

The IP Protocol

Tom Kelliher, CS 325

Apr. 13, 2011

1 Administrivia

Announcements

Written assignment due in Friday.

Assignment

Read 4.5–6.

From Last Time

Introduction to network layer protocols.

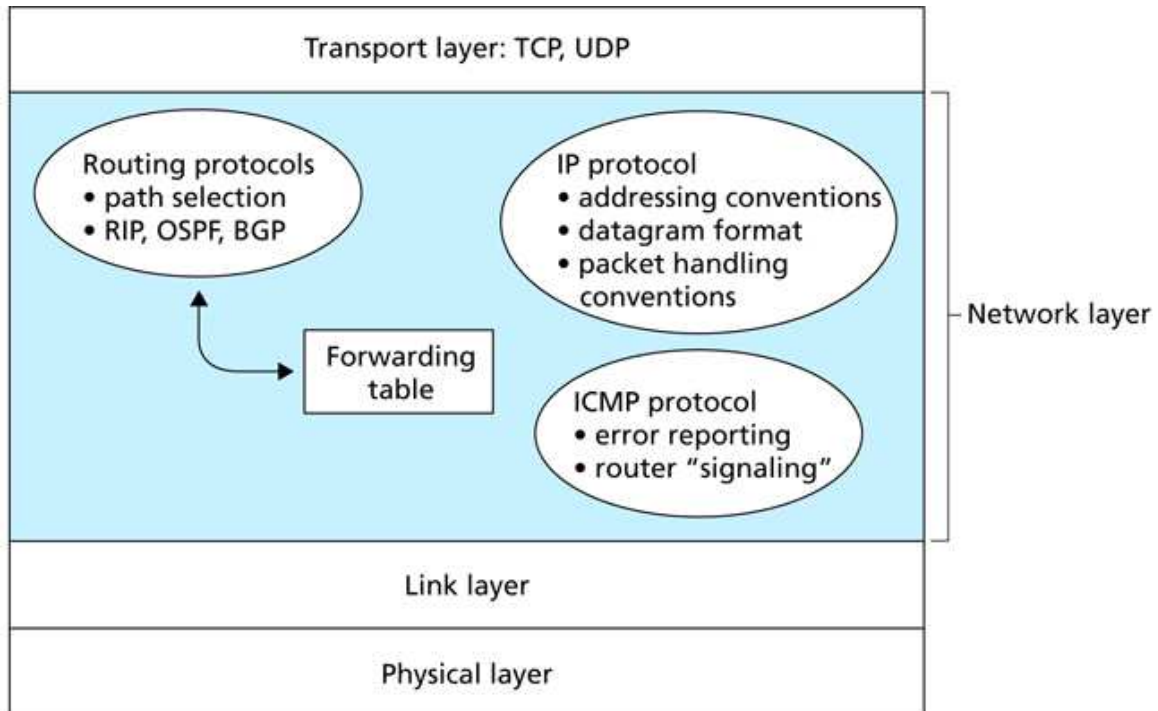
Outline

1. Introduction.
2. Network addressing.
3. IP address management: DHCP and NAT.
4. IPV6.

Coming Up

Project day; Routing.

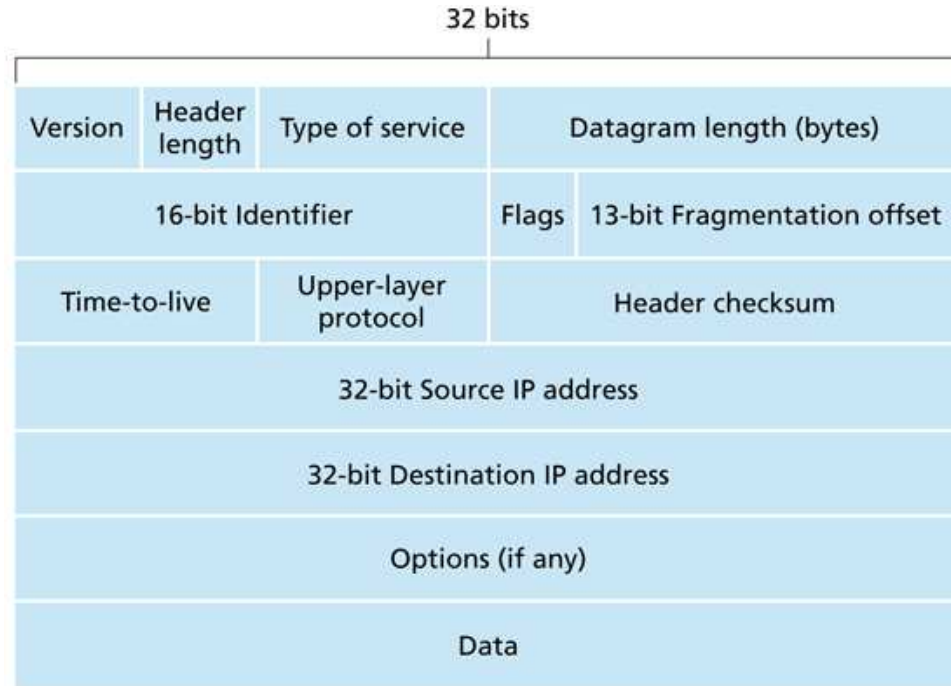
2 Introduction



Network layer protocols:

1. IP: datagram protocol.

Header:



Selected header fields:

- (a) Version — V4 or V6.
- (b) Identifier — sequence number.
- (c) Flags, fragmentation offset: for fragment reassembly.
- (d) Time-to-live (TTL) — number of router hops remaining.
(Used by `traceroute`.)
- (e) Protocol — identifies transport layer protocol to route datagram to (TCP or UDP).
- (f) Source and destination IP addresses.

2. ICMP: Internet Control Message Protocol.

Control messages sent as IP datagrams, with Protocol field set to “ICMP.”


Sample messages:

- (a) Echo request/reply (ping).
- (b) Destination host/network/protocol/port unreachable.
Port unreachable refers to UDP — a kluge?
- (c) Destination host/network unknown.
- (d) Router advertisement/discovery.
- (e) TTL expired.

3. Routing:

- (a) Intra-AS: RIP, OSPF.
- (b) Inter-AS: BGP.

Fragmentation:



1. What happens if a datagram's size exceeds a link's MTU?

Split it into fragments!

2. Downstream routers then route the fragments.
3. Destination host re-assembles fragments into the original datagram.

Example: 4,000 byte datagram (20 bytes header; 3,980 bytes data) being sent across a link with an MTU of 1,500 bytes.

| Fragment | Data Bytes | Offset | Flag |
|--------------|--------------------------------------|---|------|
| 1st Fragment | 1,480 | 0 | 1 |
| 2nd Fragment | 1,480 | 185 ($185 \times 8 = 1,480$) | 1 |
| 3rd Fragment | 1,020 ($= 3,980 - 2 \times 1,480$) | 370 ($370 \times 8 = 2 \times 1,480$) | 0 |

Notes:

- (a) Each datagram contains a 20 byte header.
 - (b) Each fragment uses the original datagram's identifier field.
 - (c) The offset is the eight byte offset from the beginning of the original datagram of the first byte of the fragment.

Hence, all fragments, save the last, must have a data length that is a multiple of 8.
 - (d) Flag = 0 means additional fragments. Flag = 1 means this is the final fragment.
 - (e) The fragmentation mechanism permits fragments to be further fragmented, with due care.
4. Obviously, only complete datagrams are passed up to transport layer.


3 Network/Host Addressing

IP addresses are in "dotted-decimal" form:

$$10.67.1.26 = 00001010\ 01000011\ 00000001\ 00011101$$

$$10.32.3.39 = 00001010\ 00100000\ 00000011\ 00100111$$

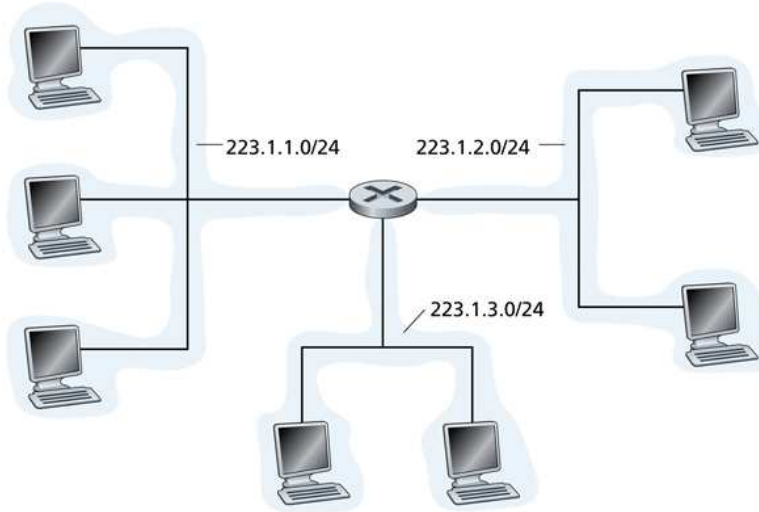
Each host/router on a network has its own IP address:



1. By definition, routers connect several networks. Hence, they have multiple IP addresses.
2. All hosts on the same network have the network portion of their IP address in common.

IP addresses are hierarchical: network/host-on-network.

The network portion of these three networks:



1. The /24 specifies the netmask, or network portion of the IP address: 255.255.255.0 — 24 leading 1's.

The network portion of the IP address is the first 24 bits. The remaining 8 bits is the host portion of the IP address.

2. A /20 corresponds to a netmask of 255.255.240.0 — 20 leading 1's.
3. /17 = 255.255.128.0; /18 = 255.255.192.0; /22 = 255.255.252.0; /23 = 255.255.254.0.
192 = 128 + 64; 252 = 128 + 64 + 32 + 16 + 8 + 4.

4. /n addressing is referred to as CIDR — Classless Interdomain Routing.

Another way of conserving IP addresses.

5. CIDR replaced classful routing — Class A (/8), Class B (/16), Class C (/24).
6. It was/is common for organizations to subnet IP address ranges to create “networks within a network.”

Example: You're given a 64 address range: 66.240.10.64/26.


- (a) Note: Two IP addresses are reserved within each network — The network IP number itself (host portion all 0's) and the broadcast address (host portion all 1's).
- (b) You can create four subnets using a /28 netmask. The subnets are 66.240.10.64, 66.240.10.80, 66.240.10.96, and 66.240.10.112.

Each subnet has 16 IP addresses; 14 usable.

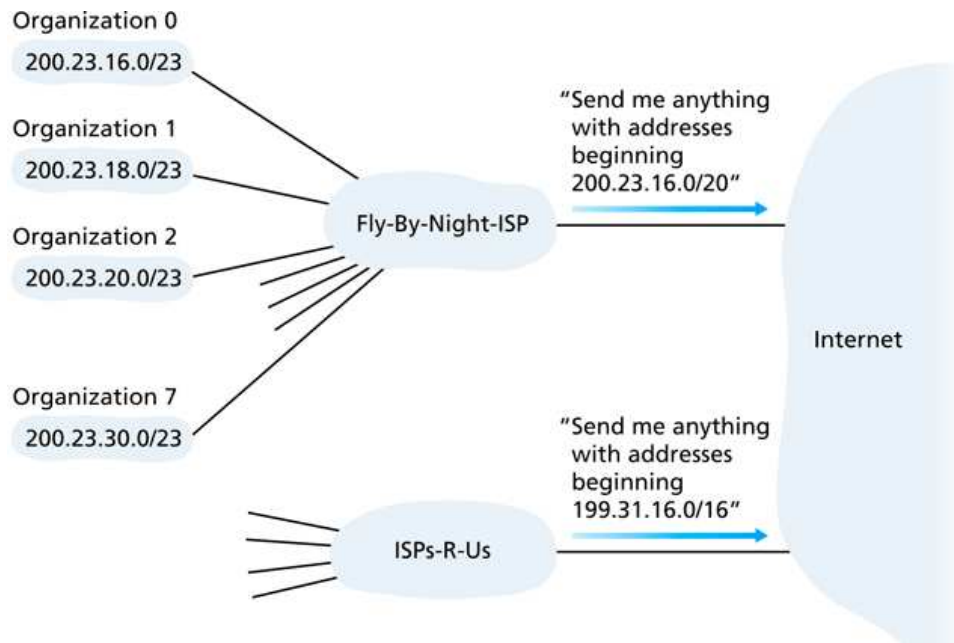
- (c) You can create eight subnets using a /29 netmask. The subnets are 66.240.10.64, 66.240.10.72, 66.240.10.80, etc.

Each subnet has 8 IP addresses; 6 usable.

The links connecting routers are considered networks themselves:



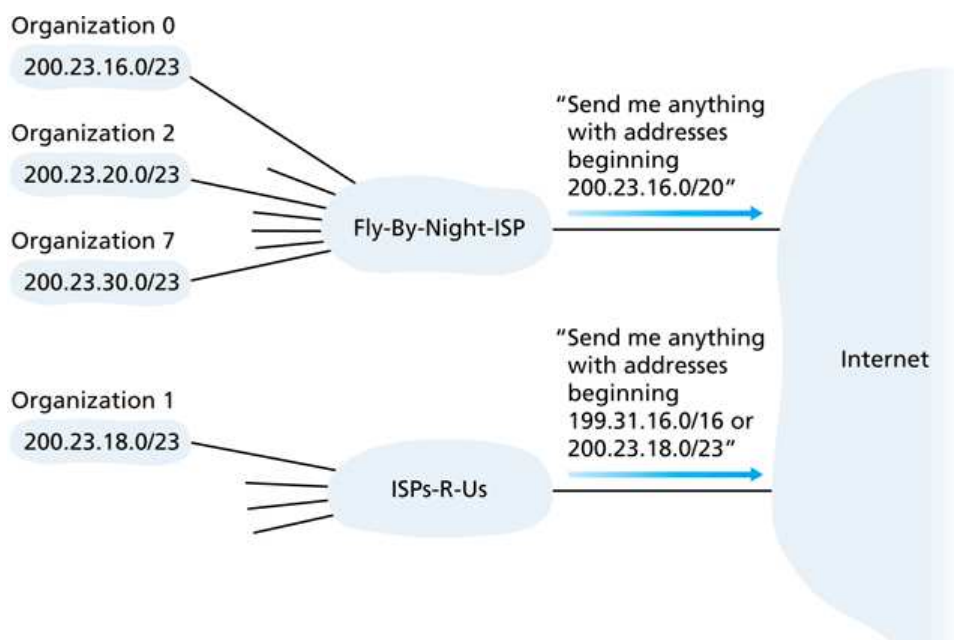
Networks can be supernetted:



Fly-by-Night has 8 /23's which have been aggregated (supernetted) into a single /20 for routing purposes.

Address aggregation helps reduce forwarding table size.

Organization 1 can change ISPs, keeping its IP addresses, so long as the new ISP advertises a "more specific" network (/23 vs. /20):



4 IP Address Management: DHCP and NAT


Where do I get an IP address(es)...

1. to create a public network? From your ISP.
2. if I'm an ISP? From ARIN. Also, to get an ASN.
3. if I need public access to a host from within an existing network? Obtain public IP address from IT.
4. If I don't need public access from within an existing network? From DHCP.

4.1 DHCP

(Dynamic Host Configuration Protocol)

1. Typically uses private networks: 10.0.0.0/8 or 192.168.0.0/16.
Must also use NAT.
2. DHCP server dynamically assigns IP addresses from pool. Assignments are temporary — leased.
3. Each network needs a DHCP server or a DHCP relay agent (usually a router).
4. DHCP clients also typically receive netmask, gateway, DNS servers, WINS servers, etc. info from DHCP server.



DHCP exchange:

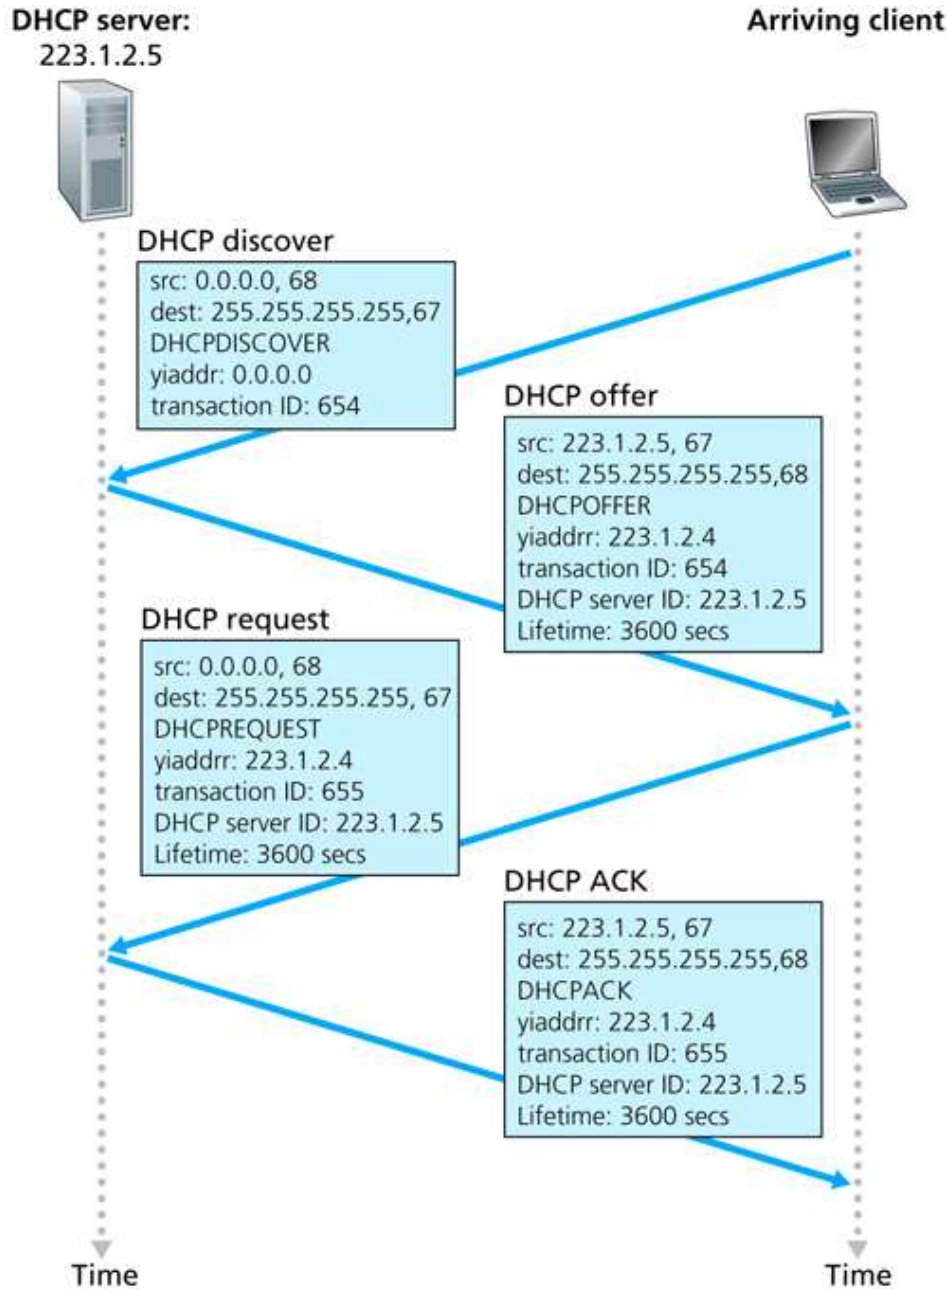
1. Client sends DHCP Discover message. Note src IP addr of 0.0.0.0 (“this host”) and broadcast dest IP addr of 255.255.255.255.

Hopefully, the transaction ID is unique.

2. Server responds with a DHCP offer message. Dest IP addr must be the broadcast addr, but sent to client’s port address, using the client’s transaction ID.

Client *might* receive several offers.

3. Client responds to one offer with a DHCP request message, repeating the original offer’s parameters.
4. Server responds with a DHCP ACK message, confirming the request.



4.2 NAT

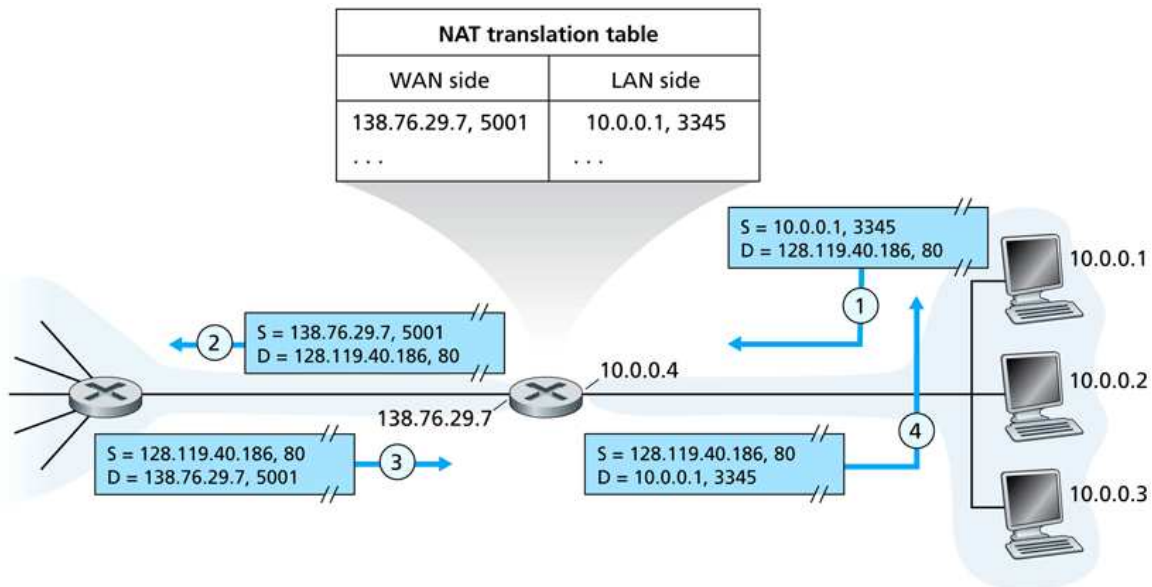
(Network Address Translation)

Comcast will give a residential customer one IP address (via DHCP). How does the customer configure a home network? With NAT, and an internal DHCP server!

NAT helps conserve IP addresses.

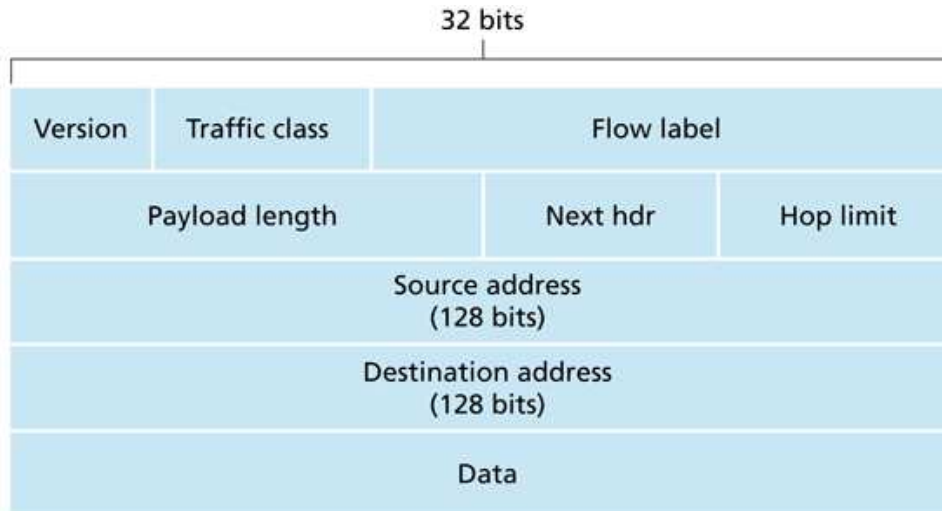
1. Externally, a NAT router appears to handle only a single IP address — It's not really a router.
2. Each device on the NAT network has its own IP address.
3. The NAT router maintains a table of (external port, (local IP, local port)) entries to handle the internal/external mapping of TCP and UDP connections.
4. As NAT routers look into the transport layer data, they're more than just network layer devices.
5. Provisions must be made for servers and P2P hosts behind NAT routers.

NAT translation example:



5 IPV6

IPV6 header:



1. Streamlined — no options.
2. No fragmentation — easing burden on routers.
3. No checksum — again, easing burden on routers.
4. Traffic class, flow — to specify type of data (video stream, interactive, etc.).
5. Next header — transport layer protocol.
6. Hop limit — TTL.
7. Source and destination IP addresses — 128 bits!!!.


Forwarding table lookups?

We're not going to be able to switch from IPV4 to IPV6 overnight.

IPV4 and IPV6 hosts/routers may have to interoperate for a period of time.

Two interoperation approaches:

1. Dual stack:



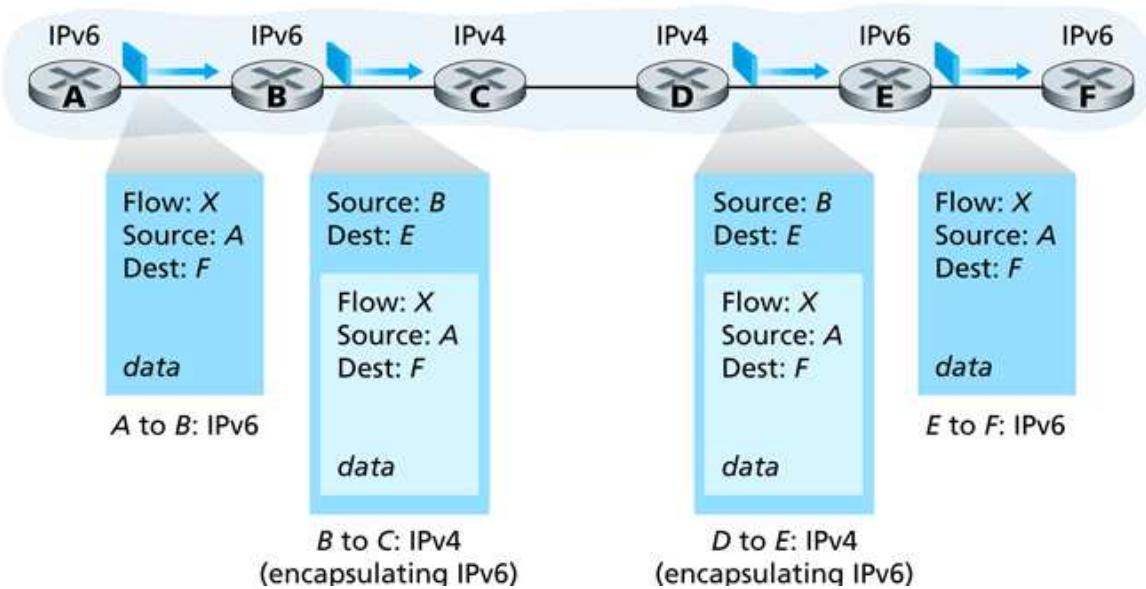
Note that IPV6 characteristics get stripped.

2. Tunnel:

Logical view



Physical view



Note that IPV6 datagrams are encapsulated in IPV4 datagrams.