NETWORKING

UNIVERSITY OF THE
FREE STATE
UNIVERSITEIT VAN DIE
VRYSTAAT
YUNIVESITHI YA
FREISTATA
UFS

OPEN SYSTEMS INTERCONNECTION (OSI) MODEL

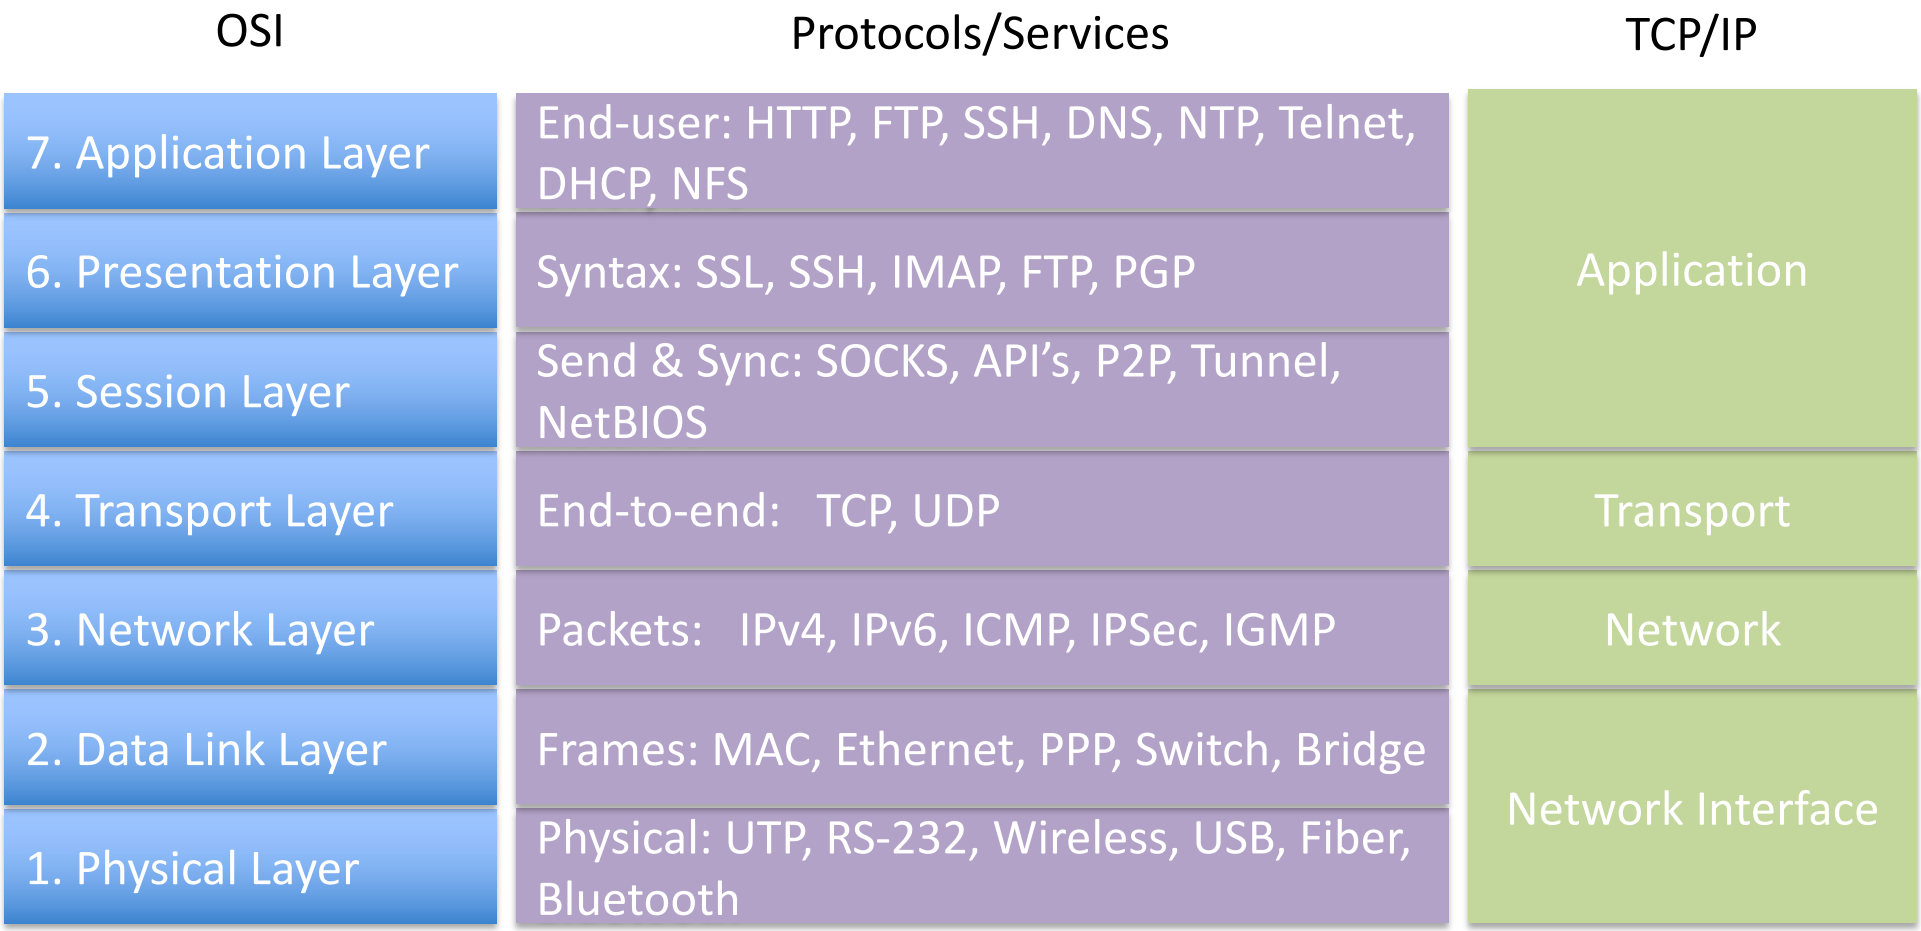

UFS

OSI 1: PHYSICAL LAYER

- Cat 5/6/7/8 Cable
- RS 232
- SFP+
- SFP28
- QSFP28
- InfiniBand
- OmniPath
- Optical Fiber (LC, SC, ST)

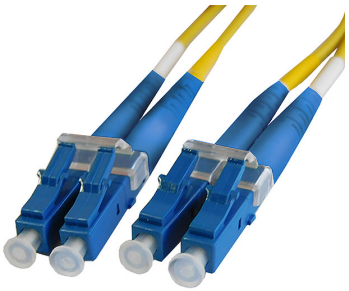

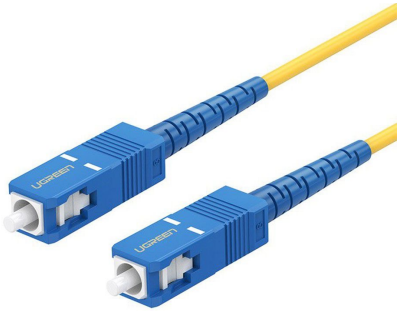

UFS

OSI 2: DATA LINK LAYER

- **M**edia **A**ccess **C**ontrol (MAC)
  - Lowest level of communication between network devices
  - IEEE 802
- MAC Address
  - Hexadecimal device identifier, unique per vendor and per device. For instance:

    Intel: **f8:16:54**:aa:bb:cc
  - Used to translate IP to MAC
  - Can be used by DHCP to provide specific IP Address
    - Broadcast is sent to ff:ff:ff:ff:ff:ff

UFS

OSI 2: DATA LINK LAYER – ARP TABLE

[hpc@node0501 ~]$ **arp -an**
? (172.31.1.44) at 0c:c4:7a:59:9e:cc [ether] on eno1
? (172.31.1.30) at 0c:c4:7a:88:fe:d9 [ether] on eno1
? (172.31.1.200) at 0c:c4:7a:1d:af:48 [ether] on eno1
? (172.31.1.254) at 00:25:90:5d:e5:37 [ether] on eno1
? (172.31.1.54) at 0c:c4:7a:94:a8:f6 [ether] on eno1

[hpc@login ~]$ **arp -an**
? (196.21.184.46) at 52:54:00:f3:7c:99 [ether] on bond0.350
? (172.31.1.74) at ac:1f:6b:8a:3b:14 [ether] on bond0.172
? (172.31.1.60) at 0c:c4:7a:94:45:a6 [ether] on bond0.172
? (172.31.1.49) at 0c:c4:7a:94:46:28 [ether] on bond0.172
? (196.21.184.22) at 00:25:90:5d:e5:37 [ether] on bond0.350
? (172.31.1.51) at 0c:c4:7a:94:46:56 [ether] on bond0.172
? (172.31.1.44) at 0c:c4:7a:59:9e:cc [ether] on bond0.172
? (172.31.1.84) at <incomplete> on bond0.172

? – Hostname (ignored due to -n option)

Q: What would you guess does the <incomplete> mean?

UFS

OSI 2 & 3: DATA LINK LAYER TO NETWORK LINK

- Layer 2 switches:
  - Can send/receive traffic only on the same subnet
  - Learn and remember address tables (ARP table)
  - Can support VLAN (802.1q)

- Layer 3 switches:
  - Can route VLAN traffic
  - Can route traffic between subnets, if the client device has a gateway set

- Q: A hub did not keep an address table.
  - How was traffic managed?

UFS

IP ADDRESSES

- Managed by Internet Assigned Numbers Authority (IANA)
- Each country/continent has a registration body
  - Regional Internet Registry

- Africa: AFRINIC
  - Situated in Mauritius

UFS

IPv4 PACKET

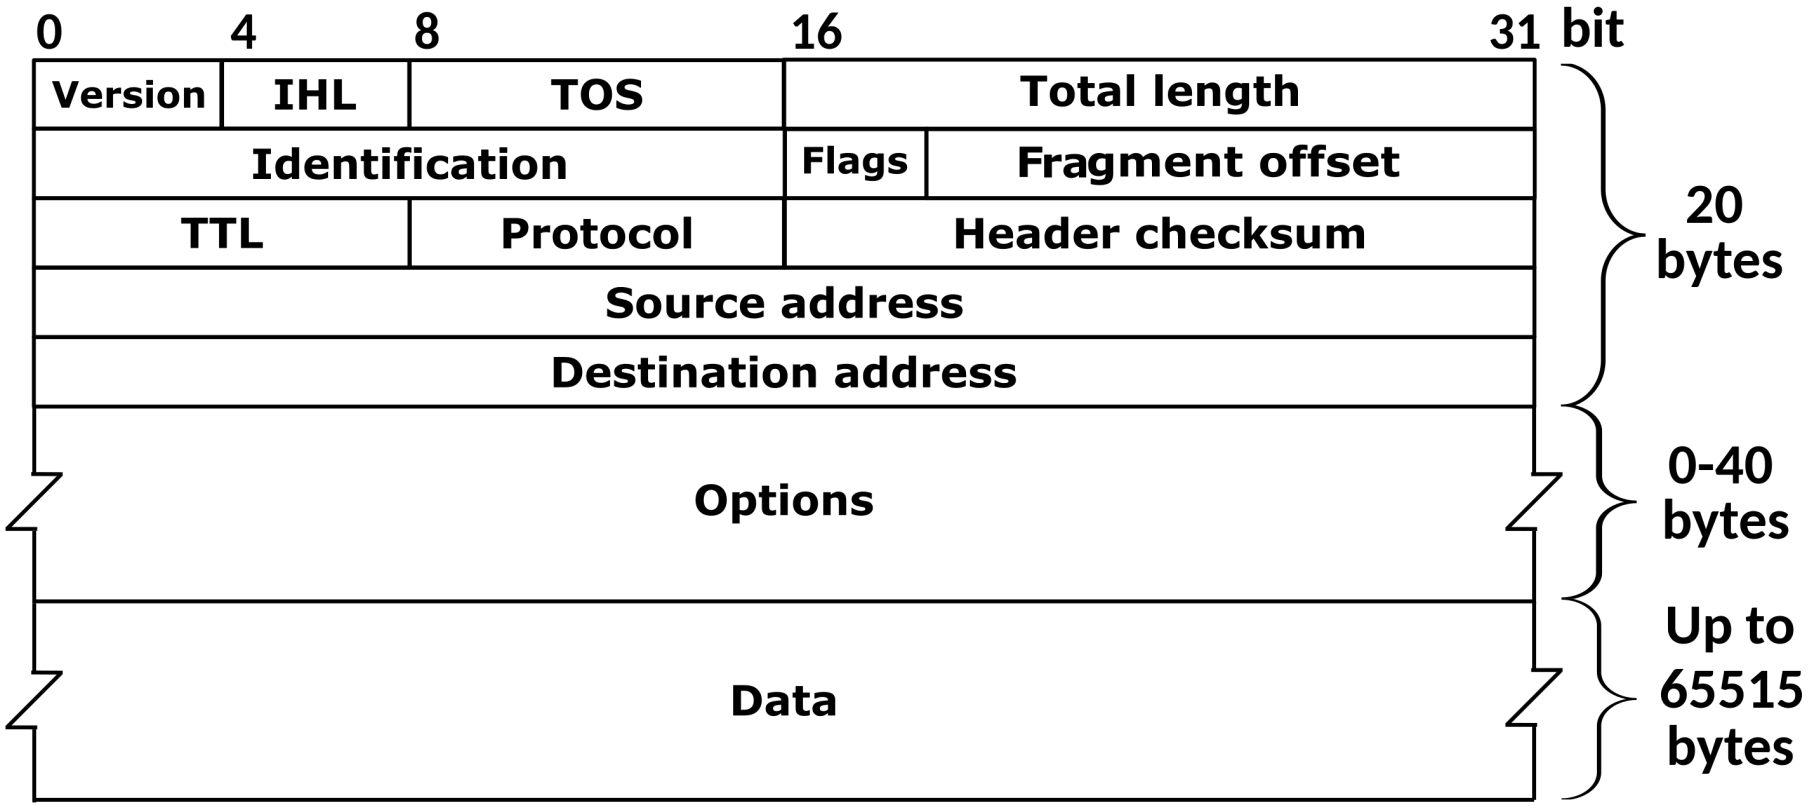

Internet Header Length (IHL): Size of the header
Type Of Service (TOS): Real-time data streaming such as VoIP
TTL – Time to live: Router subtracts one; if zero, ICMP Time-out
Protocol: RFC790 Number representing the protocol
Source & Destination: 32Bit address may be affected by NAT

UFS

IP ADDRESSING: IPv4

Format: xxx . xxx . xxx . xxx
Example: 142 . 251. 47 . 163

Four groups (octets) of numbers ranging between 0-255
First group > 0
Last group < 255 and usually > 0

Group a.k.a. Octet (referring to **eight** bits, in binary)
Binary: 11111111.11111111.**11111100**.00000010
Decimal: 255 . 255 . **252** . 2

The IPv4 (**32**bit) address space allows for a total of:
2^{32} addresses = 4 294 967 296
(That is about “only“ half of the earth’s population – if everyone had a cellphone, only half would be working at a given time)

Q: If only ± 4bn IPv4 addresses exist;
How can we all work on the Internet simultaneously?

UFS

IPv4 ADDRESS

- IP Address can be broken up into two parts: Network and a host.
- The network part remains static, and the host part uniquely changes for each host in the network.
- For instance, a network with the IP ranges:

  176.22.254.1 to **176.22.254.**255

- In this example, the first three bold octets are the **network** address, and the last (underlined) octet is the host address
- Binary representation:

176.22.254.1	**10110000 . 00010110 . 11111110 .**	00000001
176.22.254.2	**10110000 . 00010110 . 11111110 .**	00000010
	To:	
176.22.254.254	**10110000 . 00010110 . 11111110 .**	111111110
176.22.254.255	**10110000 . 00010110 . 11111110 .**	111111111
	Network	Host

UFS

IPv4 SUBNET MASK

- The subnet determines which part of the address is the network part and which is the host part.
- Let's look at an IP address in our previous IP range:

	176	**22**	**254**	42
Binary:	**10110000**	**00010110**	**11111110**	00101010
Subnet Mask:	**11111111**	**11111111**	**11111111**	**00000000**
Mask in decimal:	**255**	**255**	**255**	**0**

In this example, the first three octets are masked

(values of one in binary), and where the first zero occurs (in binary), it indicates where the host address starts (24th position)

UFS

IPv4 SUBNET MASK ...CONTINUED

	176	**22**	**254**	42
	10110000	00010110	11111110	00101010
Subnet Mask:	11111111	11111111	11111111	**00000000**
Subnet in decimal:	**255**	**255**	**255**	**0**

A subnet can be other values than 255 and 0, for instance:

	10110000	00010110	11111**10**	00101010
Subnet Mask:	**11111111**	**11111111**	**11111100**	**00000000**
Mask in decimal:	**255**	**255**	**252**	**0**

Note that the subnet mask has shifted two bits to the left, and thus the original network address will also have to shift two bits:

10110000.00010110.11111**10**.00101010 **becomes:**

10110000.00010110.11111**00**.00101010

This subnet mask gives us IP addresses in the range from:

176.22.252.1 to 176.22.255.254 (1024 hosts)

UFS

CLASSLESS INTER-DOMAIN ROUTING (CIDR)

IPv4 Classes: A . B . C . __

Class	CIDR notation	IP Range	IP Addresses
A	**123**.0.0.0/8	123.0.0.1 – 123.255.255.254	16 777 214
B	**123.44**.0.0/16	123.44.0.1 – 123.44.255.254	65 534
C	**123.44.55**.0/24	123.44.55.1 – 123.44.55.254	254

- **Classless**: Not providing an entire class A, B or C to an ISP, but rather subdividing (subnetting)
- **Inter-Domain Routing**:
  - The internal switches pass all addresses inside the subnet
  - All addresses outside the subnet are sent to the router/gateway

UFS

IPv4 CLASSLESS INTER-DOMAIN ROUTING (CIDR)

- CIDR notation is a short format for writing a subnet mask

	176	**22**	**254**	42
	10110000	00010110	11111110	00101010
Subnet Mask:	**11111111**	**11111111**	**11111111**	00000000
Subnet in decimal:	**255**	**255**	**255**	**0**

The basic concept is the number of ones (in binary) before a zero occurs in a subnet.

For the IPv4 (**32**bit) address space, the value will be 0 – **32**

If we count the above-mentioned binary ones, it equals **24**
The CIDR notation for this address will therefore be:

176.22.254.42/24

UFS

CIDR: ANOTHER EXAMPLE

	176	.	22	.	**252**	.	1
Subnet Mask:	**11111111**	**.**	**11111111**	**.**	**11111100**	**.**	**00000000**
Mask in decimal:	**255**	**.**	**255**	**.**	**252**	**.**	**0**

- Counting the ones: **22**
- CIDR notation will be: 176.22.252.1/**22**

UFS

CIDR - ROUTING

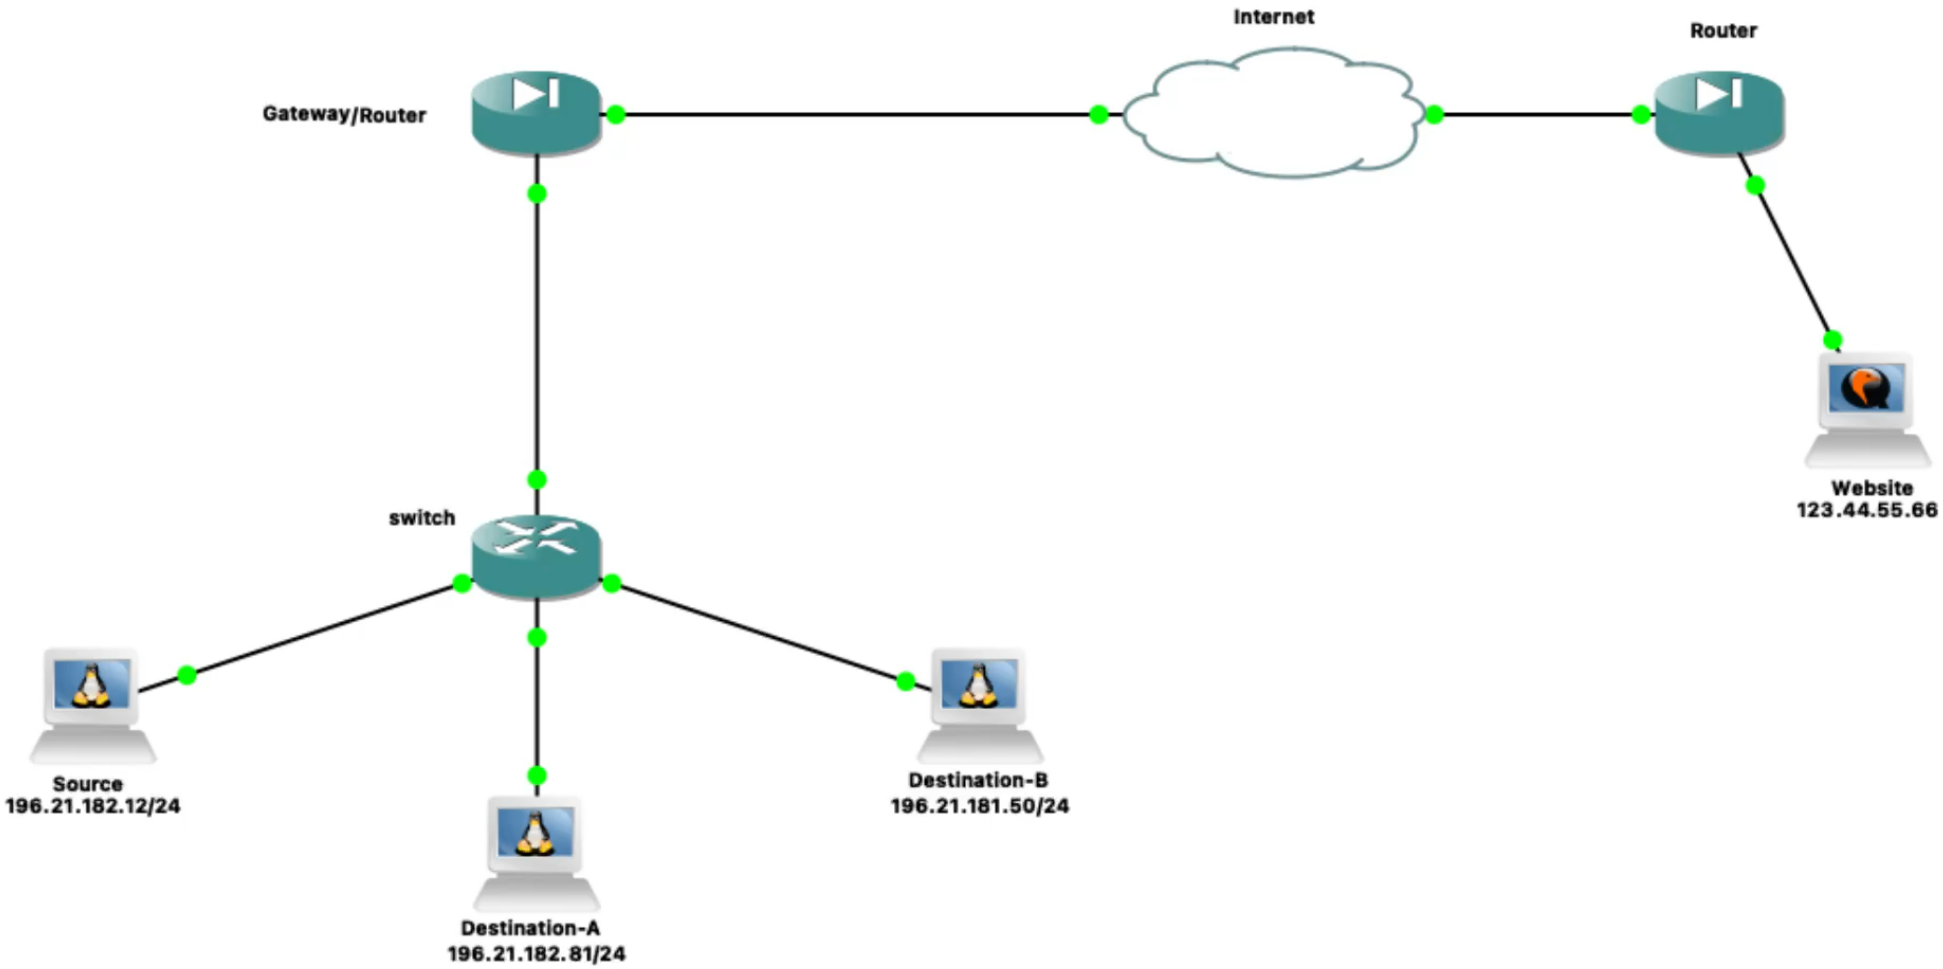

IPv4 …CONTINUED

Private Addresses (RFC1918)

Class	CIDR notation	IP Range	IP Addresses
A	10.0.0.0/8	10.0.0.1 – 10.255.255.254	16 777 214
B	192.168.0.0/16	192.168.0.1 – 192.168.255.254	65 534
C	172.16.0.0/12	172.16.0.1 – 172.31.255.254	1 048 574

Private addresses can't be routed on the Internet

Also known as Bogon networks

About 18 million IP addresses are reserved for private IPs, and another 270 million for multicast addresses

UFS

IP ADDRESSING: IPv6

Format: xxxx : xxxx : xxxx : xxxx : xxxx : xxxx : xxxx : xxxx
Example: fe80::f8f6:38ff:fe0e:505f

Eight groups of four (quartet) hexadecimal (16bit) values 0-ffff (between 0 and 65 535)
First group(s) can be 0
Last group < ffff and usually > 0

A double :: notates that one or more adjacent groups are zero
Only one double :: is allowed
If more than one group of adjacent zero groups exist,
the leftmost zeros are replaced by the double ::

Example:
2201:0df8:**0000**:**0000**:64f4:0000:0000:43df becomes

2201:0df8::64f4:0000:0000:43df or even more simplified:

2201:0df8::64f4:0:0:43df

The **128**bit address space, allows for a total of:
2^{128} addresses = 3.4 x 10^{38} = ±340 undecillion (trillion trillion trillion) addresses:
340,282,366,920,938,000,000,000,000,000,000,000,000

UFS

IPv6 SUBNETTING

	2201	:	0df8	::	64f4	:0:0:	43df					
Binary:	0010001000010001	:	0000110111111100	::	0110010011110100	:0:0:	0100001111011111					
Decimal:	8705	:	3576	::	25844	:0:0:	17375					

UFS

DON'T PANIC

UFS

IPv4 ADDRESSES THAT ARE OFTEN OBSERVED

CIDR notation	IP Range	IP Addresses	Notes
0.0.0.0/0	0.0.0.0 – 255.255.255.254	4 294 967 296	Internet (anything outside of "This" network)
127.0.0.1/8	127.0.0.1 – 127.255.255.254	16 777 214	Local host (loopback) and local networks
10.0.0.0/8	10.0.0.1 – 10.255.255.254	16 777 214	**Private IPs** Not routable
172.16.0.0/12	172.16.0.1 – 172.31.255.254	1 048 574	
192.168.0.0/16	192.168.0.1 – 192.168.255.254	65 534	
0.0.0.0/8	0.0.0.0 – 0.255.255.254	16 777 214	"This" network
224.0.0.0/4	224.0.0.1 – 239.255.255.254	268 435 456	Multicast
255.255.255.255/32	255.255.255.255	1	Limited Broadcast

Bold ones are of interest to us

UFS

IPv4 ADDRESSING SUMMARIZED

- The network part of the address is static:
  - **A.** **172**.15.81.88/8
  - **B.** **172.15**.81.88/16
  - **C.** **172.15.81**.88/24

- The subnet determines the host part of the address
- The CIDR notation is shorthand to specify the subnet
- The CIDR notation indicates the number of ones (in binary), before the first zero occurs (network) in the subnet mask
- The "smaller" the value of the CIDR, the larger the network
- We will only focus on /32, /24, /16, and /8 networks
  - Keep subnetting simple

UFS

Q&A

- In an ARP table:
  - What would you guess does the **<incomplete>** mean?
  - The host did not respond to ICMP requests and is deemed offline

- A hub did not keep an address table. How was traffic managed?
  - Traffic was not managed.
  - Each packet was sent to all ports on the hub.

- If only ± 4bn IPv4 addresses exist; how can we all work on the Internet simultaneously?
  - CIDR was designed to provide Public IPs to only ISPs.
    - An ISP would make use of NAT to provide IPs to clients.
  - IPv6 allows for enough routable IP addresses over the Internet

NETWORK EXERCISE

- Execute the following:

  dnf -y install net-utils telnet \

  bind-utils nmap traceroute curl whois

- What is your GNU Linux computer's IPv4 address?
- What is your GNU Linux computer's MAC address?
- How would you debug/check if your network is working?
- What is the IP address for www.google.com?
- Which ports are open on 8.8.8.8?
- Can you connect to 8.8.8.8 on port 80?

UFS

NETWORK EXERCISE

- What is the private/local IP address of your gateway?
- In a terminal:
  - Determine what is the Public IP of your gateway
  - Determine what is the city of your ISP's IP
  - Determine what is the country of your ISP's origin

- What are the name servers for ufs.ac.za?
- Who is the domain administrator(s) for ufs.ac.za?
- What are the mail exchange servers for ufs.ac.za?
- Does the ufs.ac.za domain have any SPF records?
- Demo: Use telnet to send an email

UFS