

Systems and Devices 2 (Network)

Lec 5: Data Link Layer

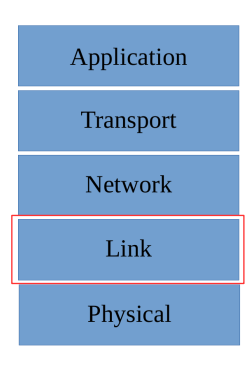
University of York : M Freeman 2024

Before we get started ...

- We have looked at how packets are transferred at the network layer i.e. host to host, using:
 - ▶ IP addresses across a WAN, a network or networks.
 - ▶ BUT, we have not considered how packets get from one host to another on a LAN e.g. how does a switch know which cable to transmit a packet down to get to a specific destination?
- Moving down to the Data link layer i.e. host to host communications on a LAN :
 - ▶ Physical addressing rather than logical addressing
 - ▶ Ethernet protocol, MAC addresses ...


University of York : M Freeman 2024

Internet protocol stack

- 
- Application
 - Transport
 - Network
 - Link (layer 2)
 - ▶ Moves packets (**frames**) from one node (PC, router, switch etc) to another, error detection / correction. Different protocols depending on transmission medium used, we will focus on:
 - ♦ Ethernet protocol
 - ▶ Also, contains house keeping protocols to identify network addresses used.
 - ♦ Address Resolution Protocol (ARP)
 - Physical

University of York : M Freeman 2024

Ethernet Protocol



Name	Description	Status
IEEE 802.1	Higher Layer LAN Protocols Working Group	Active
IEEE 802.2	LLC	Disbanded
IEEE 802.3	Ethernet	Active
IEEE 802.4	Token bus	Disbanded
IEEE 802.5	Token Ring MAC layer	Disbanded
IEEE 802.6	MANs (DQDB)	Disbanded
IEEE 802.7	Broadband LAN using Coaxial Cable	Disbanded
IEEE 802.8	Fiber Optic TAG	Disbanded
IEEE 802.9	Integrated Services LAN (ISLAN or isoEthernet)	Disbanded
IEEE 802.10	Interoperable LAN Security	Disbanded
IEEE 802.11	Wireless LAN (WLAN) & Mesh (Wi-Fi certification)	Active
IEEE 802.12	100BaseVG	Disbanded
IEEE 802.13	Unused ^[2]	reserved for Fast Ethernet development ^[3]
IEEE 802.14	Cable modems	Disbanded
IEEE 802.15	Wireless PAN	Active

- Ethernet protocol
 - ▶ IEEE 802 family : https://en.wikipedia.org/wiki/IEEE_802

University of York : M Freeman 2024

Ethernet Protocol


Ethernet standard	Date	Description
802.3a	1985	10BASE2 10 Mbit/s (1.25 MB/s) over thin Coax (a.k.a. thinnet or cheapernet)
802.3b	1985	10BROAD36
802.3c	1985	10 Mbit/s (1.25 MB/s) repeater specs
802.3-1985	1985	a revision of the base standard from 1983
802.3d	1987	Fiber-optic inter-repeater link
802.3e	1987	1BASE5 or StarLAN, first use of (voice-grade) twisted pair cabling, 1 Mbit/s, maximum reach of 250 to 500 m
802.3i	1990	10BASE-T 10 Mbit/s (1.25 MB/s) over twisted pair
802.3j	1993	10BASE-F 10 Mbit/s (1.25 MB/s) over optical fiber
802.3q	1993	GDMO (ISO 10164-4) format for Layer Managed Objects
802.3u	1995	100BASE-TX, 100BASE-T4, 100BASE-FX Fast Ethernet at 100 Mbit/s (12.5 MB/s) with autonegotiation
802.3x	1997	Full Duplex and flow control; also incorporates DIX framing, so there's no longer a DIX/802.3 split
802.3y	1998	100BASE-T2 100 Mbit/s (12.5 MB/s) over voice-grade twisted pair
802.3z	1998-07	1000BASE-X Gbit/s Ethernet over optical fiber at 1 Gbit/s (125 MB/s)
802.3-1998	1998-07	(802.3aa) A revision of base standard incorporating the above amendments and errata
802.3ab	1999-06	1000BASE-T Gbit/s Ethernet over twisted pair at 1 Gbit/s (125 MB/s)

Ethernet protocol

- IEEE 802.3 family : https://en.wikipedia.org/wiki/IEEE_802.3

University of York : M Freeman 2024

Ethernet Protocol

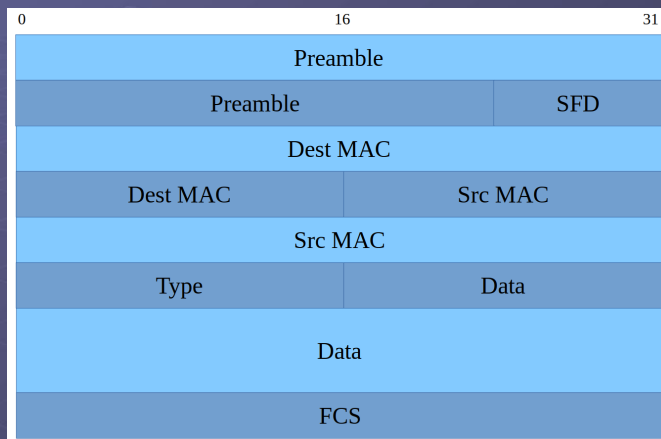


Ethernet protocol

- Create in 1983, no RFC rather an IEEE standard: 802.3
- The default protocol used in LANs.
- On a subnet hosts are identified using their Media Access Control (MAC) address, rather than their IP address.
 - A unique identifier, typically hard coded into each NIC at manufacture as a 48bit address, represented as six two-digit hexadecimal values.
- Q : why are there min / max frame sizes?

University of York : M Freeman 2024


Ethernet Protocol



- Ethernet frame : $7 + 1 + 6 + 6 + 2 + 4 = 26B$

University of York : M Freeman 2024

Ethernet Protocol

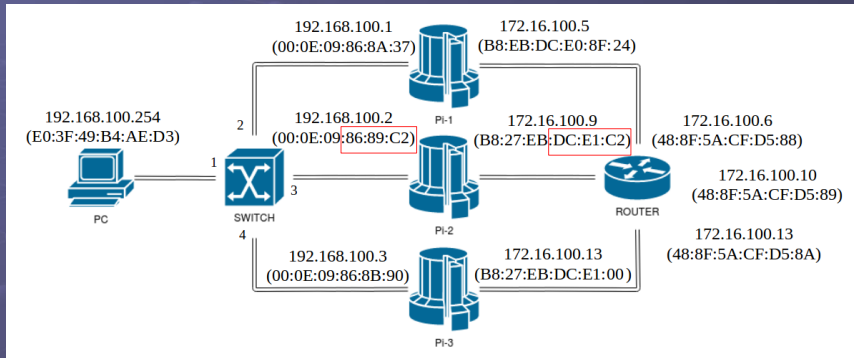


MAC addresses, two parts

- Organisationally Unique Identifier (OUI)
- Network Interface Controller Specific (NIC)

University of York : M Freeman 2024

Ethernet Protocol



- MAC addresses, two parts
 - ▶ Organisationally Unique Identifier (OUI)
 - ▶ Network Interface Controller Specific (NIC)

University of York : M Freeman 2024

Ethernet Protocol

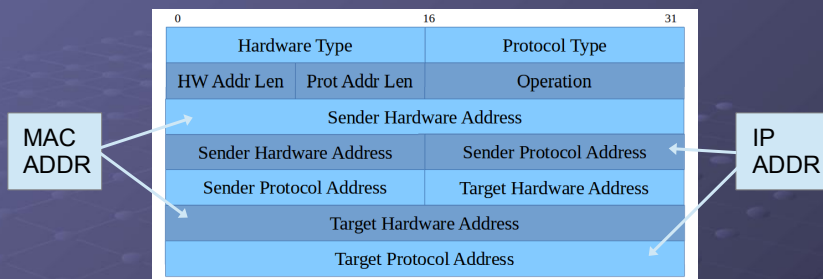
Switch port



- Key points to remember:
 - ▶ Unlike a router, switches know nothing about IP addresses, they only “see” layer 2 protocols i.e. on a router each port is assigned an IP address, on a switch they are not.
 - ▶ Communication between ports is based on MAC addresses i.e. switch learns what host (MAC address) is connected to each port.
- Q: how does a host discover what MAC address is used by another host?

University of York : M Freeman 2024

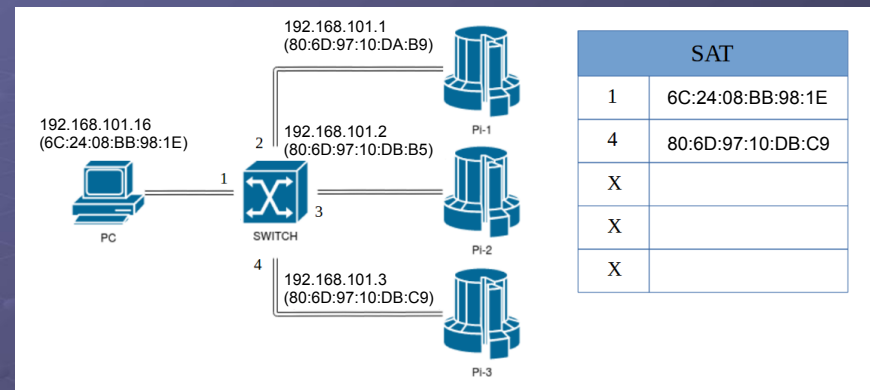
ARP



- Address Resolution Protocol
 - ▶ RFC 826 : <https://datatracker.ietf.org/doc/html/rfc826>
 - ▶ Defined in 1982, used by a number of early protocols e.g. when using IP this protocol allow a host to convert a logical address (IP) into a physical address (MAC).
 - ▶ Uses the broadcast MAC address FF:FF:FF:FF:FF:FF.

University of York : M Freeman 2024

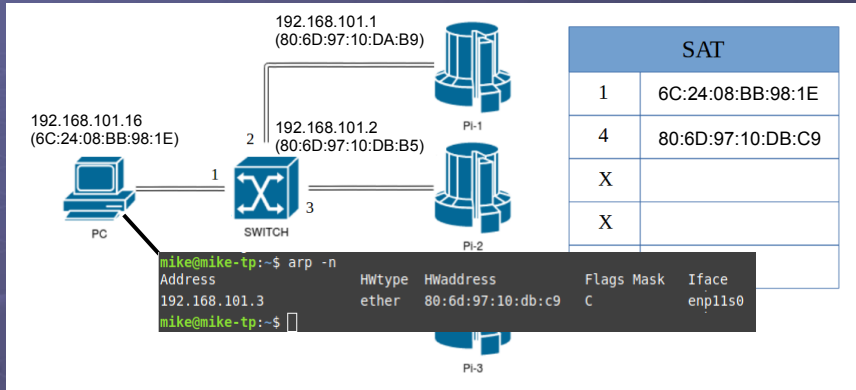
Example



- Source Address Table (SAT)
 - ▶ Mapping port number to MAC address. In this example PC and Pi-3 have communicated previously (ports 1 & 4)

University of York : M Freeman 2024

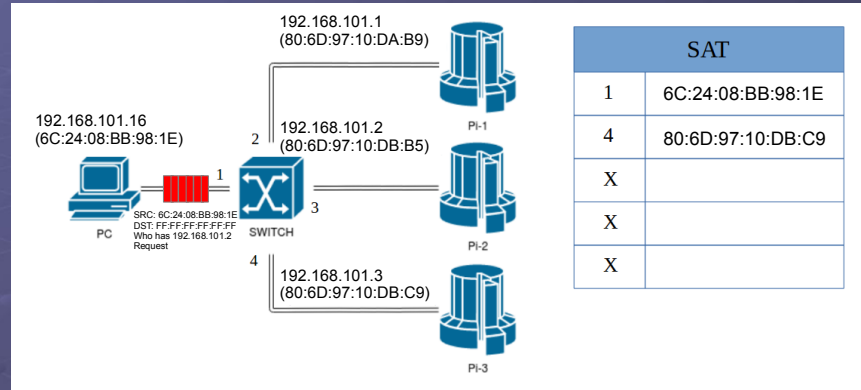
Example



- PC wants to communicate with Pi-2
 - PC wants to communicate with Pi-2, it checks for Pi-2's IP address in its ARP table, MISS :(

University of York : M Freeman 2024

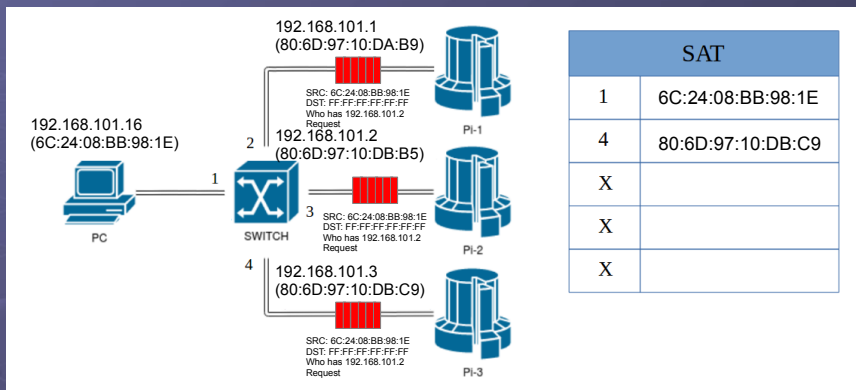
Example



- PC wants to communicate with Pi-2
 - PC knows IP address, doesn't know MAC address. Broadcasts a "Who has"

University of York : M Freeman 2024

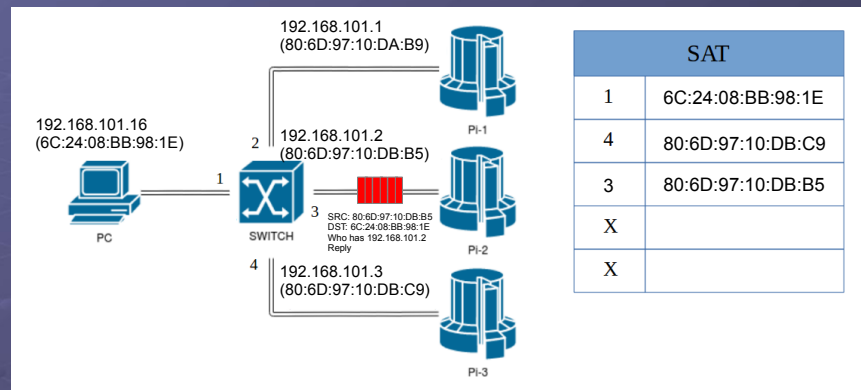
Example



- Broadcast received by switch, no matching MAC in SAT, switch re-broadcast on other ports (not src).

University of York : M Freeman 2024

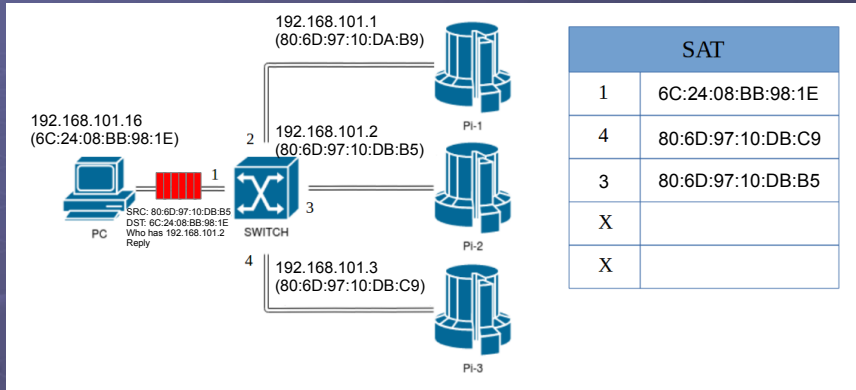
Example



- Pi-2 matches requested IP address, transmits back an ARP response packet via switch.

University of York : M Freeman 2024

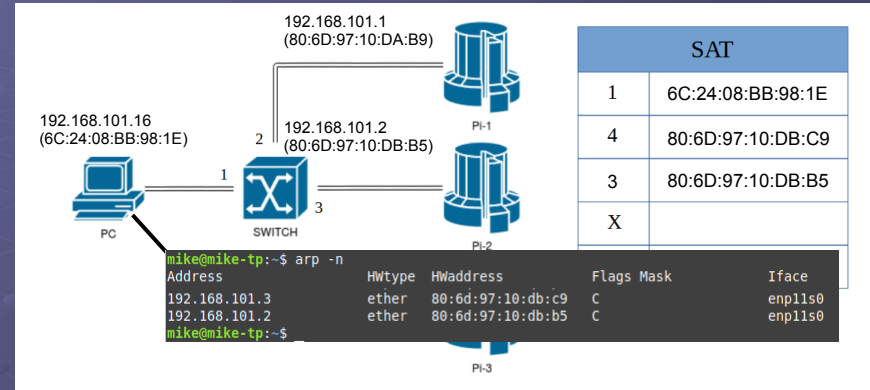
Example



- Destination MAC matched in SAT, packet forwarded to port 1.

University of York : M Freeman 2024

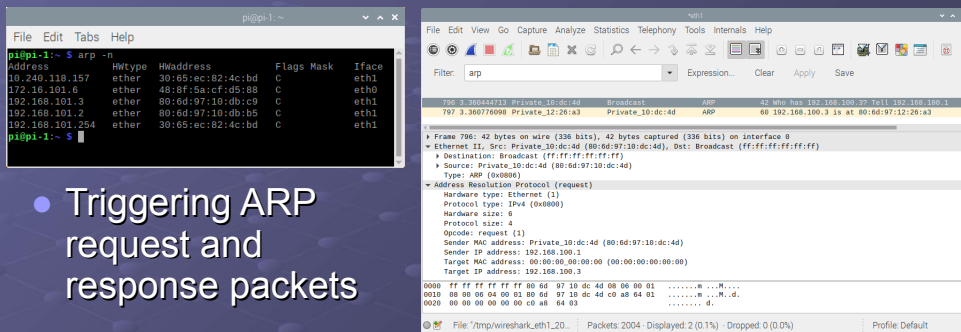
Example



- PC wants to communicate with Pi-2 again
 - Checks for Pi-2's IP address in its ARP table, HIT, use MAC address from table.

University of York : M Freeman 2024

Demo




- Triggering ARP request and response packets

- arp -n
- arp -d <IP>
- ping <IP>

University of York : M Freeman 2024


Quick Quizzz



- To test out these ideas consider what will happen to the Raspberry Pi system if we add a second switch

University of York : M Freeman 2024

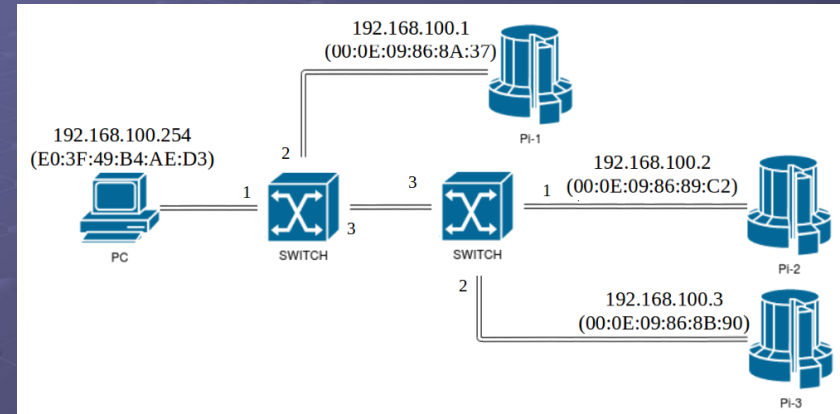
Quick Quizzz



- Move eth1 from Pi-2 and Pi-3 onto second switch + one patch cable between switches.

University of York : M Freeman 2024

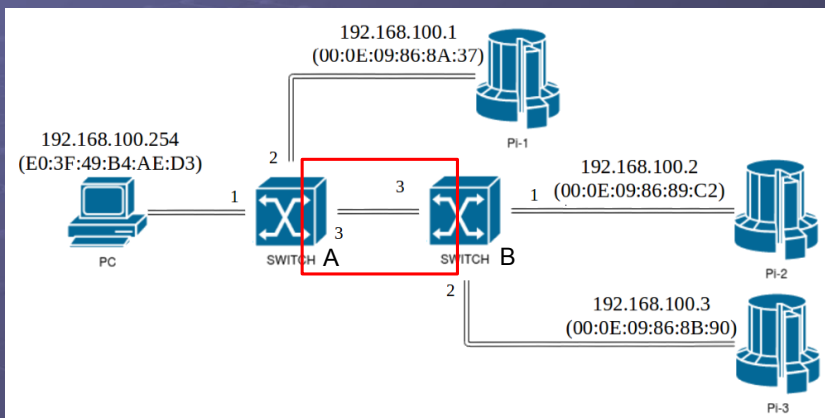
Quick Quizzz



- Q : will this network still work?

University of York : M Freeman 2024


Quick Quizzz



- Q : What will the SAT look like in the two switches (A & B) for port 3 i.e. SWA-p3 and SWB-p3?

University of York : M Freeman 2024

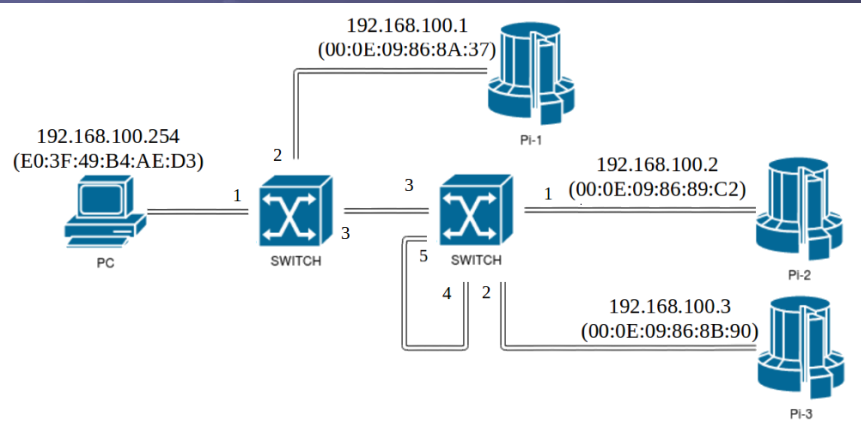
Quick Quizzz



- Move eth1 from Pi-2 and Pi-3 onto second switch + one patch cable between switches + a link between ports.

University of York : M Freeman 2024


Quick Quizzz



- Q : will this network still work?

University of York : M Freeman 2024

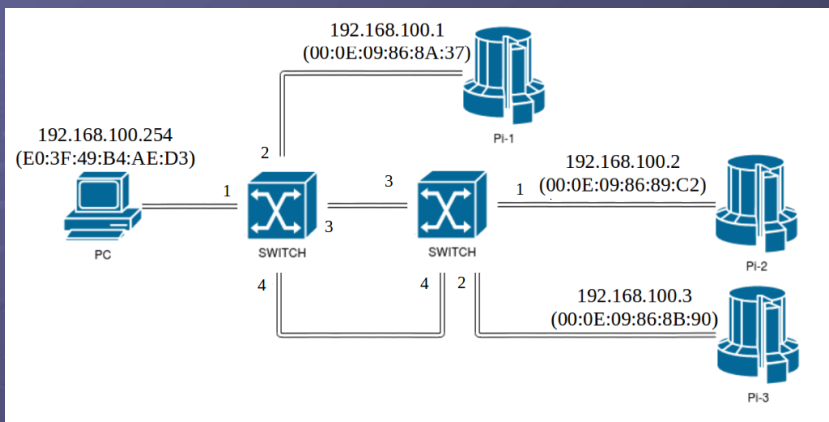
Quick Quizzz



- Move eth1 from Pi-2 and Pi-3 onto second switch + two patch cables between switches.

University of York : M Freeman 2024


Quick Quizzz



- Q : will this network still work?
 - ▶ Note, cheap dumb switches no link aggregation allowed :)

University of York : M Freeman 2024

Quick Quizzz



- Move eth0 from Pi-1, Pi-2 and Pi-3 onto second switch + patch cable between switches
- Q : will the network still work?

University of York : M Freeman 2024

Quick Quizzz

```

pi@pi-2: ~
File Edit Tabs Help
pi@pi-2:~ $ route -n
Kernel IP routing table
Destination      Gateway         Genmask         Flags Metric Ref    Use Iface
0.0.0.0          172.16.101.10  0.0.0.0         UG  0      0      0 eth0
172.16.101.8     0.0.0.0        255.255.255.252 U  0      0      0 eth0
192.168.0.0      0.0.0.0        255.255.0.0     U  0      0      0 eth1
pi@pi-2:~ $

```

- A : no, but perhaps not because of the reason you are thinking of.
- Q : what changes do you need to make to this system to allow it to work correctly?
 - ▶ What routes do you need to delete?
 - ▶ What routes do you need to add?

University of York : M Freeman 2024

Some questions to consider ...

```


pi@pi-1:~ $ ping 10.100.2.2
PING 10.100.2.2 (10.100.2.2) 56(84) bytes of data.
From 172.16.100.6 icmp_seq=1 Destination Net Unreachable
From 172.16.100.6 icmp_seq=2 Destination Net Unreachable
^C
--- 10.100.2.2 ping statistics ---
2 packets transmitted, 0 received, +2 errors, 100% packet loss, time 2ms
pi@pi-1:~ $

```

- Q : what protocol does the Ping command use?
 - ▶ How does Ping know that the destination network is unreachable?
- Q : how are errors signalled on a network?
 - ▶ Destination network, host, protocol, port is unreachable, TTL count expires etc.
- Q : do we need another protocol?

University of York : M Freeman 2024

ICMP



- Internet Control Message Protocol
 - ▶ RFC792 : <https://datatracker.ietf.org/doc/html/rfc792>
 - ▶ Defined in 1981, like UDP designed to be a very light weight protocol.
 - ▶ ICMP uses the IP protocol to transport its data, however, unlike TCP/UDP it is not used to transfer “user” data (segments). Therefore, as this protocol was designed alongside IP it is considered a layer 3 protocol (network layer), could argue layer 4, but it feels like a layer 3.5 :).

University of York : M Freeman 2024

ICMP

Type	Code	Status	Description
0 - Echo Reply ^{[6]:14}	0		Echo reply (used to ping)
1 and 2		unassigned	Reserved
	0		Destination network unreachable
	1		Destination host unreachable
	2		Destination protocol unreachable
	3		Destination port unreachable
	4		Fragmentation required, and DF flag set
	5		Source route failed
	6		Destination network unknown
	7		Destination host unknown
3 - Destination Unreachable ^{[6]:4}	8		Source host isolated
	9		Network administratively prohibited
	10		Host administratively prohibited
	11		Network unreachable for ToS
	12		Host unreachable for ToS
	13		Communication administratively prohibited
	14		Host Precedence Violation
	15		Precedence cutoff in effect

- ICMP types and codes

University of York : M Freeman 2024

ICMP

Type	Code	Status	Description
4 - Source Quench	0	deprecated	Source quench (congestion control)
	0		Redirect Datagram for the Network
5 - Redirect Message	1		Redirect Datagram for the Host
	2		Redirect Datagram for the ToS & network
	3		Redirect Datagram for the ToS & host
6		deprecated	Alternate Host Address
7		unassigned	Reserved
8 - Echo Request	0		Echo request (used to ping)
9 - Router Advertisement	0		Router Advertisement
10 - Router Solicitation	0		Router discovery/selection/solicitation
11 - Time Exceeded ^[6]	0		TTL expired in transit
	1		Fragment reassembly time exceeded
	0		Pointer indicates the error
12 - Parameter Problem: Bad IP header	1		Missing a required option
	2		Bad length
13 - Timestamp	0		Timestamp
14 - Timestamp Reply	0		Timestamp reply
15 - Information Request	0	deprecated	Information Request
16 - Information Reply	0	deprecated	Information Reply

ICMP types and codes

University of York : M Freeman 2024

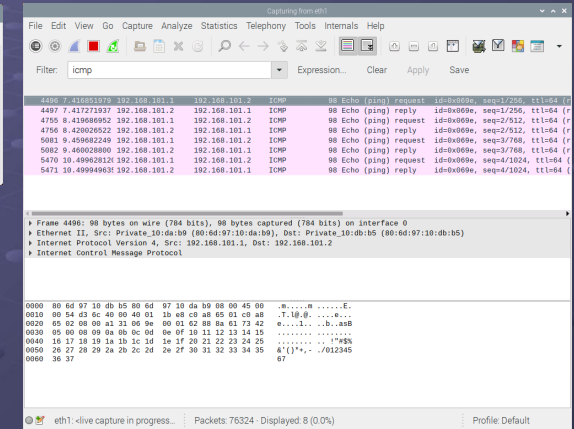
Demo

```

pi@pi-1:~$ ping -c 4 192.168.101.2
PING 192.168.101.2 (192.168.101.2) 56(84) bytes of data:
64 bytes from 192.168.101.2: icmp_seq=1 ttl=64 time=0.472 ms
64 bytes from 192.168.101.2: icmp_seq=2 ttl=64 time=0.393 ms
64 bytes from 192.168.101.2: icmp_seq=3 ttl=64 time=0.305 ms
64 bytes from 192.168.101.2: icmp_seq=4 ttl=64 time=0.352 ms
--- 192.168.101.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 80ms
rtt min/avg/max/mdev = 0.352/0.403/0.472/0.043 ms
pi@pi-1:~$

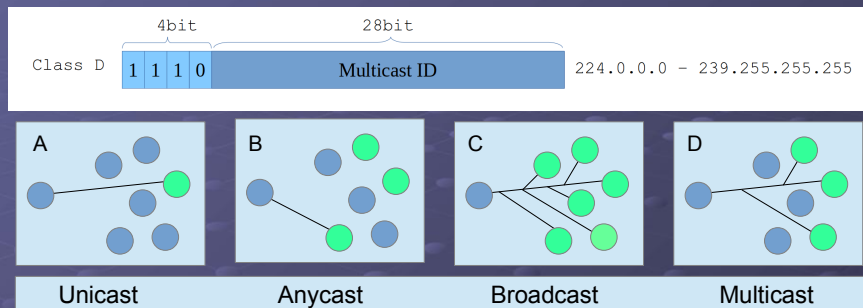
```

- ICMP ping packets
- More than just PING e.g. try running your `udpTX.py` without `udpRX.py` ICMP will tell the TX host to stop.



University of York : M Freeman 2024

Multicast



- What happens in layer 2 when we use multicast?
- Transfers are now not to a single destination MAC?
 - ▶ We need one-to-many communications.
- Note, multicast address blocks allocated to specific functions e.g. local multicast : 224.0.0.0 to 224.0.0.255

University of York : M Freeman 2024

Multicast

```

Terminal - pi@pi-1: ~
pi@pi-1:~$ ping raspberrypi-mail.local
PING raspberrypi-mail.local (192.168.100.2) 56(84) bytes of data:
64 bytes from raspberrypi-mail.local (192.168.100.2): icmp_seq=1 ttl=64 time=0.813 ms
64 bytes from raspberrypi-mail.local (192.168.100.2): icmp_seq=2 ttl=64 time=0.429 ms
64 bytes from raspberrypi-mail.local (192.168.100.2): icmp_seq=3 ttl=64 time=0.481 ms
--- raspberrypi-mail.local ping statistics ---
7 packets transmitted, 7 received, 0% packet loss, time 204ms
rtt min/avg/max/mdev = 0.421/0.509/0.813/0.126 ms
pi@pi-1:~$

```

- Easy way of generating “multicast” packets is mDNS
 - ▶ Multicast Domain Name System (mDNS)
 - ▶ RFC 6762 : <https://datatracker.ietf.org/doc/html/rfc6762>
 - ▶ Configuration free, resolves hostnames to IP addresses in small networks, common implementations : Bonjour & Avahi
 - ♦ `.local` is the local area network (local link)
 - ▶ IP address : IPv4 – 224.0.0.251 and IPv6 – ff02::fb
 - ▶ Port : 5353

University of York : M Freeman 2024

Multicast

The screenshot shows a Wireshark capture on interface eth1. The filter is 'udp.port==5353'. The selected packet is a DNS query from 192.168.100.2 to 224.0.0.251 (multicast). The query asks for the IP address of raspberrypi-mail.local. The response packet shows the answer: raspberrypi-mail.local is at 192.168.100.2.

- To resolve .local names Pi-1 puts two multicast packets onto the network i.e. an IPv4 and IPv6 multicast DNS query.

University of York : M Freeman 2024

Multicast

This screenshot is similar to the one on the left, showing a DNS query and response. It highlights the multicast address 224.0.0.251 used for local network discovery.

- The Raspberry Pi mail server matches the query to its host name and responds with its IP address.
 - Q : why would ping pi-1.local cause issues?

University of York : M Freeman 2024

Multicast

Two screenshots of Wireshark. The top one shows a query from 192.168.100.2 to 224.0.0.251. The bottom one shows the response from 192.168.100.2 back to 224.0.0.251, indicating the IP of raspberrypi-mail.local.

- To implement one-to-many connections the Ethernet protocol defines multicast MAC addresses for IPv4 and IPv6.

University of York : M Freeman 2024

Example : video streams

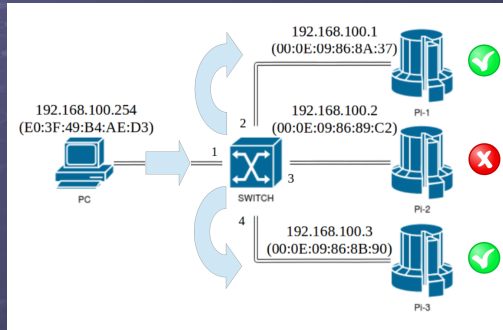
The top part shows a terminal window displaying the output of the 'ffprobe' command on a video stream. It shows details like codec (h264), resolution (640x360), and frame rate (30 fps). The bottom part shows a video player window displaying a stream of jellyfish.

- Example : in the lab, a stream of jellyfish :)

University of York : M Freeman 2024

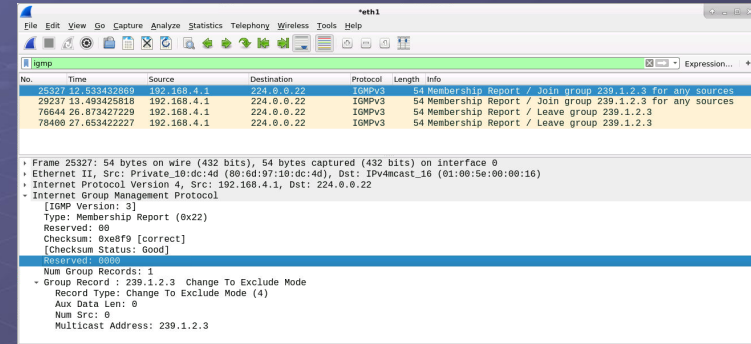
Multicast

- When using multicast we only want packets to go to hosts that want them
 - ▶ Not a broadcast.
- Clients wanting to RX multicast packets will listen on the multicast MAC address.
 - ▶ Do not TX on this MAC
- Therefore, what will the switch do? What is a switch's default behaviour when it can not find a MAC address in its SAT?



University of York : M Freeman 2024

Multicast



- Internet Group Management Protocol (IGMP)
 - ▶ Membership reports : join and leave messages
 - ▶ RFC 2236 : <https://datatracker.ietf.org/doc/html/rfc2236>
 - ▶ Need a switch that supports IGMP snooping :(

University of York : M Freeman 2024

Summary

- We have now finished looking at network protocols
 - ▶ There are lot more protocols to look at and the ones we have looked at we have only scratched the surface e.g. DNS, TCP, IPv6, routing protocols ...
 - ▶ BUT you should now have a basic understanding of the Internet protocol stack and the different protocols commonly used.
- However, we still have some unanswered questions:
 - ▶ What cables do these signals travel across, what do the electrical signals look like ...
 - ▶ Are all Ethernet cables the same?

University of York : M Freeman 2024