

Last time

□ Network layer

◆ Introduction

- forwarding vs. routing

◆ Virtual circuit vs. datagram details

- connection setup, teardown
- VC# switching
- forwarding tables, longest prefix matching

◆ IP: the Internet Protocol


- packet structure
- fragmentation & reassembly

This time

- IP Addressing
- ARP
- DHCP
- ICMP
- IPv6


IP Addressing: introduction

- IP address: 32-bit identifier for host, router *interface*
- *interface*: connection between host/router and physical link
 - ◆ routers typically have multiple interfaces
 - ◆ host typically has one interface
 - ◆ IP addresses associated with each interface



Subnets

- IP address:
 - ◆ subnet part (high order bits)
 - ◆ host part (low order bits)
- *What's a subnet ?*
 - ◆ device interfaces with same subnet part of IP address
 - ◆ can physically reach each other without intervening router




network consisting of 3 subnets

Subnets

Recipe

- To determine the subnets, detach each interface from its host or router, creating islands of isolated networks. Each isolated network is called a **subnet**.




Subnet mask: /24

Subnets

How many?


How large?



IP addressing: CIDR

CIDR: Classless InterDomain Routing

- ◆ subnet portion of address of arbitrary length
- ◆ address format: **a.b.c.d/x**, where x is # bits in subnet portion of address



IP addresses: how to get one?

Q: How does *host* get IP address?

- Hard-coded by system admin in a file
 - ◆ Wintel: control-panel->network->configuration->tcp/ip->properties
 - ◆ Linux: /etc/network/interfaces
- **DHCP: Dynamic Host Configuration Protocol:** dynamically get address from as server
 - ◆ “plug-and-play”
 - ◆ later

IP addresses: how to get one?

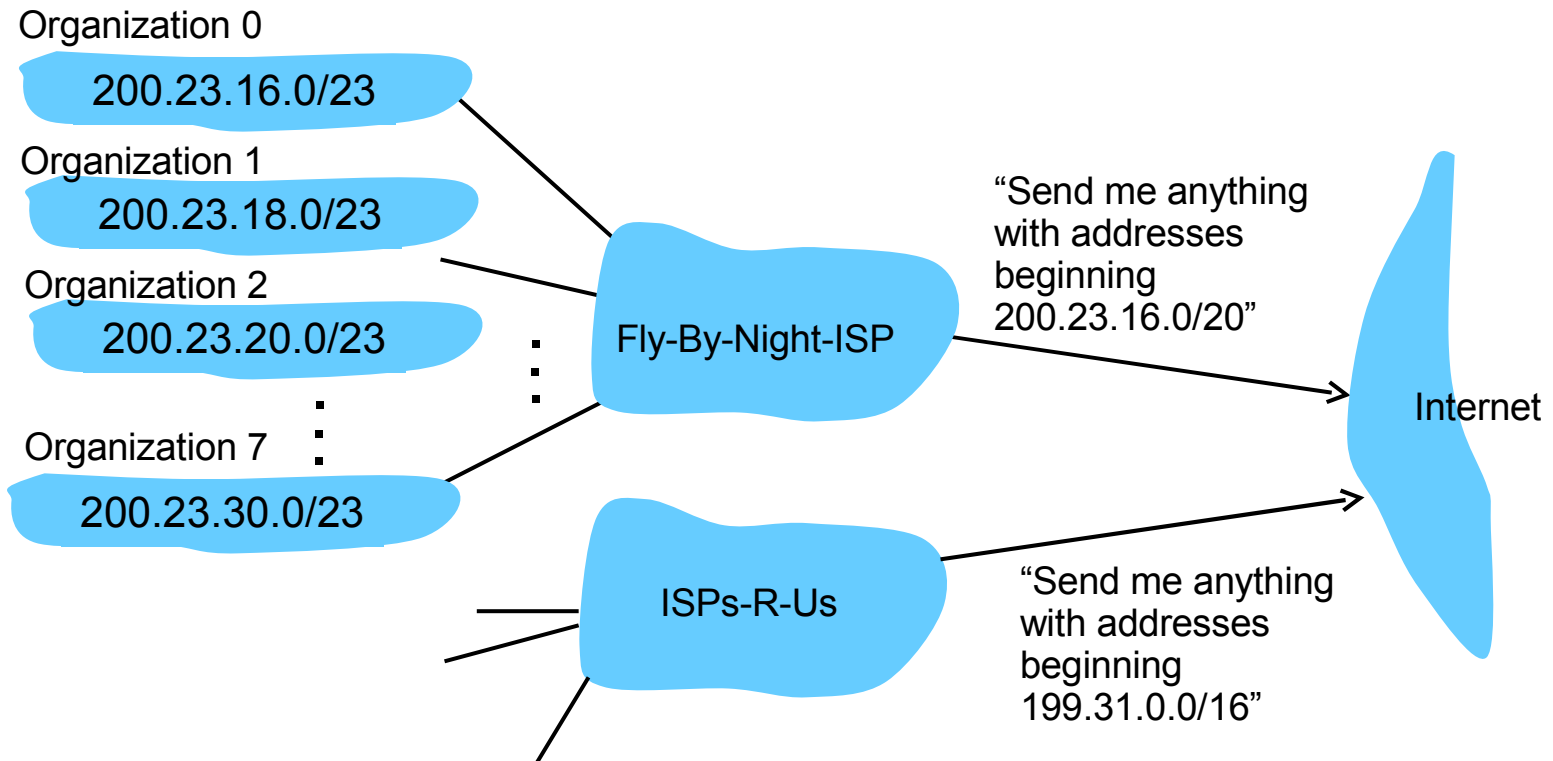
Q: How does *network* get subnet part of IP addr?

A: gets allocated portion of its provider ISP's address space

ISP's block	<u>11001000</u>	<u>00010111</u>	<u>00010000</u>	00000000	200.23.16.0/20
Organization 0	<u>11001000</u>	<u>00010111</u>	<u>00010000</u>	00000000	200.23.16.0/23
Organization 1	<u>11001000</u>	<u>00010111</u>	<u>00010010</u>	00000000	200.23.18.0/23
Organization 2	<u>11001000</u>	<u>00010111</u>	<u>00010100</u>	00000000	200.23.20.0/23
...
Organization 7	<u>11001000</u>	<u>00010111</u>	<u>00011110</u>	00000000	200.23.30.0/23

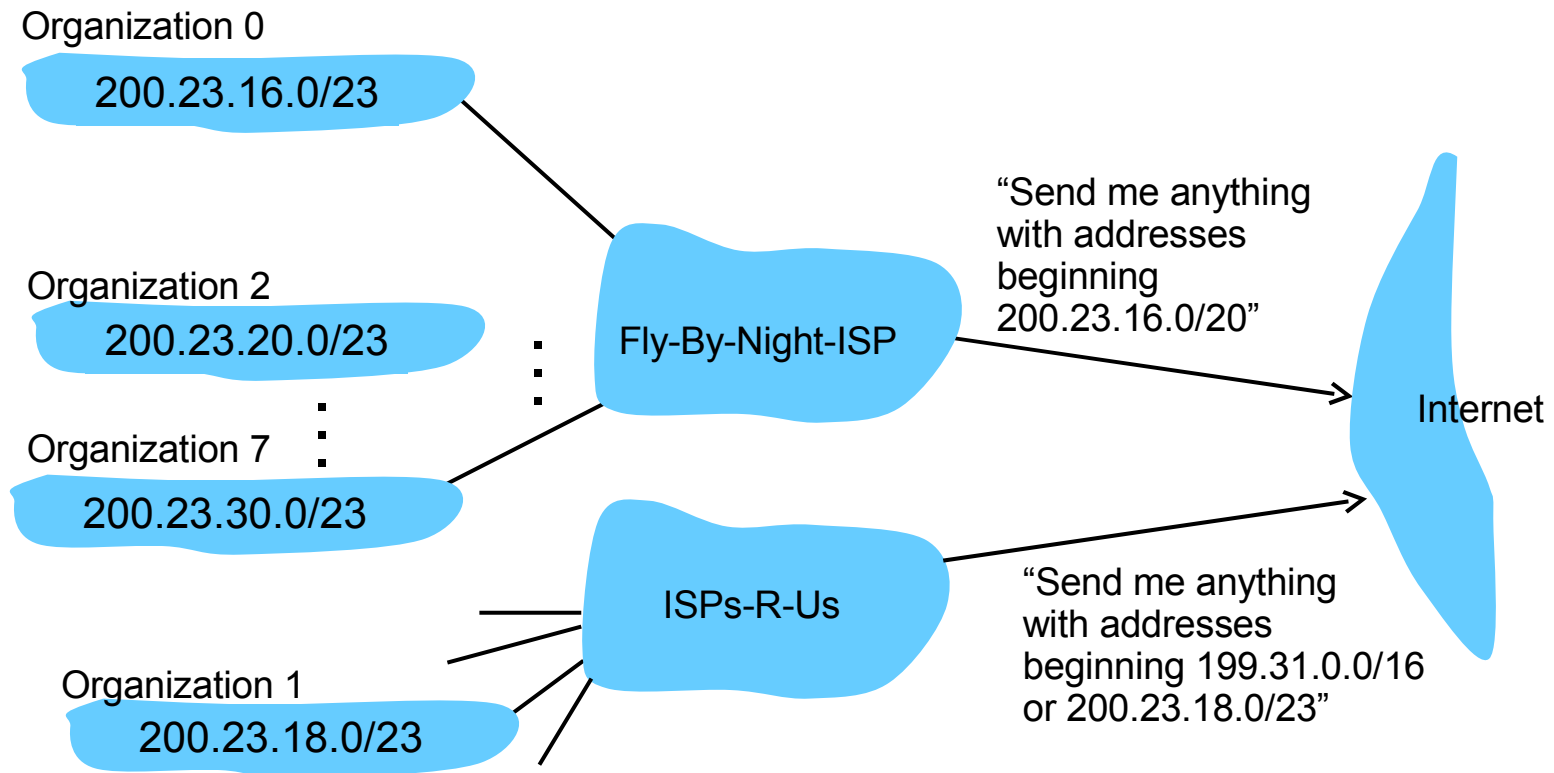
Hierarchical addressing: route aggregation

Hierarchical addressing allows efficient advertisement of routing information:



Hierarchical addressing: more specific routes

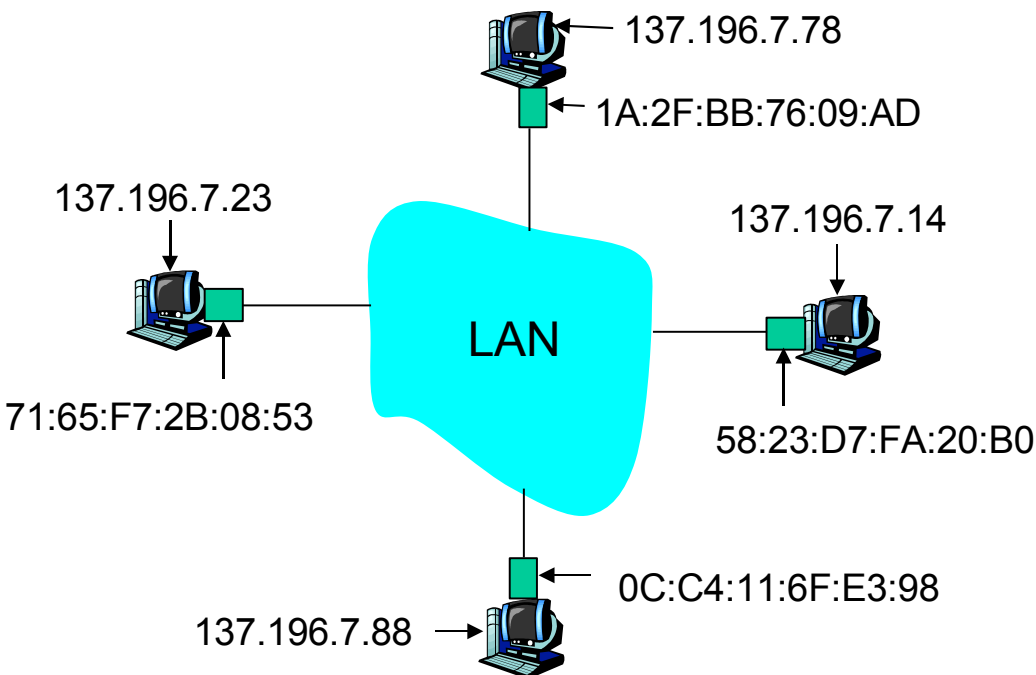
ISPs-R-Us has a more specific route to Organization 1



ARP: Address Resolution Protocol

Question: how to determine MAC address of B knowing B's IP address?

- Each IP node (Host, Router) on LAN has **ARP** table
 - ARP Table: IP/MAC address mappings for some LAN nodes
- < IP address; MAC address; TTL >
- ◆ TTL (Time To Live): time after which address mapping will be forgotten (typically 20 min)




ARP protocol: Same LAN (network)

- A wants to send datagram to B, and B's MAC address not in A's ARP table.
- A **broadcasts** ARP query packet, containing B's IP address
 - ◆ Dest MAC address = FF:FF:FF:FF:FF:FF
 - ◆ all machines on LAN receive ARP query
- B receives ARP packet, replies to A with its (B's) MAC address
 - ◆ frame sent to A's MAC address (unicast)
- A caches (saves) IP-to-MAC address pair in its ARP table until information becomes old (times out)
 - ◆ soft state: information that times out (goes away) unless refreshed
- ARP is “plug-and-play”:
 - ◆ nodes create their ARP tables without intervention from net administrator

Routing to another LAN


walkthrough: **send datagram from A to B via R**

assume A know's B IP address



- Two ARP tables in router R, one for each IP network (LAN)

- A creates datagram with source A, destination B
- A uses ARP to get R's MAC address for 111.111.111.110
- A creates link-layer frame with R's MAC address as dest, frame contains A-to-B IP datagram
- A's adapter sends frame
- R's adapter receives frame
- R removes IP datagram from Ethernet frame, sees its destined to B
- R uses ARP to get B's MAC address
- R creates frame containing A-to-B IP datagram sends to B



DHCP: Dynamic Host Configuration Protocol

Goal: allow host to *dynamically* obtain its IP address from network server when it joins network

Can renew its lease on address in use


Allows reuse of addresses (only hold address while connected and “on”)

Support for mobile users who want to join network (more soon)

DHCP overview:

- ◆ host broadcasts “DHCP discover” msg
- ◆ DHCP server responds with “DHCP offer” msg
- ◆ host requests IP address: “DHCP request” msg
- ◆ DHCP server sends address: “DHCP ack” msg

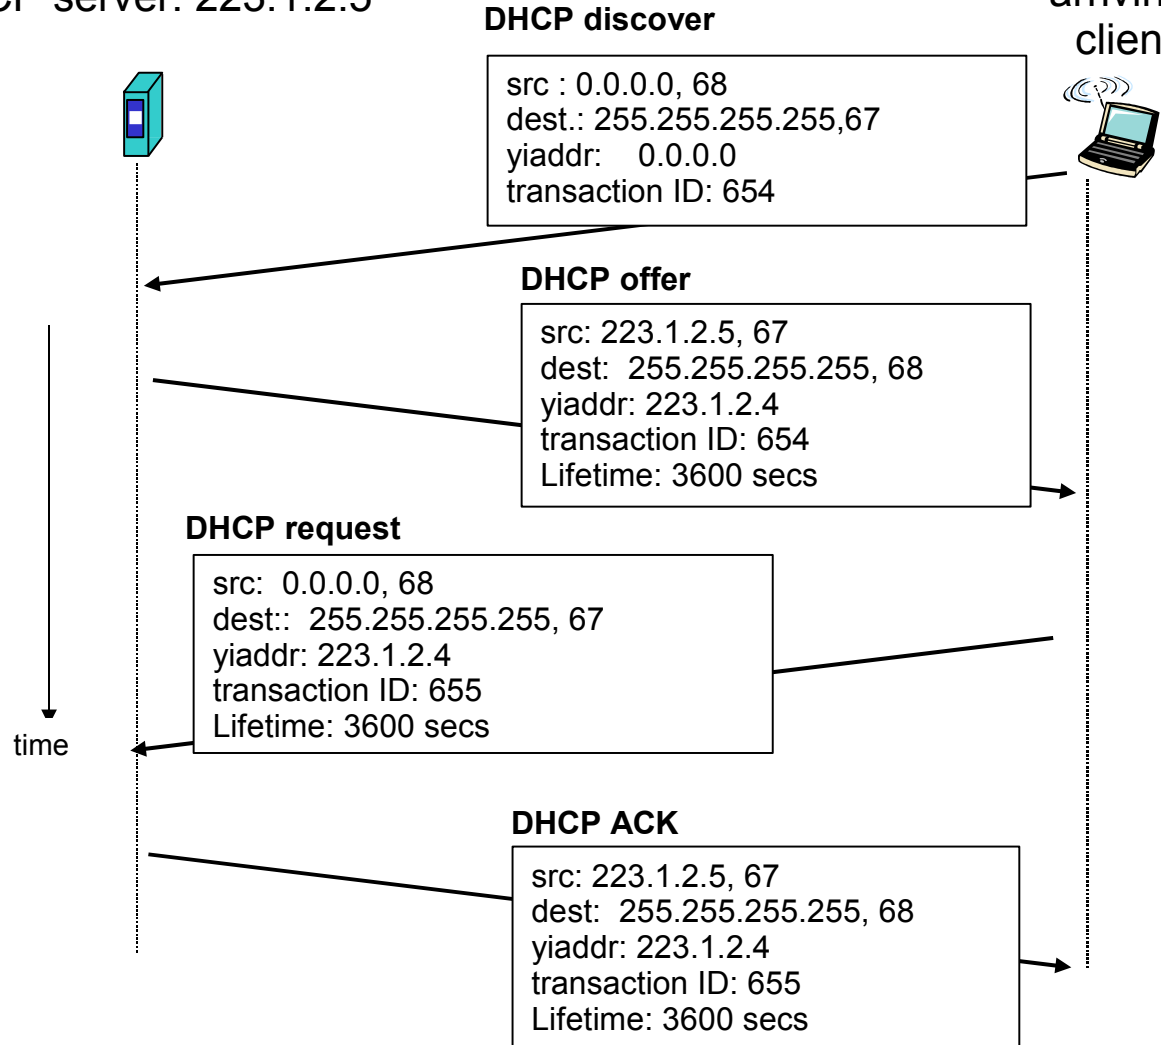
DHCP client-server scenario



DHCP client-server scenario

DHCP server: 223.1.2.5

arriving client



Chapter 4: Network Layer

- 4. 1 Introduction
- 4.2 Virtual circuit and datagram networks
- 4.3 What's inside a router
- **4.4 IP: Internet Protocol**
 - ◆ Datagram format
 - ◆ IPv4 addressing
 - ◆ **ICMP**
 - ◆ IPv6
- 4.5 Routing algorithms
 - ◆ Link state
 - ◆ Distance Vector
 - ◆ Hierarchical routing
- 4.6 Routing in the Internet
 - ◆ RIP
 - ◆ OSPF
 - ◆ BGP
- 4.7 Broadcast and multicast routing

ICMP: Internet Control Message Protocol

- Used by hosts & routers to communicate network-level information
 - ◆ error reporting: unreachable host, network, port, protocol
 - ◆ echo request/reply (used by ping)
- Network-layer “above” IP:
 - ◆ ICMP msgs carried in IP datagrams
- **ICMP message:** type, code plus first 8 bytes of IP datagram causing error

<u>Type</u>	<u>Code</u>	<u>description</u>
0	0	echo reply (ping)
3	0	dest. network unreachable
3	1	dest host unreachable
3	2	dest protocol unreachable
3	3	dest port unreachable
3	6	dest network unknown
3	7	dest host unknown
4	0	source quench (congestion control - not used)
8	0	echo request (ping)
9	0	route advertisement
10	0	router discovery
11	0	TTL expired
12	0	bad IP header

Traceroute and ICMP

- Source sends series of UDP segments to dest
 - ◆ First has TTL =1
 - ◆ Second has TTL=2, etc.
 - ◆ Unlikely port number
 - When nth datagram arrives to nth router:
 - ◆ Router discards datagram
 - ◆ And sends to source an ICMP message (type 11, code 0)
 - ◆ Message includes name of router & IP address
 - When ICMP message arrives, source calculates RTT
 - Traceroute does this 3 times
- Stopping criterion
- UDP segment eventually arrives at destination host
 - Destination returns ICMP “port unreachable” packet (type 3, code 3)
 - When source gets this ICMP, stops.

Chapter 4: Network Layer

- 4. 1 Introduction
- 4.2 Virtual circuit and datagram networks
- 4.3 What's inside a router
- **4.4 IP: Internet Protocol**
 - ◆ Datagram format
 - ◆ IPv4 addressing
 - ◆ ICMP
 - ◆ **IPv6**
- 4.5 Routing algorithms
 - ◆ Link state
 - ◆ Distance Vector
 - ◆ Hierarchical routing
- 4.6 Routing in the Internet
 - ◆ RIP
 - ◆ OSPF
 - ◆ BGP
- 4.7 Broadcast and multicast routing

IPv6

- **Initial motivation:** 32-bit address space soon to be completely allocated.
- Additional motivation:
 - ◆ header format helps speed processing/forwarding
 - ◆ header changes to facilitate QoS

IPv6 datagram format:

- fixed-length 40 byte header
- no fragmentation allowed

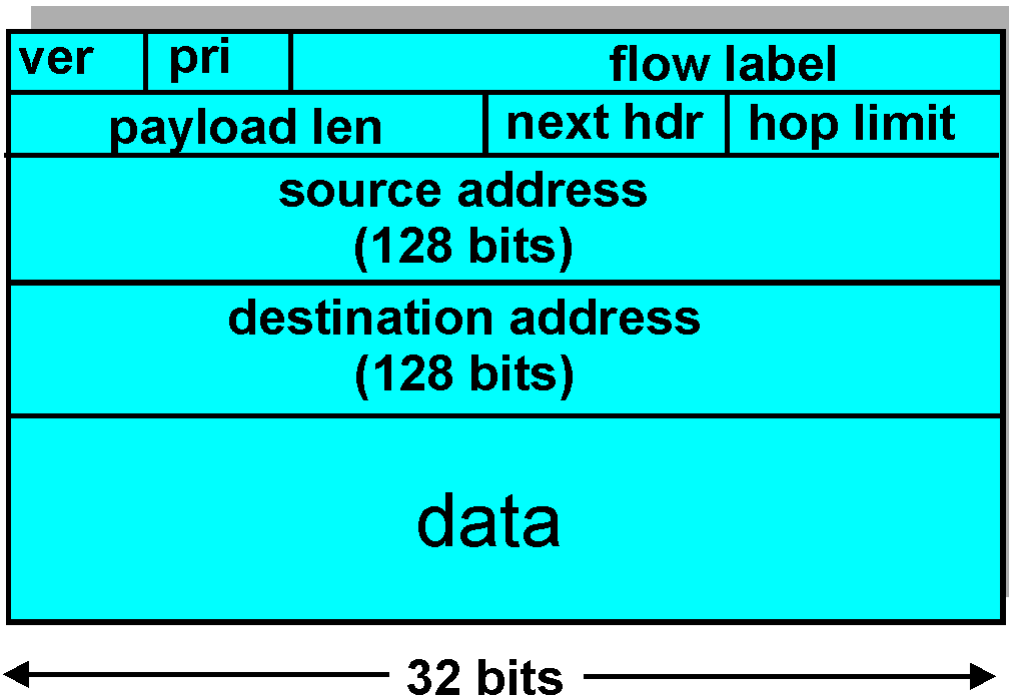
IPv6 Header (Cont)

Priority: identify priority among datagrams in flow

Flow Label: identify datagrams in same “flow.”

(concept of “flow” not well defined).

Next header: identify upper layer protocol for data



Other Changes from IPv4

- *Checksum*: removed entirely to reduce processing time at each hop
- *Options*: allowed, but outside of header, indicated by “Next Header” field
- *ICMPv6*: new version of ICMP
 - ◆ additional message types, e.g. “Packet Too Big”
 - ◆ multicast group management functions

Recap

- IP addressing
 - ◆ addressing, subnets, CIDR
 - ◆ address aggregation
- ARP
 - ◆ Learning other hosts' MAC addresses
 - ◆ Same LAN only
- DHCP
 - ◆ Learning your own IP address
- ICMP
 - ◆ Internet “error messages”
 - ◆ How traceroute works
- IPv6
 - ◆ Differences from IPv4

Next time

- Transitioning to IPv6

- Routing
 - ◆ Link-state routing

 - ◆ Distance-vector routing