

JOB-READY SERIES · COMPLETE EDITION

PACKET EDITION

NET WORKS

7 Application

6 Presentation

5 Session

4 Transport


3 Network

2 Data Link

1 Physical

Computer Networks · From OSI to Offer Letter

Everything interviews ask and real jobs use: OSI, TCP/IP, subnetting, routing, DNS, HTTP, TLS, security and troubleshooting, with worked numericals and 100+ questions with answers.



OSI Model

TCP / UDP

Subnetting

DNS

HTTP/3

TLS

BGP

Routing

ARP

NAT

Wireshark

Firewalls

WRITTEN BY

Shreyash Tripathi

15 Chapters · Interview Corner in every chapter

Edition 2026

JOB-READY SERIES • COMPLETE EDITION



Computer Networks

Job-Ready Guide for developers preparing for their next job



WRITTEN BY

Shreyash Tripathi

Complete Edition, 2026

© Shreyash Tripathi. All rights reserved.

Preface

You already have one job. You write code, you push to production, and sometimes things break in ways nobody explained in college. Now you want the next job, and the interviews ask the same "basic" Computer Networks questions again, except now the interviewer also expects you to connect them to real work. This book is written for exactly that moment.

It follows the style of a school textbook on purpose: numbered chapters, simple language, worked examples step by step, a check with answers after every topic, and full exercises with an answer key at the end of every chapter. Every chapter also has an **Interview Corner** with questions interviewers really ask, and an **At Work** box that shows where the concept appears in a real job. There is humour too, because a joke is easier to remember than a definition.

Read one chapter a day, solve the questions before you look at the answers, and revise the Interview Corners the night before your interview. You will walk in prepared.

Shreyash Tripathi

Features of Every Chapter

DO YOU KNOW?

Real facts and history.

FUNNY SIDE

A joke that makes the idea stick.

KEY POINT

The line to remember.

INTERVIEW CORNER

Questions interviewers really ask, with model answers.

AT WORK

How the concept shows up in a real job.

COMMON MISTAKE

Traps that cost marks and cause bugs.

CHECK YOUR PROGRESS

Short questions with answers after every topic.

Each chapter ends with **What Have We Learnt?**, **Key Terms**, **Exercises** and a complete **Answer Key**. The last chapter is a full **Interview Question Bank**.

Contents

PART I: FOUNDATIONS

1	Introduction to Computer Networks	6
2	Network Models: OSI and TCP/IP	28
3	Physical Layer, Transmission Media and Switching	46
4	Data Link Layer	67

PART II: NETWORK LAYER

5	Network Layer I: IP Addressing and Subnetting	89
6	Network Layer II: Routing and Supporting Protocols	114

PART III: TRANSPORT AND APPLICATIONS

7	Transport Layer: TCP and UDP	137
8	Application Layer I: DNS and the Web	160
9	Application Layer II: Email, File Transfer and Remote Access	183

PART IV: SECURITY AND MODERN NETWORKS

10	Network Security	207
11	Wireless, Mobile and Modern Networks	229
12	Data Center and Cloud Networking	251

PART V: PRACTICE AND INTERVIEWS

13	Network Troubleshooting and Tools	275
14	Putting It Together: What Happens When You Type a URL, and Socket Programming	298
15	Computer Networks Interview Question Bank	320

APPENDICES

A	Ports and Protocols Cheat Sheet	346
B	Subnetting Quick Reference	350
C	Glossary	355
D	Sample Question Paper with Answers	360

PART I

Foundations

What a network is, the layered models, and how bits travel.

CHAPTER 1

Introduction to Computer Networks

"It works on my machine" is a networking problem waiting to be discovered.

After studying this chapter, you will be able to

- Define a computer network and explain why networks exist, with everyday examples.
- Classify networks as PAN, LAN, MAN and WAN, and tell internet, intranet and extranet apart.
- Compare bus, star, ring, mesh, tree and hybrid topologies with their pros and cons.
- Name common network devices, the OSI layer each works at, and how they split collision and broadcast domains.
- Explain simplex, half duplex and full duplex, and client-server versus peer-to-peer.
- Use the terms bandwidth, throughput, latency and jitter correctly, with numbers.
- Explain who makes networking standards (IEEE, IETF, ISO, ITU) and what an RFC is.

Think about the dabbawalas of Mumbai. Thousands of lunch boxes leave homes every morning, travel by bicycle, Mumbai local and handcart, get sorted at stations using simple coded marks, and reach the right office desk before lunch. Then the empty boxes travel back. Nobody in the chain knows the whole route. Each person only knows the next hop. That, in one paragraph, is how a computer network works.

As a developer, you already use networks every minute: `git push`, an API call from your React app, a database connection from your backend, a Slack message. When something breaks ("connection refused", "timeout", "DNS resolution failed"), the people who fix it fast are the ones who understand what is happening underneath. This chapter builds the vocabulary. Later chapters go deep into each layer.

1.1 What Is a Computer Network?

A **computer network** is a set of two or more devices (called **nodes** or **hosts**) connected by communication links so that they can share data and resources. The links can be wired (copper cable, optical fibre) or wireless (Wi-Fi, mobile radio, satellite).

Every network has five basic parts:

1. **Sender:** the device that sends the message (your laptop).
2. **Receiver:** the device that receives it (a web server).

3. **Message:** the data itself (an HTTP request, an email, a video frame).
4. **Transmission medium:** the physical path (Wi-Fi radio waves, fibre, copper).
5. **Protocol:** the set of rules both sides agree on (HTTP, TCP, IP, Ethernet).



Figure 1.1: The basic communication model. Protocols govern every arrow.

1.1.1 Why do we need networks?

Before networks, data moved on floppy disks carried by hand (people jokingly called this "sneakernet"). Networks solve real problems:

- **Resource sharing:** one printer for a whole office, one database for many app servers.
- **Communication:** email, chat, video calls, WhatsApp.
- **Information access:** the web, search, online banking.
- **Reliability:** data copied to several machines (replicas) survives when one dies.
- **Scalability and cost:** cloud computing lets you rent servers in a data centre instead of buying them.
- **Commerce and services:** UPI payments, IRCTC bookings, food delivery apps.

💡 DO YOU KNOW?

The ancestor of the Internet, ARPANET, sent its first message in October 1969 between UCLA and Stanford Research Institute. The plan was to type "LOGIN". The system crashed after "LO", so the first message on the network that became the Internet was "LO".

✅ CHECK YOUR PROGRESS 1.1

1. Name the five components of a data communication system.

Answer: Sender, receiver, message, transmission medium and protocol.

2. Give two reasons why companies build networks.

Answer: To share resources (printers, databases, storage) and to communicate (email, chat). Reliability through replication is another strong reason.

3. What is a protocol?

Answer: A protocol is an agreed set of rules about message format, order and actions, so that two devices can understand each other. HTTP, TCP and IP are examples.

1.2 Types of Networks by Size

Networks are classified mainly by the geographical area they cover. Think of it as the difference between talking inside your room, inside your building, across your city and across the country.

1.2.1 PAN (Personal Area Network)

A PAN connects devices around one person, usually within about 10 metres. Your phone talking to Bluetooth earbuds, a smartwatch syncing with your phone, or a laptop tethered over USB are PANs. Technologies: Bluetooth, USB, NFC, Zigbee.

1.2.2 LAN (Local Area Network)

A LAN covers a small area such as a home, office floor, school or one building. It is usually owned and managed by one organisation. LANs are fast (1 Gbps to the desk is common, 10 Gbps or more between switches), have low latency (well under a millisecond inside a building) and low error rates. Technologies: Ethernet (IEEE 802.3) and Wi-Fi (IEEE 802.11). A Wi-Fi LAN is often called a **WLAN**.

1.2.3 MAN (Metropolitan Area Network)

A MAN covers a city or a large campus, roughly up to a few tens of kilometres. Examples: a cable TV operator's network across a city, a bank connecting all its branches in Pune with fibre, or a city-wide Wi-Fi project. It is usually run by a service provider or a large organisation.

1.2.4 WAN (Wide Area Network)

A WAN covers a country, a continent or the whole world. It connects many LANs and MANs. WAN links are often leased from telecom companies (MPLS circuits, leased lines, undersea cables). WANs have higher latency (tens to hundreds of milliseconds) and are more expensive per bit. The Internet is the largest WAN. A company with offices in Bengaluru, Mumbai and Singapore connecting them is building a WAN. Modern companies often use **SD-WAN** (software-defined WAN) to manage these links.

Feature	PAN	LAN	MAN	WAN
Area	About 10 m	Room to building/campus	City	Country to world
Ownership	One person	One organisation	Provider or large org	Many providers
Speed	Low to moderate	High	Moderate to high	Varies, costly
Latency	Very low	Very low	Low	Higher
Example	Bluetooth earbuds	Office Ethernet	City cable network	The Internet

1.2.5 Internet, intranet and extranet

- **Internet** (capital I): the global public network of networks that uses the TCP/IP protocol suite. Anyone can join through an ISP.
- **internet** (small i): any interconnection of two or more networks. The Internet is the biggest one.
- **Intranet**: a private network inside an organisation that uses Internet technologies (web servers, browsers, HTTP) but is accessible only to employees. Example: your company's HR portal, internal wiki or Jira that only opens on the office network or VPN.
- **Extranet**: an intranet that is partly opened to selected outsiders such as vendors, partners or dealers, with login. Example: a car maker's portal where dealers check stock and place orders.

✓ CHECK YOUR PROGRESS 1.2

1. Your office has 200 PCs on three floors of one building. What type of network is this?

Answer: A LAN, because it covers a single building and is owned by one organisation.

2. What is the difference between an intranet and an extranet?

Answer: An intranet is for internal employees only. An extranet extends part of that intranet to trusted outsiders like suppliers or partners, usually with login and limited access.

3. Is a VPN a type of network by size?

Answer: No. A VPN is a technique that creates an encrypted private tunnel over a public network such as the Internet. It can connect a user to a LAN or two LANs over a WAN.

1.3 Network Topologies

A **topology** is the arrangement of nodes and links. **Physical topology** is how cables actually run. **Logical topology** is how data actually flows. These can differ: a hub-based network looks like a star physically but behaves like a bus logically, because every frame reaches every port.

1.3.1 Bus topology

All devices connect to one shared cable called the **backbone**, with terminators at both ends to stop signal reflection. Old 10BASE2 and 10BASE5 coaxial Ethernet used this.




Figure 1.2: Bus topology

- **Pros:** cheap, little cable, easy to add one device.
- **Cons:** one cable break takes down the whole network; collisions grow with more devices; hard to find faults.
- **Analogy:** one shared corridor in a hostel. If someone blocks it, nobody moves.

1.3.2 Star topology

Every device connects to a central device (hub or, today, a switch) with its own cable. This is how almost every office and home LAN is built.




Figure 1.3: Star topology

- **Pros:** one cable failure affects only one device; easy to add, remove and troubleshoot; works well with switches.

- **Cons:** the central device is a single point of failure; more cable than bus.
- **Analogy:** a chai stall at the centre and every office worker walks directly to it.

1.3.3 Ring topology

Each device connects to exactly two neighbours, forming a closed loop. Data travels in one direction (or both, in a dual ring). Token Ring (IEEE 802.5) and FDDI used rings; a special frame called a **token** is passed around and only the holder may transmit, so there are no collisions. SONET/SDH telecom rings still use the idea for fast failover.

- **Pros:** orderly access, no collisions with token passing, predictable performance under load.
- **Cons:** in a single ring, one broken node or link can break the loop; adding a node disturbs the ring; delay grows with ring size.
- **Analogy:** passing the microphone around in a college function. Only the person holding it speaks.

1.3.4 Mesh topology

Devices have many links to each other. In a **full mesh**, every device connects to every other device. For n devices, the number of links is $n(n-1)/2$ and each device needs $n-1$ ports.

Worked example: A full mesh of 6 routers needs $6 \times 5 / 2 = 15$ links, and each router needs 5 ports. For 20 routers: $20 \times 19 / 2 = 190$ links. You can see why full mesh becomes costly fast.

A **partial mesh** connects only important nodes redundantly. The Internet core and data centre spine-leaf fabrics are partial meshes.

- **Pros:** highly reliable (many paths), no traffic bottleneck, good privacy on dedicated links.
- **Cons:** very expensive in cables and ports, complex to set up.

1.3.5 Tree topology

A tree is a hierarchy of stars: a root switch connects to floor switches, which connect to desk devices. Large campus networks use a tree (core, distribution and access layers).




Figure 1.4: Tree (hierarchical) topology

- **Pros:** scalable, easy to manage in segments, faults isolated to a branch.
- **Cons:** if the root or a branch link fails, everything under it is cut off; more cabling.

1.3.6 Hybrid topology

A hybrid mixes two or more topologies. Real networks are almost always hybrid: star LANs on each floor, a tree across the building, and a partial mesh between data centres.

Topology	Cost	Fault tolerance	Single point of failure	Typical use
Bus	Lowest	Poor	Backbone cable	Old coax Ethernet
Star	Moderate	Good per link	Central switch	Home and office LANs
Ring	Moderate	Poor (single), good (dual)	Any link in a single ring	Token Ring, telecom rings
Mesh	Highest	Excellent	None (full mesh)	Internet core, data centres
Tree	Moderate to high	Medium	Root node	Campus networks
Hybrid	Varies	Designed as needed	Depends	Almost all real networks

ACTIVITY 1.1

Draw the network in your home. Mark the ISP fibre or cable, the ONT or modem, the Wi-Fi router, and every device. Which topology is it? (Hint: most homes are a star around the Wi-Fi router, connected to the ISP's tree.) Now count: if one cable from the router to your smart TV breaks, who is affected?

CHECK YOUR PROGRESS 1.3

1. How many links are needed for a full mesh of 10 devices?

Answer: $n(n-1)/2 = 10 \times 9 / 2 = 45$ links, and each device needs 9 ports.

2. Why is star topology the most common for LANs?

Answer: A broken cable affects only one device, devices are easy to add and troubleshoot, and switches at the centre give each port its own bandwidth. The central switch is a single point of failure, but it is easy to replace or make redundant.

3. What is the difference between physical and logical topology?

Answer: Physical topology is the actual layout of cables; logical topology is how data flows. A hub network is a physical star but a logical bus.

1.4 Network Devices

Network devices are the "dabbawalas" of the network: each one receives data, looks at some label and decides what to do. The key to understanding them is to ask: **what label does this device read?** A repeater reads nothing, a switch reads the MAC address, a router reads the IP address.

1.4.1 NIC (Network Interface Card)

The NIC (network adapter) connects a computer to the network. It converts data into signals and back, and it has a unique 48-bit **MAC address** burned in by the manufacturer, such as `3C:22:FB:1A:9E:04`. The first 24 bits identify the vendor (OUI). A NIC works at **Layer 1 and Layer 2** (Physical and Data Link). Your laptop has at least two: Wi-Fi and often Ethernet.

1.4.2 Repeater

Signals weaken (attenuate) over distance. A repeater receives a weak signal, regenerates it to full strength and sends it on. It does not understand addresses at all. Works at **Layer 1**. Analogy: the person in a long queue who shouts the announcement again for the people at the back.

1.4.3 Hub

A hub is a multi-port repeater. Whatever comes in on one port is copied out on **all** other ports. All devices share the bandwidth and are in one collision domain. Works at **Layer 1**. Hubs are obsolete; switches replaced them. Analogy: a loudspeaker in a wedding hall. Everyone hears every announcement, even "the owner of car MH-12 AB 1234, please move it".

1.4.4 Bridge

A bridge connects two LAN segments and filters traffic using MAC addresses. It learns which MAC addresses live on which side and forwards a frame only if the destination is on the other side. Works at **Layer 2**. A bridge splits a network into separate collision domains.

1.4.5 Switch

A switch is a multi-port bridge. It keeps a **MAC address table** (also called CAM table) mapping MAC addresses to ports. When a frame arrives, it sends it only out the port where the destination lives. If the destination is unknown, it floods the frame to all ports except the incoming one. Each port is its own collision domain and can run full duplex. A standard switch works at **Layer 2**. A **Layer 3 switch** can also route between VLANs using IP addresses.

Example MAC table:

MAC address	Port	VLAN
3C:22:FB:1A:9E:04	Gi0/1	10
A4:83:E7:55:10:2B	Gi0/7	10
00:1B:44:11:3A:B7	Gi0/24	20

1.4.6 Router

A router connects **different networks** and forwards packets based on the destination **IP address** using a **routing table**. It picks the best next hop, decrements the TTL, and blocks broadcasts from crossing to other networks. Works at **Layer 3**. Your home "Wi-Fi router" is really a combo box: router plus switch plus access point plus NAT plus DHCP server plus firewall. Analogy: the railway junction officer who reads the destination station and sends the train to the right track.

1.4.7 Gateway

A gateway connects networks that use **different protocols** and translates between them, for example an email gateway, a VoIP gateway converting SIP calls to the phone network, or an IoT gateway converting Zigbee to IP. A full protocol-translating gateway can work at **all layers, up to**

Layer 7. Careful: in everyday IP settings, "default gateway" simply means the router your machine sends non-local traffic to (for example 192.168.1.1). In cloud, an "API gateway" is a Layer 7 reverse proxy. Same word, different meanings.

1.4.8 Modem

Modem stands for **MODulator-DEModulator**. It converts digital data to analog signals suitable for the line (telephone, cable, radio) and back. DSL modems, cable modems and the fibre ONT (Optical Network Terminal) your ISP installs play this role. Works at **Layer 1** (some also have Layer 2 functions).

1.4.9 Access Point (AP)

A wireless access point lets Wi-Fi devices join a wired LAN. It acts as a bridge between 802.11 wireless and 802.3 Ethernet. Works at **Layer 2** (with radio at Layer 1). Offices use many APs controlled by a wireless controller so you can walk around without dropping your call.

1.4.10 Firewall

A firewall filters traffic according to security rules. A **packet-filtering** firewall checks IP addresses and ports (Layers 3 and 4). A **stateful** firewall also tracks connections, allowing reply traffic for connections started from inside. A **next-generation firewall** or **WAF** (web application firewall) inspects application data (Layer 7), for example blocking SQL injection patterns. In AWS, Security Groups and Network ACLs are firewalls you configure in code.

Device	OSI layer	Reads which address?	Splits collision domain?	Splits broadcast domain?
Repeater	1	None	No	No
Hub	1	None	No	No
Modem	1 (some 2)	None	Not applicable	Not applicable
NIC	1 and 2	Own MAC	Not applicable	Not applicable
Bridge	2	MAC	Yes	No
Switch	2 (L3 switch: 3)	MAC (and IP for L3)	Yes, every port	No (yes with VLANs or L3)
Access point	2	MAC	Shared radio channel	No
Router	3	IP	Yes	Yes, every interface
Firewall	3, 4, up to 7	IP, port, app data	Yes	Yes (when routing)
Gateway	Up to 7	Protocol-specific	Yes	Yes

KEY POINT

Hub: copies to everyone (Layer 1). Switch: sends to one MAC (Layer 2). Router: sends to another network by IP (Layer 3). If you remember "nothing, MAC, IP", you will answer 80% of device questions correctly.

1.4.11 Collision domain and broadcast domain

A **collision domain** is the part of a network where two devices sending at the same time will have their signals collide. On a hub, all ports share one collision domain. On a switch, each port is a separate collision domain, and with full duplex there are no collisions at all.

A **broadcast domain** is the set of devices that receive a broadcast frame (destination MAC `FF:FF:FF:FF:FF:FF`), such as an ARP request or DHCP discover. Switches forward broadcasts to all ports in the same VLAN, so a switch is one broadcast domain (per VLAN). **Routers stop broadcasts**, so every router interface is a separate broadcast domain.

Worked example: A router has two interfaces. Interface 1 connects to a 24-port switch with 10 PCs. Interface 2 connects to an 8-port hub with 5 PCs.

- Collision domains: the switch side has 10 PC links + 1 link to the router = 11. The hub side is 1 (all hub ports plus the link to the router share it). Total = **12 collision domains**.
- Broadcast domains: one per router interface = **2 broadcast domains**.

COMMON MISTAKE

Saying "a switch breaks broadcast domains". It does not, unless you configure VLANs (each VLAN is a separate broadcast domain) or use a Layer 3 switch to route. A plain switch breaks **collision** domains. A router breaks **broadcast** domains.

INTERVIEW CORNER

Q: What is the difference between a hub, a switch and a router?

A: A hub is Layer 1 and repeats every bit to all ports, so everyone shares one collision domain. A switch is Layer 2, learns MAC addresses and forwards frames only to the right port, giving each port its own collision domain. A router is Layer 3, connects different IP networks using a routing table and stops broadcasts.

Q: What does a switch do when it receives a frame for an unknown MAC address?

A: It floods the frame out all ports except the one it arrived on (within the same VLAN). When the destination replies, the switch learns its port and sends future frames only there.

Q: What is a default gateway?

A: The IP address of the router a host sends packets to when the destination is not in its own subnet. Without it, a host can talk only to its local network.

AT WORK

A classic production issue: a new Docker container cannot reach the internet but can reach other containers. The cause is often a missing default route or a wrong gateway inside the container network. Commands like `ip route` (Linux) or `route print` (Windows) show the default gateway. In AWS, the equivalent is a subnet route table missing a route `0.0.0.0/0` to an Internet Gateway or NAT Gateway. Knowing "who forwards my packet next" saves hours.

```
$ ip route
default via 192.168.1.1 dev wlan0 proto dhcp metric 600
192.168.1.0/24 dev wlan0 proto kernel scope link src 192.168.1.42
```

✓ CHECK YOUR PROGRESS 1.4

1. At which layer does a bridge work and what address does it use?

Answer: Layer 2 (Data Link). It uses MAC addresses to decide whether to forward or filter a frame.

2. Why are hubs obsolete?

Answer: Hubs send every frame to every port, wasting bandwidth, causing collisions and offering no privacy. Switches give each port dedicated full-duplex bandwidth at a similar cost.

3. Which device separates broadcast domains?

Answer: A router (or a Layer 3 switch). VLANs on a switch also create separate broadcast domains.

1.5 Transmission Modes

Transmission mode (or **data flow direction**) describes which way data can travel between two connected devices, and whether both can talk at once.

1.5.1 Simplex

Data flows in **one direction only**. One side always sends, the other always receives. Examples: a keyboard to a CPU (traditionally), TV broadcast, a radio station, a PA system at a railway station. Pro: the whole channel capacity is used in one direction. Con: no way to reply or acknowledge.

1.5.2 Half duplex

Both sides can send, but **not at the same time**. They take turns. Examples: walkie-talkies ("over"), old hub-based Ethernet, and classic Wi-Fi on a single channel (a radio cannot easily transmit and receive on the same frequency simultaneously). Con: time is lost while switching direction, and collisions are possible if both talk.

1.5.3 Full duplex

Both sides send and receive **simultaneously**. Examples: a phone call, switched Ethernet (separate wire pairs or fibre strands for each direction), a TCP connection carrying data both ways. A 1 Gbps full-duplex link can carry 1 Gbps in each direction at once.

Mode	Direction	Simultaneous?	Example
Simplex	One way	Not applicable	TV broadcast, FM radio
Half duplex	Both ways	No, one at a time	Walkie-talkie, hub Ethernet
Full duplex	Both ways	Yes	Phone call, switched Ethernet

⚠ COMMON MISTAKE

A **duplex mismatch** (one side forced to full duplex, the other auto-negotiating and falling back to half duplex) is a real-world bug. The link "works" but is painfully slow with many errors called late collisions. If a server's network suddenly crawls after someone changed switch settings, check speed and duplex on both ends (`ethtool eth0` on Linux).

✓ CHECK YOUR PROGRESS 1.5

1. Classify: a walkie-talkie, a phone call, a TV broadcast.

Answer: Walkie-talkie is half duplex, phone call is full duplex, TV broadcast is simplex.

2. Why do collisions not happen on a full-duplex switched Ethernet link?

Answer: Each direction has its own path (separate wire pairs or fibre strands) and only two devices share the link, so sending and receiving never compete for the same medium.

1.6 Network Architectures: Client-Server and Peer-to-Peer

1.6.1 Client-server model

In the **client-server** model, a dedicated, always-on machine (the **server**) provides a service, and many **clients** request it. The client starts the conversation; the server waits on a well-known port (HTTP on 80, HTTPS on 443, PostgreSQL on 5432). Examples: web browsing, email, your mobile app calling a REST API, a database accessed by an app server.



Figure 1.5: Client-server request and response

- **Pros:** central control of data and security, easy backup, easy to update in one place.
- **Cons:** the server can become a bottleneck or single point of failure; scaling costs money (more servers, load balancers).
- **Analogy:** IRCTC. Millions of clients, one booking system. At 10 a.m. Tatkal time, you see exactly what happens when every client hits the server at once.

1.6.2 Peer-to-peer (P2P) model

In **peer-to-peer**, every node can be both client and server. Peers share resources directly without a central server (or with a server only for discovery). Examples: BitTorrent file sharing, blockchain networks like Bitcoin, and parts of video calling (WebRTC can send media directly between peers when NAT allows).

- **Pros:** scales naturally (more peers bring more capacity), no single point of failure, cheaper.
- **Cons:** hard to secure and control, data may be unreliable, peers come and go (churn), NAT traversal is tricky.

- **Analogy:** a group of friends sharing notes before exams. Everyone gives and takes. No teacher in the middle.

Aspect	Client-server	Peer-to-peer
Central server	Yes	No (or only for discovery)
Scalability	Add servers (costs money)	Grows with peers
Control and security	Easy, centralised	Difficult
Failure	Server failure hurts all	No single point of failure
Examples	Web, email, banking apps	BitTorrent, Bitcoin

💡 DO YOU KNOW?

Early Skype (2003 onwards) was built largely as a peer-to-peer system, where some well-connected users' computers acted as "supernodes" helping route calls for others. After Microsoft acquired Skype, it moved this work to servers it ran itself.

✅ CHECK YOUR PROGRESS 1.6

1. Who starts communication in the client-server model?

Answer: The client. The server passively listens on a known port and responds.

2. Give one advantage and one disadvantage of P2P.

Answer: Advantage: capacity grows as more peers join and there is no single point of failure.
Disadvantage: it is hard to control security and guarantee data availability since peers leave at any time.

1.7 Measuring Network Performance

"The network is slow" is the most common complaint in IT. To debug it, you must separate four ideas: bandwidth, throughput, latency and jitter.

1.7.1 Bandwidth

In networking, **bandwidth** means the maximum data rate a link can carry, measured in bits per second (bps, Kbps, Mbps, Gbps). A "100 Mbps broadband plan" means up to 100 megabits per second. (In signal theory, bandwidth also means a range of frequencies in hertz; Chapter 3 covers that.)

Remember: **1 byte = 8 bits**, and network speeds use small "b". $100 \text{ Mbps} = 100/8 = 12.5 \text{ MB/s}$ at best. So a 1 GB file (8,000 megabits, using $1 \text{ GB} = 1,000 \text{ MB}$) needs at least $8000 / 100 = 80 \text{ seconds}$ on a perfect 100 Mbps link.

1.7.2 Throughput

Throughput is the actual rate of successful data delivery. It is always less than or equal to bandwidth, because of protocol overheads (headers), congestion, retransmissions, slow servers and the slowest link in the path (the **bottleneck**).

Worked example: A file of 25 MB downloads in 4 seconds. Throughput = $25 \times 8 / 4 = 50$ Mbps, even though your plan says 100 Mbps. The path's throughput equals the minimum of all link rates on it: if your Wi-Fi gives 300 Mbps, the ISP link 100 Mbps and the server is limited to 50 Mbps, you get at most 50 Mbps.

Analogy: bandwidth is the number of lanes on the Mumbai-Pune Expressway; throughput is how many cars actually reach Pune per hour, which drops when there is a jam at the toll booth.

1.7.3 Latency (delay)

Latency is the time taken for data to travel from source to destination, usually in milliseconds. It has four parts (covered with numbers in Chapter 3): transmission delay, propagation delay, queuing delay and processing delay. **RTT** (round-trip time) is the time to go and come back, which is what ping shows.

```
$ ping -c 3 google.com
64 bytes from 142.250.183.46: icmp_seq=1 ttl=117 time=14.2 ms
64 bytes from 142.250.183.46: icmp_seq=2 ttl=117 time=13.8 ms
64 bytes from 142.250.183.46: icmp_seq=3 ttl=117 time=31.5 ms
```

Light in fibre travels at about 2×10^8 m/s (about two-thirds of its speed in vacuum). Mumbai to London is roughly 7,200 km in a straight line, and real cable paths are longer. Even with a perfectly straight fibre, one-way propagation would be $7.2 \times 10^6 / 2 \times 10^8 = 36$ ms, so an RTT below about 72 ms is impossible. No amount of money buys you faster light, which is why companies use CDNs to put servers near users.

1.7.4 Jitter

Jitter is the variation in latency between packets. In the ping output above, delays were 14.2, 13.8 and 31.5 ms: the jumpiness is jitter. Jitter hurts real-time traffic (voice, video calls, online gaming) more than average latency does, because audio must be played at a steady rate. Applications use a **jitter buffer** that holds packets for a short time and plays them out smoothly, trading a little extra delay for smoothness.

Simple calculation: for delays 20, 25, 18, 30 ms, the differences between consecutive packets are 5, 7, 12 ms, so average jitter $\approx (5 + 7 + 12) / 3 = 8$ ms. (Real protocols like RTP use a smoothed formula, but the idea is the same.)

Metric	Question it answers	Unit	Tool
Bandwidth	How wide is the pipe?	bps	ISP plan, link speed
Throughput	How much actually flows?	bps	iperf3, speedtest

Metric	Question it answers	Unit	Tool
Latency / RTT	How long does one trip take?	ms	ping, traceroute
Jitter	How steady is the latency?	ms	ping series, VoIP stats
Packet loss	How much is lost?	%	ping, mtr

😄 FUNNY SIDE

Your manager: "Our users in Chennai see 2-second load times. Let us buy a bigger bandwidth plan." The 2 seconds came from 40 sequential API calls of 50 ms each. Upgrading bandwidth for a latency problem is like buying a wider plate at a wedding buffet when the real problem is the 45-minute queue.

🎯 INTERVIEW CORNER

Q: What is the difference between bandwidth and latency?

A: Bandwidth is how much data can be sent per second (width of the pipe). Latency is how long one bit takes to reach the other end (length of the pipe). A satellite link can have high bandwidth but high latency.

Q: Why is throughput lower than bandwidth?

A: Header overhead, congestion, packet loss with retransmissions, TCP window limits, slow endpoints and the bottleneck link on the path all reduce it.

Q: Why does jitter matter for video calls?

A: Audio and video must be played at a constant rate. Varying delays make packets arrive too late to play, causing choppy audio, so apps use jitter buffers.

✅ CHECK YOUR PROGRESS 1.7

1. A 200 MB file takes 40 seconds to download. What is the throughput?

Answer: $200 \times 8 = 1600$ megabits; $1600 / 40 = 40$ Mbps.

2. How long, at minimum, to send 500 MB over a 50 Mbps link?

Answer: $500 \times 8 = 4000$ Mb; $4000 / 50 = 80$ seconds, ignoring overheads.

3. What is jitter?

Answer: The variation in packet delay over time. It matters most for real-time traffic like voice and video.

1.8 Protocols and Standards

1.8.1 What a protocol defines

A protocol has three elements:

- **Syntax:** the format of the data (for example, the first 4 bits of an IPv4 header are the version).

- **Semantics:** the meaning of each field and what action to take (a TCP RST flag means "abort this connection").
- **Timing:** when and how fast to send (timeouts, retransmission rules, speed matching).

Analogy: making a phone call in India. Syntax: dial the 10-digit number. Semantics: "Hello?" means the call is connected. Timing: if nobody answers in about 30 seconds, hang up and try again.

1.8.2 Why standards?

Without standards, a Dell laptop could not talk to a Cisco switch or an Apple phone to a Google server. Standards guarantee **interoperability**, create competitive markets and let vendors build compatible products. There are two kinds: **de jure** (by law or formal body, like IEEE 802.3) and **de facto** (by wide adoption, like TCP/IP originally became popular before formal acceptance).

1.8.3 The main standards bodies

Body	Full name	Known for
IEEE	Institute of Electrical and Electronics Engineers	802 family: 802.3 Ethernet, 802.11 Wi-Fi, 802.1Q VLANs, 802.15.1 (original Bluetooth basis)
IETF	Internet Engineering Task Force	Internet protocols published as RFCs: IP, TCP, UDP, HTTP, DNS, TLS
ISO	International Organization for Standardization	The OSI reference model
ITU-T	International Telecommunication Union, Telecommunication sector	Telecom standards: G.652 fibre, H.264 video (with ISO), X.509 certificates
ICANN / IANA	Internet Corporation for Assigned Names and Numbers / Internet Assigned Numbers Authority	Domain names, IP address blocks, port numbers

1.8.4 RFCs

An **RFC** (Request for Comments) is a numbered document published by the IETF (and related groups) that describes Internet protocols and practices. Not every RFC is a standard; some are informational, experimental or best current practice. Once published, an RFC is never changed; corrections come as new RFCs that "obsolete" or "update" old ones. Famous RFCs: RFC 791 (IPv4), RFC 9293 (TCP, which replaced the original RFC 793), RFC 1034/1035 (DNS), RFC 9110 (HTTP semantics), RFC 8446 (TLS 1.3).

DO YOU KNOW?

Some RFCs published on 1 April are jokes. RFC 1149 describes sending IP packets by carrier pigeon ("IP over Avian Carriers"), and in 2001 a group in Norway actually tested it, with significant packet loss. RFC 2324 defines the Hyper Text Coffee Pot Control Protocol, which gave the world HTTP status 418 "I'm a teapot".

 ACTIVITY 1.2

Open <https://www.rfc-editor.org> and look up RFC 1918. What address ranges does it reserve for private networks? Now run `ipconfig` (Windows) or `ip addr` (Linux/Mac). Is your laptop's IP in one of those ranges? (It almost certainly is: 10.0.0.0/8, 172.16.0.0/12 or 192.168.0.0/16.)

 CHECK YOUR PROGRESS 1.8**1. Which body publishes RFCs?**

Answer: The IETF (through the RFC Editor), along with related groups like the IAB and IRTF.

2. Which IEEE standard defines Wi-Fi and which defines Ethernet?

Answer: IEEE 802.11 defines Wi-Fi and IEEE 802.3 defines Ethernet.

What Have We Learnt?

- A computer network connects devices to share data and resources, using a medium and agreed protocols.
- The five components of communication are sender, receiver, message, medium and protocol.
- By size: PAN (about 10 m), LAN (building), MAN (city), WAN (country or world). The Internet is the largest WAN.
- Intranet is private for employees, extranet extends access to partners, Internet is public.
- Topologies: bus, star, ring, mesh, tree and hybrid. Star is the most common LAN layout; full mesh needs $n(n-1)/2$ links.
- Repeaters and hubs work at Layer 1, bridges, switches and access points at Layer 2, routers at Layer 3, firewalls at Layers 3 to 7, gateways up to Layer 7.
- Switches split collision domains; routers (and VLANs) split broadcast domains.
- Transmission modes: simplex (one way), half duplex (both ways, one at a time), full duplex (both ways together).
- Client-server centralises the service; peer-to-peer lets every node act as both client and server.
- Bandwidth is capacity, throughput is actual delivery, latency is delay, jitter is variation in delay.
- A protocol has syntax, semantics and timing.
- IEEE makes LAN standards (802.x), IETF makes Internet standards (RFCs), ISO made the OSI model, ITU-T makes telecom standards.

Key terms: node, host, protocol, PAN, LAN, MAN, WAN, WLAN, VPN, Internet, intranet, extranet, topology, bus, star, ring, mesh, tree, hybrid, NIC, MAC address, repeater, hub, bridge, switch, MAC table, router, routing table, gateway, default

gateway, modem, access point, firewall, collision domain, broadcast domain, simplex, half duplex, full duplex, client-server, peer-to-peer, bandwidth, throughput, latency, RTT, jitter, packet loss, IEEE, IETF, ISO, ITU-T, RFC.

Exercises

A. Multiple Choice Questions

1. Which device works at the Network layer?
 - a. Hub
 - b. Switch
 - c. Router
 - d. Repeater
2. How many links are needed for a full mesh of 8 devices?
 - a. 16
 - b. 28
 - c. 56
 - d. 64
3. Which topology has a central device as a single point of failure?
 - a. Mesh
 - b. Ring
 - c. Star
 - d. Bus
4. A walkie-talkie is an example of:
 - a. Simplex
 - b. Half duplex
 - c. Full duplex
 - d. Multiplex
5. Which device breaks broadcast domains?
 - a. Hub
 - b. Bridge
 - c. Layer 2 switch without VLANs
 - d. Router
6. A company portal that suppliers can log into is a(n):
 - a. Intranet
 - b. Extranet
 - c. PAN
 - d. MAN
7. Which organisation publishes RFCs?
 - a. IEEE
 - b. ISO
 - c. IETF
 - d. W3C
8. A 40 MB file downloads in 8 seconds. Throughput is:
 - a. 5 Mbps
 - b. 40 Mbps
 - c. 320 Mbps
 - d. 8 Mbps
9. The variation in packet delay is called:
 - a. Latency
 - b. Jitter
 - c. Attenuation
 - d. Throughput
10. IEEE 802.11 is the standard for:
 - a. Ethernet
 - b. Wi-Fi
 - c. Token Ring
 - d. VLAN tagging

B. Fill in the Blanks

1. A hub is a multi-port _____.
2. A switch forwards frames using the _____ address table.
3. The broadcast MAC address is _____.
4. Light in optical fibre travels at about _____ m/s.
5. The three elements of a protocol are syntax, semantics and _____.
6. A network covering a city is called a _____.

C. True or False

1. Throughput can be greater than bandwidth.
2. Each port of a switch is a separate collision domain.
3. A modem converts between digital data and signals suitable for the carrier line.
4. In peer-to-peer networks, a node can act as both client and server.
5. Every RFC is an Internet standard.

D. Short Answer Questions

1. Differentiate between LAN and WAN on four points.
2. What is the difference between physical and logical topology? Give an example.
3. Explain collision domain and broadcast domain.
4. Why do we need networking standards?
5. Explain why a "faster internet plan" may not fix a slow application.
6. What is a default gateway and how do you find it on your machine?

E. Long Answer Questions

1. Describe the six network topologies with diagrams, advantages and disadvantages. Which would you choose for a new office and why?
2. Explain the working of hub, switch and router. Include the layer, the address used and their effect on collision and broadcast domains.
3. Compare client-server and peer-to-peer architectures with examples. When would you use a hybrid?
4. Define bandwidth, throughput, latency and jitter. Explain how each affects a video call, with a numerical example.

F. Design Challenge

A small startup in Bengaluru has 30 employees on one floor, 2 printers, a meeting room with a video conferencing system, a guest Wi-Fi requirement, and a few internal tools (a wiki and a staging server). Design the network: choose the topology, list the devices with quantities, decide how guests are kept away from internal tools, and explain where the single points of failure are and how you would reduce them.

Answer Key: Chapter 1

A. MCQ

1. (c) A router forwards by IP address at Layer 3.
2. (b) $8 \times 7 / 2 = 28$.
3. (c) Star depends on the central hub or switch.
4. (b) Both can talk, but one at a time.
5. (d) Routers do not forward broadcasts between interfaces.
6. (b) Extranet extends access to outsiders.
7. (c) IETF.
8. (b) $40 \times 8 / 8 = 40$ Mbps.
9. (b) Jitter.
10. (b) Wi-Fi.

B. Fill in the Blanks

1. repeater
2. MAC (CAM)
3. FF:FF:FF:FF:FF:FF
4. 2×10^8
5. timing
6. MAN

C. True or False

1. False. Throughput is the actual delivered rate and cannot exceed the link capacity.
2. True. A switch gives each port its own collision domain.
3. True. That is modulation and demodulation.
4. True. Peers both request and serve resources.
5. False. Many RFCs are informational, experimental or even April Fools' jokes.

D. Short Answers

1. A LAN covers a small area like a building while a WAN covers a country or the world. A LAN is owned by one organisation, while WAN links are often leased from telecom providers. LANs have higher speed and lower latency; WANs have higher latency and higher cost per bit. LAN errors are rare; WAN links are more prone to congestion and loss.
2. Physical topology is how cables and devices are laid out; logical topology is how data actually flows. A hub network is physically a star (all cables go to the hub) but logically a bus, because every frame reaches every device.
3. A collision domain is the area in which simultaneous transmissions collide; each switch port is a separate one while a hub shares one across all ports. A broadcast domain is the set of devices that receive a Layer 2 broadcast; it is bounded by routers or VLANs.
4. Standards ensure devices from different vendors interoperate, create open markets, reduce cost and allow the global Internet to function. Without them every vendor's equipment would need custom adapters.

5. A faster plan increases bandwidth, but if slowness comes from latency (many sequential requests, a far-away server), slow server processing or a bottleneck elsewhere, extra bandwidth does not help. You must measure first: TTFB, RTT, and the number of round trips.
6. The default gateway is the router's IP address to which a host sends packets for destinations outside its own subnet. Find it with `ipconfig` on Windows or `ip route` on Linux ("default via ...").

E. Long Answers

1. Model answer: Bus uses one backbone with terminators, cheap but one break kills all and collisions grow. Star connects all to a central switch, robust per link and easy to manage but the centre is a single point of failure. Ring passes data (often a token) around a loop, collision-free but a single break can stop a single ring. Mesh gives multiple paths; full mesh needs $n(n-1)/2$ links, very reliable but costly. Tree is a hierarchy of stars, scalable and segmentable but the root is critical. Hybrid combines these. For a new office, choose a star (switches) per floor connected in a tree to a core, with redundant uplinks for the important links: that is a hybrid design.
2. Model answer: A hub (Layer 1) repeats bits to all ports, no addresses, one collision domain and one broadcast domain. A switch (Layer 2) learns source MACs into a MAC table, forwards known unicasts to one port, floods unknown and broadcast frames; each port is a collision domain, all ports in a VLAN share one broadcast domain. A router (Layer 3) reads the destination IP, looks up the routing table, forwards to the next hop, decrements TTL, and each interface is a separate collision and broadcast domain. Routers connect different networks; switches connect devices in the same network.
3. Model answer: Client-server has dedicated servers and requesting clients, e.g. web, email, banking; central control, easy security and backup, but costs grow with scale and the server can be a bottleneck. P2P makes every node both, e.g. BitTorrent, Bitcoin; it scales with peers and has no single point of failure but is hard to secure and peers churn. Hybrid: a central server for login and discovery plus P2P for heavy data, e.g. WebRTC calls with a signalling server, or OS updates shared between LAN peers.
4. Model answer: Bandwidth is maximum capacity in bps; throughput is the actual delivered rate; latency is one-way or round-trip delay; jitter is variation in delay. For a video call needing about 2 Mbps, a 100 Mbps plan has plenty of bandwidth, but if latency is 300 ms conversation feels awkward, and jitter of 50 ms forces a larger jitter buffer adding delay. Example: delays 40, 90, 45, 100 ms give differences 50, 45, 55, average jitter 50 ms, which causes choppy audio even though bandwidth is ample.

F. Design Challenge: Model Solution

Use a star topology around one 48-port managed switch (PoE capable to power access points and phones) in a small rack, connected to a business router or firewall that links to the ISP. Devices: 1 router/firewall, 1 managed switch (48-port), 3 to 4 ceiling access points, 30 wired desk ports, 2 printers on wired ports, 1 port for the video conferencing unit. Create VLANs: VLAN 10 for staff, VLAN 20 for printers and the meeting room, VLAN 30 for guests. The guest SSID maps to VLAN 30, and firewall rules allow it only to reach the Internet, not the wiki or staging server (which sit in VLAN 10 or a separate server VLAN). Single points of failure: the ISP link, the router and the switch. Reduce risk

with a second ISP (even a 4G/5G backup) with failover, a spare switch kept on the shelf, saved configurations, and a UPS for the rack.

CHAPTER 2

Network Models: OSI and TCP/IP

Seven layers in theory, four in practice, and one more (Layer 8: the user) in every incident report.

After studying this chapter, you will be able to

- Explain why networks are designed in layers.
- Describe all seven OSI layers: function, PDU, devices, protocols and examples.
- Describe the TCP/IP model (4-layer and 5-layer views) and compare it with OSI.
- Trace encapsulation and decapsulation of an HTTP request step by step.
- Identify the address used at each layer: port, IP and MAC.
- Answer the most common interview questions on network models.

Imagine sending a Diwali gift from Pune to your cousin in Kolkata. You pack the gift (your job). The courier company puts it in their own box with a tracking number. The hub sorts it into a bag for Kolkata. The bag goes into a truck or a plane. In Kolkata, everything happens in reverse: plane, bag, box, and finally your cousin opens the gift. You never worried about which flight it took. The airline never cared what was inside the box.

This is **layering**. Each layer does one job, uses the service of the layer below and serves the layer above. Networks are built exactly like this. When you call `fetch("https://api.example.com/users")`, your code does not think about Wi-Fi radio waves. Layers below handle that.

2.1 Why Layering?

Networking is a huge problem: electrical signals, error recovery, routing across the world, reliability, encryption, application formats. Solving it as one giant program would be impossible to build or change. Layering splits it into smaller problems.

2.1.1 Benefits of layering

- **Modularity**: each layer can be designed and tested separately.
- **Abstraction**: a layer only needs to know the interface of the layer below, not its internals. HTTP works the same over Wi-Fi, Ethernet or 5G.
- **Interoperability**: vendors build products for one layer (a switch, a NIC, a browser) that work with others.

- **Easy change:** you can replace a layer's technology (copper to fibre, IPv4 to IPv6) without rewriting applications.
- **Easy troubleshooting:** engineers debug "bottom up": cable, link, IP, port, application.

2.1.2 Costs of layering

- **Overhead:** every layer adds a header, so a small message carries many bytes of headers.
- **Duplication:** some functions repeat at several layers (error checking in Ethernet and TCP).
- **Hidden information:** a layer may lack details it could use (TCP cannot directly tell whether loss was due to congestion or a bad Wi-Fi signal).

2.1.3 Key vocabulary

- **Service:** what a layer offers to the layer above (for example, "reliable byte stream").
- **Interface:** how the layer above accesses that service (for example, the sockets API).
- **Protocol:** the rules two **peer** layers on different machines use to talk. TCP on your laptop "talks" to TCP on the server, logically, even though the bits physically go down and up the stack.
- **PDU (Protocol Data Unit):** the unit of data at a layer, including its header (segment, packet, frame).
- **SDU (Service Data Unit):** the data handed down from the layer above, before this layer adds its header.

KEY POINT

Each layer talks **logically** to the same layer on the other machine (peer-to-peer communication), but **physically** data only travels down the sender's stack, across the wire, and up the receiver's stack.

CHECK YOUR PROGRESS 2.1

1. Give three advantages of a layered design.

Answer: Modularity, abstraction (a layer can change without affecting others) and interoperability between vendors. It also makes troubleshooting systematic.

2. What is the difference between a service and a protocol?

Answer: A service is what a layer offers vertically to the layer above. A protocol is the set of rules horizontally between peer layers on two machines.

2.2 The OSI Reference Model

The **OSI** (Open Systems Interconnection) model was published by ISO (as ISO 7498, in 1984). It divides communication into seven layers. OSI protocols themselves never became popular, but the **model** became the common language of networking. When someone says "it is a Layer 7 load balancer", they are using OSI terms.

No.	Layer	One-line job	PDU
7	Application	Network services to user applications	Data (message)
6	Presentation	Format, encrypt, compress	Data
5	Session	Start, manage, end dialogues	Data
4	Transport	Process-to-process delivery	Segment (TCP) / Datagram (UDP)
3	Network	Host-to-host delivery across networks	Packet
2	Data Link	Node-to-node delivery on one link	Frame
1	Physical	Bits as signals on the medium	Bits

2.2.1 Mnemonics

From Layer 1 to 7: **"Please Do Not Throw Sausage Pizza Away"** (Physical, Data Link, Network, Transport, Session, Presentation, Application). From Layer 7 to 1: **"All People Seem To Need Data Processing"**. Desi version from 1 to 7: **"Please Don't Need To Sip Pepsi Always"**. PDUs from top: **"Don't Some People Fear Birthdays"** (Data, Segment, Packet, Frame, Bits).

2.2.2 Layer 1: Physical

Job: transmit raw bits (0s and 1s) as electrical, light or radio signals over a medium.

- **Defines:** cables and connectors (RJ-45, fibre LC), voltage levels, pin layouts, bit rate, bit synchronisation, line coding, transmission mode (simplex, duplex), physical topology.
- **PDU:** bit.
- **Devices:** hubs, repeaters, cables, modems, transceivers.
- **Standards/examples:** Ethernet physical layers like 1000BASE-T, USB, DSL, the radio part of 802.11, Bluetooth radio.
- **Problems seen here:** loose cable, damaged fibre, Wi-Fi interference, link light off.

2.2.3 Layer 2: Data Link

Job: deliver frames reliably from one node to the **next directly connected node** (hop-to-hop).

- **Functions:** framing (marking start and end of data), physical addressing using **MAC addresses**, error detection (CRC in the frame trailer), flow control on the link, and access control on shared media (who may transmit, e.g. CSMA/CD, CSMA/CA).
- **Sublayers (IEEE):** LLC (Logical Link Control, talks to the network layer) and MAC (Media Access Control, talks to the physical layer).
- **PDU:** frame (header + payload + trailer).
- **Devices:** switches, bridges, NICs, access points.
- **Protocols:** Ethernet (802.3), Wi-Fi MAC (802.11), PPP, VLAN tagging (802.1Q), ARP (sits between Layers 2 and 3).

2.2.4 Layer 3: Network

Job: deliver packets from the **source host to the destination host**, possibly across many networks.

- **Functions:** logical addressing (**IP addresses**), routing (choosing the path), forwarding (moving packet to the correct output), fragmentation when a packet is too large for a link, TTL to stop endless loops.
- **PDU:** packet (or datagram).
- **Devices:** routers, Layer 3 switches.
- **Protocols:** IPv4, IPv6, ICMP (used by ping and traceroute), routing protocols like OSPF and BGP (OSPF runs directly over IP; BGP actually runs over TCP, but both serve the network layer's routing job), IPsec.

Analogy: Layer 2 is the local auto-rickshaw from your house to the station; Layer 3 is the full train route from Pune to Kolkata across many stations.

2.2.5 Layer 4: Transport

Job: deliver data from a **process on one host to a process on another** (end-to-end). Many apps on your laptop share one IP; port numbers tell them apart.

- **Functions:** port addressing, segmentation and reassembly, connection management, reliability (acknowledgements, retransmission), flow control (do not overwhelm the receiver), congestion control (do not overwhelm the network), multiplexing and demultiplexing.
- **PDU:** segment (TCP), datagram (UDP).
- **Devices:** none specific to the path; it runs on end hosts. Stateful firewalls and Layer 4 load balancers inspect it.
- **Protocols:** TCP (reliable, ordered, connection-oriented), UDP (fast, connectionless, no guarantees), QUIC (runs over UDP, used by HTTP/3), SCTP.

2.2.6 Layer 5: Session

Job: establish, maintain, synchronise and terminate **sessions** (dialogues) between applications.

- **Functions:** dialogue control (who talks when, half or full duplex), synchronisation using **checkpoints** so a long transfer can resume from the last checkpoint instead of from the start, session recovery.
- **Examples:** NetBIOS, RPC, PPTP, SQL session setup. In practice, these jobs are often done by applications or libraries (for example, TLS session resumption, or a resumable file upload).

2.2.7 Layer 6: Presentation

Job: make sure data from one system can be **understood** by another. It is the "translator" layer.

- **Functions:** translation of data formats and character encodings (ASCII, UTF-8, EBCDIC), serialisation (JSON, XML, Protocol Buffers, ASN.1), encryption and decryption (TLS is often mapped here), compression (gzip).
- **Examples:** JPEG, PNG, MPEG, MIME types, TLS encryption, gzip.

2.2.8 Layer 7: Application

Job: provide network services directly to user applications. Note: Chrome itself is not Layer 7; the **protocol** Chrome uses (HTTP) is.

- **Functions:** resource sharing, remote login, mail, file transfer, directory services, web.
- **Protocols:** HTTP/HTTPS, DNS, SMTP, POP3, IMAP, FTP, SSH, DHCP, SNMP, Telnet, MQTT.
- **Devices:** Layer 7 load balancers, API gateways, reverse proxies (Nginx, Envoy), web application firewalls.

2.2.9 Complete OSI summary table

Layer	Address / key info	Devices	Protocols	Typical issue
7 Application	URLs, hostnames	L7 load balancer, proxy, WAF	HTTP, DNS, SMTP, SSH	HTTP 500, wrong URL
6 Presentation	Encoding, cipher	(in software)	TLS, JPEG, UTF-8, JSON	Garbled text, cert error
5 Session	Session ID	(in software)	RPC, NetBIOS	Session timeout
4 Transport	Port number	L4 load balancer, firewall	TCP, UDP, QUIC	Connection refused, port blocked
3 Network	IP address	Router, L3 switch	IP, ICMP, OSPF	No route to host
2 Data Link	MAC address	Switch, bridge, AP	Ethernet, Wi-Fi, PPP, ARP	VLAN misconfig, MAC flap
1 Physical	Signals	Hub, repeater, cable	1000BASE-T, DSL	Cable unplugged

😄 FUNNY SIDE

Network engineers joke about "Layer 8": the user. "The issue was at Layer 8" is the polite way of saying the laptop's Wi-Fi switch was off. Some add Layer 9 (politics) and Layer 10 (budget), which block more projects than any firewall.

💡 DO YOU KNOW?

In the 1980s many governments, including the US through its GOSIP policy, officially backed OSI protocols. Yet TCP/IP won in practice, because it had free, working implementations (notably in BSD Unix) years earlier. This contest is often called the "protocol wars".

✓ CHECK YOUR PROGRESS 2.2

1. Which OSI layer handles routing, and which address does it use?

Answer: The Network layer (Layer 3), using IP addresses.

2. Which layer provides process-to-process delivery?

Answer: The Transport layer (Layer 4), using port numbers to identify processes.

3. What does the Presentation layer do?

Answer: It translates data formats and encodings, and handles encryption and compression, so that both sides understand the data.

4. Name the two sublayers of the Data Link layer.

Answer: LLC (Logical Link Control) and MAC (Media Access Control).

2.3 The TCP/IP Model

The **TCP/IP model** (also called the Internet protocol suite) came from ARPANET research in the 1970s. Unlike OSI, the protocols came first and the model was written to describe them. It is the model the real Internet runs on. ARPANET switched to TCP/IP on 1 January 1983.

💡 DO YOU KNOW?

The design of TCP/IP was first described by Vint Cerf and Bob Kahn in their 1974 paper "A Protocol for Packet Network Intercommunication". Originally TCP did everything; it was later split into TCP and IP so that applications like voice, which prefer speed over perfect reliability, could use IP without TCP. That split is why UDP exists.

2.3.1 Four-layer view (RFC 1122)

TCP/IP layer	Matching OSI layers	Protocols
Application	7, 6, 5	HTTP, DNS, SMTP, FTP, SSH, TLS (commonly placed here)
Transport (Host-to-host)	4	TCP, UDP, QUIC
Internet	3	IP, ICMP, IGMP
Link (Network Access)	2, 1	Ethernet, Wi-Fi, ARP, PPP

2.3.2 Five-layer (hybrid) view

Many textbooks (for example, Kurose and Ross, and Forouzan) teach a five-layer model: **Application, Transport, Network, Data Link, Physical**. It keeps TCP/IP's top but separates the link and physical layers, because they are very different technologies. Interviewers accept either; just say which one you are using.

OSI (7)	TCP/IP (4)	Hybrid (5)
+-----+	+-----+	+-----+
Application		
Presentation	Application	Application
Session		
+-----+	+-----+	+-----+
Transport	Transport	Transport
+-----+	+-----+	+-----+
Network	Internet	Network
+-----+	+-----+	+-----+
Data Link		Data Link
Physical	Link	Physical
+-----+	+-----+	+-----+

Figure 2.1: OSI, TCP/IP and the five-layer model side by side

2.3.3 The hourglass shape

The TCP/IP suite is often drawn as an **hourglass**: many applications at the top, many link technologies at the bottom, and just one protocol in the narrow waist: **IP**. Everything runs over IP and IP runs over everything. That narrow waist is why a WhatsApp message can cross Wi-Fi, 4G, fibre and undersea cable without the app knowing.




Figure 2.2: The hourglass: IP is the narrow waist

✓ CHECK YOUR PROGRESS 2.3

1. Name the four layers of the TCP/IP model.

Answer: Application, Transport, Internet and Link (Network Access).

2. Where do OSI's Session and Presentation functions go in TCP/IP?

Answer: They are absorbed into the Application layer. Applications and libraries (such as TLS, JSON serialisers) handle them.

3. Why is IP called the "narrow waist"?

Answer: Because all applications and transports use IP, and IP can run over any link technology. One common protocol in the middle connects diverse technologies above and below.

2.4 OSI versus TCP/IP

Point	OSI	TCP/IP
Developed by	ISO	US DoD research (ARPANET), later IETF
Number of layers	7	4 (or 5 in the hybrid view)
Approach	Model first, protocols later	Protocols first, model later
Nature	Generic reference model, used for teaching and troubleshooting	Practical, the Internet runs on it

Point	OSI	TCP/IP
Session and Presentation	Separate layers	Part of Application
Network layer service	Both connection-oriented and connectionless	Connectionless only (IP)
Transport layer service	Connection-oriented mainly	Both (TCP connection-oriented, UDP connectionless)
Distinction of service, interface, protocol	Very clear	Less clear
Flexibility	Rigid, protocols fit the model	Flexible, new protocols added easily

⚠ COMMON MISTAKE

Saying "the OSI model is used on the Internet". The Internet uses the TCP/IP suite. OSI is the **vocabulary** we use to describe it ("Layer 4 load balancer", "Layer 2 issue"). Also, do not say UDP is a Layer 3 protocol: it is Transport (Layer 4), just like TCP.

✓ CHECK YOUR PROGRESS 2.4

1. Give two differences between OSI and TCP/IP.

Answer: OSI has 7 layers and TCP/IP has 4. OSI was designed as a model before protocols existed, while TCP/IP's model was written after the protocols were in use.

2. Which model offers only connectionless service at the network layer?

Answer: TCP/IP. IP is connectionless; reliability, if needed, is added by TCP at the transport layer.

2.5 Encapsulation and Decapsulation

Encapsulation is the process in which each layer, while sending, takes data from the layer above and wraps it with its own **header** (and, at Layer 2, a **trailer**). **Decapsulation** is the reverse on the receiver: each layer reads and removes its own header and passes the rest up. Analogy: Russian dolls, or the dabbawala's tiffin inside a cloth bag inside a crate.

2.5.1 Step by step: sending an HTTP request

Your laptop (IP 192.168.1.42, MAC AA:AA:AA:11:11:11) sends `GET /users` to a server at 93.184.216.34 on port 443. Assume the request (after TLS encryption) is 200 bytes.

- Application (7/6/5):** the browser builds the HTTP request; TLS encrypts it. Data = 200 bytes.
- Transport (4):** TCP adds a 20-byte header (minimum) with source port (say 52344, a random ephemeral port), destination port 443, sequence number, ACK number, flags, window, checksum. Segment = 220 bytes.
- Network (3):** IP adds a 20-byte header (minimum) with source IP 192.168.1.42, destination IP 93.184.216.34, TTL (often 64), protocol = 6 (TCP), checksum. Packet = 240 bytes.

4. **Data Link (2):** the destination is outside the local subnet, so the laptop uses ARP to find the MAC of the default gateway (router 192.168.1.1). Ethernet adds a 14-byte header (destination MAC = router's MAC, source MAC = laptop's MAC, EtherType 0x0800 for IPv4) and a 4-byte trailer (FCS, a CRC-32). Frame = 258 bytes.
5. **Physical (1):** the NIC sends the frame as bits on the wire or radio, preceded by a 7-byte preamble and 1-byte start frame delimiter for synchronisation.

```

[ HTTP data 200 B ]
[ TCP 20 B ] [ HTTP data 200 B ] = segment 220 B
[ IP 20 B ] [ TCP 20 B ] [ HTTP data 200 B ] = packet 240 B
[ ETH 14 ] [ IP 20 B ] [ TCP 20 B ] [ HTTP data 200 B ] [ FCS 4 ] = frame 258 B

```

Figure 2.3: Encapsulation of a 200-byte request

Overhead calculation: headers + trailer = 20 + 20 + 14 + 4 = 58 bytes. Overhead = $58 / 258 \approx 22.5\%$ of the frame. For a full 1460-byte payload, the frame is 1518 bytes and overhead drops to $58 / 1518 \approx 3.8\%$. This is why many tiny messages are inefficient, and why batching matters.

Why 1460? The standard Ethernet **MTU** (maximum payload of a frame) is 1500 bytes. Subtract 20 bytes IP and 20 bytes TCP and you get a **MSS** (maximum segment size) of 1460 bytes.

2.5.2 What happens at each hop

1. The **switch** reads only the Ethernet header (destination MAC) and forwards the frame to the router's port. It does not touch IP or TCP.
2. The **router** receives the frame, checks the FCS, removes the Ethernet header and trailer (decapsulates to Layer 3), reads the destination IP, decrements TTL, updates the IP checksum, looks up its routing table, then **builds a new frame** for the next link with new source and destination MAC addresses. If it does NAT (as home routers do), it also rewrites the source IP and port.
3. This repeats at every router across the Internet.
4. The **server** decapsulates fully: Ethernet checks the MAC and FCS, IP checks the destination IP, TCP uses port 443 to hand the data to the web server process, TLS decrypts it, and HTTP finally sees `GET /users`.




Figure 2.4: The layers each device processes along the path

KEY POINT

MAC addresses change at every hop (they are hop-to-hop). IP addresses stay the same end to end (except where NAT rewrites them). Port numbers identify the process at each end.

 INTERVIEW CORNER

Q: Does a switch look at the IP header? Does a router look at the TCP header?

A: A plain Layer 2 switch reads only the Ethernet header. A router reads up to the IP header to route; it does not need TCP, though routers doing NAT or firewalling also read and change ports.

Q: When your laptop sends a packet to a server on the Internet, whose MAC address is in the destination field?

A: The default gateway's (router's) MAC, found by ARP. The destination IP is the server's, but the frame only travels to the next hop.

 ACTIVITY 2.1

Install Wireshark, start a capture on your Wi-Fi interface, and open any website. Click one packet. The middle pane shows the layers exactly as in this chapter: "Frame" (physical info), "Ethernet II" or "802.11" (MAC addresses), "Internet Protocol Version 4" (IP addresses, TTL), "Transmission Control Protocol" (ports, sequence numbers) and finally "Transport Layer Security" or "Hypertext Transfer Protocol". Note the source and destination MAC: is the destination your router?

 CHECK YOUR PROGRESS 2.5

1. Which layer adds a trailer, and what does it contain?

Answer: The Data Link layer. The Ethernet trailer contains the FCS (Frame Check Sequence), a CRC-32 used to detect errors.

2. A 1000-byte application message is sent over TCP/IPv4/Ethernet with minimum headers. What is the frame size?

Answer: $1000 + 20 \text{ (TCP)} + 20 \text{ (IP)} + 14 \text{ (Ethernet header)} + 4 \text{ (FCS)} = 1058 \text{ bytes}$.

3. What does a router change in a packet as it forwards it?

Answer: It replaces the Layer 2 header (new source and destination MAC), decrements TTL and recomputes the IP header checksum. With NAT, it also changes IP addresses and ports.

2.6 Addressing at Each Layer

Four kinds of addresses help a message find its target. Think of the courier: city and street (IP), the house's door (MAC, used for the last step), the person's name inside the house (port), and a friendly name everyone remembers (domain name).

Address	Layer	Size	Example	Scope	Changes per hop?
Physical (MAC)	2	48 bits	3C:22:FB:1A:9E:04	Local link only	Yes
Logical (IP)	3	32 bits (IPv4), 128 bits (IPv6)	192.168.1.42, 2001:db8::1	Global (or private)	No (except NAT)

Address	Layer	Size	Example	Scope	Changes per hop?
Port	4	16 bits (0 to 65535)	443, 5432, 52344	Within a host	No (except NAT/PAT)
Application-specific	7	Variable	api.example.com, user@mail.com	Global names	No

2.6.1 Port ranges

- **Well-known ports (0 to 1023):** 20/21 FTP, 22 SSH, 23 Telnet, 25 SMTP, 53 DNS, 67/68 DHCP, 80 HTTP, 110 POP3, 143 IMAP, 443 HTTPS.
- **Registered ports (1024 to 49151):** 3306 MySQL, 5432 PostgreSQL, 6379 Redis, 27017 MongoDB, 8080 alternative HTTP.
- **Dynamic/ephemeral ports (49152 to 65535):** chosen temporarily by clients. (Linux by default actually uses 32768 to 60999.)

A TCP connection is uniquely identified by the **5-tuple**: protocol, source IP, source port, destination IP, destination port. That is how one server on port 443 handles thousands of clients at once: each connection has a different client IP or port.

2.6.2 How the layers find each other: DNS and ARP

- **DNS** maps a name to an IP: api.example.com → 93.184.216.34.
- **ARP** maps an IP to a MAC on the local network: "Who has 192.168.1.1? Tell 192.168.1.42" → "192.168.1.1 is at 9C:53:22:..."

AT WORK

Debug by layer, bottom up. Service unreachable? (1) Is the link up (`ip link` , cable, Wi-Fi)? (2) Can you reach the gateway (`arp -a` , `ping 192.168.1.1`)? (3) Can you reach the server IP (`ping` , `traceroute`)? (4) Is the port open (`nc -vz host 5432` , `telnet host 443`)? "Connection refused" means the host answered but nothing listens on that port (L4). "Timeout" often means a firewall silently drops packets (L3/L4). (5) Does the app respond (`curl -v https://host/health`)? A certificate error is L6/L7. This one habit makes you the go-to person in on-call.

```
$ nc -vz db.internal 5432
nc: connect to db.internal port 5432 (tcp) failed: Connection refused
$ curl -v https://api.example.com/health
* Connected to api.example.com (93.184.216.34) port 443
* SSL certificate problem: certificate has expired
```

DO YOU KNOW?

Port numbers are assigned by IANA (Internet Assigned Numbers Authority), which also coordinates IP address allocation and DNS root zone management. The full official list of port assignments is public on the IANA website.

✓ CHECK YOUR PROGRESS 2.6

1. How many bits are in a MAC address, an IPv4 address and a port number?

Answer: 48 bits, 32 bits and 16 bits respectively.

2. Which protocol maps an IP address to a MAC address?

Answer: ARP (Address Resolution Protocol), used on IPv4 local networks. IPv6 uses Neighbor Discovery instead.

3. What uniquely identifies a TCP connection?

Answer: The 5-tuple: protocol, source IP, source port, destination IP and destination port.

2.7 Interview Favourites

These questions appear in almost every networking round, from campus placements to backend interviews.

INTERVIEW CORNER

Q: What happens when you type google.com in the browser and press Enter?

A: The browser checks its cache, then asks DNS to resolve the name to an IP. It opens a TCP connection to port 443 (three-way handshake), does a TLS handshake, sends an HTTP GET request, and the server responds with HTML. The browser parses it and fetches CSS, JS and images, often reusing the connection. Underneath, each packet is encapsulated with TCP, IP and Ethernet/Wi-Fi headers, ARP finds the gateway MAC, and routers forward by IP.

Q: At which layer does a Layer 4 load balancer work versus a Layer 7 one?

A: A Layer 4 balancer routes using IP and port only (fast, does not read the content), like AWS NLB. A Layer 7 balancer reads HTTP (URL path, headers, cookies) and can route /api to one service and /images to another, like AWS ALB or Nginx.

Q: Is TLS in the Transport layer or the Application layer?

A: It sits between them. It runs on top of TCP and provides encryption to the application, so in OSI it is often mapped to Presentation (Layer 6); in TCP/IP it is part of the Application layer. Saying "between L4 and L7, mapped to L6" is a safe answer.

Q: In which layer does ARP work?

A: It is debated: ARP is carried directly in Ethernet frames (Layer 2) but serves Layer 3 by resolving IP to MAC. Most answers place it at Layer 2 or call it a "Layer 2.5" protocol.

Protocol / device	Layer (OSI)
HTTP, DNS, SMTP, FTP, SSH, DHCP	7
TLS/SSL, JPEG, ASCII/UTF-8	6
RPC, NetBIOS	5
TCP, UDP, QUIC	4

Protocol / device	Layer (OSI)
IP, ICMP, IPsec, router	3
Ethernet, PPP, ARP (usually), switch, bridge	2
Hub, repeater, cables, modem	1

THINK AND DISCUSS

QUIC runs over UDP, yet it provides reliability, ordering, congestion control and built-in TLS 1.3 encryption. Is QUIC a Transport protocol or an Application protocol? What does this tell you about the limits of neat layer models?

CHECK YOUR PROGRESS 2.7

1. Which layer does DHCP belong to?

Answer: Application layer (Layer 7). It runs over UDP ports 67 (server) and 68 (client).

2. Which type of load balancer can route by URL path?

Answer: A Layer 7 load balancer, because it reads the HTTP request.

What Have We Learnt?

- Layering breaks networking into independent parts, giving modularity, abstraction and interoperability at the cost of header overhead.
- Peer layers communicate logically using protocols; adjacent layers communicate through interfaces and services.
- OSI has seven layers: Physical, Data Link, Network, Transport, Session, Presentation, Application.
- PDUs: bits (L1), frames (L2), packets (L3), segments/datagrams (L4), data (L5 to L7).
- Layer 2 delivers node to node using MAC; Layer 3 delivers host to host using IP; Layer 4 delivers process to process using ports.
- TCP/IP has four layers (Application, Transport, Internet, Link); the five-layer hybrid splits Link into Data Link and Physical.
- IP is the narrow waist of the Internet hourglass.
- OSI is a reference model and vocabulary; TCP/IP is what the Internet actually runs.
- Encapsulation adds headers going down (plus an Ethernet trailer); decapsulation removes them going up.
- Minimum TCP + IPv4 + Ethernet overhead is 58 bytes; Ethernet MTU is 1500 bytes and TCP MSS is typically 1460.
- MAC addresses change at each hop; IP addresses stay the same end to end unless NAT rewrites them.
- A connection is identified by the 5-tuple; troubleshoot bottom up, layer by layer.

Key terms: layering, service, interface, protocol, peer, PDU, SDU, OSI, Physical, Data Link, LLC, MAC sublayer, Network, Transport, Session, Presentation, Application, frame, packet, segment, datagram, TCP/IP model, Internet layer, Link layer, hourglass, encapsulation, decapsulation, header, trailer, FCS, MTU, MSS, port, ephemeral port, 5-tuple, ARP, DNS, NAT, Layer 4 and Layer 7 load balancer.

Exercises

A. Multiple Choice Questions

1. The PDU of the Network layer is:
 - a. Frame
 - b. Segment
 - c. Packet
 - d. Bit
2. Which layer is responsible for process-to-process delivery?
 - a. Network
 - b. Transport
 - c. Session
 - d. Data Link
3. Encryption and compression belong to which OSI layer?
 - a. Application
 - b. Session
 - c. Presentation
 - d. Transport
4. Which address changes at every router hop?
 - a. Destination IP
 - b. Destination port
 - c. Destination MAC
 - d. Domain name
5. How many layers does the TCP/IP model (RFC 1122) have?
 - a. 4
 - b. 5
 - c. 6
 - d. 7
6. Which layer adds a trailer?
 - a. Transport
 - b. Network
 - c. Data Link
 - d. Application
7. UDP belongs to which layer?
 - a. Network
 - b. Transport
 - c. Application
 - d. Data Link
8. The size of a port number is:
 - a. 8 bits
 - b. 16 bits
 - c. 32 bits
 - d. 48 bits
9. A 500-byte message over TCP/IPv4/Ethernet (minimum headers) gives a frame of:
 - a. 540 bytes
 - b. 554 bytes
 - c. 558 bytes
 - d. 520 bytes
10. Checkpointing for resuming long transfers is a function of:
 - a. Session layer
 - b. Physical layer
 - c. Network layer
 - d. Data Link layer

B. Fill in the Blanks

1. The OSI model was developed by _____.
2. The two sublayers of the Data Link layer are LLC and _____.
3. The standard Ethernet MTU is _____ bytes.
4. _____ resolves an IPv4 address to a MAC address.
5. Ports 0 to 1023 are called _____ ports.
6. In the TCP/IP model, the layer corresponding to OSI Layer 3 is called the _____ layer.

C. True or False

1. A switch reads IP addresses to forward frames.
2. The source IP address of a packet normally stays the same across all routers, unless NAT is used.
3. OSI protocols are what the modern Internet runs on.
4. TCP/IP's network layer offers only connectionless service.
5. HTTP is an Application layer protocol.

D. Short Answer Questions

1. What is encapsulation? Why does each layer add its own header?
2. Differentiate between the Data Link layer and the Network layer.
3. Why did TCP/IP win over OSI?
4. What is the difference between MTU and MSS?
5. Explain "connection refused" versus "connection timed out" in terms of layers.
6. List the four types of addresses used in TCP/IP networks with examples.

E. Long Answer Questions

1. Explain all seven layers of the OSI model with function, PDU, devices and protocols.
2. Compare the OSI and TCP/IP models in a table and explain the mapping between them.
3. Trace an HTTPS request from your laptop to a web server, describing encapsulation at the sender, processing at the switch and router, and decapsulation at the server.
4. "Layering is good, but not free." Discuss the advantages and disadvantages of layered network architecture with examples.

F. Design Challenge

Your team runs a public API behind a load balancer. Product wants to send all requests for `/v2/*` to a new service, add HTTPS, and block a misbehaving client IP. For each requirement, state the OSI layer at which it must be handled, and choose between a Layer 4 and Layer 7 load balancer, with reasons.

Answer Key: Chapter 2

A. MCQ

- | | |
|---|---|
| 1. (c) Packet. | 6. (c) Data Link adds the FCS trailer. |
| 2. (b) Transport uses ports. | 7. (b) Transport. |
| 3. (c) Presentation. | 8. (b) 16 bits, so 0 to 65535. |
| 4. (c) MAC addresses are rewritten per hop. | 9. (c) $500 + 20 + 20 + 14 + 4 = 558$. |
| 5. (a) 4. | 10. (a) Session layer synchronisation. |

B. Fill in the Blanks

1. ISO
2. MAC (Media Access Control)
3. 1500
4. ARP
5. well-known
6. Internet

C. True or False

1. False. A Layer 2 switch uses MAC addresses.
2. True. Routers forward by IP without changing it; NAT is the exception.
3. False. The Internet runs TCP/IP; OSI is a reference model.
4. True. IP is connectionless.
5. True.

D. Short Answers

1. Encapsulation is wrapping data from the upper layer with a header (and at Layer 2 a trailer) as it moves down the stack. Each header carries control information that only the peer layer needs, such as ports for transport, IP addresses for routing and MAC addresses for the local link.
2. The Data Link layer moves frames between two directly connected nodes using MAC addresses, handling framing, error detection and media access. The Network layer moves packets from source host to destination host across multiple networks using IP addresses and routing.
3. TCP/IP had free, working implementations (e.g. in BSD Unix) and was already running ARPANET by 1983, while OSI protocols were complex and late. Its simple, flexible design and IETF's "running code" culture helped it spread.
4. MTU is the largest payload a link-layer frame can carry (1500 bytes for Ethernet), i.e. the maximum IP packet size. MSS is the largest TCP data payload in one segment, typically MTU minus IP and TCP headers: $1500 - 40 = 1460$ bytes.
5. "Connection refused" means the packet reached the host but no process listens on that port, so the host's TCP replied with RST (Layer 4). "Timed out" means no reply came at all, often due to a firewall dropping packets, a wrong route, or the host being down (Layer 3/4).

6. MAC address (Layer 2, e.g. 3C:22:FB:1A:9E:04), IP address (Layer 3, e.g. 192.168.1.42), port number (Layer 4, e.g. 443), and application-specific addresses (Layer 7, e.g. api.example.com or an email address).

E. Long Answers

1. Model answer: Physical: bits as signals, cables, hubs, repeaters, 1000BASE-T. Data Link: frames, MAC addressing, CRC error detection, media access, switches, Ethernet, Wi-Fi, PPP; sublayers LLC and MAC. Network: packets, IP addressing, routing, fragmentation, TTL, routers, IP, ICMP, OSPF. Transport: segments/datagrams, ports, reliability, flow and congestion control, TCP, UDP. Session: dialogue control, checkpoints, RPC, NetBIOS. Presentation: translation, encryption, compression, TLS, JPEG, UTF-8. Application: services to apps, HTTP, DNS, SMTP, FTP, SSH; L7 load balancers and proxies.
2. Model answer: Use the table in Section 2.4. Mapping: OSI Application, Presentation and Session map to TCP/IP Application; Transport maps to Transport; Network maps to Internet; Data Link and Physical map to Link. Key differences: number of layers, model-first versus protocol-first, connectionless-only network layer in TCP/IP, clearer separation of service, interface and protocol in OSI, practical use of TCP/IP.
3. Model answer: DNS resolves the name. Application builds the HTTP request, TLS encrypts. TCP adds ports (ephemeral to 443), sequence numbers. IP adds source and destination IP, TTL. ARP finds gateway MAC; Ethernet adds MAC header and FCS. Bits go on the wire. Switch reads destination MAC and forwards. Router strips L2, checks IP, decrements TTL, does NAT (home), builds a new frame for the next hop. Routers repeat. Server checks FCS and MAC, IP address, uses port 443 to deliver to the web server, TLS decrypts, HTTP processes the request.
4. Model answer: Advantages: modularity, abstraction, interoperability, independent evolution (IPv6 without changing HTTP), systematic troubleshooting, vendor specialisation. Disadvantages: header overhead (58 bytes on a small 200-byte message is about 22.5%), duplicated functions (checksums in L2 and L4), loss of information across layers (TCP misreading Wi-Fi loss as congestion), and real protocols that do not fit neatly (TLS, ARP, QUIC).

F. Design Challenge: Model Solution

Routing `/v2/*` needs reading the HTTP path, which is Layer 7, so a Layer 4 balancer cannot do it. HTTPS (TLS) is Layer 6/7 in OSI terms: terminate TLS at a Layer 7 balancer (it must decrypt to read the path) and optionally re-encrypt to the backends. Blocking a client IP is Layer 3, which can be done at a firewall, security group or the load balancer itself. Choose a Layer 7 load balancer (e.g. AWS ALB, Nginx, Envoy) with path-based rules, TLS certificates installed on it, and an IP deny rule (or a WAF rule) in front. Use Layer 4 only if you needed raw speed for non-HTTP protocols.