

UNIT 4

The Network Layer

(8 Hours) — Computer Networking (CACS303)

The Network Layer is the third layer of the OSI model and the layer at which the Internet Protocol (IP) operates. Its central job is to move packets from a source host to a destination host, possibly across many intermediate routers and multiple networks, choosing a suitable path along the way. This unit covers the core responsibilities of the network layer, the addressing schemes (IPv4 and IPv6) used to identify hosts, the algorithms routers use to build forwarding tables, and the helper protocols that keep the network layer functioning.

4.1 Functions of the Network Layer

The network layer is responsible for the end-to-end delivery of packets across one or more networks. Unlike the data link layer, which only handles delivery between adjacent nodes on the same physical link, the network layer is aware of the entire path a packet must travel. Its key functions are:

- **Logical Addressing:** Assigning a unique logical address (such as an IP address) to every device so that packets can be delivered across different networks, independent of the underlying physical (MAC) addressing scheme.
- **Routing:** Determining the best path a packet should take from source to destination among the many possible paths in a network, using routing algorithms and routing tables maintained by routers.
- **Packet Forwarding:** Moving packets from an input interface to the appropriate output interface of a router, based on the routing table, so the packet progresses toward its destination.
- **Fragmentation and Reassembly:** Breaking a packet into smaller fragments when it is too large for the Maximum Transmission Unit (MTU) of an underlying network, and reassembling the fragments at the destination.
- **Congestion Control:** Monitoring and controlling the amount of traffic injected into the network so that routers and links are not overwhelmed.
- **Internetworking:** Providing a uniform addressing and delivery mechanism across heterogeneous physical networks (Ethernet, Wi-Fi, etc.), which is what allows the Internet to function as a single logical network.
- **Error Handling and Diagnostics:** Reporting problems in packet delivery (e.g. unreachable destination, time exceeded) using protocols such as ICMP.

Key Takeaways — 4.1

- ✓ Network layer provides host-to-host delivery across multiple networks.
- ✓ Main jobs: logical addressing, routing, forwarding, fragmentation, congestion control.
- ✓ It hides the differences between physical networks, enabling internetworking.

4.2 Virtual Circuits and Datagram Subnets

Inside the network layer, a subnet can be organized in one of two fundamentally different ways: as a virtual-circuit subnet (connection-oriented) or as a datagram subnet (connectionless).

Virtual Circuit Subnets

In a virtual-circuit subnet, before any data can be sent, a path (the virtual circuit) must be established between the source and destination. This path is chosen once, and all packets belonging to that connection follow the same route through the network, in order. Each router along the path stores an entry in a table for every open virtual circuit passing through it, so it knows where to forward packets carrying a particular circuit identifier. When the connection is released, the table entries are removed.

Datagram Subnets

In a datagram subnet, no advance setup is required. Every packet (called a datagram) is routed independently, carrying the full destination address, and different packets belonging to the same message may travel by different routes and can arrive out of order. Routers make a fresh forwarding decision for every packet by consulting their routing tables.




Figure 4.1: Virtual circuit subnet (fixed path for all packets) vs. Datagram subnet (each packet routed independently)

Aspect	Virtual Circuit Subnet	Datagram Subnet
Connection setup	Required before data transfer	Not required
Path used	Same fixed path for all packets	May differ for each packet
Packet ordering	Preserved (in-order delivery)	Not guaranteed
Router state	Keeps per-circuit table entries	Stateless — no per-connection state
Failure recovery	Circuit fails if a router on path fails	Easily re-routed around failures
Overhead	Setup/teardown overhead	Per-packet addressing overhead
Example	ATM, MPLS	IP (Internet)

Note: The Internet's IP layer is a classic example of a datagram (connectionless) subnet — this is why packets can arrive out of order and why IP itself does not guarantee delivery.

4.3 IPv4 Addresses

Address Space and Notation

An IPv4 address is a 32-bit number that uniquely identifies a host's connection to a network. Since it is 32 bits long, the total address space contains 2^{32} (about 4.29 billion) distinct addresses. IPv4 addresses are almost always written in dotted-decimal notation: the 32 bits are divided into four 8-bit groups (octets), and each octet is written as a decimal number from 0 to 255, separated by dots — for example, 192.168.1.10. The equivalent binary form of this address is 11000000.10101000.00000001.00001010.

Classful Addressing

In classful addressing, the IPv4 address space was originally divided into five classes (A through E), distinguished by the leading bits of the address. This scheme fixed how many bits were used for the network portion versus the host portion of the address.




Figure 4.2: Bit layout of the five classes of classful IPv4 addressing

Class	First Byte Range	Default Mask	Purpose
A	1 – 126	255.0.0.0 (/8)	Very large networks (few networks, many hosts)
B	128 – 191	255.255.0.0 (/16)	Medium-sized networks
C	192 – 223	255.255.255.0 (/24)	Small networks (many networks, few hosts)
D	224 – 239	— (not subnetted)	Multicast addressing
E	240 – 255	— (not subnetted)	Reserved for experimental/future use

Note: 127.x.x.x is reserved for loopback testing and does not belong to any usable class range in practice, even though it falls within the Class A range.

Classless Addressing (CIDR)

Classful addressing wastes large numbers of addresses (e.g. a Class C network with only 5 hosts still consumes 254 host addresses). Classless Inter-Domain Routing (CIDR) solves this by allowing the network/host boundary to fall on any bit position, not just at 8, 16, or 24 bits. An address is written with a slash notation indicating how many leading bits form the network prefix — for example, 192.168.1.0/24 means the first 24 bits are the network part and the remaining 8 bits are for hosts. CIDR allows efficient, flexible allocation of address blocks and underlies modern Internet routing.

Example 4.1 — Interpreting CIDR Notation

Given the address block 172.16.32.0/20, determine the number of host addresses available.

Solution

A /20 prefix uses 20 bits for the network, leaving $32 - 20 = 12$ bits for hosts.

Total addresses = $2^{12} = 4096$.

Usable host addresses = $4096 - 2 = 4094$ (subtracting the network address and the broadcast address).

Subnetting

Subnetting is the process of further dividing a single network (or address block) into several smaller sub-networks, by borrowing bits from the host portion of the address to extend the network prefix. This allows an organization to structure its internal network into departments or physical segments, improves security by isolating traffic, and reduces broadcast domain size.

Subnetting: Dividing 192.168.1.0/24 into Four /26 Subnets (borrowing 2 host bits)




Figure 4.3: Subnetting 192.168.1.0/24 into four /26 subnets

Example 4.2 — Subnetting Calculation

A network administrator has the block 192.168.1.0/24 and needs to create 4 equal-sized subnets. Find the subnet mask and the range of each subnet.

Solution

To create 4 subnets, we need 2 additional bits (since $2^2 = 4$), so the new prefix length is $/24 + 2 = /26$.

New subnet mask: 255.255.255.192.

Each subnet has $2^6 = 64$ addresses (62 usable hosts).

Subnet 1: 192.168.1.0 – 192.168.1.63

Subnet 2: 192.168.1.64 – 192.168.1.127

Subnet 3: 192.168.1.128 – 192.168.1.191

Subnet 4: 192.168.1.192 – 192.168.1.255

Network Address Translation (NAT)

Because the pool of public IPv4 addresses is limited, most private networks (homes, offices) use private address ranges internally (such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16) and rely on Network Address Translation (NAT) to communicate with the public Internet. A NAT-enabled router or gateway translates the private source address and port of outgoing packets into its own single public IP address (and a unique port number), keeping a translation table so that returning packets can be mapped back to the correct internal host. NAT conserves public IPv4 addresses and adds a degree of privacy, since internal addressing is hidden from the outside world.

Key Takeaways — 4.3

- ✓ IPv4 address is 32 bits, written in dotted-decimal notation.
- ✓ Classful addressing (A–E) fixes the network/host boundary at 8, 16, or 24 bits.
- ✓ CIDR (classless addressing) allows any prefix length, improving address efficiency.
- ✓ Subnetting borrows host bits to create smaller internal sub-networks.
- ✓ NAT lets many private hosts share one public IPv4 address.

4.4 IPv4 Datagram Format and Fragmentation

Every packet sent using IP is wrapped in an IPv4 datagram, consisting of a header (normally 20 bytes, up to 60 bytes with options) followed by the data payload.




Figure 4.4: IPv4 datagram header format

Field	Size	Purpose
Version	4 bits	IP version number (4 for IPv4)
IHL	4 bits	Internet Header Length, in 32-bit words
Type of Service	8 bits	Requested quality of service / priority
Total Length	16 bits	Total datagram length (header + data) in bytes
Identification	16 bits	Uniquely identifies fragments of the same datagram
Flags	3 bits	Controls fragmentation (DF = don't fragment, MF = more fragments)
Fragment Offset	13 bits	Position of this fragment within the original datagram
Time to Live (TTL)	8 bits	Decrement per hop; datagram discarded when it reaches 0
Protocol	8 bits	Identifies the upper-layer protocol (e.g. TCP=6, UDP=17)

Field	Size	Purpose
Header Checksum	16 bits	Error-checking for the header only
Source Address	32 bits	IPv4 address of the sender
Destination Address	32 bits	IPv4 address of the receiver
Options + Padding	variable	Optional fields (e.g. security, routing records)


Fragmentation

Different physical networks impose different Maximum Transmission Unit (MTU) limits on frame size. When a datagram is larger than the MTU of the network it must cross, the router splits it into smaller fragments, each becoming its own IP datagram with its own header. The Identification field is copied into every fragment so the destination knows they belong to the same original datagram; the Fragment Offset field records each fragment's position (in units of 8 bytes) so they can be reassembled in the correct order; and the More Fragments (MF) flag is set on every fragment except the last one. Reassembly happens only at the final destination host, not at intermediate routers, because fragments of the same datagram may travel by different paths.

Note: If the Don't Fragment (DF) flag is set and a datagram is too large for a link, the router discards it and returns an ICMP 'Destination Unreachable — Fragmentation Needed' message instead of fragmenting it.

4.5 IPv6 Address Structure and Advantages over IPv4

IPv6 was designed to replace IPv4, primarily to solve the exhaustion of the IPv4 address space. An IPv6 address is 128 bits long, giving an address space of 2^{128} — vastly larger than IPv4's 2^{32} . IPv6 addresses are written as eight groups of four hexadecimal digits, separated by colons, for example: 2001:0db8:85a3:0000:0000:8a2e:0370:7334. Leading zeros in a group may be omitted, and one contiguous run of all-zero groups may be replaced with a double colon (::), used only once per address.



Example: 2001:0db8:85a3:0000:0000:8a2e:0370:7334 (128 bits, 8 groups of 16 bits in hex)

Figure 4.5: Structure of a 128-bit IPv6 address

Field	Size	Purpose
Global Routing Prefix	48 bits	Identifies a site; assigned by a provider/registry
Subnet ID	16 bits	Identifies a subnet within the site
Interface ID	64 bits	Identifies a specific interface/host on the subnet

Advantages of IPv6 over IPv4

- **Larger address space:** 128-bit addressing eliminates address exhaustion, removing the practical need for NAT.
- **Simplified header:** The IPv6 header is fixed-length and simpler than IPv4's, which speeds up router processing.
- **No header checksum:** Removed since it is handled at other layers, reducing per-hop processing overhead.
- **Built-in security:** IPsec support was designed in as an integral part of IPv6 (though it is now also available for IPv4).
- **Better support for extensions and options:** Implemented through a chain of extension headers rather than embedded option fields.
- **No fragmentation by routers:** Only the sending host fragments packets in IPv6, using Path MTU Discovery, reducing router workload.
- **Improved support for auto-configuration and mobility.**

Key Takeaways — 4.5

- ✓ IPv6 uses 128-bit addresses, written as 8 groups of hex digits.
- ✓ Structure: Global Routing Prefix + Subnet ID + Interface ID (48+16+64 bits).
- ✓ IPv6 offers a vastly larger address space, simpler header, and built-in security support.

4.6 Routing Algorithms

A routing algorithm is the part of network-layer software responsible for deciding which output line an incoming packet should be forwarded on, and for building and maintaining the routing tables that make that decision fast. The two classical approaches used inside an autonomous system are Distance Vector Routing and Link State Routing.




Figure 4.6: Distance Vector routing vs. Link State routing

Distance Vector Routing

In distance vector routing, each router maintains a table (a vector) listing, for every destination, the best known distance (e.g. hop count) and the next-hop router to use. Periodically, each router exchanges its entire distance vector with its directly connected neighbors. On receiving a neighbor's vector, a router updates its own table if the neighbor offers a shorter path to some destination. This algorithm is based on the Bellman-Ford algorithm and is used in protocols such as RIP (Routing Information Protocol).

- Each router knows only the distances reported by its immediate neighbors, not the full network topology.
- Convergence can be slow, and the algorithm suffers from the well-known count-to-infinity problem when a link fails.
- Simple to implement and requires little computation per router.

Link State Routing

In link state routing, each router first discovers its neighbors and measures the cost (e.g. delay or bandwidth) to reach each of them. It then constructs a small packet, the Link State Advertisement (LSA), describing these costs, and floods this packet to every other router in the network. Because every router receives every other router's LSA, each one can independently build an identical, complete map of the network topology. Each router then runs Dijkstra's shortest-path algorithm on this map to compute the best route to every destination. This approach is used in protocols such as OSPF (Open Shortest Path First).

- Every router has full knowledge of the network topology, allowing faster and more accurate convergence than distance vector routing.
- Requires more memory and computation (Dijkstra's algorithm) at each router.
- Reacts to topology changes more quickly since LSAs are flooded immediately on a change.

Aspect	Distance Vector Routing	Link State Routing
Information shared	Own entire routing table, to neighbors only	Link costs only, flooded to all routers
Topology knowledge	Partial (indirect, via neighbors)	Complete (every router builds a full map)
Algorithm used	Bellman-Ford	Dijkstra's shortest path
Convergence speed	Slower; count-to-infinity issue	Faster and more reliable
Example protocol	RIP	OSPF

Key Takeaways — 4.6

- ✓ Distance vector: routers share full tables with neighbors only; uses Bellman-Ford.
- ✓ Link state: routers flood link costs to everyone; each builds a full map and runs Dijkstra.
- ✓ Link state converges faster but needs more router resources than distance vector.

4.7 Internet Control Protocols: ARP, RARP, ICMP

Alongside IP itself, several helper protocols operate at (or just above) the network layer to resolve addresses and report errors.




Figure 4.7: Roles of ARP, RARP, and ICMP

Address Resolution Protocol (ARP)

A host that wants to send a frame on a local network (e.g. Ethernet) needs the recipient's physical (MAC) address, but it typically only knows the recipient's IP address. ARP solves this by broadcasting a request containing the target IP address onto the local network, asking 'who has this IP address?'. The host that owns that IP address replies directly with its MAC address, which the sender then caches in an ARP table for future use, avoiding repeated broadcasts.

Reverse Address Resolution Protocol (RARP)

RARP performs the opposite mapping: a host that knows its own MAC address but not its IP address (for example, a diskless workstation booting up) broadcasts a RARP request containing its MAC address, and a RARP server on the network replies with the IP address assigned to that MAC address. RARP has largely been superseded by more capable protocols such as BOOTP and DHCP, which provide the IP address plus other configuration information (subnet mask, gateway, DNS server) in one exchange.

Internet Control Message Protocol (ICMP)

ICMP is used by routers and hosts to report errors and exchange diagnostic information about IP packet processing — it is not used to carry user data. Common ICMP message types include:

- **Destination Unreachable:** Sent when a packet cannot be delivered (e.g. no route to host, port unreachable).
- **Time Exceeded:** Sent when a packet's TTL reaches zero before reaching its destination — the basis of the traceroute utility.
- **Echo Request / Echo Reply:** Used by the ping utility to test whether a host is reachable and measure round-trip time.
- **Redirect:** Sent by a router to tell a host to use a better route for future packets to a given destination.
- **Source Quench (historical):** Previously used to signal congestion and ask a sender to slow down (now largely deprecated in favor of other congestion control mechanisms).

Key Takeaways — 4.7

- ✓ ARP maps a known IP address to an unknown MAC address (used before sending on a LAN).
- ✓ RARP maps a known MAC address to an unknown IP address (largely replaced by DHCP).
- ✓ ICMP reports errors and diagnostics for IP (e.g. ping, traceroute) — not user data.

4.8 Routing Protocols: OSPF, BGP, Unicast, Multicast and Broadcast

OSPF (Open Shortest Path First)

OSPF is an interior gateway protocol (IGP) — it is used for routing within a single autonomous system (AS), such as a single organization's network. It is a link state protocol: every router floods information about its directly connected links, builds a complete topology map of the AS, and computes shortest paths using Dijkstra's algorithm. OSPF supports hierarchical design through areas, allowing large networks to be split into manageable regions, and it converges quickly after topology changes.

BGP (Border Gateway Protocol)

BGP is the exterior gateway protocol (EGP) used to route between different autonomous systems — it is, in effect, the routing protocol that holds the global Internet together. Rather than being purely a shortest-path (distance/cost based) protocol, BGP is a path-vector protocol: each BGP router advertises the entire sequence (path) of autonomous systems needed to reach a destination network, which allows routing decisions to be based on policy (business agreements, avoiding certain ASes, etc.) rather than purely on distance.

Aspect	OSPF	BGP
Scope	Within a single autonomous system (IGP)	Between autonomous systems (EGP)
Type	Link state protocol	Path vector protocol
Metric basis	Cost (bandwidth-based)	Policy-based path selection
Convergence	Fast	Slower, but very stable at Internet scale

Unicast, Multicast and Broadcast Routing

Network-layer delivery can also be classified by how many destinations a single packet is meant to reach:

- **Unicast:** A packet is sent from one source to exactly one specific destination. This is the most common form of delivery (e.g. loading a web page).
- **Broadcast:** A packet is sent from one source to all hosts on a network or subnet (e.g. an ARP request, or a DHCP discovery message). Broadcasts are typically limited to a local network and are not forwarded across routers by default.
- **Multicast:** A packet is sent from one source to a specific group of interested receivers, identified by a multicast group address (Class D in IPv4). Multicast is efficient for applications like video conferencing or live streaming, since a single copy of the data can be replicated by routers only along the branches of the network that actually have subscribed group members, rather than being duplicated by the source for every receiver.

Key Takeaways — 4.8

- ✓ OSPF is a link-state IGP used for routing within one autonomous system.
- ✓ BGP is a path-vector EGP used to route between autonomous systems across the Internet.
- ✓ Unicast = one-to-one, Broadcast = one-to-all (local), Multicast = one-to-group.

— End of Unit 4 —