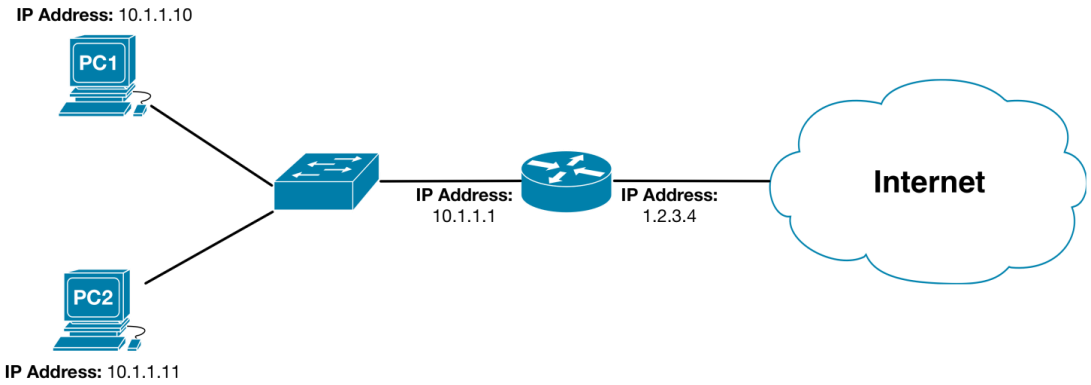

Internal IP	Internal Port	External IP	External Port
10.1.1.10	32453	1.2.3.4	56732
10.1.1.11	54532	1.2.3.4	54532

Murdoch UNIVERSITY

NAT Operation (cont.)

- To the outside world, only the external (public) IP address and port numbers are visible
- Your local router will translate headers back to internal IP address and port number

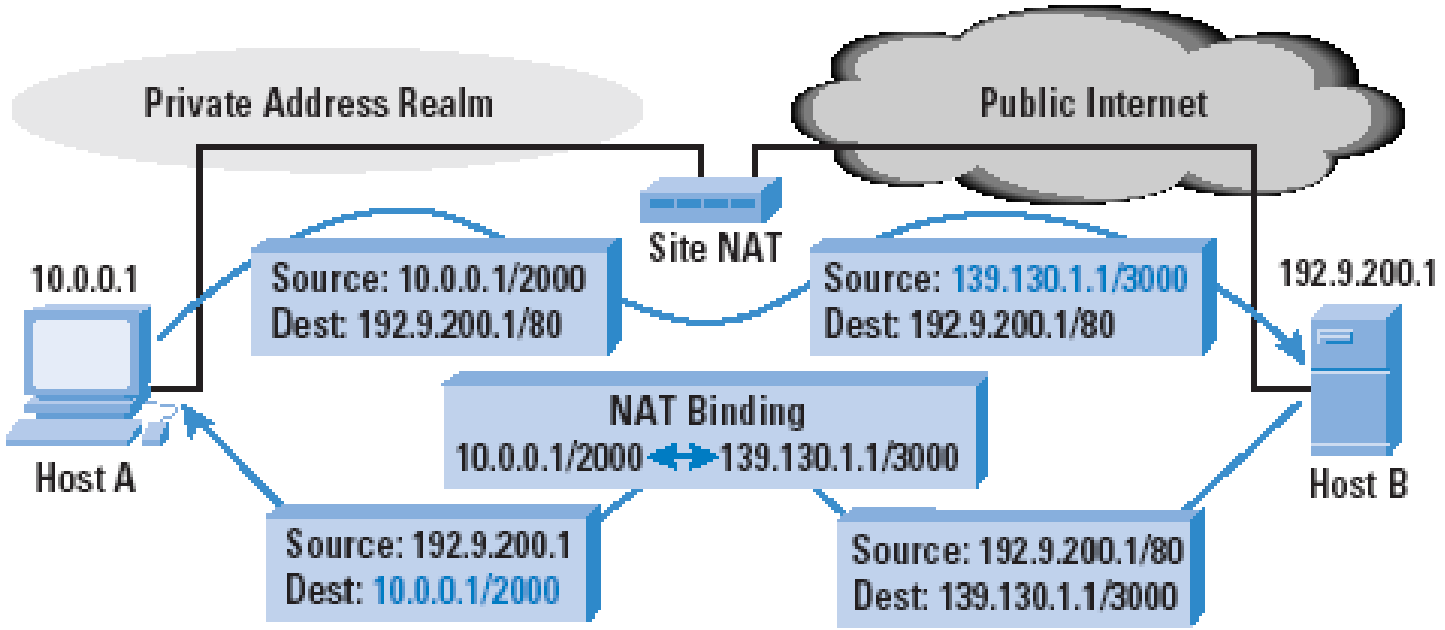

G. Huston, "Anatomy: A Look Inside Network Address Translators", The Internet Protocol Journal - Volume 7, Number 3, http://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-29/anatomy.html

Murdoch UNIVERSITY

Carrier-grade NAT (CGNAT)

- Due to IPv4 address exhaustion, NAT is also sometimes used at the Service Provider level
- The basic concept remains the same but the scale increases dramatically
- Usually additional to NAT on your local router

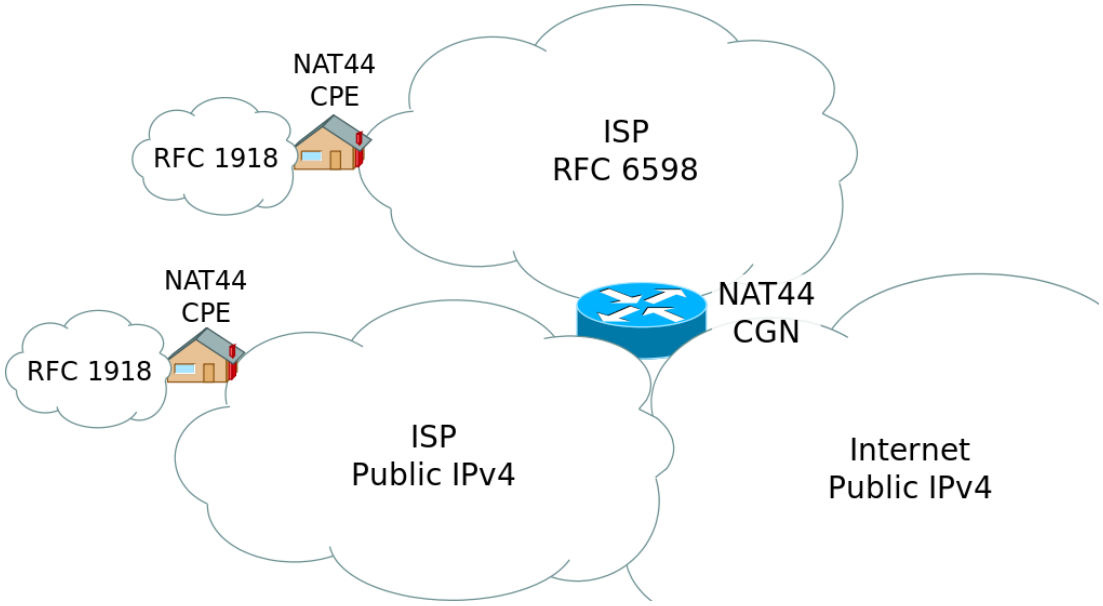

https://en.wikipedia.org/wiki/Carrier-grade_NAT#/media/File:CGN_IPv4.svg

Challenges with NAT – Application Connectivity

- NAT mappings are normally created on-demand, so communication can't be initiated by an external host
  - Leads NAT to sometimes be thought of as a firewall
  - Makes it difficult to run servers
  - Peer-to-peer communications might also be effected
- Can be addressed using static NAT (aka. port forwarding)
  - Create static NAT mapping, allowing communication to be initiated by external hosts

Murdoch UNIVERSITY

Challenges with NAT – Double NAT

- Because CGNAT is often additional to NAT on your local router it increases the complexity of the network
- This scenario is referred to as ‘Double NAT’
    - Also possible within local network (if you have two routers)
- Applications that rely on connections established from outside the local network may not work
    - Local port forwarding is no longer sufficient
    - ISP must also port forward (difficult to setup)

Murdoch UNIVERSITY

Challenges with NAT – Geolocation

- Identifying the location of individual hosts becomes challenging when multiple devices share a single IP address
- ISP will still be able to determine location, but a content provider (eg. Google, Yahoo) may not
- Interferes with content personalisation (eg. optimising search results based on location, location data for maps)
- Can still use user-reported location information though

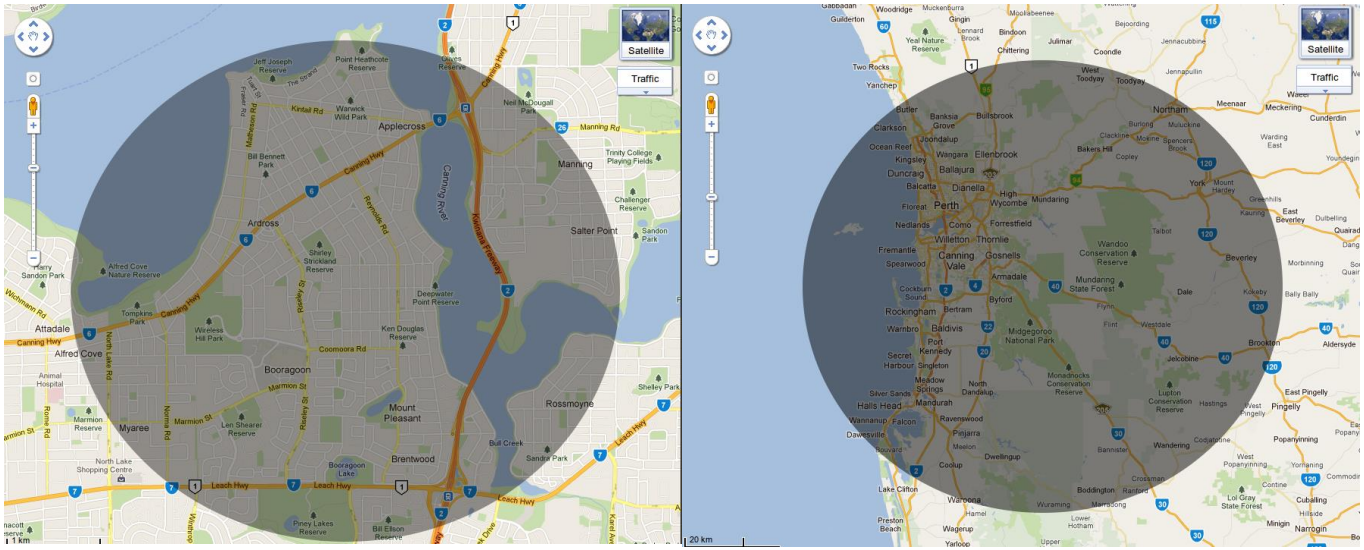

Challenges with NAT – User Identification

- NAT also makes it difficult to identify individual users on the network
- If a user behind NAT misbehaves (eg. cheating in a game, starts a cyberattack), investigation would point to all users behind the same NAT
- Blocking the IP address would ban all users

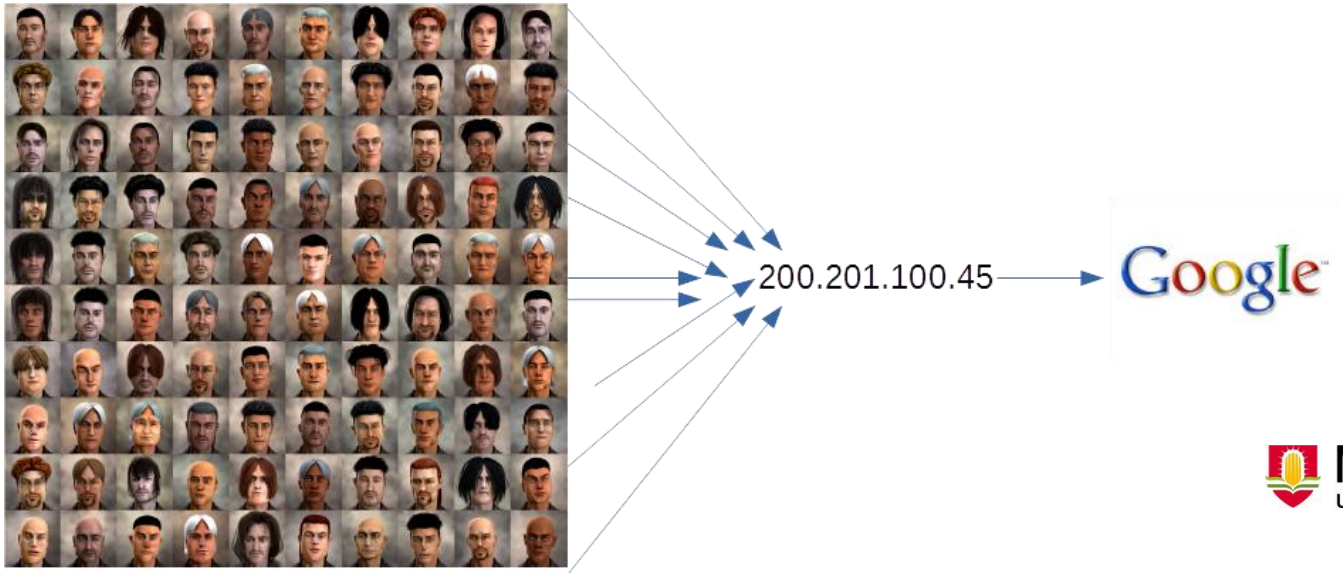

Murdoch UNIVERSITY

Upsize to Internet Protocol v6

- NAT was always intended to be an interim fix to IPv4 address exhaustion; IPv6 was designed to be the long-term solution
- IPv6 increases the length of IP addresses from 32 bits to 128 bits
    - Provides 2^{128} (3.4×10^{38}) addresses; said to be enough to provide an IPv6 address to every atom on Earth and have enough left over to service 50—100 additional Earths
- But IPv6 has struggled with adoption
    - Transitioning to IPv6 was seen as unnecessary and difficult (more on this later)

Internet Protocol version 6 (IPv6)

- Standardised in 1998 to supersede IPv4
- Increase address length to 128 bits represented as eight 16 bit hexadecimal numbers separated by colons (:)
    - Example: 2001:0db8:0031:0001:020a:95ff:fef5:246e
    - Each hexadecimal digit is 4 bits long
- Simplified packet headers
- Automatic address configuration (without requiring DHCP)
- Mobility; devices to be always reachable
- Integrated IP Security (IPSec)

Murdoch UNIVERSITY

IPv6 Addresses

- IPv6 addresses are divided into eight 16 bit hexadecimal numbers separated by colons
  - Example: 1080:0db8:0031:0000:0008:0800:200c:417a
- Subnet mask is now **always** in slash notation (eg. /64)
  - Usually referred to as the prefix length
  - Includes routing prefix and subnet ID
- Any part of the address not included in the prefix is the interface identifier

bits	48 (or more)	16 (or fewer)	64
field	*routing prefix*	*subnet id*	*interface identifier*

https://en.wikipedia.org/wiki/IPv6_address

Murdoch UNIVERSITY

Writing IPv6 Addresses

- 128 bit addresses are not especially convenient to write, so we have a few methods of shortening them
    - Start with the address: 1080:0db8:0031:0000:0008:0800:200c:417a
- First, leading zeroes can be omitted:
    - Shorten to: 1080:db8:31:0000:8:800:200c:417a
- Largest consecutive set of zeroes can also be substituted with '::'
    - Shorten again to: 1080:db8:31::8:800:200c:417a

Exercise: Condensing IPv6 Addresses

- Write the following IPv6 addresses in their most condensed form:
  - fe80:0000:0000:0000:0202:b3ff:fe1e:8329
  - aaaa:0000:0450:0000:0000:0000:acdc:0210
- Solutions:
  - fe80:~~0000:0000:0000~~:~~0~~202:b3ff:fe1e:8329
  - **fe80::202:b3ff:fe1e:8329**
  - aaaa:~~000~~0:~~0~~450:~~0000:0000:0000~~:acdc:~~0~~210
  - **aaaa:0:450::acdc:210**

Murdoch UNIVERSITY

Condensing IPv6 Addresses (cont.)

- Notice that in the second address, we could only use the :: once, even though it would be shorter if we could use it multiple times?
    - aaaa:0:450::acdc:210 would become **aaaa::450::acdc:210**
- There would be three 'valid' expanded versions of this address:
    - aaaa:0000:0450:0000:0000:0000:acdc:0210
    - aaaa:0000:0000:0450:0000:0000:acdc:0210
    - aaaa:0000:0000:0000:0450:0000:acdc:0210
- How could you disambiguate which one was being actually being used?

Murdoch UNIVERSITY

Exercise: Expanding IPv6 Addresses

Write the following IPv6 addresses in their fully expanded form:

- 2001:ba0::1234
- aaaa::

Solutions:

- **2001:0ba0:0000:0000:0000:0000:0000:1234**
- **aaaa:0000:0000:0000:0000:0000:0000:0000**

Murdoch UNIVERSITY

Break

When we return: More on IPv6 and transitioning from IPv4

Murdoch UNIVERSITY

Subnetting in IPv6

- As with IPv4, we can subdivide address blocks in IPv6
- ISPs usually allocate users and organisations /64 networks, so VLSM isn't necessary; subnets can be of equal size
- Simplifies process of subnetting, as well as network management
- You'll probably want a calculator for this..

http://www.fix6.net/shop/

Murdoch UNIVERSITY

Subnetting in IPv6 – An Example

- Subnet the 2aa1:ea0::/32 network into /48 subnets
- Write the network addresses
  - 2aa1:ea0:0::/48
  - 2aa1:ea0:1::/48
  - 2aa1:ea0:2::/48
  - 2aa1:ea0:3::/48
  - 2aa1:ea0:4::/48

http://www.fix6.net/shop/

Murdoch UNIVERSITY

Assigning IPv6 Addreses

- DHCP provided automated address configuration in IPv4
- IPv6 introduces an additional option for automated address assignment: **Stateless Address Autoconfiguration** (SLAAC)
    - DHCPv6 is still an option
- Routers periodically advertise network prefix, with hosts generating the remaining 64 bits

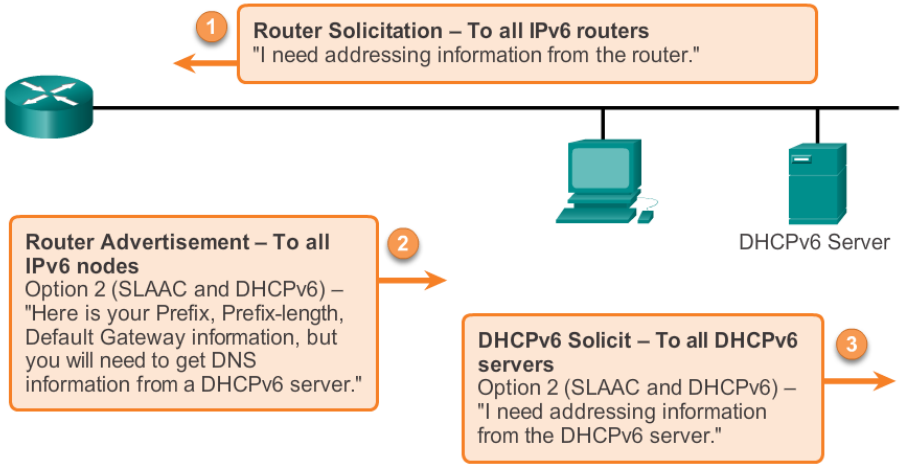

Murdoch UNIVERSITY

Assigning IPv6 Addresses (cont.)

- Host portion can be automatically generated based on MAC address
- MAC addresses are only 48 bits long, so FE:EE is usually inserted
- Host portion can also be randomly generated

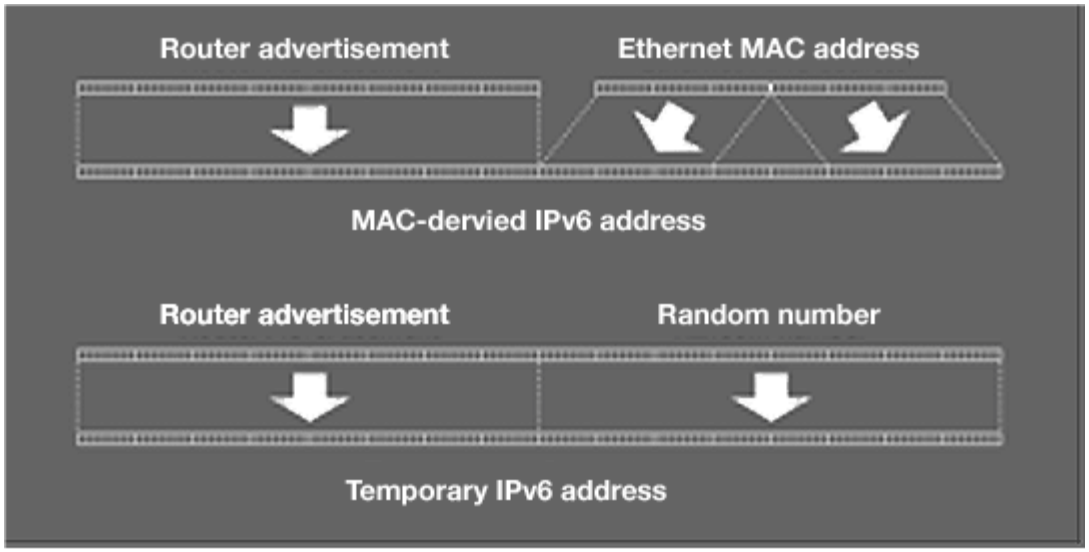

http://arstechnica.com/hardware/news/2007/03/IPv6.ars/2

Murdoch UNIVERSITY

DHCP in IPv6 (DHCPv6)

- SLAAC is not a complete replacement to DHCP
    - DHCP also provides information about DNS servers, and optionally other information (eg. NTP servers)
- Router advertisements can direct hosts to DHCPv6 server
- DNS servers can also be configured using router advertisements

Murdoch UNIVERSITY

IPv6 Headers

- IPv6 headers are only double the size of IPv4, despite having addresses 4x as long
- Some IPv4 header fields have been removed, others have been renamed

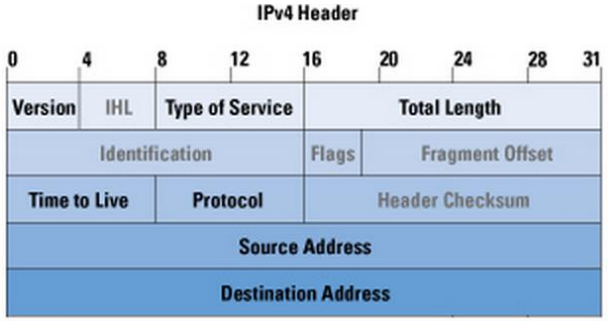

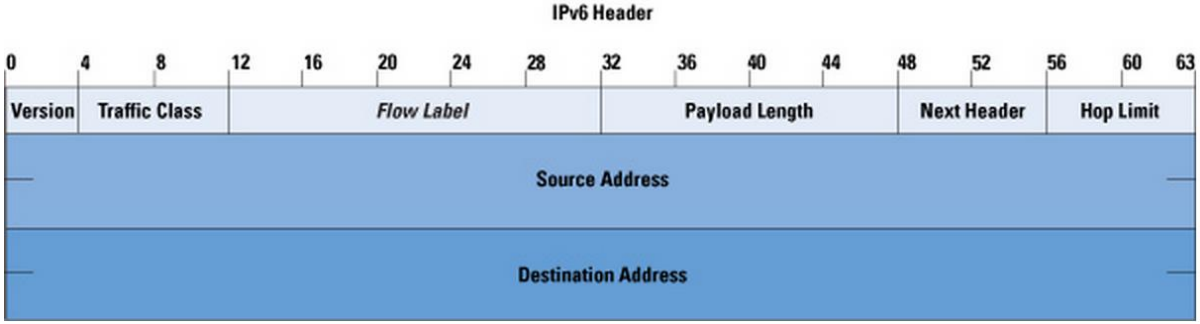

Murdoch UNIVERSITY

IPv6 Headers (cont.)

- IPv4 Time to Live → IPv6 Hop Count
- IPv4 Type of Service → IPv6 Traffic Class
- **Removed:** IPv4 Identification, Fragment Offset, Header Checksum
  - Fragmentation removed for performance and security
  - IPv6 probes for maximum packet size for the path
  - Lower and higher layer protocols already have checksums

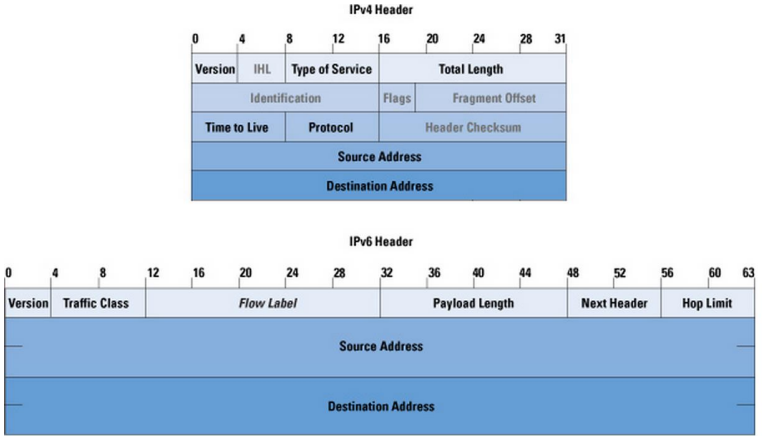

Murdoch UNIVERSITY

IPv6 Headers (cont.)

- Flow label field added to IPv6
  - Pseudo-random identifier used in with the source and destination addresses to uniquely identify a packet stream
  - Replaces use of source and destination ports, as well as protocol field for this purpose
  - Allows for non-standard Quality of Service (eg. all packets should be transmitted via the same path)

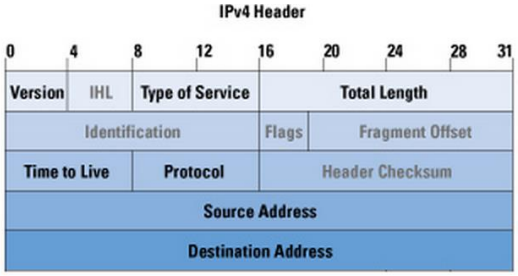

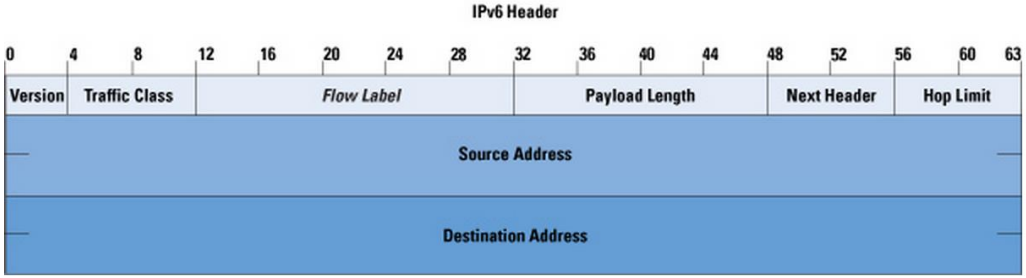

Murdoch UNIVERSITY

Types of IPv6 Transmissions

- Supports Unicasts and Multicasts from IPv4
  - Multicasts used in place of broadcasts (more efficient)
- Adds a new type of transmission: Anycast
  - One-to-one(-of-many) transmission
  - Facilitates load balancing and use of Content Delivery Networks (CDNs)
  - Any member of the anycast group can respond

Contrasting Multicast and Anycast Transmission

- Recall that in multicast transmission, a packet is delivered to all members of the multicast group
- Anycast alters this behaviour by instead sending a packet to **one member of the anycast group**
- Responding anycast group member should be closest to the sender, so anycast can also be considered as one-to-nearest

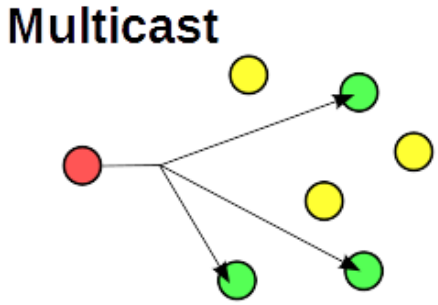

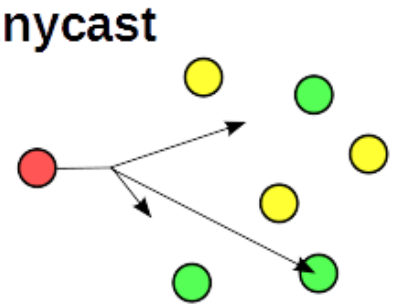

https://en.wikipedia.org/wiki/Routing

Murdoch UNIVERSITY

Neighbour Discovery

- Neighbour Discovery Protocol (NDP) replaces ARP in IPv6
- NDP is a Layer 3 protocol
- Uses multicast in place of broadcasts (which no longer exist in IPv6)

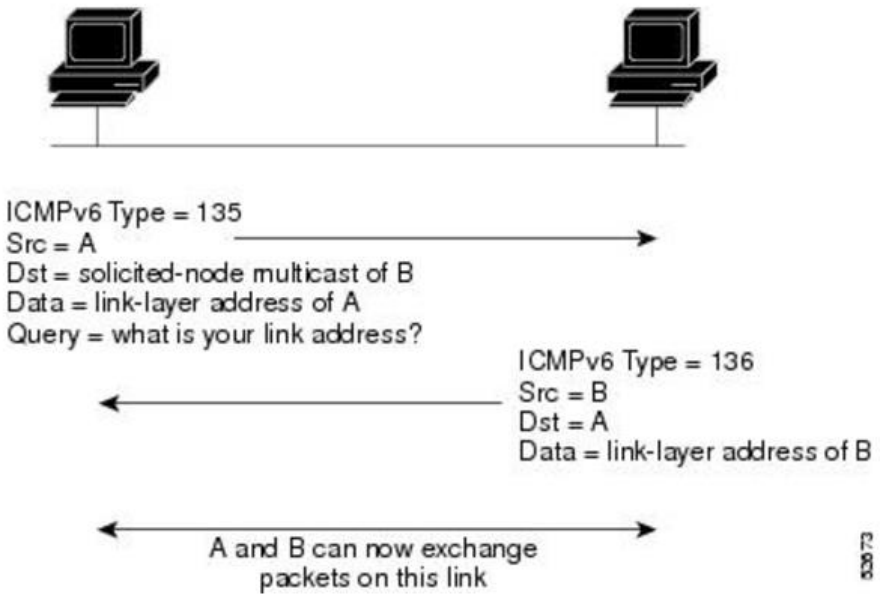

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-15-2mt-book/ip6-neighb-disc.html

The Road to IPv6

- So IPv6 sounds great on a technical level, so why are we still stuck with IPv4?
  - Lack of economic incentives for organisations
  - Lack of backwards compatibility with IPv4
  - Lack of perceived urgency (NAT and IPv4 markets exist)

http://blogs.technet.com/b/ipv6/archive/2007/11/13/windows-vista-earns-ipv6-ready-logo-phase-2.aspx

Murdoch UNIVERSITY

The Road to IPv6 (cont.)

- Slow transition hasn't been through lack of trying
- Government initiatives; big push in the early 2000s
  - Japan set a deadline to adopt IPv6 by 2005
  - European Union developed the eEurope Action Plan
  - India established a target for all ISPs to transition by 2006
  - China developed the China Next Generation Internet plan ahead of 2008 Beijing Olympics
- Industry-driven initiatives
  - World IPv6 Day and World IPv6 Launch Day
  - ISPs adopting IPv6

Murdoch UNIVERSITY

Transitioning to IPv6 is an Externality

- In economics, an externality is a cost or benefit that affects parties who didn't choose to incur it
- Without incentives, no clear reason exists to transition to IPv6 until IPv4 addresses are exhausted
    - Early adopters incur the highest costs with no benefits
    - Therefore, it's sensible to wait for others to transition first
    - Cheaper to adopt temporary solutions (like buying more IPv4 addresses) if needed
    - Organisations that can't get IPv4 addresses might not have a choice..

Murdoch UNIVERSITY

Lack of Backwards Compatibility

- Remember that the Internet currently runs on IPv4
- IPv6 lacks true backwards compatibility
    - All hosts and routers would need to transition simultaneously; not very likely
- Gradual transition is the only real option
    - Supported by dual stack, tunnelling and address translation

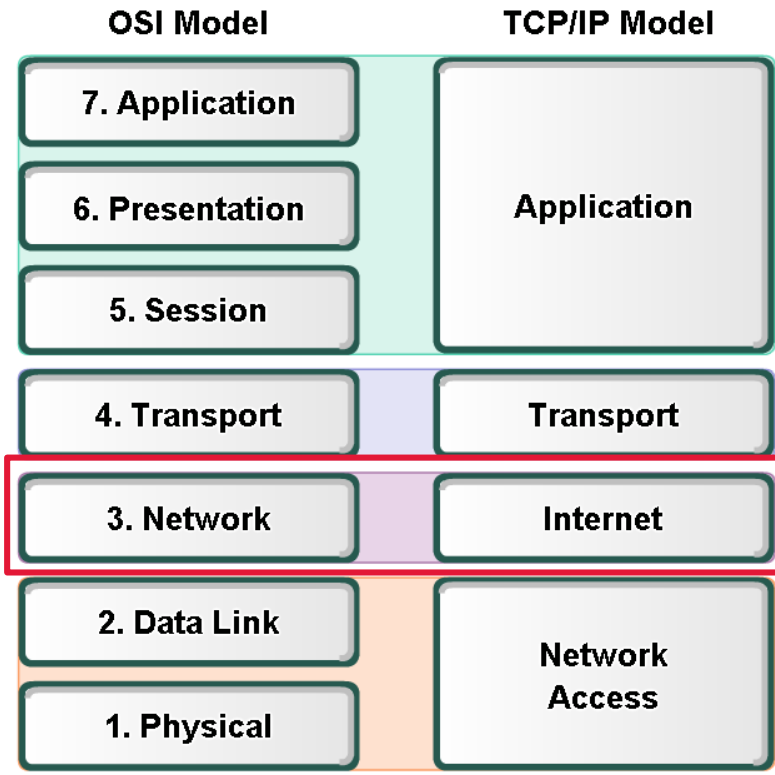

Murdoch UNIVERSITY

Dual Stack – Supporting IPv4 and IPv6 Concurrently

- IPv4 and IPv6 must coexist for some time to allow for a gradual transition
- Devices must be able to support both protocols, and have IPv4 and IPv6 addresses
    - This arrangement is known as **dual stack**
    - Commonly used operating systems support this arrangement
    - Also requires IPv6 support from ISP
- Protocol used will be determined by the host and the response of its DNS server
    - DNS server may return IPv4 or IPv6 addresses (or both)
    - Many systems will prefer IPv6 if available and configured

Murdoch UNIVERSITY

Tunnelling

- In cases where the ISP doesn't support a dual stack arrangement, IPv6 can still be enabled using tunnelling
- Encapsulate IPv6 packets inside of IPv4 ones
- Two types of tunnels available:
    - Static tunnelling requires that endpoint is manually specified
    - Automatic tunnelling identifies and connects to IPv6 capable relays

Murdoch UNIVERSITY

Tunnelling Technologies

- Many tunnelling mechanisms exist, but most now obsolete
- Some examples:
  - **6in4**
  - **6to4**
  - 6rd
  - Teredo
  - 6over4
  - 4in6
  - NAT64 / DNS64
  - TSP
  - TRT
  - SIIT
  - AYITA
  - DS-Lite

Murdoch UNIVERSITY

Tunnelling – 6in4

- IPv6 packets encapsulated in IPv4 packets
- IPv4 header is followed by IPv6 header
- IP protocol number set to 41
- Two tunnel endpoints must be configured statically
  - Host
  - Tunnel broker: Hurricane Electric, SixXS, Freenet6

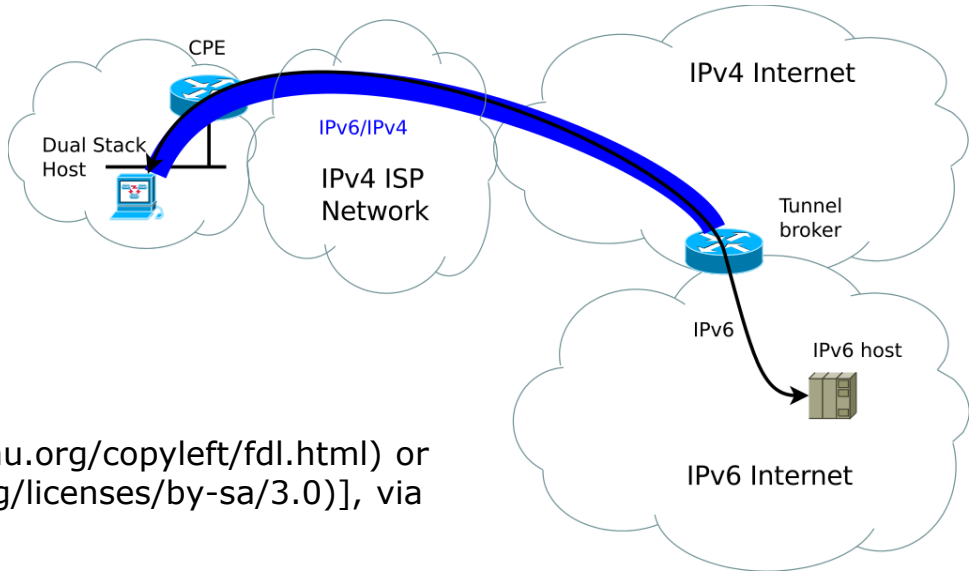

Tunnelling & Translation – 6to4

- Similar to 6in4: IPv6 packet encapsulated inside IPv4
- Address translation
  - Special IPv6 addresses map to IPv4 addresses, they start with 2002: and include host's IPv4 address
  - Hosts must have public IPv4 addresses
- Can derive IPv4 address from 6to4 address
  - Similar to NAT but translating between IPv4 and IPv6 packets
- 6to4 hosts can communicate among each other without network support
- Relay routers needed to communicate with native IPv6
  - Arriving 6to4 packets → Forward inner IPv6 packet
  - Arriving IPv6 packets → Construct outer IPv4 header

Tunnelling & Translation – 6to4 (cont.)

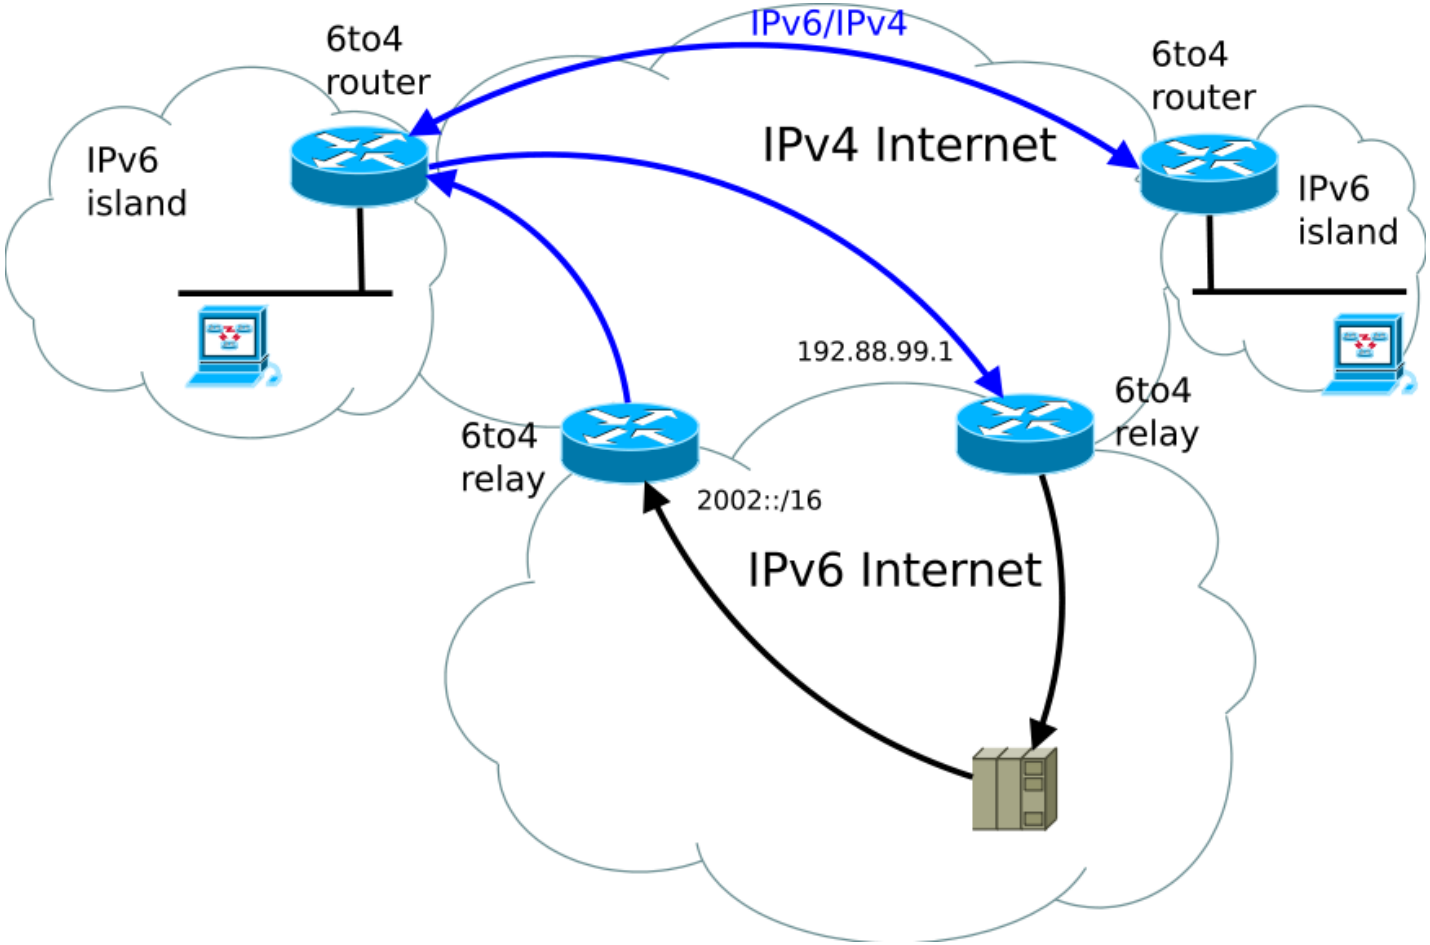

https://en.wikipedia.org/wiki/6to4

Current IPv6 Adoption

- Google monitors the amount of IPv6 traffic it receives
- Clear upward trend since 2012; now up to 21.5%
- Still a long way to go..

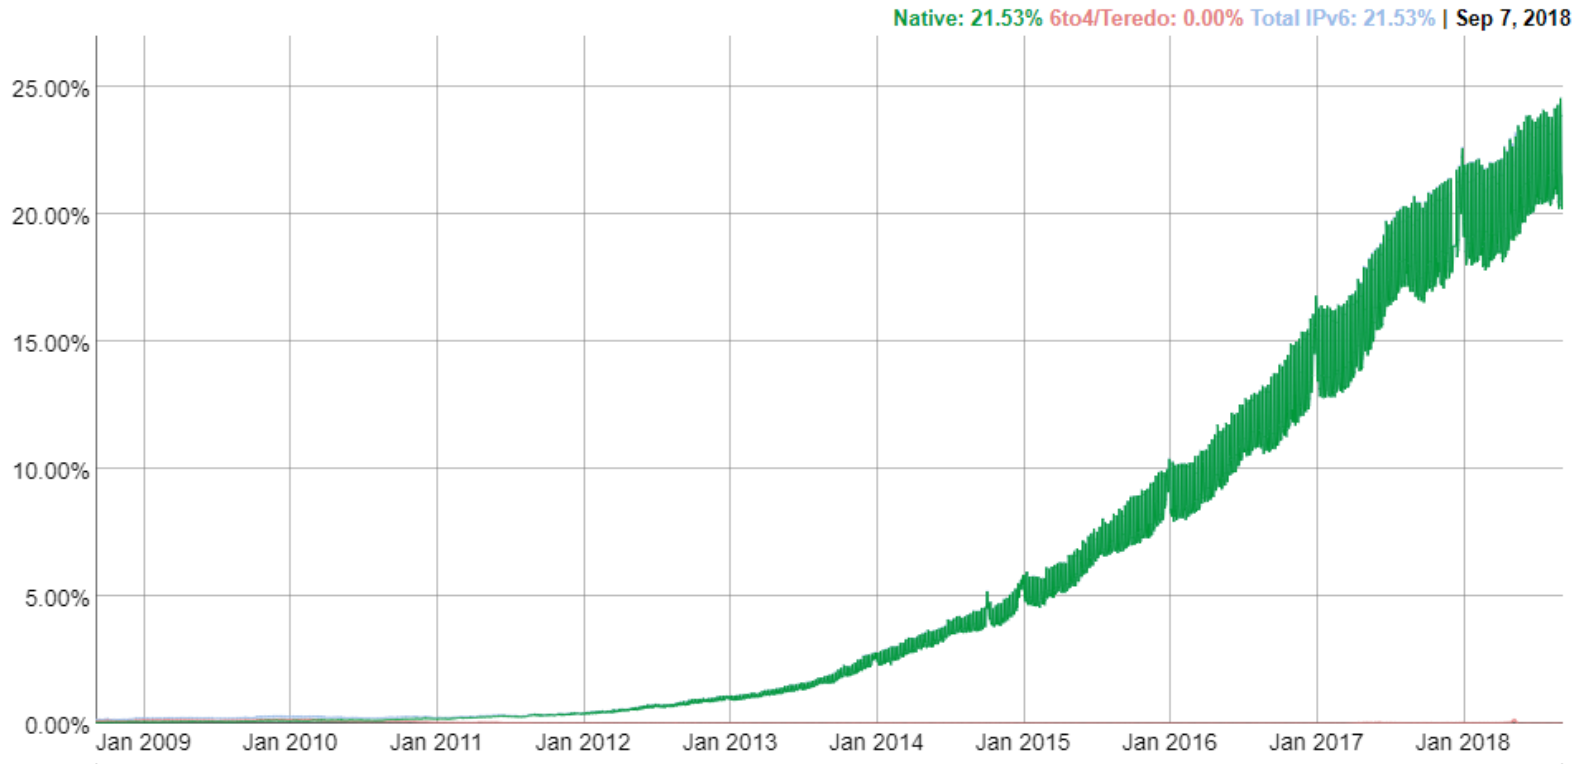

https://www.google.com/intl/en/ipv6/statistics.html#tab=ipv6-adoption

Lecture Objectives

You should now be able to:

- Describe IP version 4 address exhaustion
- Discuss the causes and consequences of IP version 4 address exhaustion
- Describe the purpose of Network Address Translation
- Describe the operation of Network Address Translation
- Identify challenges posed by the use of Network Address Translation
- Differentiate between IP version 4 and version 6
- Identify an IPv6 address
- Describe components of an IP version 6 address
- Describe IPv6 anycast transmission
- Identify approaches to transitioning from IP version 4 to IP version 6

Murdoch UNIVERSITY

Lecture Summary and the Week Ahead

- IPv4 addresses have been quickly depleted as more people connect to the Internet
- NAT was to be a temporary solution to IP address exhaustion, but persisted for much longer than expected
- IPv6 is the ultimate solution to the problem, but adoption has taken time (due to lack of backwards compatibility)
- The readings for this week are Introduction to Networks – Chapters 7 and 8, and Routing and Switching Essentials – Chapter 11
- Mid-Semester Test this week – check which lab you're enrolled in!

Murdoch UNIVERSITY

Next Week

- We'll look at routing within an organisation's network
- Static and dynamic routing
    - Router Information Protocol
    - Open Shortest Path First

Murdoch UNIVERSITY