Introduction to Service Provider networks

(Everything you wanted to know about networks, but were afraid to ask…)

Rob Lister

training@lonap.net

7 July 2025

netUK, London

Housekeeping

- Toilets
- Fire alarm
- …

LONAP

Who Am I?

- Robert Lister

  LONAP

- rob@lonap.net / training@lonap.net

LONAP

Who Are You? Who is this workshop for?

- … A student / academic?
- You are already in the industry – not in a “technical” role?
- You have some technical experience already?
- Doing something completely different?
- Just starting… changing careers?
- In this workshop because you thought it was the BNG Blaster Workshop?

- Even if you're not working in networks, a good understanding is a good foundation…

LONAP

“…because as we know, there are known knowns; there are things *we know we know*. We also *know* there are known unknowns; that is to say, *we know there are some things we do not know*.
But there are also unknown unknowns—the ones *we don't know we don't know*…”

- Donald Rumsfeld

“Rumsfeld Matrix”

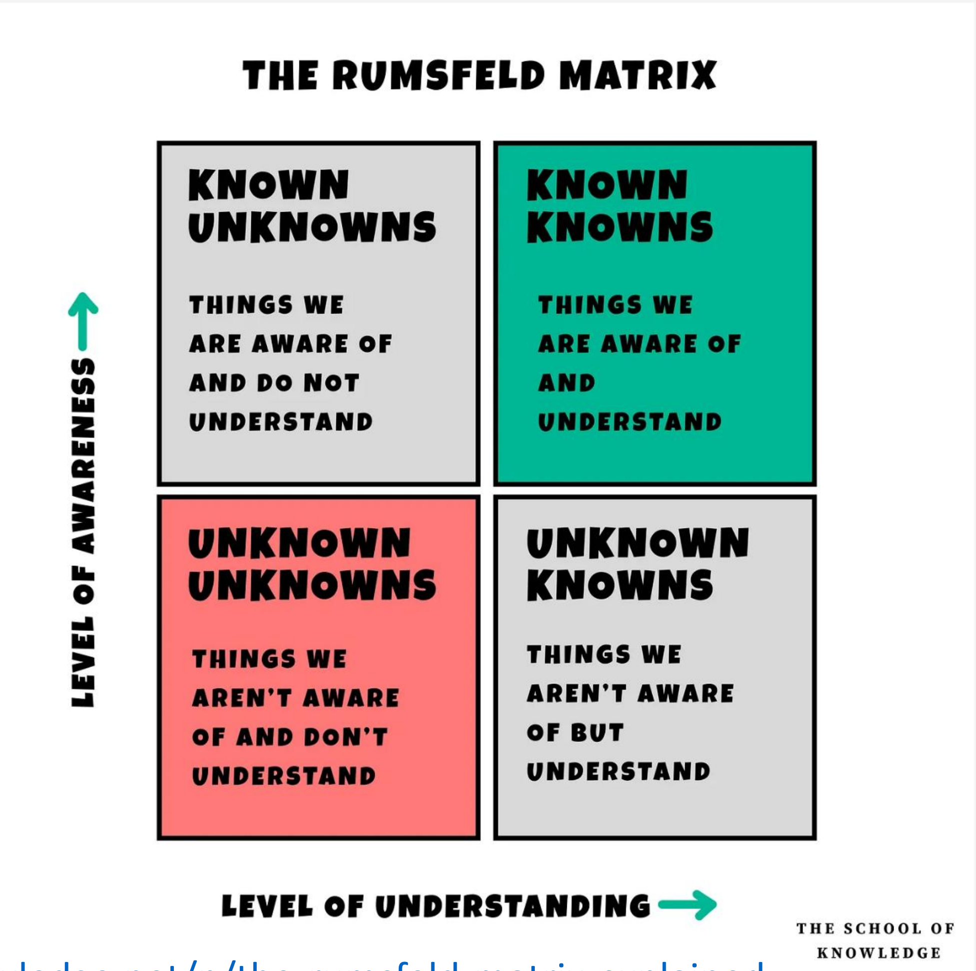

- Source: https://www.theschoolofknowledge.net/p/the-rumsfeld-matrix-explained

LONAP

Peering

PEERING:
What do you understand it to mean?

Discuss peering…

Networks

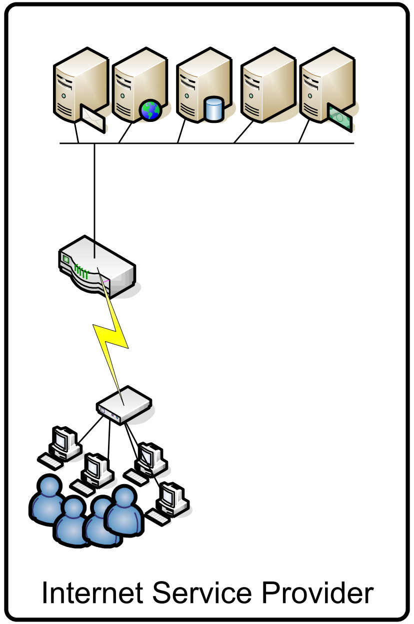

LONAP

Networks

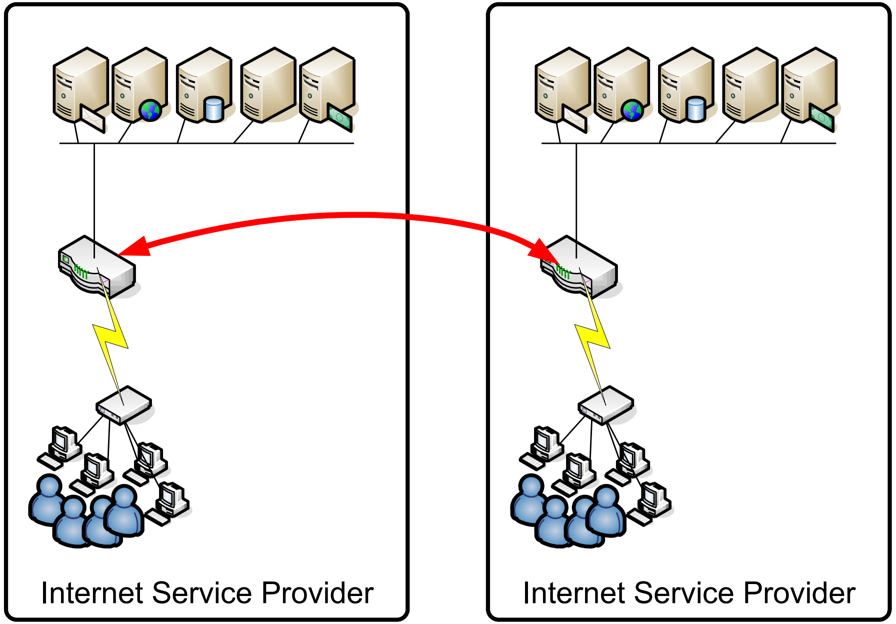

LONAP

INTERconnected NETworks

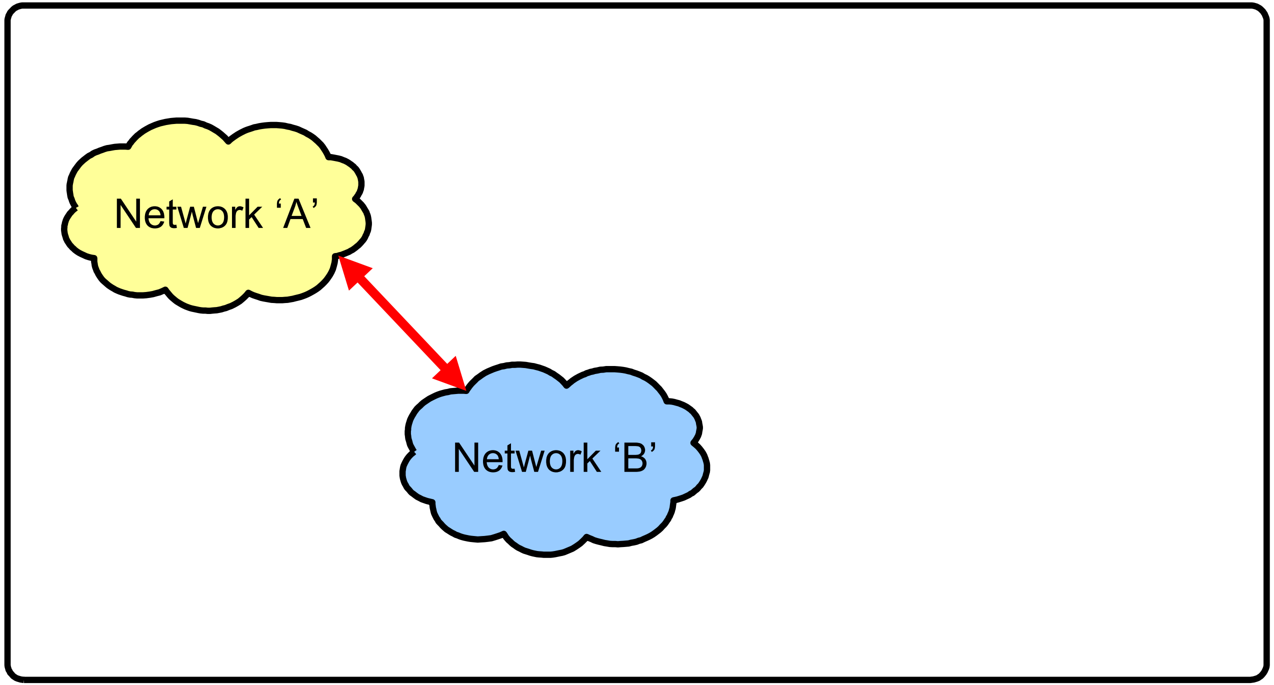

LONAP

Interconnected Networks

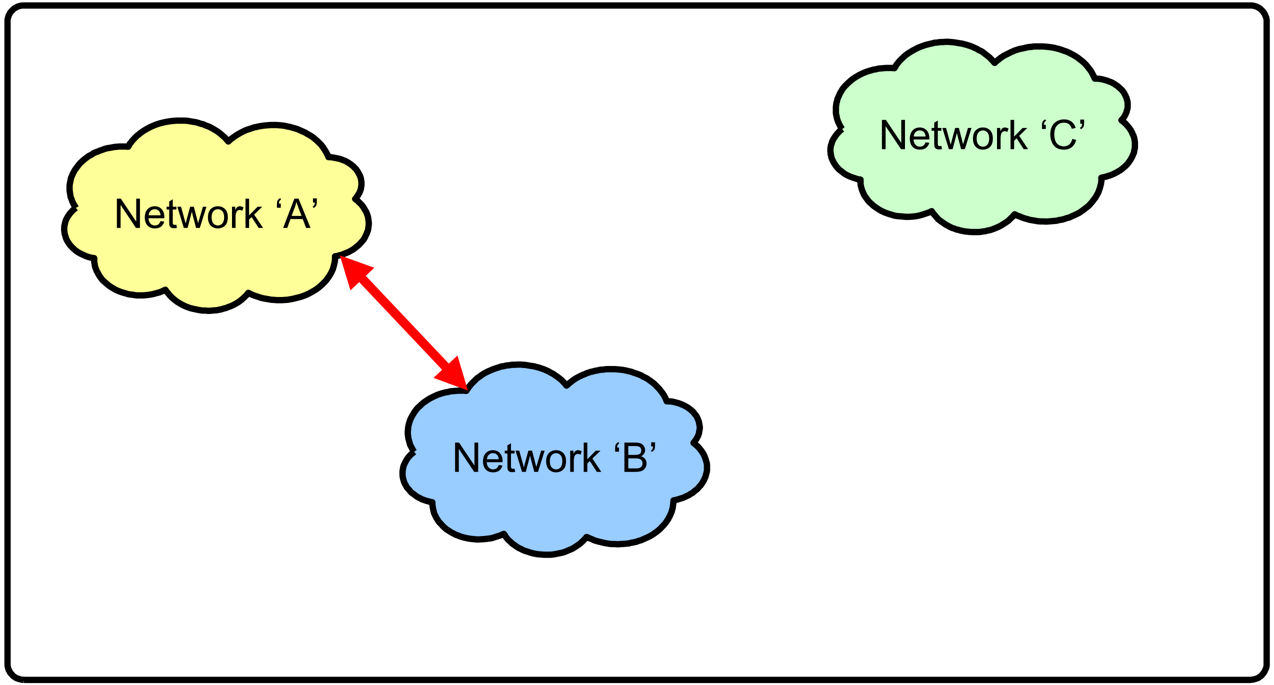

LONAP

Interconnected Networks

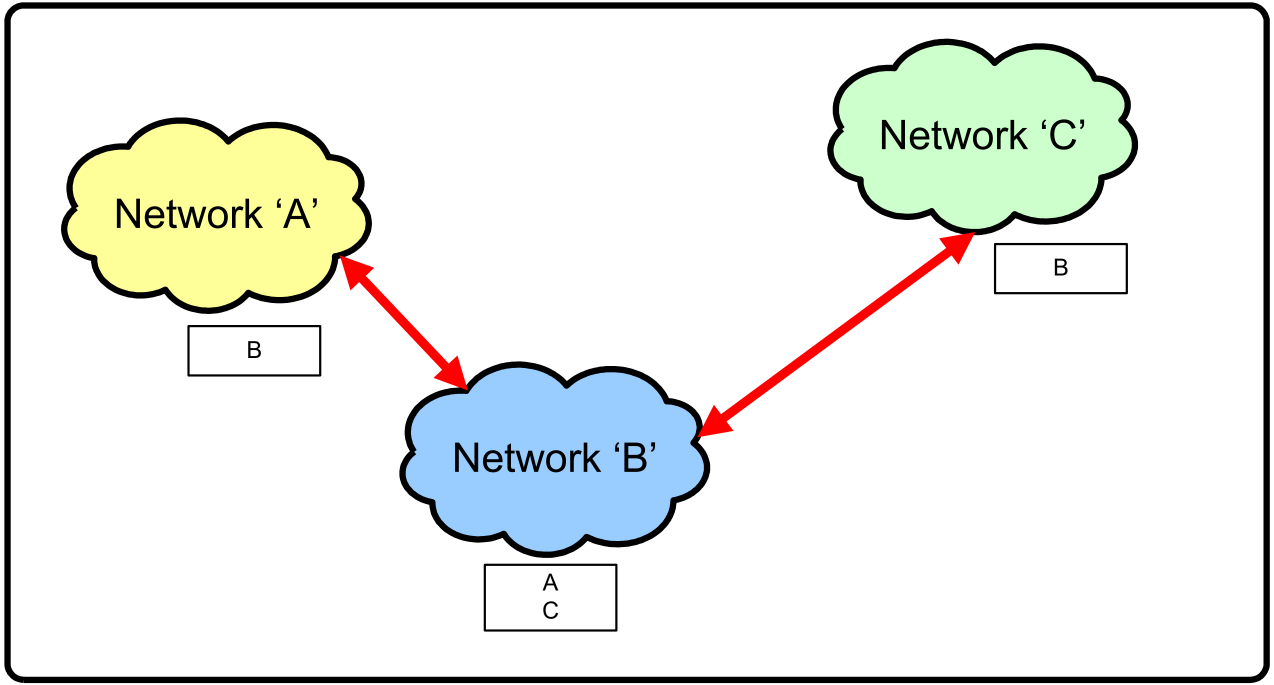

LONAP

Peers and transit

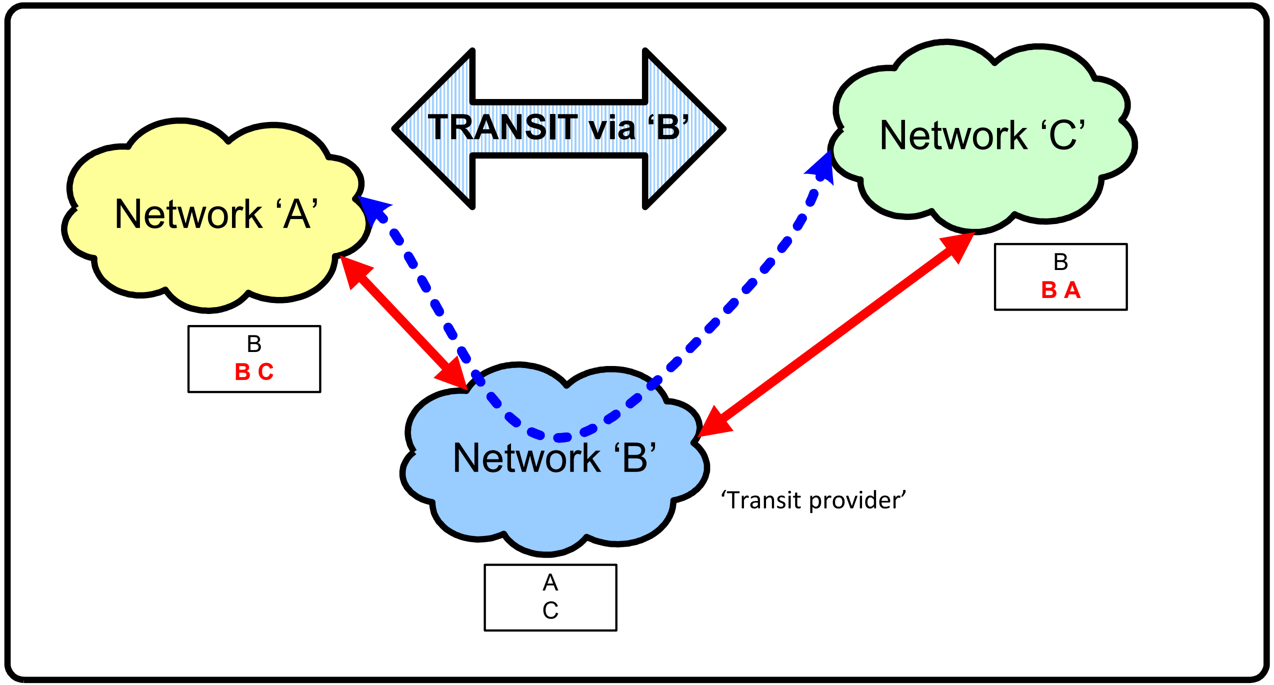

LONAP

1.3.1 Relationships between ISPs: peering and transit (p.18)

Peers and transit

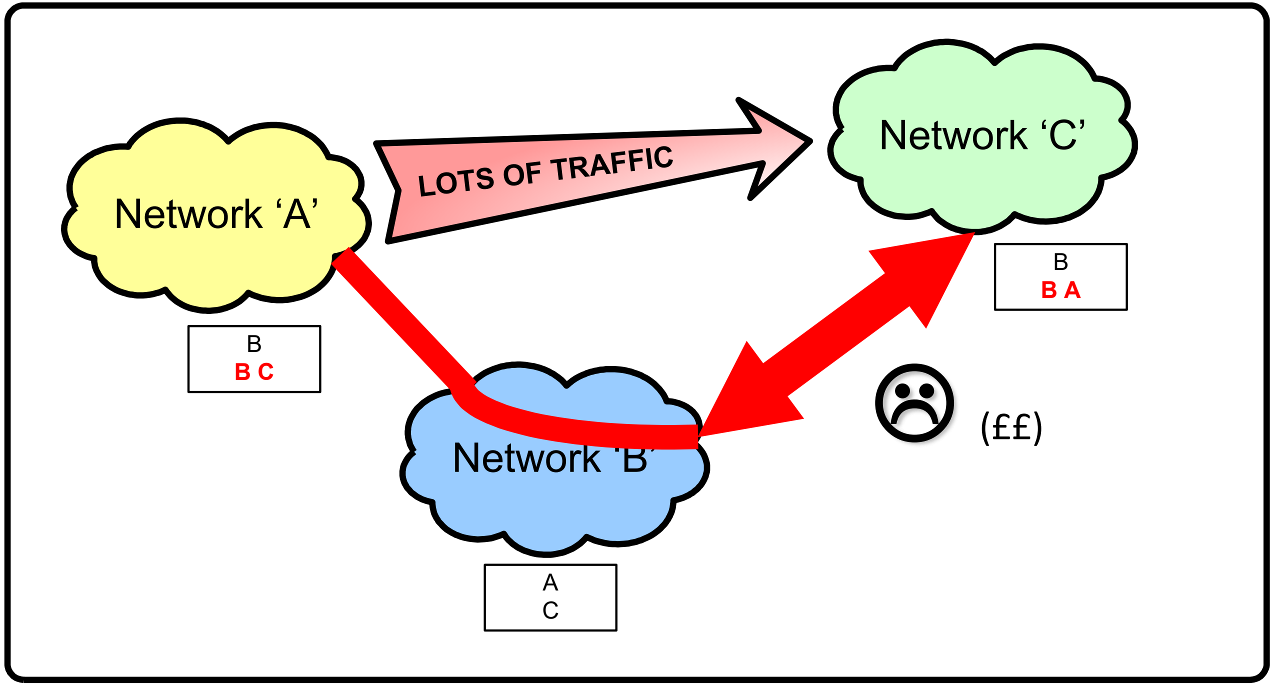

LONAP

Direct Peering

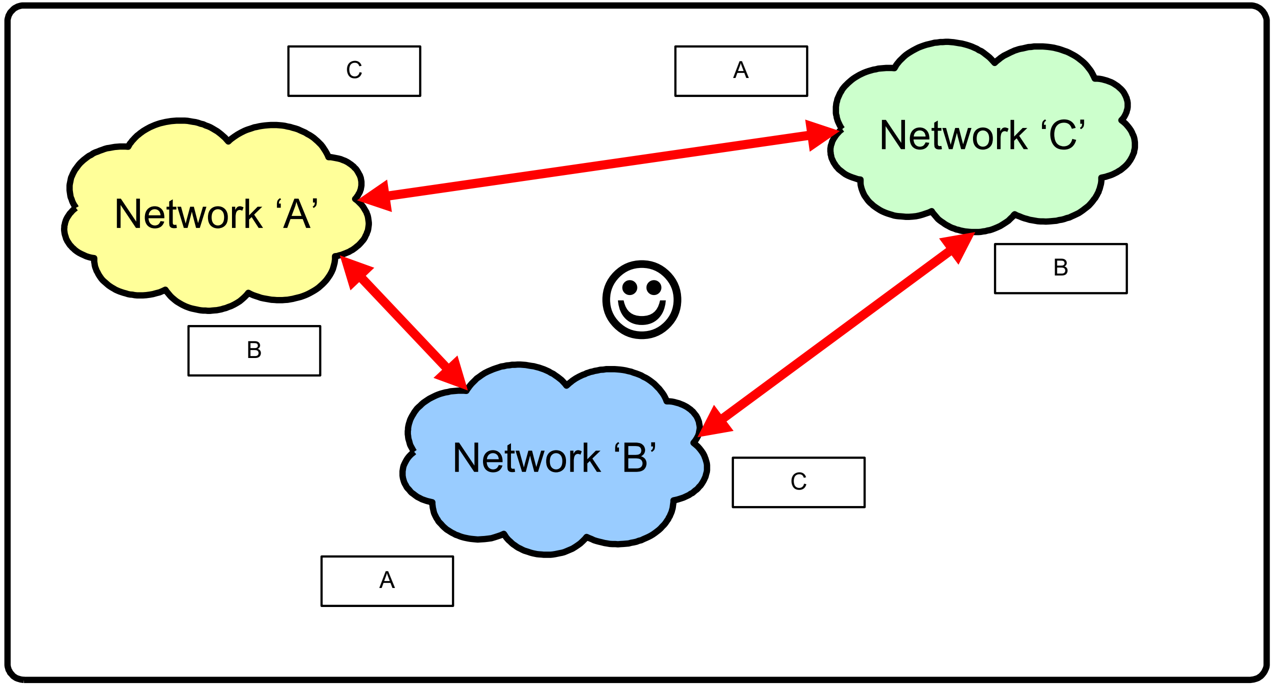

LONAP

The new network…

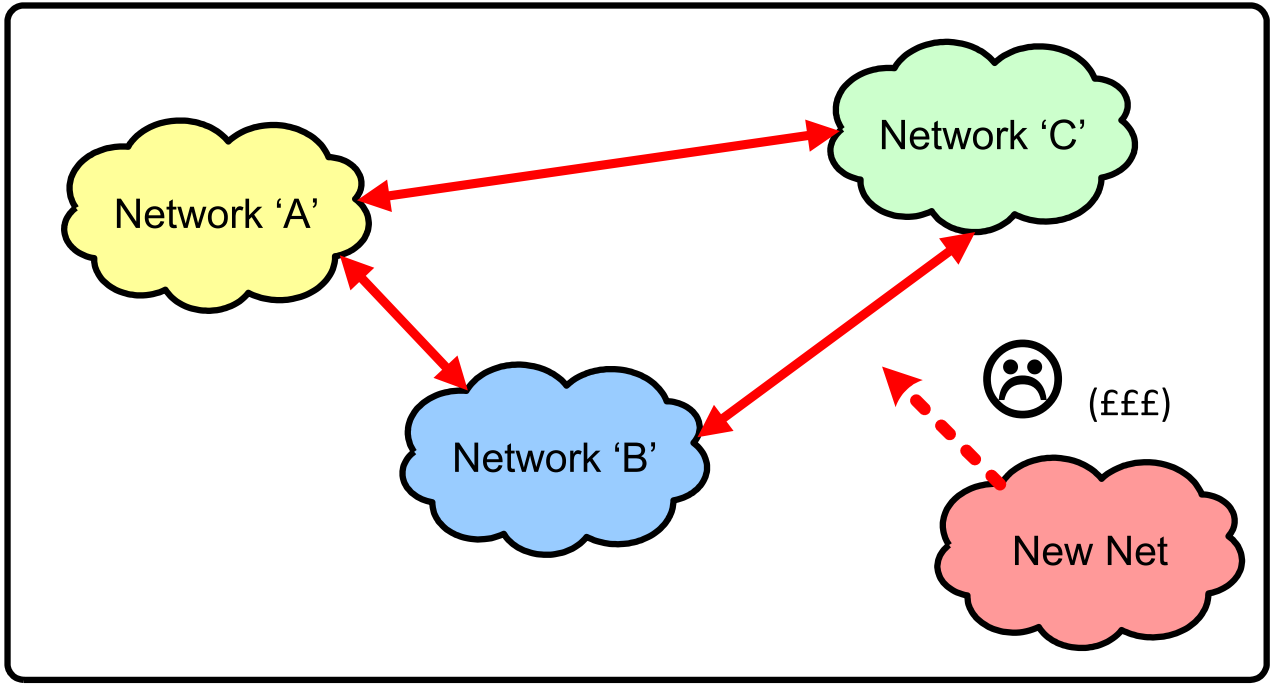

LONAP

Transit only: new net pays C

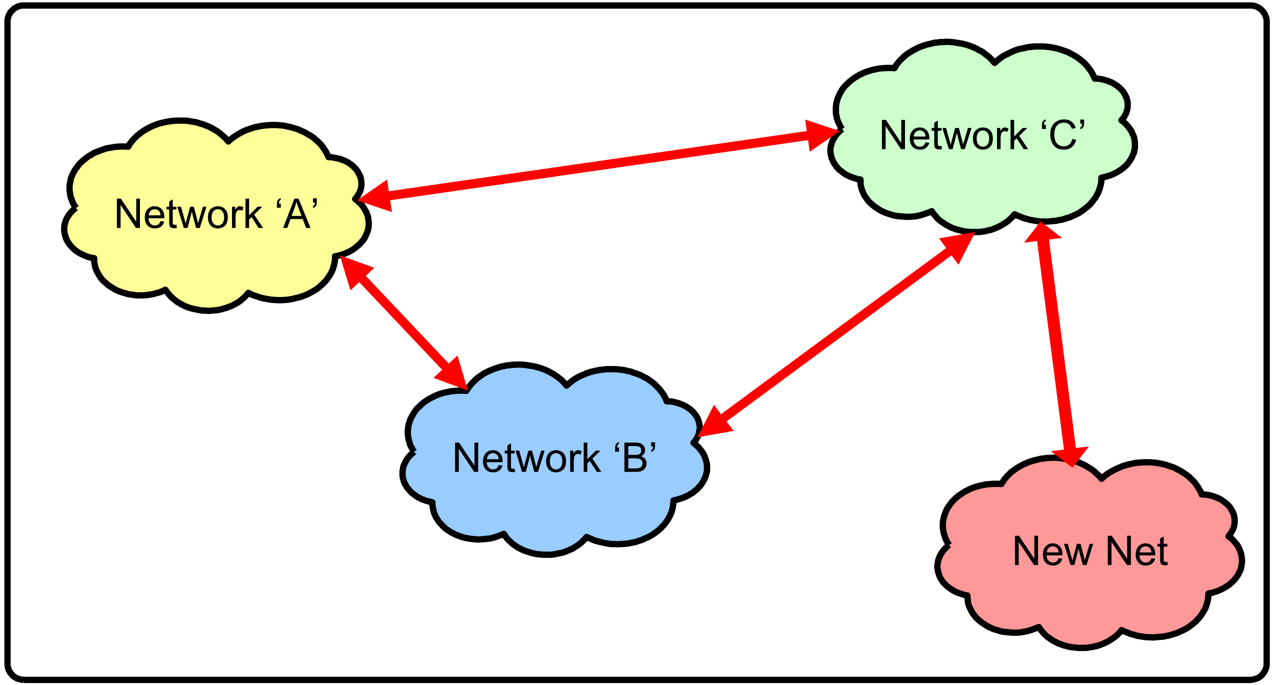

LONAP

Transit vs peering cost

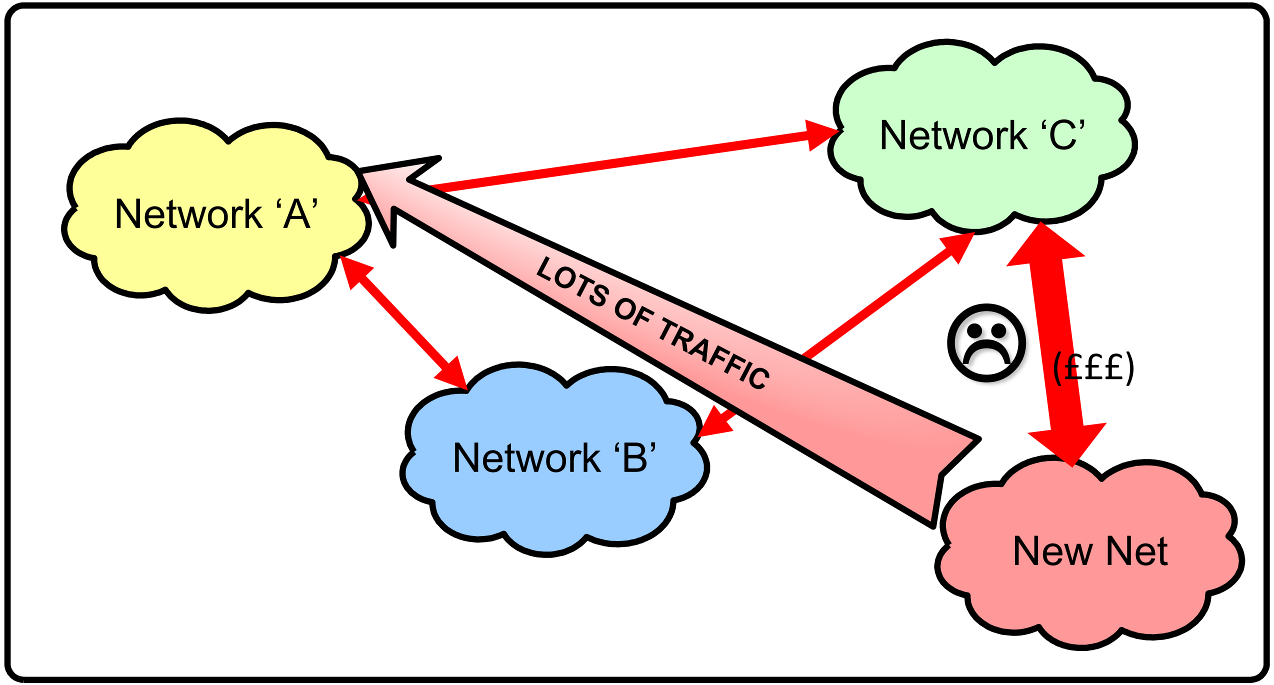

Direct peering can save money

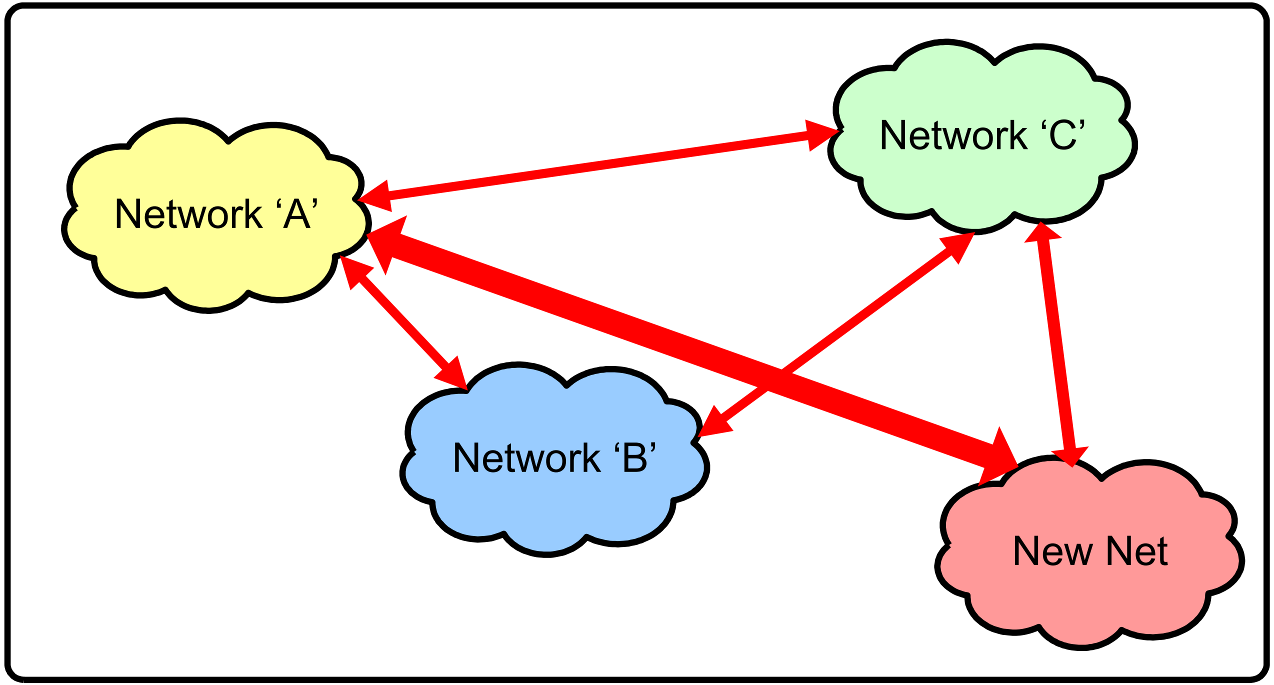

LONAP

Peering mesh

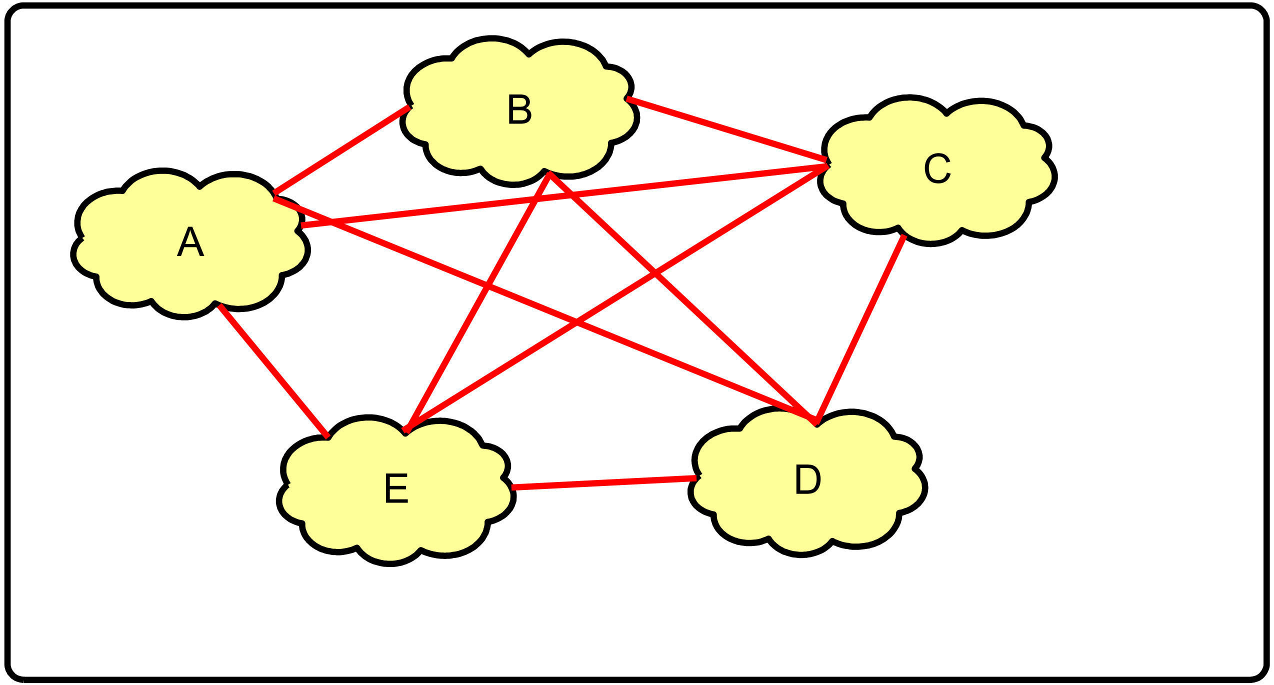

LONAP

Peering mesh

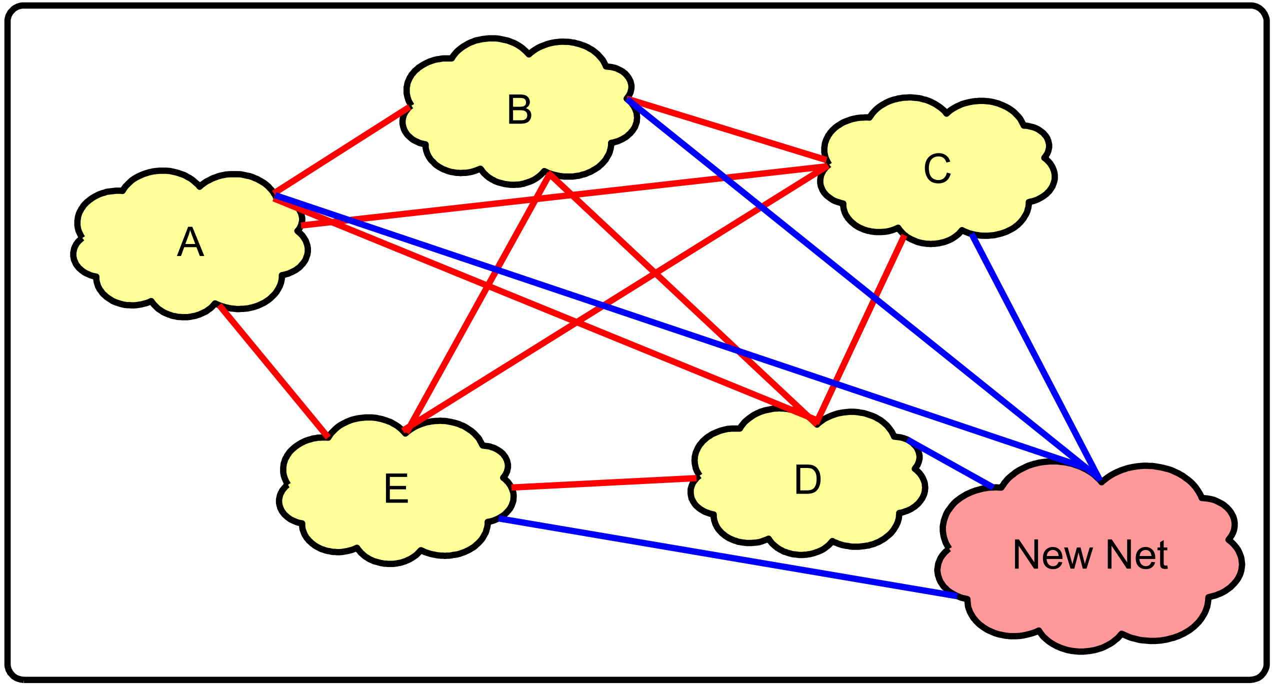

LONAP

Internet Exchange

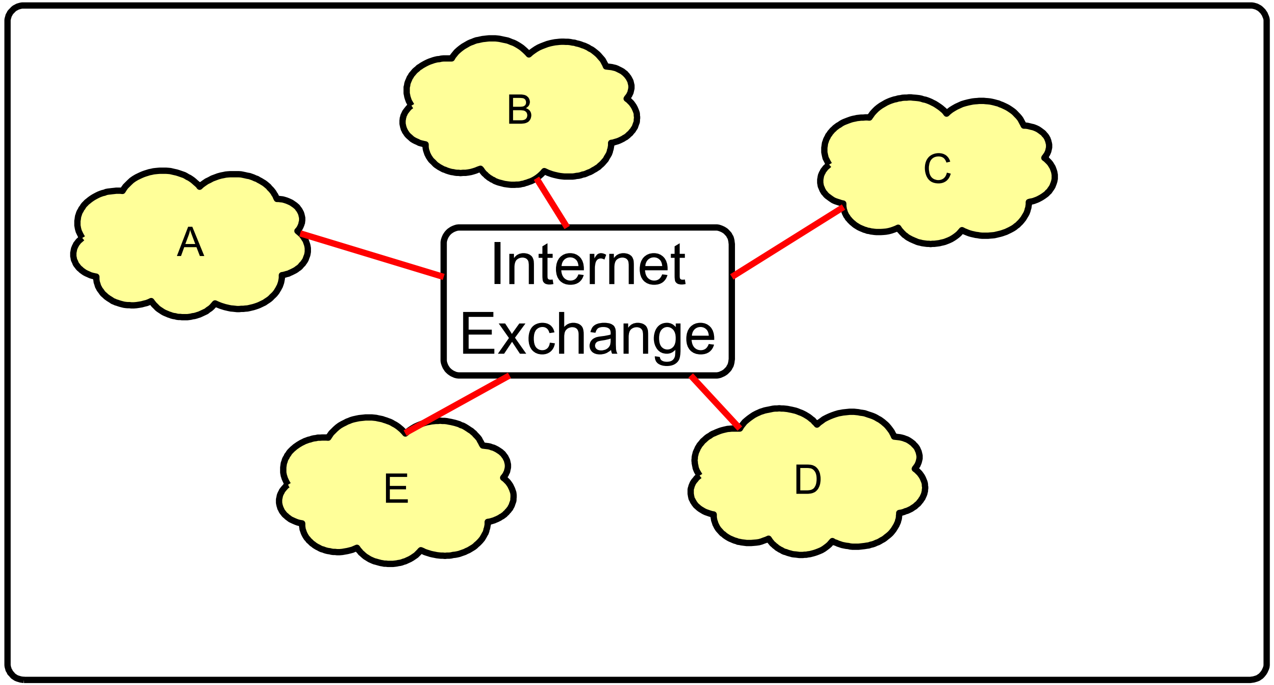

1.3.2 Internet Exchange Point (p.19)

LONAP

Bi-lateral peering

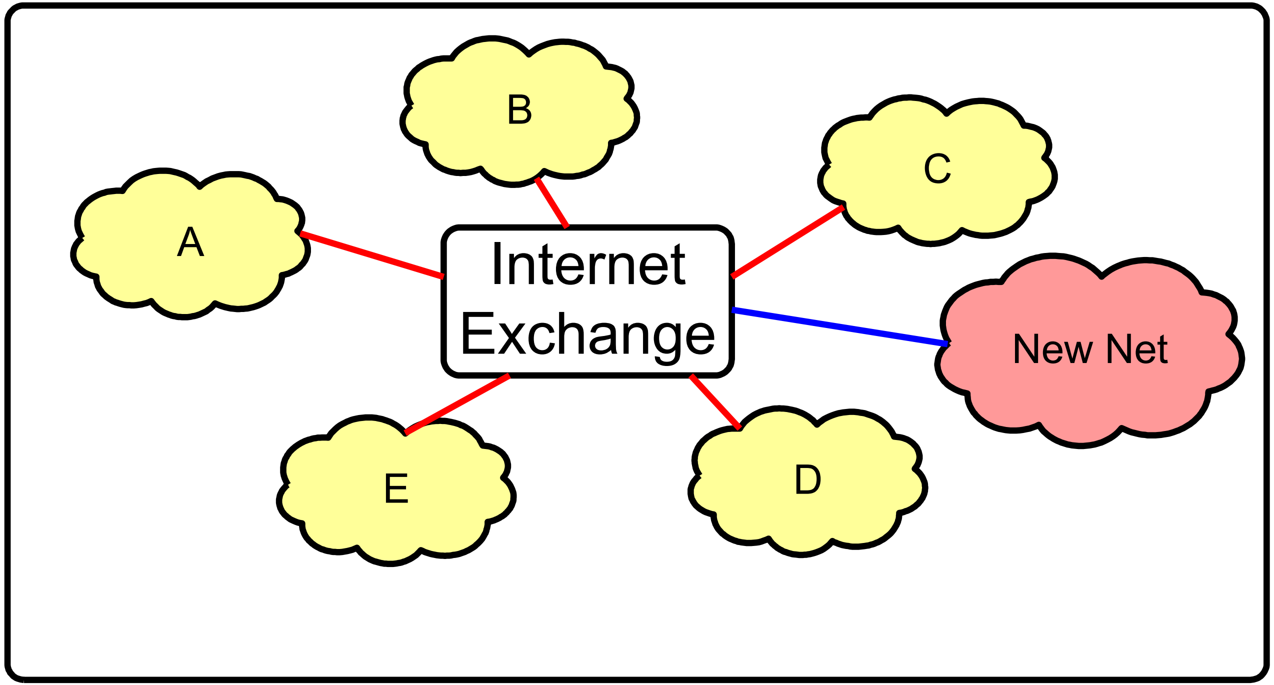

LONAP

‘Tiers’ of networks

- Some networks are larger than others
- Peering is not universal
- A mixture of transit and peering market
- Varied peering policies
- Transit can cost more than peering…
- …although this is not always the case (depends on traffic levels ‘price per Mbit’)
- Content networks as well as ISPs

‘Tiers’ of networks

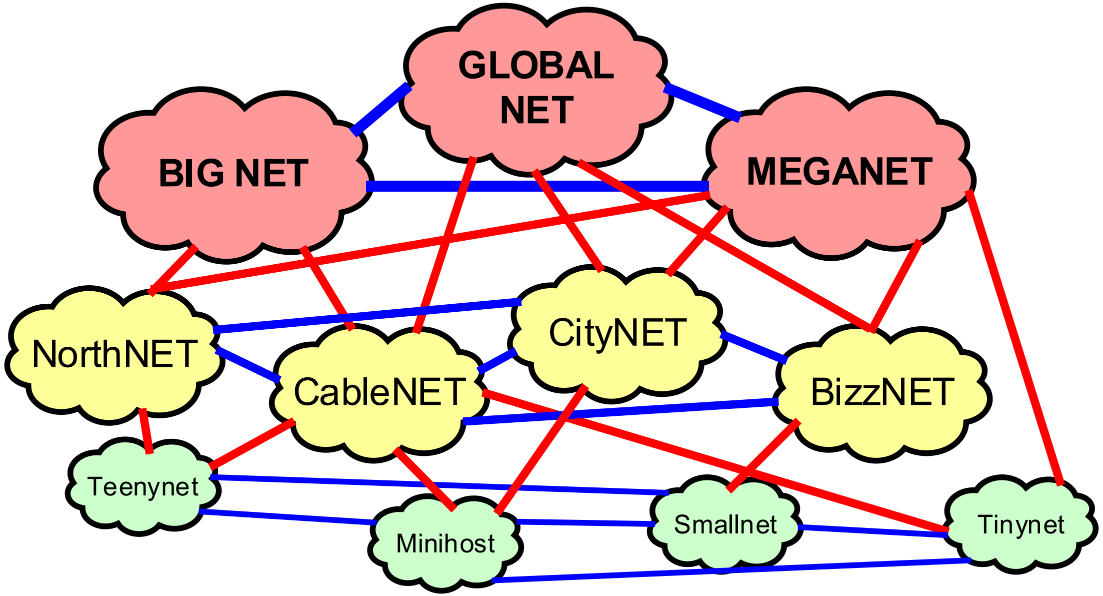

LONAP

1.3.3 ISP Classification (p.20)

‘Tiers’ of networks

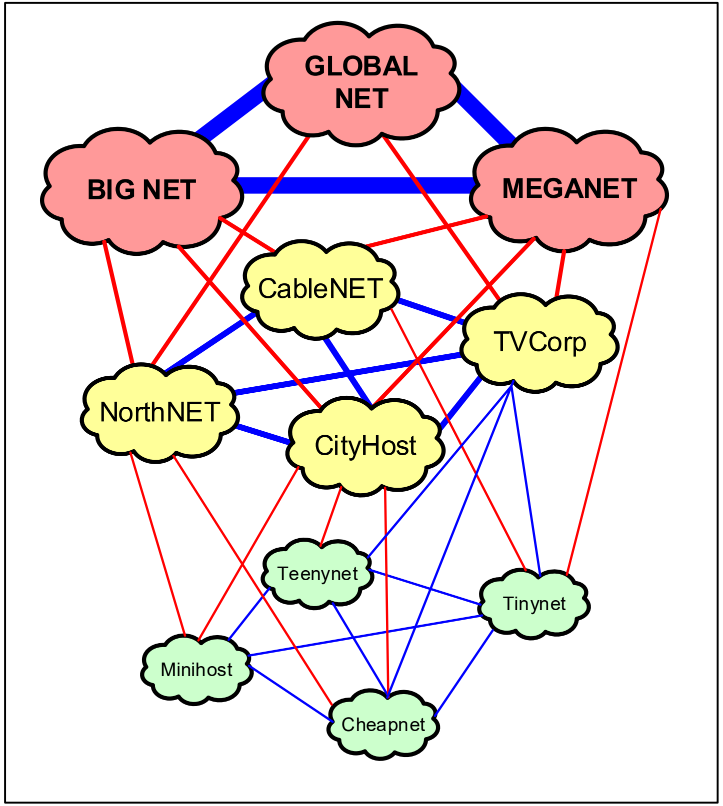

- Big networks won’t usually peer with their customers.
- Different peering policies
- What is TV Corp’s peering policy?
- Broad hierarchy but not strict.

LONAP

Peering via an IXP

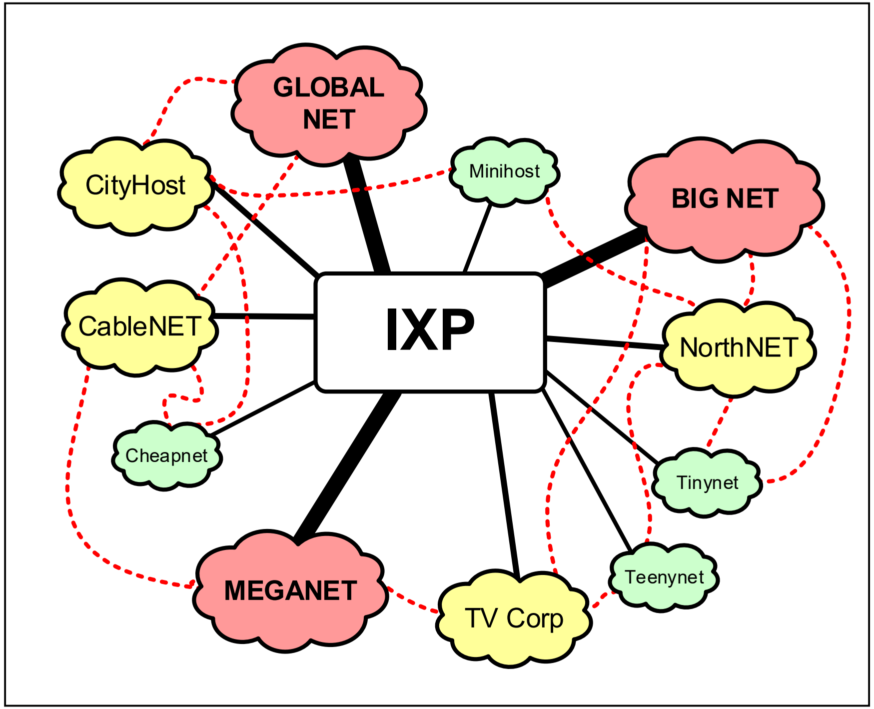

- It’s all the same to the IXP!
- Any network could peer with any other via the IXP.
- Who peers with who is controlled by member networks. (Not IXP)
- They don’t all peer with each other.

LONAP

IEEE

- IEEE (Institute of Electrical and Electronics Engineers)
- IEEE 802 Family of Working Groups
- Defines the standards for Ethernet
  - IEEE 802.3 – Ethernet
  - IEEE 802.11 – Wireless LANS
  - IEEE 802.3ad – Link aggregation (Now IEEE 802.1AX)
  - IEEE 802.3ba – 40/100Gb
  - IEEE 802.1D – LAN Switches

IETF

- The Internet Engineering Task Force
- Oversees Internet standards development
  - Internet Drafts/Standards
  - Request for Comment (RFCs)
    - RFC 2822/2821 – E-Mail (SMTP)
    - RFC 2126 – Hypertext Transfer Protocol (HTTP)
    - RFC 1771 / RFC 4271 – Border Gateway Protocol (BGP)
    - RFC 3261 – Session Initiation Protocol (SIP)
  - Best Current Practice (BCP)
    - BCP 38 - Network Ingress Filtering/Anti-Spoofing
    - BCP 214 - Mitigating the Negative Impact of Maintenance through BGP Session Culling
- http://www.ietf.org/

LONAP

7 Layer OSI Model

7	?	
6	?	
5	?	
4	?	
3	?	
2	?	
1	?	

LONAP

7 Layer OSI Model

7	Application	
6	Presentation	
5	Session	
4	Transport	
3	Network	
2	Data Link	
1	Physical	

LONAP

7 Layer OSI Model

7	**Application**	Web browser, DNS, E-Mail, Database software...	
6	**Presentation**	HTML, GIF, TIFF, JPEG, ASCII, MPEG, Encryption...	Data
5	**Session**	RPC, Named Pipes, NETBIOS	
4	**Transport**	TCP, UDP	Segments
3	**Network**	Path determination and logical addressing (IP)	Packets
2	**Data Link**	Physical addressing (Ethernet)	Frames
1	**Physical**	Media (cable) (fibre/copper) signal and binary transmission	Symbol Stream (00010110)

LONAP

Model Comparison

	OSI Model	IP Model	IP Protocol Suite
7	Application		HTTP, SMTP, SSH, FTP, DNS, RTP, SNMP
6	Presentation	Application Layer	
5	Session		
4	Transport	Transport Layer	TCP, UDP
3	Network	Internet Layer	IP, ARP, IGMP, ICMP
2	Data Link	Network Access Layer	Ethernet, Token Ring, ATM, Frame Relay
1	Physical		

1980 – 1984 ISO 7498

1978 - 1989 https://datatracker.ietf.org/doc/html/rfc1122

LONAP

Layers

- Different stages of transmitting and receiving data
- For example, sending an e-mail from e-mail application, across a network and receiving at the other end.
- Allows for different hardware and software components to be developed separately.
- Makes new development and compatibility possible without having to change everything.
- Open Systems Interconnection 7 layer model (OSI) is a **conceptual** model.
- Not everything fits “neatly” into the 7 layers
- … because IP was in use before OSI Model

LONAP

7 Layer OSI Model

- Why bother thinking about Layers?
- What’s the point?

- … discuss!

LONAP

Encapsulation

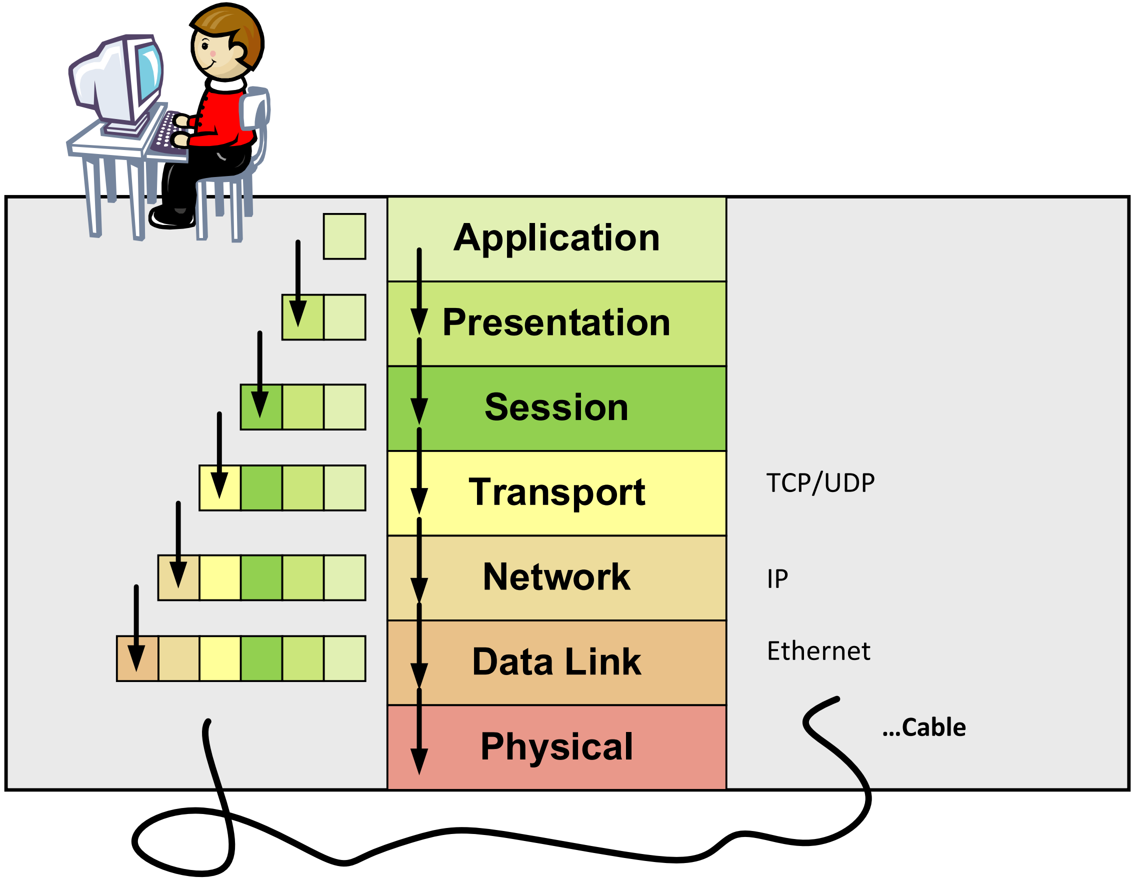

LONAP

Decapsulation

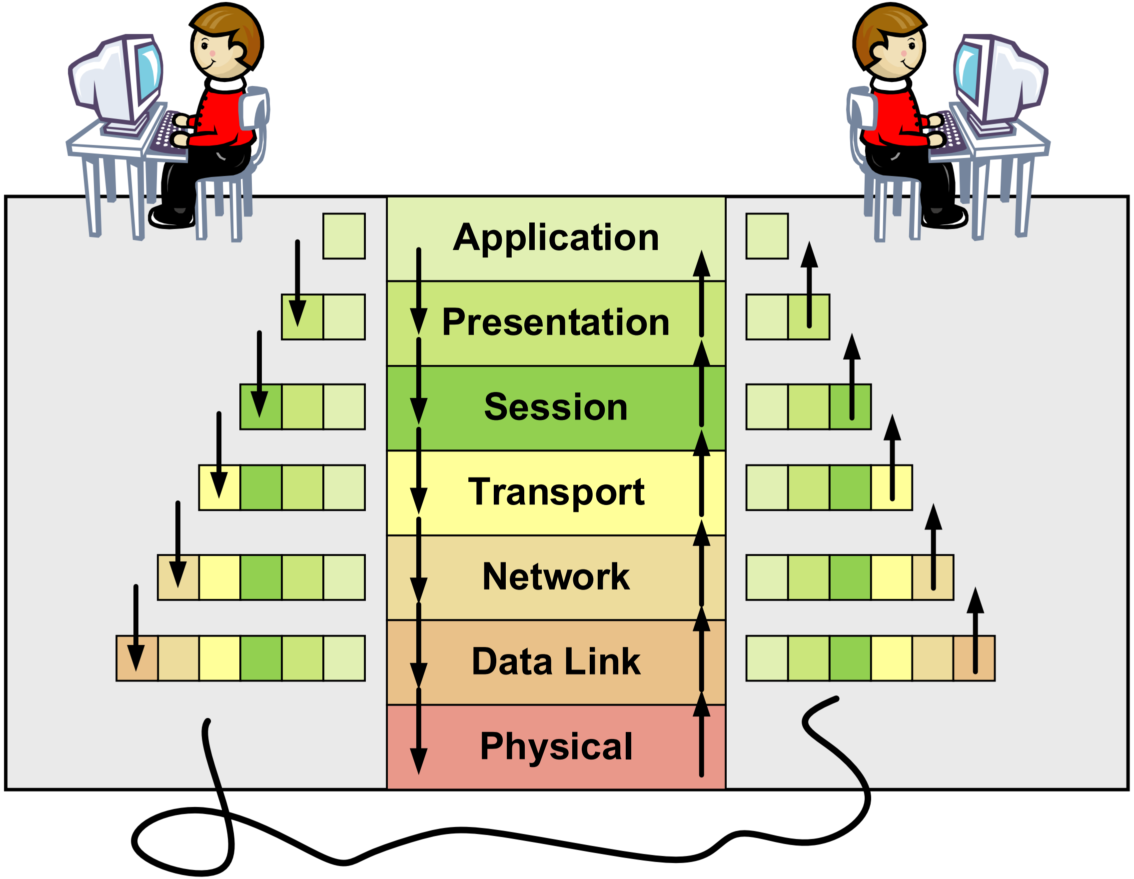

LONAP

Layer 2 / 3

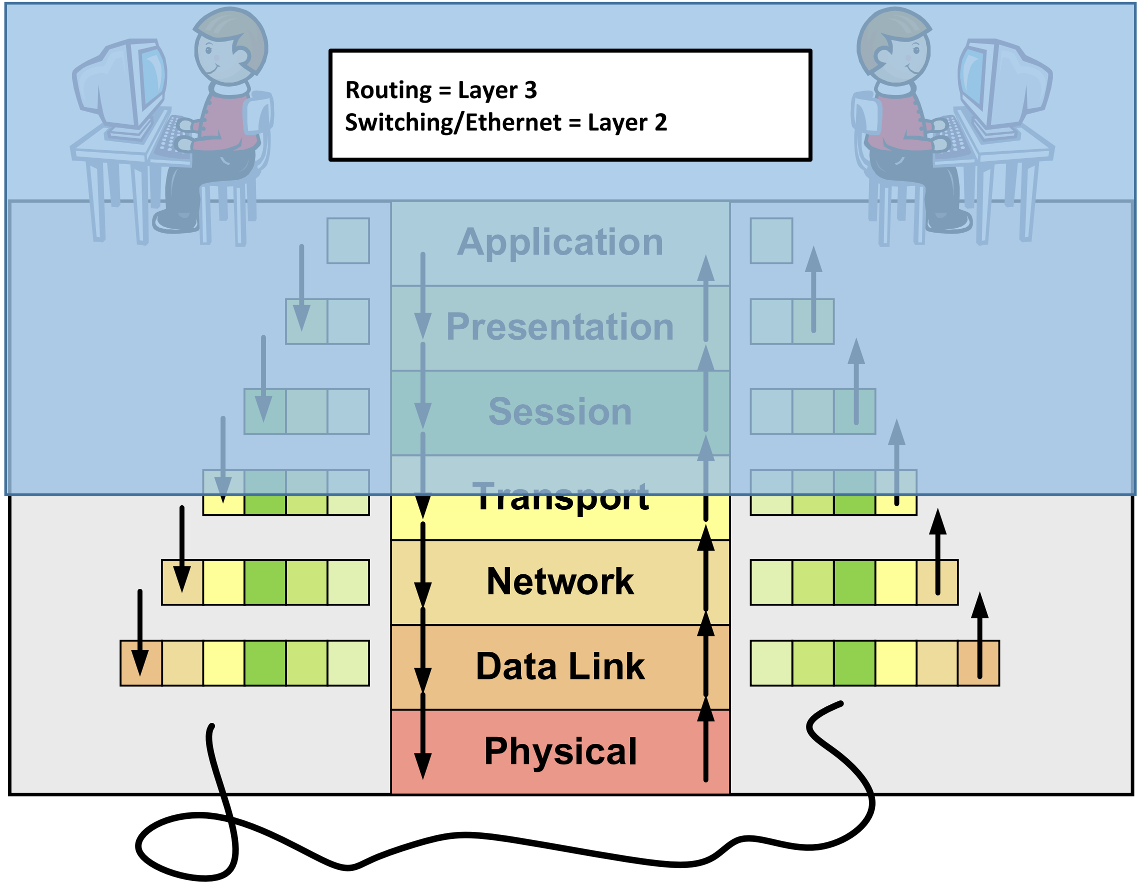

LONAP

Morse Code (1837)

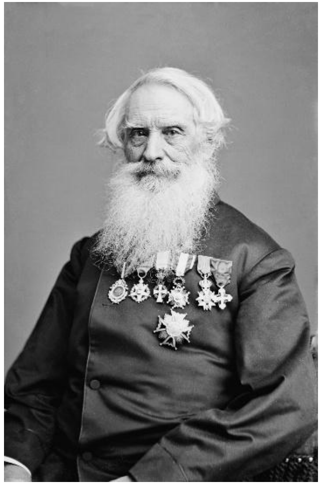

Samuel Morse, (April 27, 1791 – April 2, 1872)

International Morse Code

1. The length of a dot is one unit.
2. A dash is three units.
3. The space between parts of the same letter is one unit.
4. The space between letters is three units.
5. The space between words is seven units.

A	• —
B	— • • •
C	— • — •
D	— • •
E	•
F	• • — •
G	— — •
H	• • • •
I	• •
J	• — — —
K	— • —
L	• — • •
M	— —
N	— •
O	— — —
P	• — — •
Q	— — • —
R	• — •
S	• • •
T	—
U	• • —
V	• • • —
W	• — —
X	— • • —
Y	— • — —
Z	— — • •
1	• — — — —
2	• • — — —
3	• • • — —
4	• • • • —
5	• • • • •
6	— • • • •
7	— — • • •
8	— — — • •
9	— — — — •
0	— — — — —

LONAP

Bits/Bytes

- A ‘bit’ represents a single binary digit
(0 or 1)
- A byte is a fixed-length binary ‘word’ of 8 characters.
- 256 possible combinations from:
- 00000000 to 11111111

LONAP

Layer 1: Physical

Cables

- Copper cabling
- Fibre Cabling

- Digital signals are transmitted as a stream of binary. (1 and 0s)
- Either electrical impulses or light pulses.

Copper

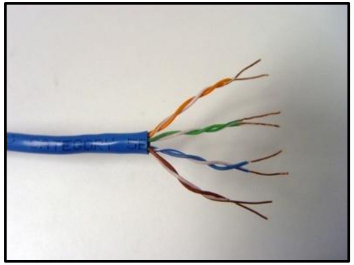

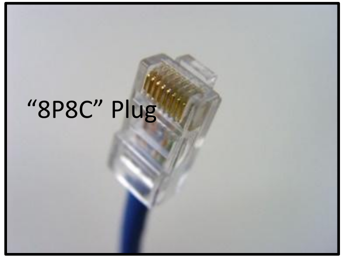

T568B
Wiring order

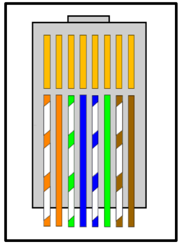

(aka RJ45)

- Category 5/CAT5e/CAT6 cables
- 10/100BASE-T only uses 2 pairs. (orange + green)
- Gigabit over copper signalling is radically different to 10/100
  - Uses all 4 pairs. This can mean a cable that was fine on 10/100 may not work when used for Gigabit!
- Newer kit is now "Auto-MDIX" but old "cross-over" cables will not work at Gigabit: only 2 pairs crossed!
- Some equipment attempts to "auto-correct" this situation.
- Auto-Negotiation is mandatory for 1000BASE-T interfaces, but NOT for older Gigabit only (SFP/GBIC) interfaces which don't support 10/100 – Limited or no auto neg support!

LONAP

Copper Channel Length

- Ethernet spec up to 100 Meters total channel length
- This includes all building wiring and patch cords!
- Can be forgotten when cabling in shared data centres
- Installer’s cable tester only indicates “FAIL” at > 100M

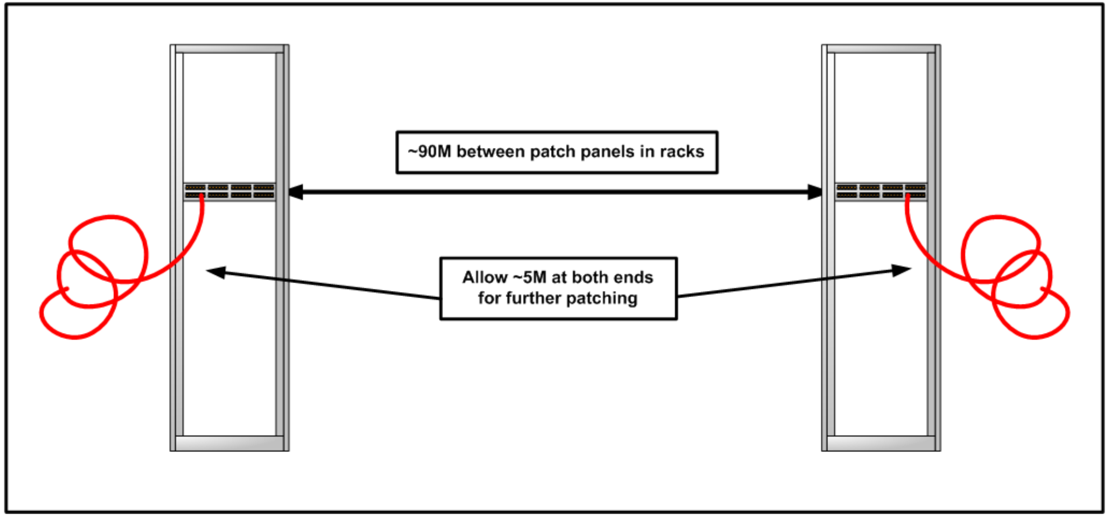

LONAP

Ethernet Copper Faults

- Faulty cable install/crimping
- Auto-Negotiation or duplex mismatches
- Interference from power cables, lighting, other kit
- Physical damage to cabling, plugs or interfaces
- Kinks or sharp bends in the cable
- “old cables”

LONAP

Fibre Optics

- Glass strands
- Generally installed in pairs of strands, one for transmit, another for receive. (TX/RX)
- Connects to an optic at either end
- SX (short haul) optic = Multi Mode
- LX (long haul) optic = Single Mode

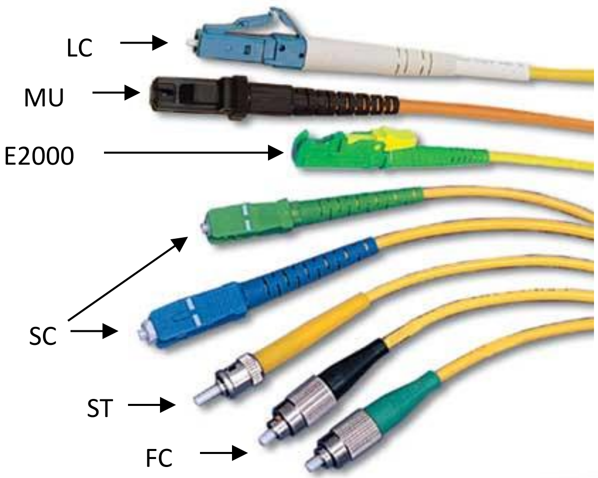

LONAP

Fibre Optic Patch Leads

- Colour is *usually* significant! (TIA-598C)

Almost all new installations are Single Mode. Use this one!

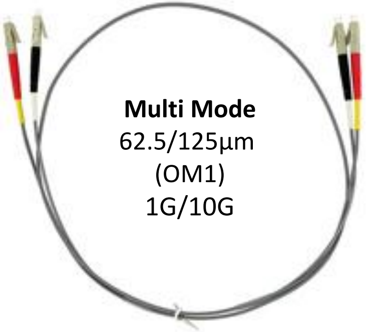

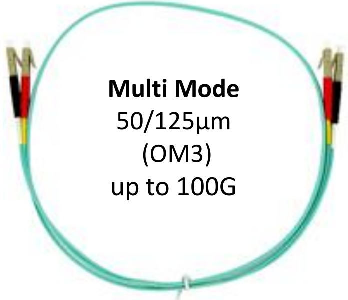

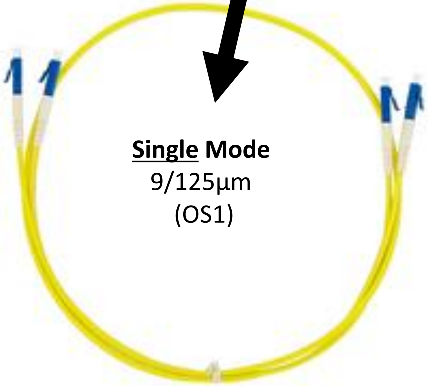

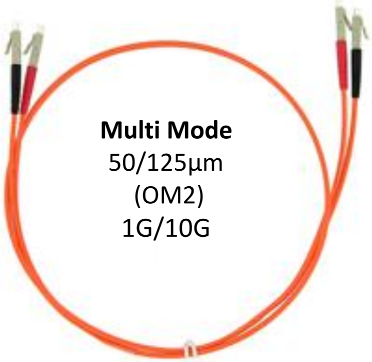

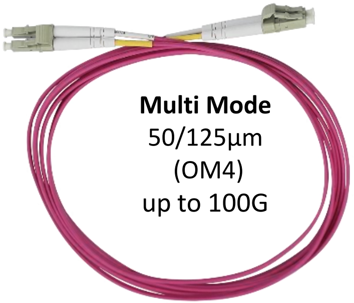

LONAP

Common Fibre Connectors

LC		**ST**	
SC	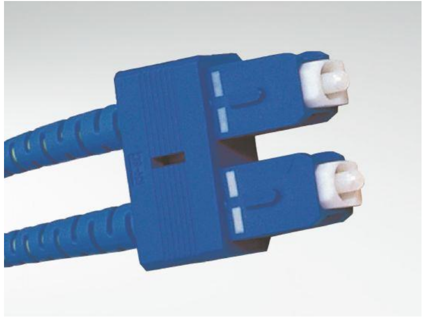	**FC**	

LONAP

Optics (aka transceivers)

Optic Form Factor		Connector
SFP (1G) SFP+ (1G/10G)		**LC**
XFP (10G)		**LC**

Optic Form Factor		Connector
GBIC (Old 1G)		**SC**
XENPAK (Old 10G)		**SC**

Optics (aka transceivers)

Optic Form Factor		Connector	Optic Form Factor		Connector
SFP (1G) SFP+ (1G/10G)	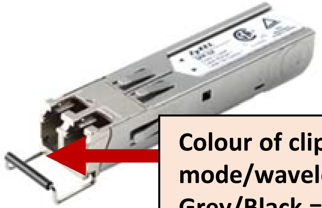	LC	**GBIC**		SC
XFP (10G)	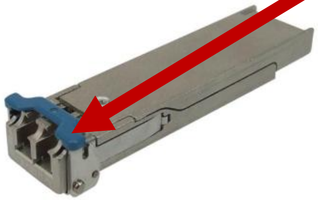	LC	**XENPAK** (Old 10G)	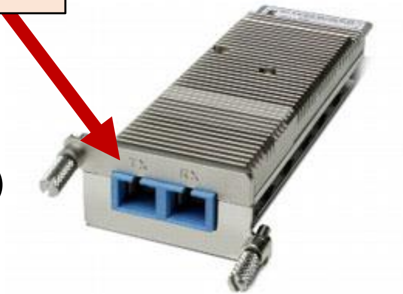	SC

Colour of clips or aperture indicates mode/wavelength
Grey/Black = Multi Mode
BLUE = Single Mode
Others = Single Mode (WDM)

LONAP

Optics – 40G and 100G

Optic Form Factor	
CFP 10×10G or 4×25G lanes	
CFP2 10×10G, 4×25G, 8×25G, or 8×50G lanes	

Optic Form Factor	
CFP4 4×10G or 4×25G lanes	
QSFP+ (40G) **QSFP28 (100G)**	

LONAP

Optics – 400G

Optic Form Factor	Optic Form Factor
OSFP Octal Small Formfactor Pluggable 	**CFP8** C form-factor pluggable 
QSFP-DD Quad Small Form Factor Pluggable-Double Density 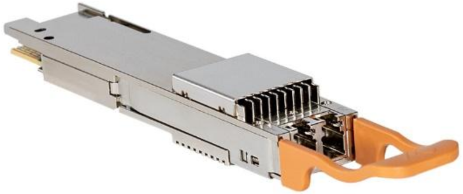	**CDFP** (Older standard) 

LONAP

Layer 2:
Data Link

Switches and Hubs

- What is a **hub**?
- What is a **switch**?

“Original” Ethernet LAN

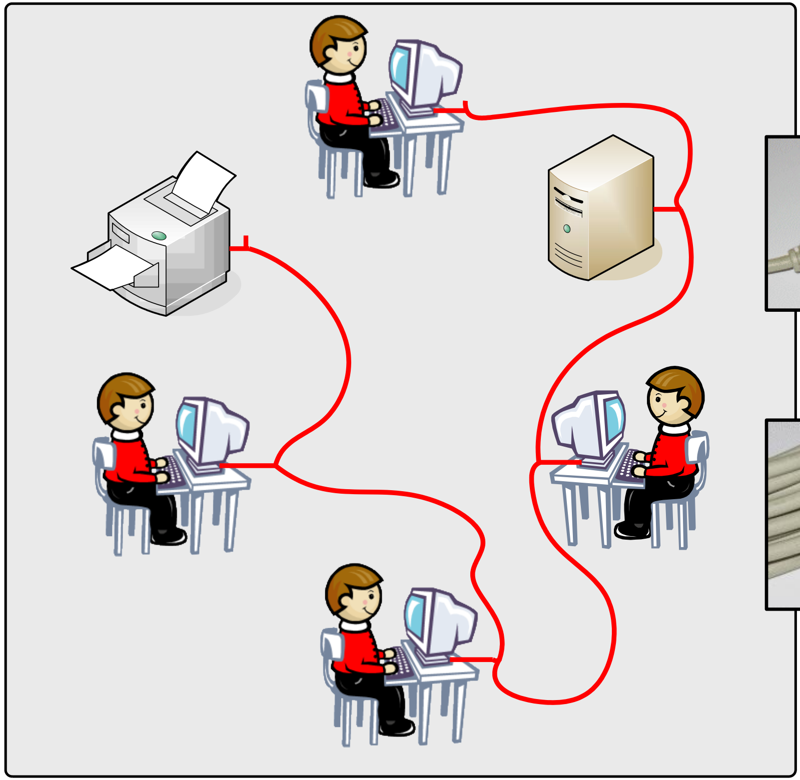

10BASE2

“Original” Ethernet LAN - MAC

CSMA/CD

Carrier sense multiple access with collision detection

ff:ff:ff:ff:ff:ff
= broadcast

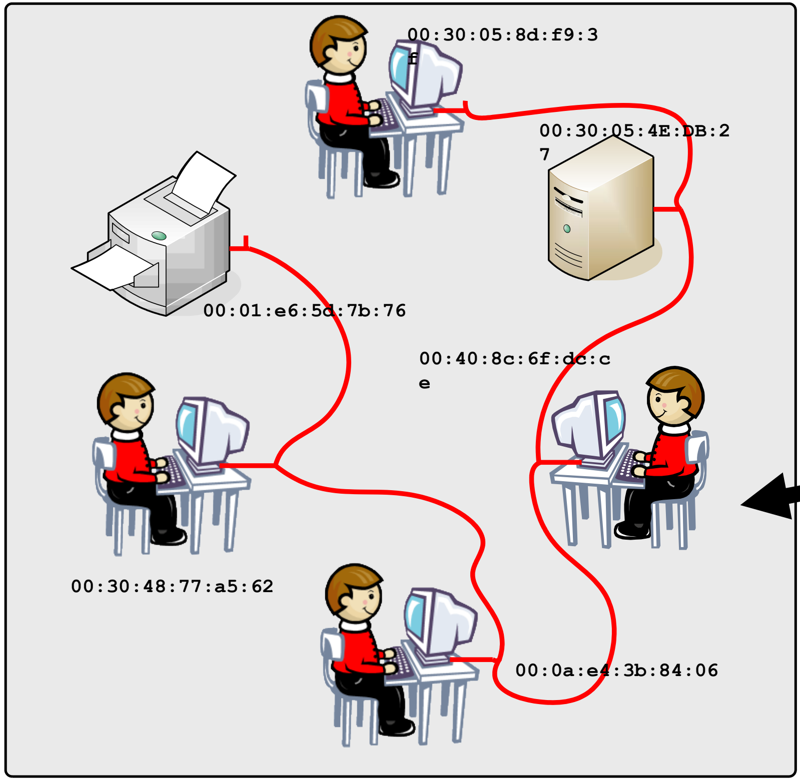

Media
Access
Control

Single **Collision Domain**
Single **Broadcast Domain**

LONAP

Hub

CSMA/CD

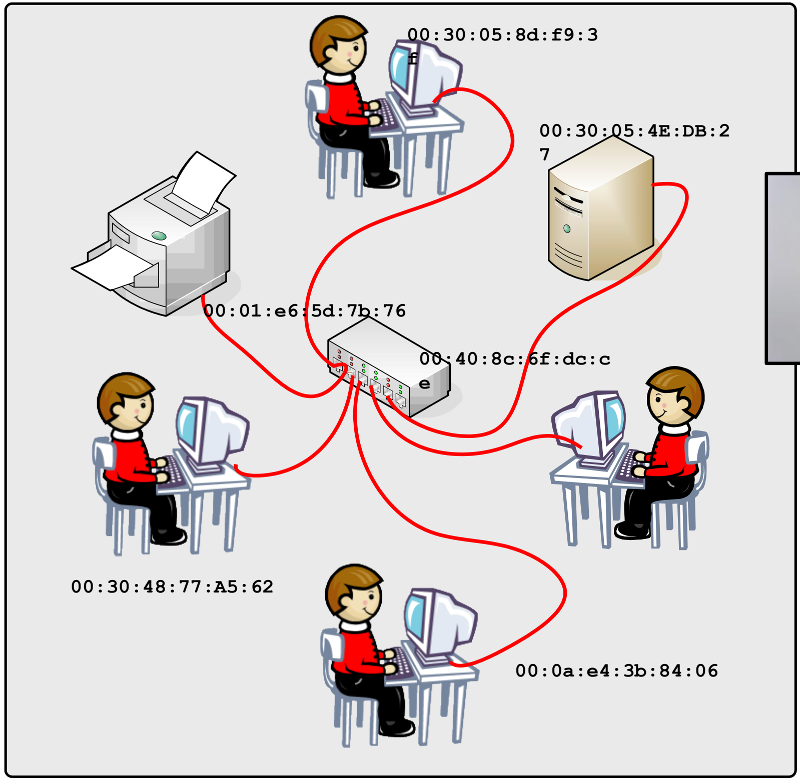

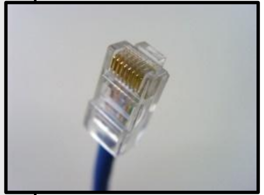

LONAP

Switch – Forwarding table

Port	MAC Address
1	00:01:e6:5d:7b:76
2	00:30:48:77:A5:62
3	00:0a:e4:3b:84:06
4	00:40:8c:6f:dc:ce
5	00:30:05:4E:DB:27
6	00:30:05:8d:f9:3f
7	
8	
…	

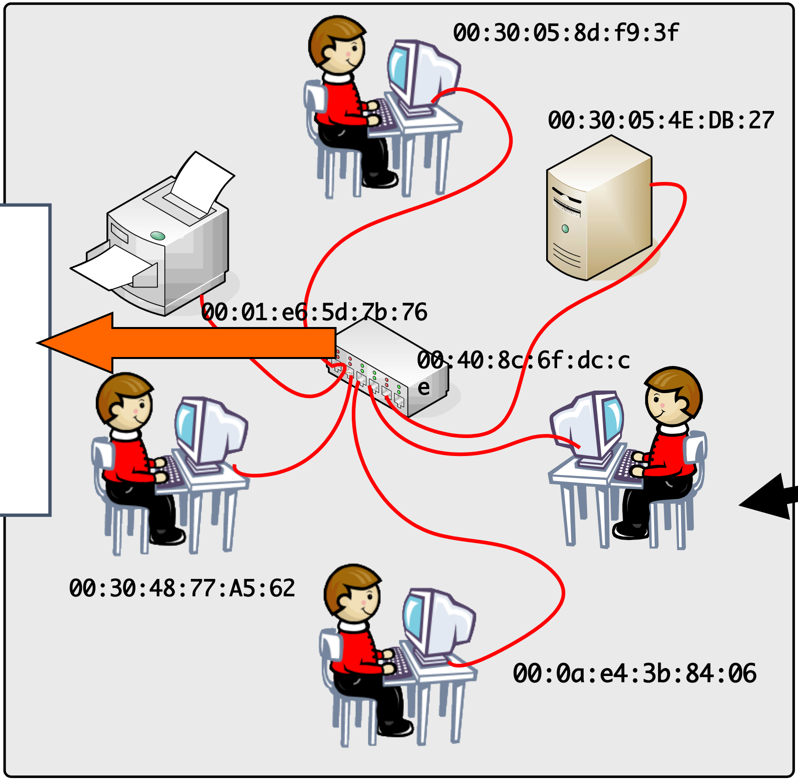

Half Duplex ports each have a Single **Collision Domain**

Full Duplex ports have **No Collision Domain**

Single **Broadcast Domain**

LONAP

Layer 3: Network

LONAP

Routers

- **What is a router?**

Router

- A router is basically a computer with a bunch of interfaces!
- Each interface is connected to a different (Layer 2) network
- Routers look at the destination of packets and decide an interface to forward it out on, **using IP** (Layer 3 **Subnets**)

- Modern (decent!) routers perform many Layer 2 and Layer 3 forwarding functions in hardware
- Also possibility for hardware redundancy (PSUs, processors etc.)

LONAP

Subnets

What is a subnet?

…

LONAP

Subnets

- Subnets allow larger networks to be split up
- Routers ‘route’ between different subnets
- Subnets are groups of IP addresses that represent IP networks. (Layer 3)
- Subnets can be various sizes, defined by a subnet mask

LONAP

IPv4 Subnets

- IP: 192.168.1.1
- Mask: 255.255.255.0
- Gateway: 192.168.1.200

32 bits long in binary

- IP: 11000000.10101000.00000001.00000001
- Mask: 11111111.11111111.11111111.00000000

- Network: 192.168.1.0/24

IPv4 Subnets

- IP: 10.50.9.1
- Mask: 255.0.0.0
- Gateway: 10.0.0.2

- IP: 00001010.00110010.00001001.00000001
- Mask: 11111111.00000000.00000000.00000000

- Network: 10.0.0.0/8

LONAP

IPv4 Subnets - CIDR

- **Classless Inter-Domain Routing** (CIDR)
- IP Allocations managed by Regional Internet Registries (RIRs) such as RIPE.
- CIDR **replaced old** "Class A, B, C" system in mid '90s to conserve address space.

- `Class A: (0.0.0.0 – 126.255.255.255)` `10.0.0.0` (/8 prefix) (~16,777,216 addresses)
- `Class B: (128.0.0.0 – 191.255.255.255)` `172.16.0.0` (/16 prefix) (~65,536 addresses)
- `Class C: (192.0.0.0 – 223.255.255.255)` `192.168.1.0` (/24 prefix) (~256 addresses)

- LONAP Allocation: `5.57.80.0/20` (4096 addresses)
- Office: **5.57.95.0/25** (~128 addresses)

LONAP

IPv4 Subnets - CIDR

- LONAP Allocation: 5.57.80.0**/20** (4096 addresses)
- Total range: 5.57.80.0 - 5.57.95.255
-    Office: 5.57.95.0**/25** (~128 addresses – half a /24)

- **IP:** 5.57.95.10
- **Mask:** 255.255.255.128

- **IP:**   **00000101.00111001.01011111.00001010**
- **Mask:**   ?

LONAP

IPv4 Subnets - CIDR

- LONAP Allocation: 5.57.80.0**/20** (4096 addresses)
- **Office: 5.57.95.0/25 (~128 addresses – half a /24)**

- **IP:** 5.57.95.10
- **Mask:** 255.255.255.128

- **IP:**   **00000101.00111001.01011111.00001010**
- **Mask:** **11111111.11111111.11111111.10000000**

LONAP

Determine what is local / remote

My IP: 5.57.95.10
Mask: 255.255.255.128 (/25)
Destination: 5.57.95.11

My IP:	00000101.00111001.01011111.00001010
Mask:	11111111.11111111.11111111.10000000
	00000101.00111001.01011111.00000000

Dest IP:	00000101.00111001.01011111.00001011
Mask:	
	00000101.00111001.01011111.00000000

Logical AND

0 AND 0 = 0
0 AND 1 = 0
1 AND 0 = 0
1 AND 1 = 1

Compare Logical AND results:

If they are **identical**, Destination is on **same subnet**. ARP for this host.

LONAP

ARP (Address Resolution Protocol)

- When an IPv4 host is **on the same subnet**, send an ARP broadcast.
- -> "Hi Everyone. Will the host at **5.57.95.11** please respond with your MAC address"
- <- "Hi. I am **5.57.95.11**. My MAC address is 00:50:c2:46:60:0b"

```
mon0:~# arp -n -i eth0
5.57.95.37               ether   00:50:c2:46:60:25   C      eth0
5.57.95.31               ether   00:22:19:cd:52:1f   C      eth0
5.57.95.25               ether   00:50:c2:46:60:19   C      eth0
5.57.95.10               ether   00:50:c2:46:60:0a   C      eth0
5.57.95.7                        (incomplete)               eth0
5.57.95.32               ether   00:50:c2:46:60:01   C      eth0
5.57.95.26               ether   00:50:c2:46:60:1a   C      eth0
5.57.95.20               ether   44:1e:a1:61:cc:16   C      eth0
5.57.95.11               ether   00:50:c2:46:60:0b   C      eth0
```

LONAP

Determine what is local / remote

My IP: 5.57.95.10
Mask: 255.255.255.128 (/25)
Destination: 5.57.80.1

	5		57		95		10	
My IP:	00000101	.	00111001	.	01011111	.	00001010	
Mask:	11111111	.	11111111	.	11111111	.	10000000	
	00000101	.	00111001	.	01011111	.	00000000	= (Net: 5.57.95.0/25)

	5		57		80		1	
Dest IP:	00000101	.	00111001	.	01010000	.	00000001	
Mask:								
	00000101	.	00111001	.	01010000	.	00000000	= (Net: 5.57.80.0)

Logical AND

0 AND 0 = 0
0 AND 1 = 0
1 AND 0 = 0
1 AND 1 = 1

Compare Logical AND results:

If they **differ**, destination is on **remote** subnet. Send to router.

LONAP

Determine what is local / remote

My IP: 5.57.95.10
Mask: 255.255.240.0 (/20)
Destination: 5.57.80.1

	5		57		95		10
My IP:	00000101	.	00111001	.	01011111	.	00001010
Mask:	11111111	.	11111111	.	11110000	.	00000000
	00000101	.	00111001	.	01010000	.	00000000

= (Net: 5.57.80.0/20)

	5		57		80		1
Dest IP:	00000101	.	00111001	.	01010000	.	00000001
Mask:							
	00000101	.	00111001	.	01010000	.	00000000

Logical AND

0 AND 0 = 0
0 AND 1 = 0
1 AND 0 = 0
1 AND 1 = 1

Compare Logical AND results:

If they are the **same**, destination is on the **same** subnet. ARP and send directly to destination host.

LONAP

IPv4 CIDR Subnetting

Addresses	Prefix	Subnet Mask	Cisco Wildcard
4	/30	255.255.255.252	0.0.0.3
8	/29	255.255.255.248	0.0.0.7
16	/28	255.255.255.240	0.0.0.15
32	/27	255.255.255.224	0.0.0.31
64	/26	255.255.255.192	0.0.0.63
128	/25	255.255.255.128	0.0.0.127
256	/24	255.255.255.0	0.0.0.255
512	/23	255.255.254.0	0.0.1.255
1,024	/22	255.255.252.0	0.0.3.255
2,048	/21	255.255.248.0	0.0.7.255
4,096	/20	255.255.240.0	0.0.15.255
8,192	/19	255.255.224.0	0.0.31.255
16,384	/18	255.255.192.0	0.0.63.255
32,768	/17	255.255.128.0	0.0.127.255
65,536	/16	255.255.0.0	0.0.255.255
131,072	/15	255.254.0.0	0.1.255.255
262,144	/14	255.252.0.0	0.3.255.255
524,288	/13	255.248.0.0	0.7.255.255
1,048,576	/12	255.240.0.0	0.15.255.255
2,097,152	/11	255.224.0.0	0.31.255.255
4,194,304	/10	255.192.0.0	0.63.255.255
8,388,608	/9	255.128.0.0	0.127.255.255
16,777,216	/8	255.0.0.0	0.255.255.255

IPv6

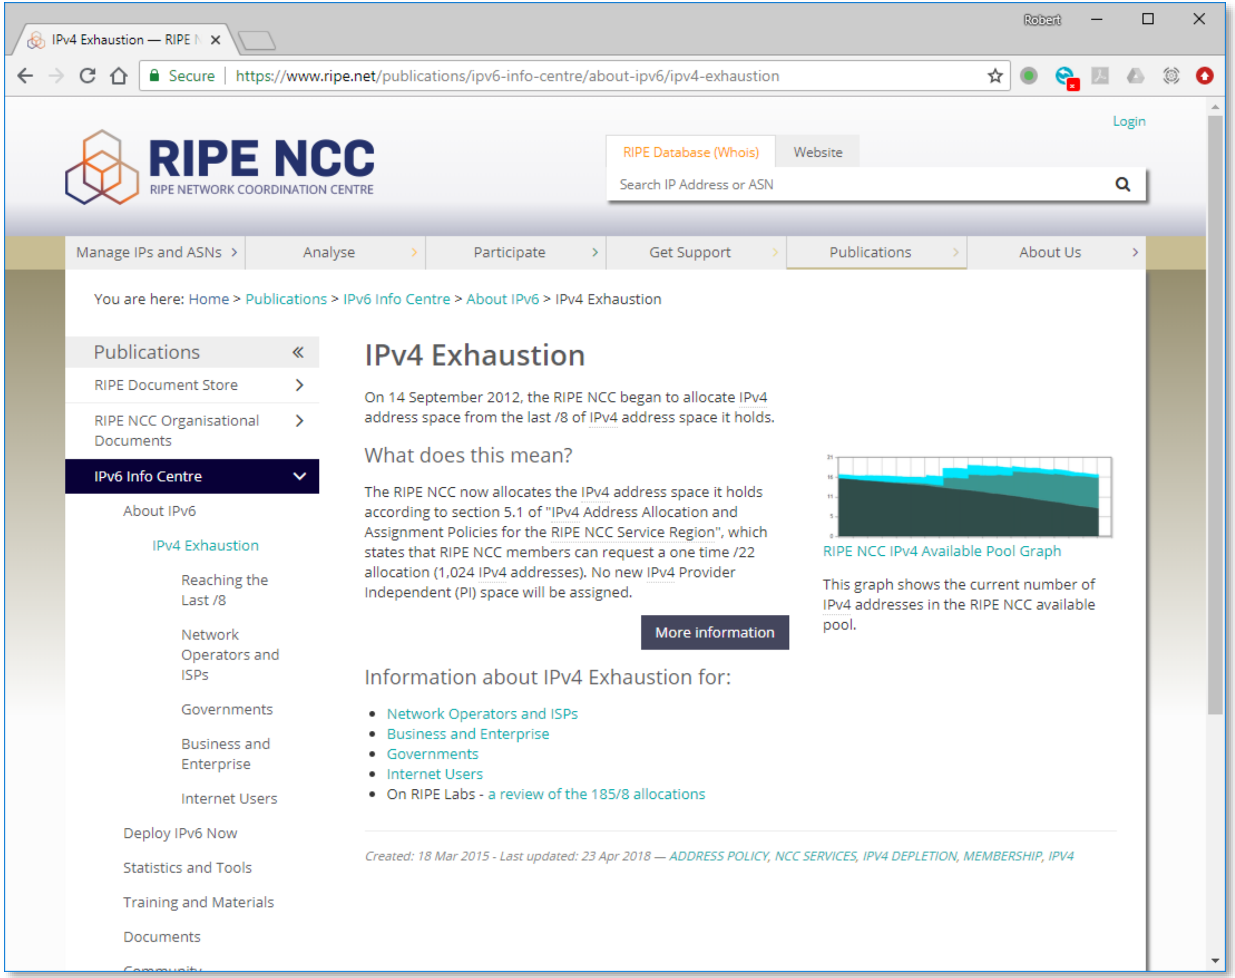

IPv4 Exhaustion — RIPE N
Robert
Secure
https://www.ripe.net/publications/ipv6-info-centre/about-ipv6/ipv4-exhaustion
Login
RIPE NCC
RIPE NETWORK COORDINATION CENTRE
RIPE Database (Whois)
Website
Search IP Address or ASN
Manage IPs and ASNs
Analyse
Participate
Get Support
Publications
About Us
You are here: Home > Publications > IPv6 Info Centre > About IPv6 > IPv4 Exhaustion
Publications
RIPE Document Store
RIPE NCC Organisational Documents
IPv6 Info Centre
About IPv6
IPv4 Exhaustion
Reaching the Last /8
Network Operators and ISPs
Governments
Business and Enterprise
Internet Users
Deploy IPv6 Now
Statistics and Tools
Training and Materials
Documents
IPv4 Exhaustion
On 14 September 2012, the RIPE NCC began to allocate IPv4 address space from the last /8 of IPv4 address space it holds.
What does this mean?
The RIPE NCC now allocates the IPv4 address space it holds according to section 5.1 of "IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region", which states that RIPE NCC members can request a one time /22 allocation (1,024 IPv4 addresses). No new IPv4 Provider Independent (PI) space will be assigned.
More information
RIPE NCC IPv4 Available Pool Graph
This graph shows the current number of IPv4 addresses in the RIPE NCC available pool.
Information about IPv4 Exhaustion for:
Network Operators and ISPs
Business and Enterprise
Governments
Internet Users
On RIPE Labs - a review of the 185/8 allocations
Created: 18 Mar 2015 - Last updated: 23 Apr 2018 — ADDRESS POLICY, NCC SERVICES, IPV4 DEPLETION, MEMBERSHIP, IPV4

LONAP

IPv6

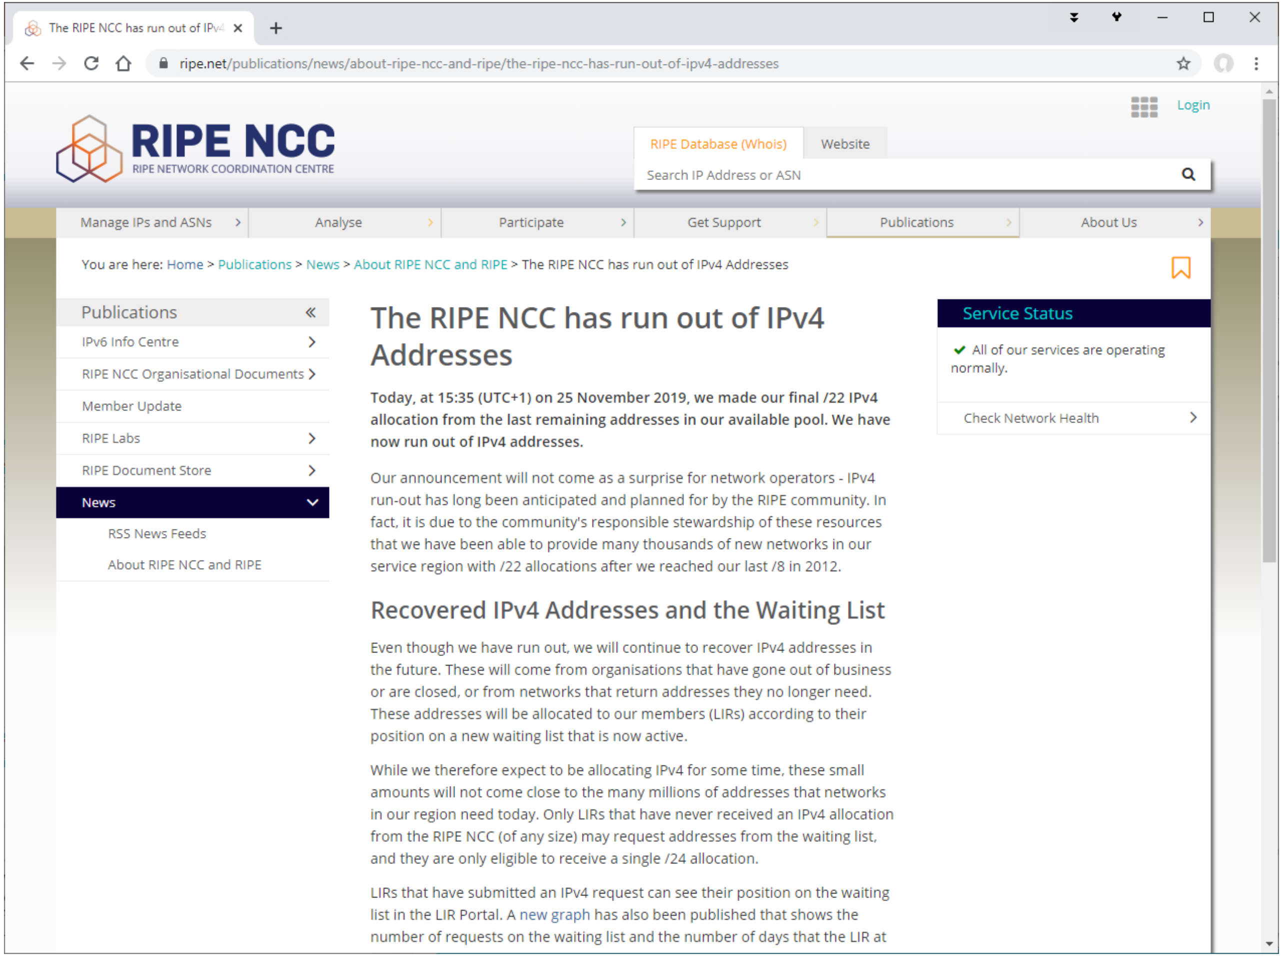

The RIPE NCC has run out of IPv4

ripe.net/publications/news/about-ripe-ncc-and-ripe/the-ripe-ncc-has-run-out-of-ipv4-addresses

Login

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

RIPE Database (Whois) | Website

Search IP Address or ASN

Manage IPs and ASNs | Analyse | Participate | Get Support | Publications | About Us

You are here: Home > Publications > News > About RIPE NCC and RIPE > The RIPE NCC has run out of IPv4 Addresses

Publications
- IPv6 Info Centre
- RIPE NCC Organisational Documents
- Member Update
- RIPE Labs
- RIPE Document Store
- News
  - RSS News Feeds
  - About RIPE NCC and RIPE

The RIPE NCC has run out of IPv4 Addresses

Today, at 15:35 (UTC+1) on 25 November 2019, we made our final /22 IPv4 allocation from the last remaining addresses in our available pool. We have now run out of IPv4 addresses.

Our announcement will not come as a surprise for network operators - IPv4 run-out has long been anticipated and planned for by the RIPE community. In fact, it is due to the community's responsible stewardship of these resources that we have been able to provide many thousands of new networks in our service region with /22 allocations after we reached our last /8 in 2012.

Recovered IPv4 Addresses and the Waiting List

Even though we have run out, we will continue to recover IPv4 addresses in the future. These will come from organisations that have gone out of business or are closed, or from networks that return addresses they no longer need. These addresses will be allocated to our members (LIRs) according to their position on a new waiting list that is now active.

While we therefore expect to be allocating IPv4 for some time, these small amounts will not come close to the many millions of addresses that networks in our region need today. Only LIRs that have never received an IPv4 allocation from the RIPE NCC (of any size) may request addresses from the waiting list, and they are only eligible to receive a single /24 allocation.

LIRs that have submitted an IPv4 request can see their position on the waiting list in the LIR Portal. A new graph has also been published that shows the number of requests on the waiting list and the number of days that the LIR at

Service Status

All of our services are operating normally.

Check Network Health

IPv6

- The ‘Internet’ has run out of IPv4 address space (~4.2 billion addresses)
- CIDR, NAT, Policy and other things have conserved IPv4 address space… but we have to address it some day!
- IPv6 – 128 bit address space, 2^{128} (~340 undecillion) unique addresses!
- …That’s a lot.

- 340,282,366,920,938,463,463,374,607,431,768,211,456

LONAP

IPv6 Subnets

- LONAP Allocation: 2a00:eb20::/32 (lots!)
- Total Range: 2a00:eb20:0000:0000:0000:0000:0000:0000 – 2a00:eb20:ffff:ffff:ffff:ffff:ffff:ffff
- Office: 2a00:eb20:100::/64 (18,446,744,073,709,551,616 addresses)

- IP: 2a00:eb20:0100:0000:0000:0000:0000:0011
- Mask: /64

- IP: 0010101000000000:1110101100100000:0000000100000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000:0000000000010001
- Mask: 1111111111111111:1111111111111111:1111111111111111:1111111111111111:0000000000000000:0000000000000000:0000000000000000:0000000000000000

LONAP

IPv6 Addresses

- Consecutive zeros between colons can be ‘compressed’ with double colons “::”
- **Leading zeros** in hextets can be stripped:
  - `:0def: = :def:`
  - ‘`:0000:`’ can be shortened to **`:0:`**

- `2001:0db8:0000:cd30:0000:0000:0000:0000/60`
- Can also be compressed as:
  - `2001:db8:0:cd30::/60`
  - `2001:db8::cd30:0:0:0:0/60` (…Don’t!)
- …But you can only do “`::`” **ONCE** in an address:
  - `2001:db8:0:cd30::/60` not `2001:db8::cd30::/60`
- “`::`” can also be used to compress ‘all zeros’ addresses - `::1`

LONAP

RFC 5952: https://datatracker.ietf.org/doc/html/rfc5952

IPv6 Address Examples

Long IPv6 Address	Compressed	Description
2001:0db8:0:0:8:800:200c:417a	2001:db8::8:800:200c:417a	A unicast address
ff01:0:0:0:0:0:0:101	ff01::101	A multicast address
0:0:0:0:0:0:0:1	::1	The loopback address
0:0:0:0:0:0:0:0	::	The unspecified address
0:0:0:0:0:0:0:0/0	::/0	“Default gateway”

- RFC 4291: http://tools.ietf.org/html/rfc4291 (Here’s lots of ways to write IPv6 addresses…)
- RFC 5952: http://tools.ietf.org/html/rfc5952 (This got annoying. Write them like **this** please!)

LONAP

IPv6 Address Space

IPv6 Prefix	Allocation	Reference	Note
0000::/8	Reserved by IETF	[RFC4291]	[1][5][6]
0100::/8	Reserved by IETF	[RFC4291]	
0200::/7	Reserved by IETF	[RFC4048]	[2]
0400::/6	Reserved by IETF	[RFC4291]	
0800::/5	Reserved by IETF	[RFC4291]	
1000::/4	Reserved by IETF	[RFC4291]	
2000::/3	**Global Unicast**	**[RFC4291]**	**[3]**
4000::/3	Reserved by IETF	[RFC4291]	
6000::/3	Reserved by IETF	[RFC4291]	
8000::/3	Reserved by IETF	[RFC4291]	
A000::/3	Reserved by IETF	[RFC4291]	
C000::/3	Reserved by IETF	[RFC4291]	
E000::/4	Reserved by IETF	[RFC4291]	
F000::/5	Reserved by IETF	[RFC4291]	
F800::/6	Reserved by IETF	[RFC4291]	
FC00::/7	Unique Local Unicast	[RFC4193]	
FE00::/9	Reserved by IETF	[RFC4291]	
FE80::/10	Link Local Unicast	[RFC4291]	
FEC0::/10	Reserved by IETF	[RFC3879]	[4]
FF00::/8	Multicast	[RFC4291]	

- [3] The IPv6 Unicast space encompasses the entire IPv6 address range except for FF00::/8. [RFC4291] IANA unicast address assignments are currently limited to the IPv6 unicast address range of 2000::/3. IANA assignments from this block are registered in the IANA registry: ipv6-unicast-address-assignments.

IPv6 Subnetting

Prefix	/48s	/56s	/64s
/24	16M	4G	1T
/25	8M	2G	512G
/26	4M	1G	256G
/27	2M	512M	128G
/28	1M	256M	64G
/29	**512K**	**128M**	**32G**
/30	256K	64M	16G
/31	128K	32M	8G
/32	**64K**	**16M**	**4G**
/33	32K	8M	2G
/34	16K	4M	1G
/35	8K	2M	512M
/36	4K	1M	256M
/37	2K	512K	128M
/38	1K	256K	64M
/39	512	128K	32M
/40	256	64K	16M
/41	128	32K	8M
/42	64	16K	4M
/43	32	8K	1M
/44	16	4K	1M

Prefix	/48s	/56s	/64s
/45	8	2K	512K
/46	4	1K	256K
/47	2	512	128K
/48	**1**	**256**	**64K**
/49		128	32K
/50		64	16K
/51		32	8K
/52		16	4K
/53		8	2K
/54		4	1K
/55		2	512
/56		1	256
/57			128
/58			64
/59			32
/60			16
/61			8
/62			4
/63			2
/64			1

LONAP

Service Provider Networks

- Network hardware (!!)
- Location with good connectivity
- IP Addresses (IPv4 and IPv6)
- AS Number
- Connectivity (transit and maybe peering)

LONAP

Regional Internet Registries (RIRs)

- RIRs manage resources (IPs, ASNs) in a service region

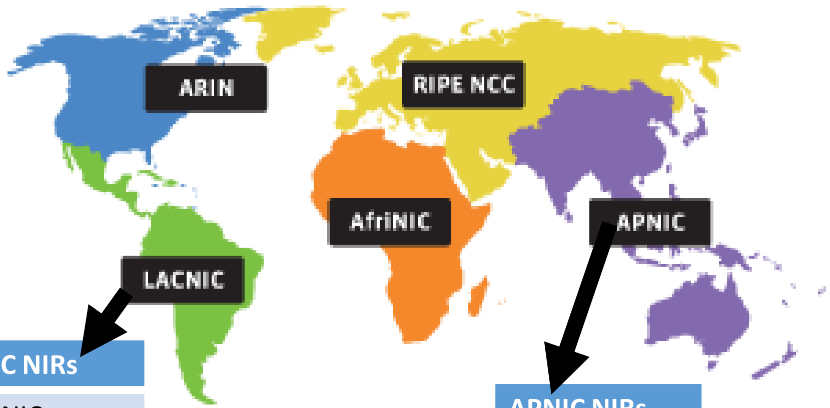

LACNIC NIRs
nic.br NIC Mexico nic.mx NIC Brazil

APNIC NIRs
APJII, CNNIC, JPNIC, KRNIC, TWNIC, VNNIC, IRINN

Registry	Area Covered
AfriNIC	Africa Region
APNIC	Asia/Pacific Region
ARIN	Canada, United States, some Caribbean nations
LACNIC	Latin America and some Caribbean nations
RIPE NCC	Europe, Russia, Middle East, and Central Asia

LONAP

Local Internet Registries (LIRs)

- An **LIR** is a member of an **RIR** (or NIR)
- Generally an LIR is a service provider wanting their own IP address space and AS number
- Provider Aggregatable (PA) IP space

- RIR makes an ALLOCATION to the LIR
- LIR makes ASSIGNMENTS to customers

- RIPE region is Europe, the Middle East, and Central Asia

LONAP

Single Homed – 'Before'

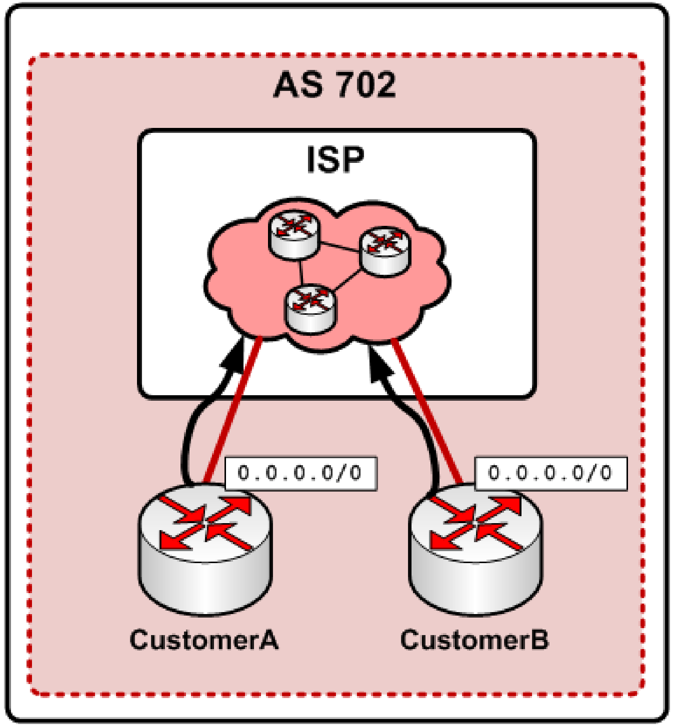

- Customer gets IPs from ISP
- Customer is **inside** the ISP's AS
- Customer points default gateway at ISP
- Most businesses, end-users/residential customers
- Cannot usually use same IP range with two ISPs.

LONAP

Multi Homed – 'After'

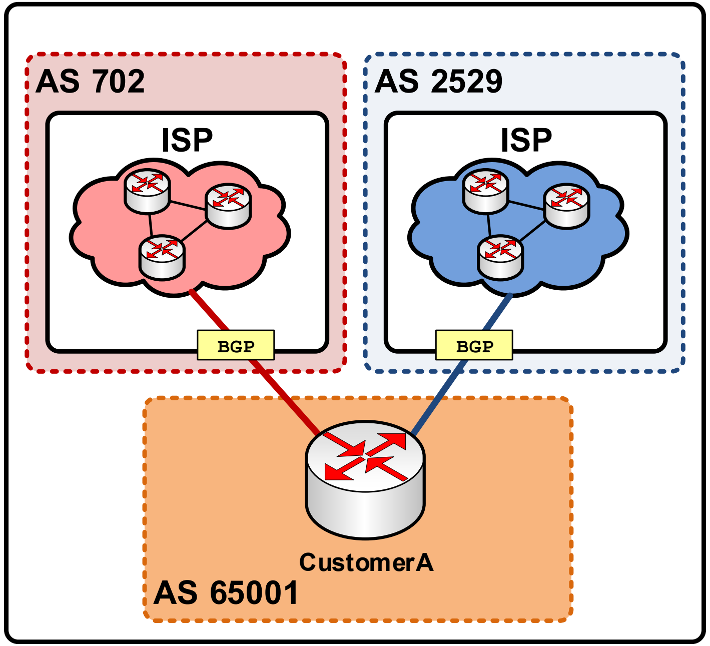

- Customer has **own IPs**
- Customer has **own AS**
- **Multiple** upstream transit providers
- BGP sessions to each, uses full table. BGP selects 'best' path
- Can move and add ISPs and peers without changing IP ranges

Why bother?

- **Advantages**
  - ‘independent’ network can be truly multi-homed
  - ‘own’ IP address space
  - Multiple transit providers, seamless to switch
  - Increase network resilience
  - Ready for peering at exchanges
  - More control where traffic goes and why (policy)
  - Easy expansion

LONAP

Why bother?

- **Disadvantages…**
  - Re-number IPs on existing network
  - Requires big routers
  - Complex procedures to get started
  - Increases configuration complexity
  - Requires specialist knowledge/skills
    - Knowledge: Learn BGP and lots more!
    - Skills: Apply knowledge, gain experience, do it well

  - Not cost-effective for every network

LONAP

BGP ROUTING

Routing Protocols

What is a routing protocol?

Discuss…

Routing Protocols

- Routers use **routing protocols** to learn and distribute which networks are connected via which interfaces.
- Routing protocols also help make decisions about the ‘best’ path to take.
- **BGP4** (Border Gateway Protocol) mostly used **between networks** on the Internet. (**EGP**)
- **OSPF, RIP, EIGRP, ISIS**.. are all Interior routing protocols. (**IGP**) that operate within an AS
- Most commonly OSPF or ISIS in ISPs
- *Can* use *only* BGP to do everything. Most have IGP+BGP for scaling.

LONAP

BGP4

- Every network has an **Autonomous System Number** (AS Number)
- AS numbers assigned by **RIR**
- Every router inside each network has the **same AS number**.
  - BGP sessions configured in the same ASN = **iBGP** (Internal)
- BGP Sessions established between ASNs = **eBGP** (External)
- No ‘magic router discovery’: every session must be configured.

- **Path Vector** routing protocol

LONAP

1.4 Basic operation (p.21)

BGP4

- Networks can be ‘**multi-homed**’ – not dependent on a single ‘upstream’ provider.
- …providing they have their own IP addresses
- Strict control of **routing policy**
- Control which prefixes are advertised where
  - Peers
  - Transit
  - Downstream customers
- Set preferences AS-wide: preferred routes

LONAP

1.4.4 Routing Policies (p.25)

Autonomous Systems

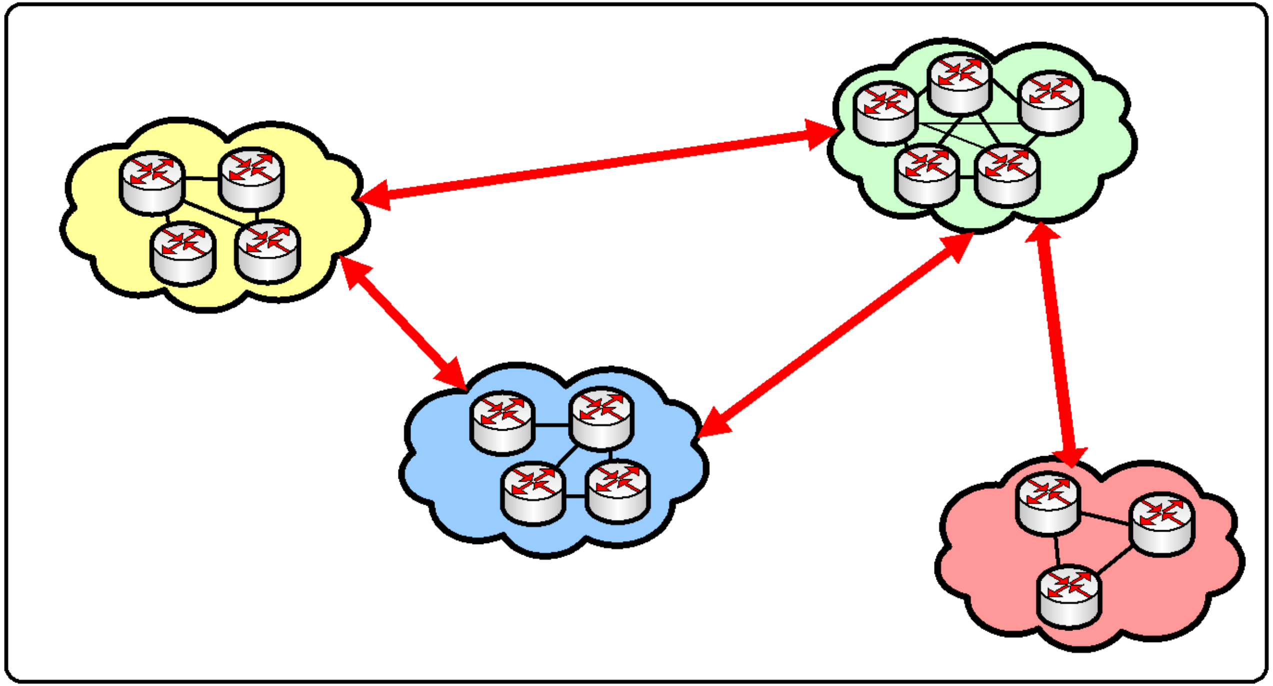

AS Path

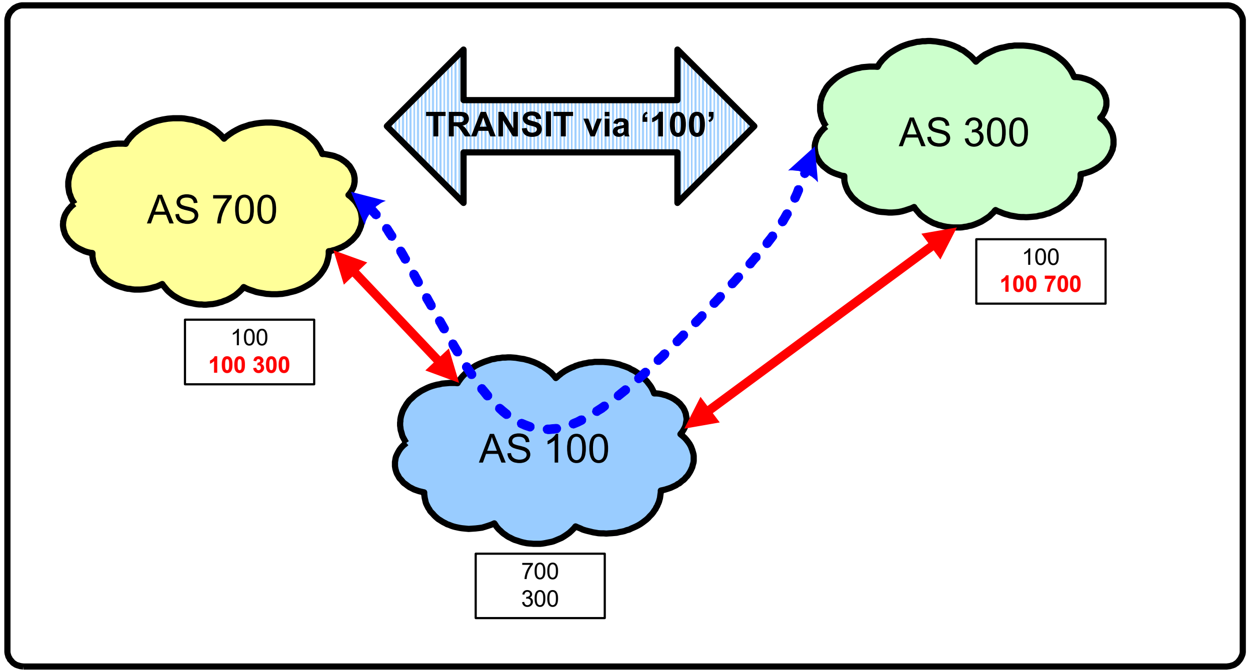

LONAP

BGP Routing Table

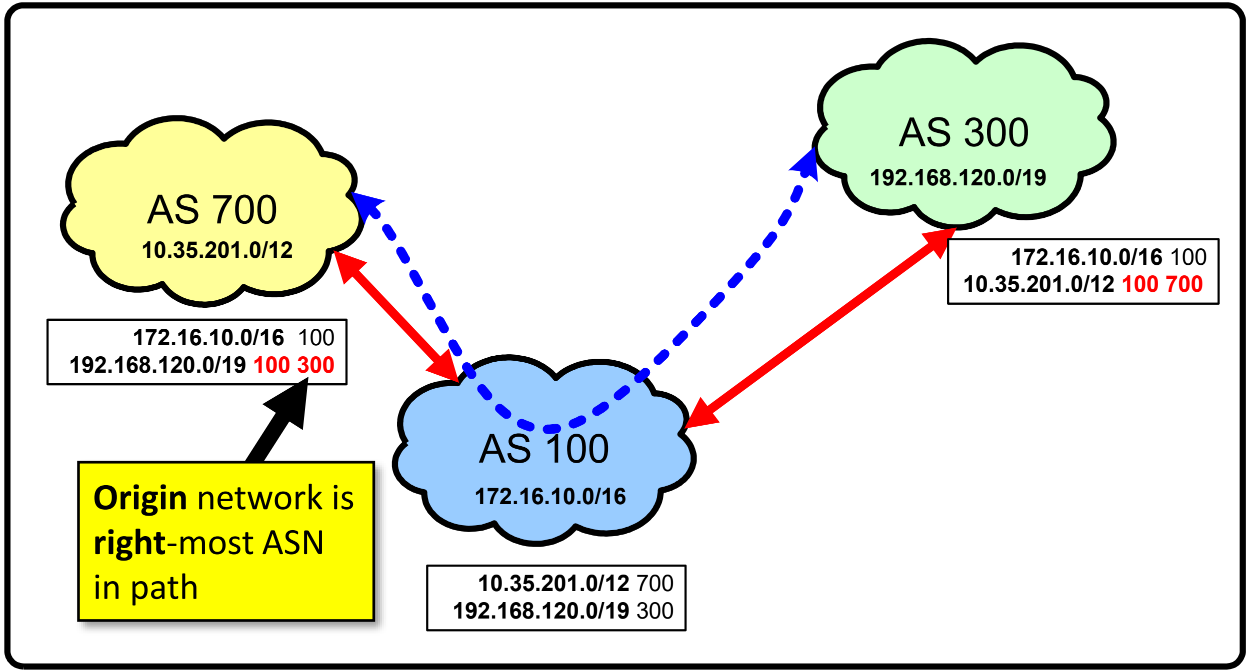

A career in networks… (my tips…)

- Making a name for yourself
- (hopefully a good one)

- Get BGP Skills!
- Do some great presentations. Tell a story, not just dull technical facts
- Get to know the community
- Help others if you can
- Help and acknowledge those who have helped you
- Build your own labs and tinker...
- Be willing to learn.
- - ..and to teach..

Skills Frameworks

- https://sfia-online.org/

LONAP

Done