IPv4

Christian Grothoff
christian@grothoff.org
http://grothoff.org/christian/

"As the builders say, the larger stones do not lie well without the lesser." – Plato

The Network Layer

- Transports datagrams from sending to receiving host
- Network layer protocols are implemented on *every* host and router

The network layer is commonly referred to as **layer 3**.

The Internet Network Layer

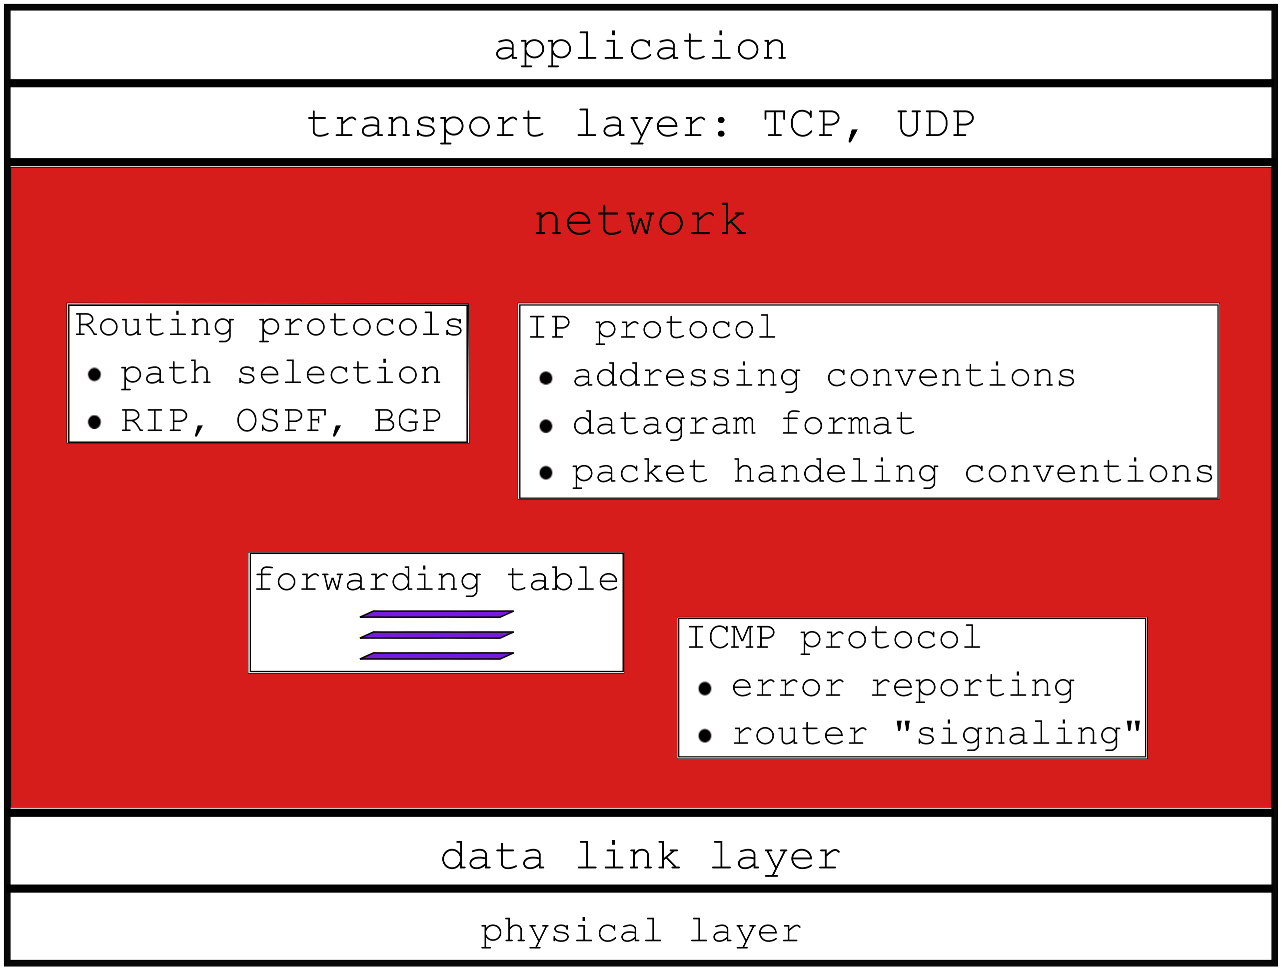

Routing and Forwarding

routing: determine route taken by datagrams from source to destination

forwardig: move datagram from input device to output device

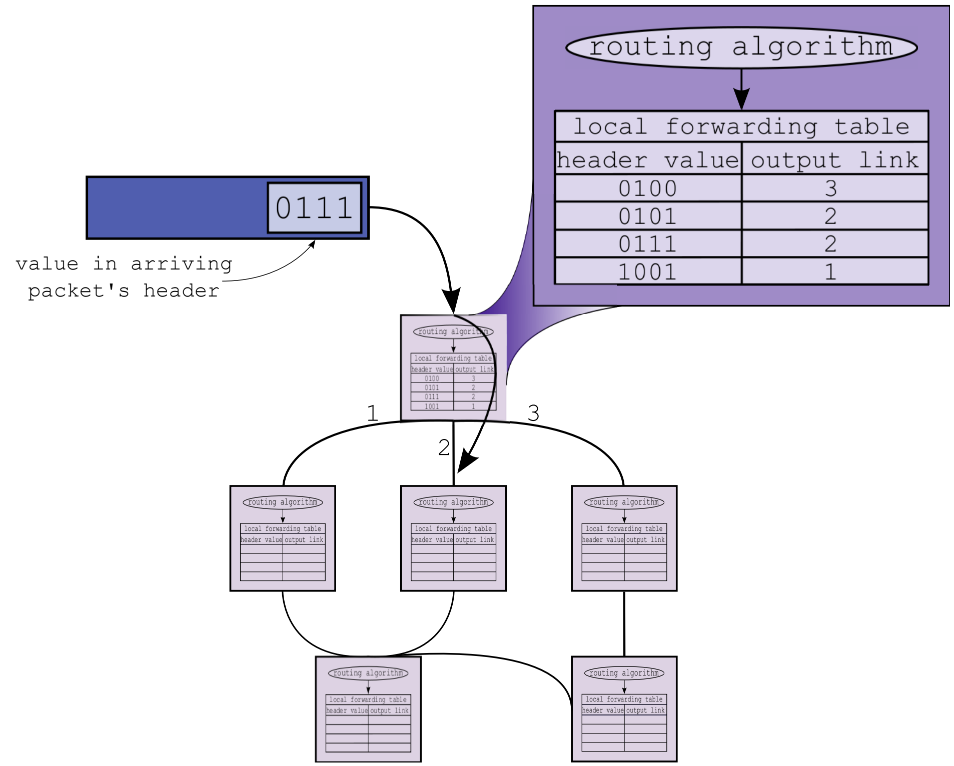

Datagram networks

- No call setup at network layer
- Routers keep *no state* about end-to-end connections
- Packets between the same source-destination pair may take different paths

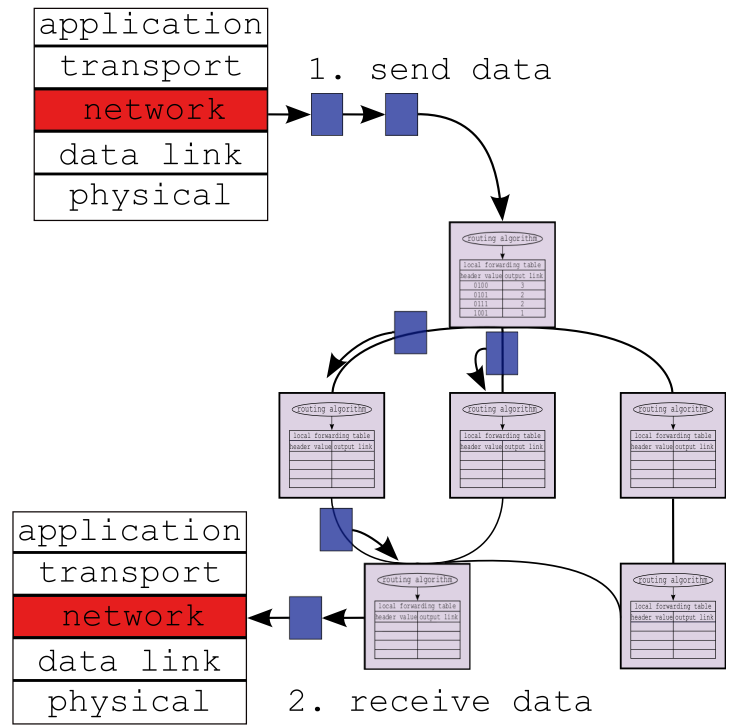

IPv4 Address Format

- 32 bits
- 4 billion[1] possible values
- Notation is dotted decimal, big-endian: 132.149.42.193
- Many values have special meanings

[1] 1 billon = 1,000,000,000

IP Addressing

- IP addresses identify interfaces
- Routers have multiple interfaces
- Hosts typically have only one network interface
- One interface can have multiple IPs

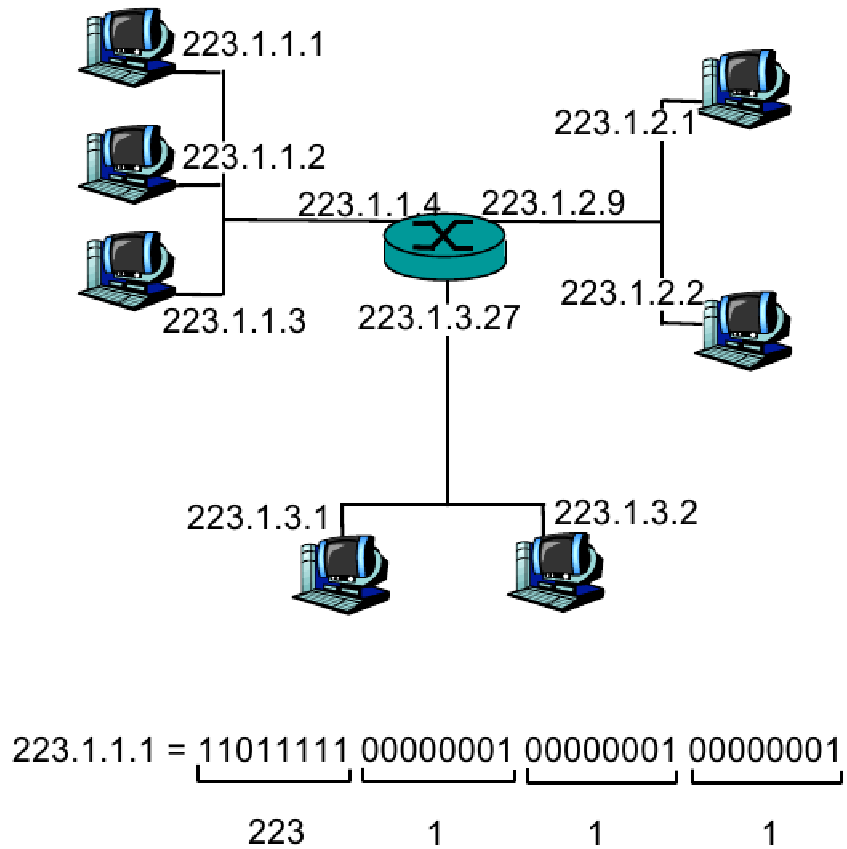

Subnets (1/3)

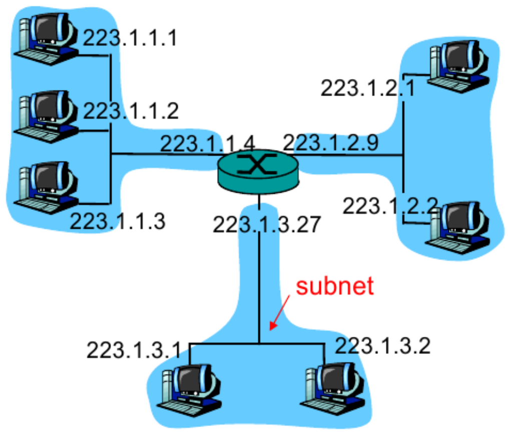

network consisting of 3 subnets

Devices in a **subnet** share the same subnet part (higher-order bits) of the IP address and can physically reach each other without the help of a router.

Subnets (2/3)

To determine the subnets, detach each interface from its host or router, creating islands of isolated networks. Those are the subnets.

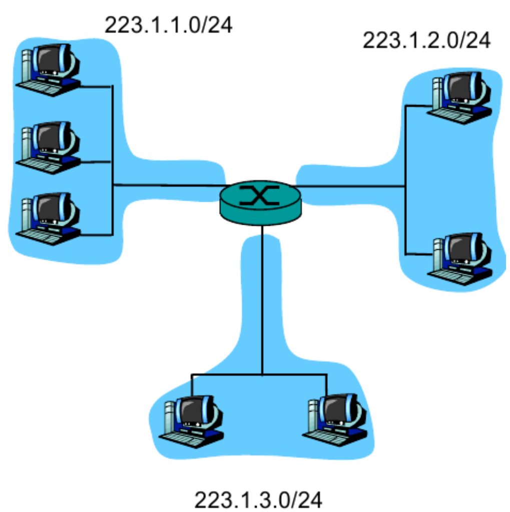

Subnet mask: /24

Question: which **devices** count as routers?

Subnets (3/3)

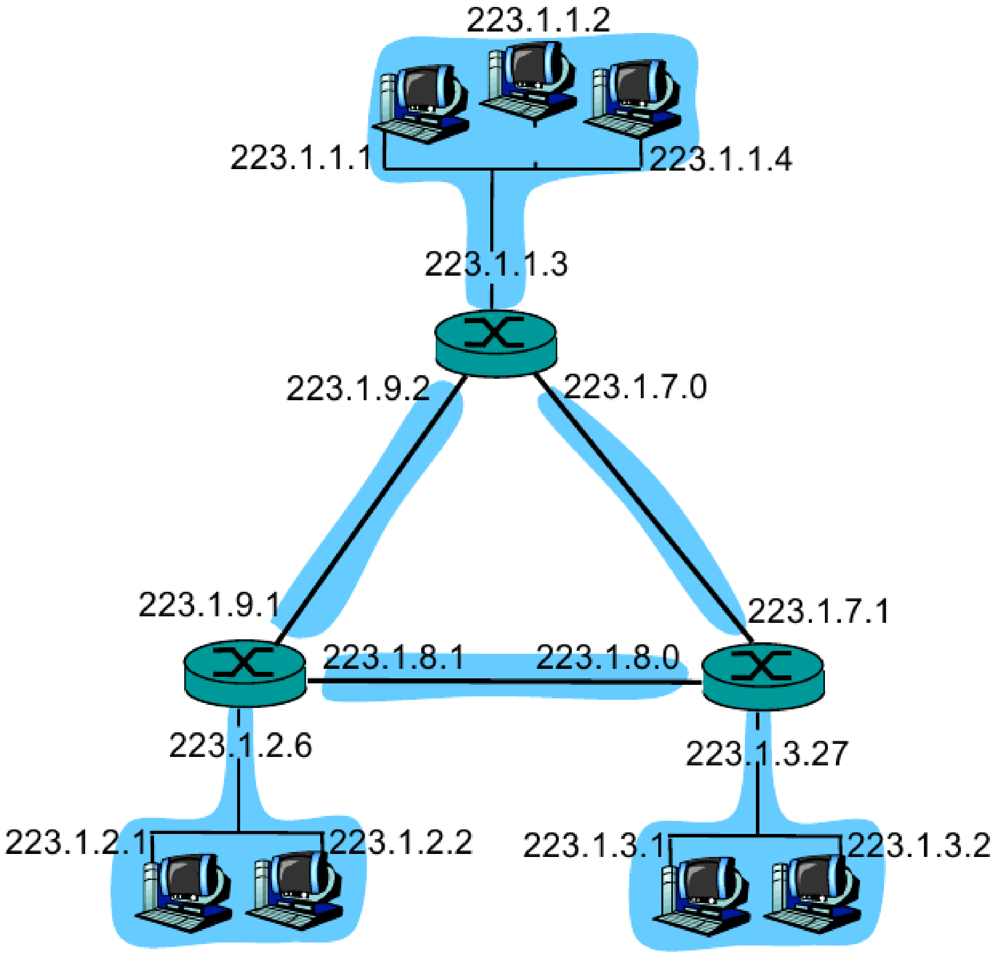
223.1.1.2
223.1.1.1
223.1.1.4
223.1.1.3
223.1.9.2
223.1.7.0
223.1.9.1
223.1.7.1
223.1.8.1
223.1.8.0
223.1.2.6
223.1.3.27
223.1.2.1
223.1.2.2
223.1.3.1
223.1.3.2

Classless InterDomain Routing (CIDR)

- Subnet part of IP address are the higher-order bits
- Address format `a.b.c.d/x` specifies that subnet part has `x` bits

How are subnet identifiers assigned?

Administrators get a portion of their providers ISP's address space:

- ISP has `200.23.16.0/20`
- ISP gives Org 1 `200.23.16.0/23`: *11001000 00010111 0001*0000
- ISP gives Org 2 `200.23.18.0/23`: *11001000 00010111 0001*0010
- ISP gives Org 3 `200.23.20.0/23`: *11001000 00010111 0001*0100
- ...
- ISP gives Org 7 `200.23.30.0/23`: *11001000 00010111 0001*1110

Longest Prefix Matching

Prefix match	Interface
11001000 00010111 00010	0
11001000 00010111 00011000	1
11001000 00010111 00011	2
otherwise	3

Example: *11001000 00010111 00011000 10101010*?

Route aggregation

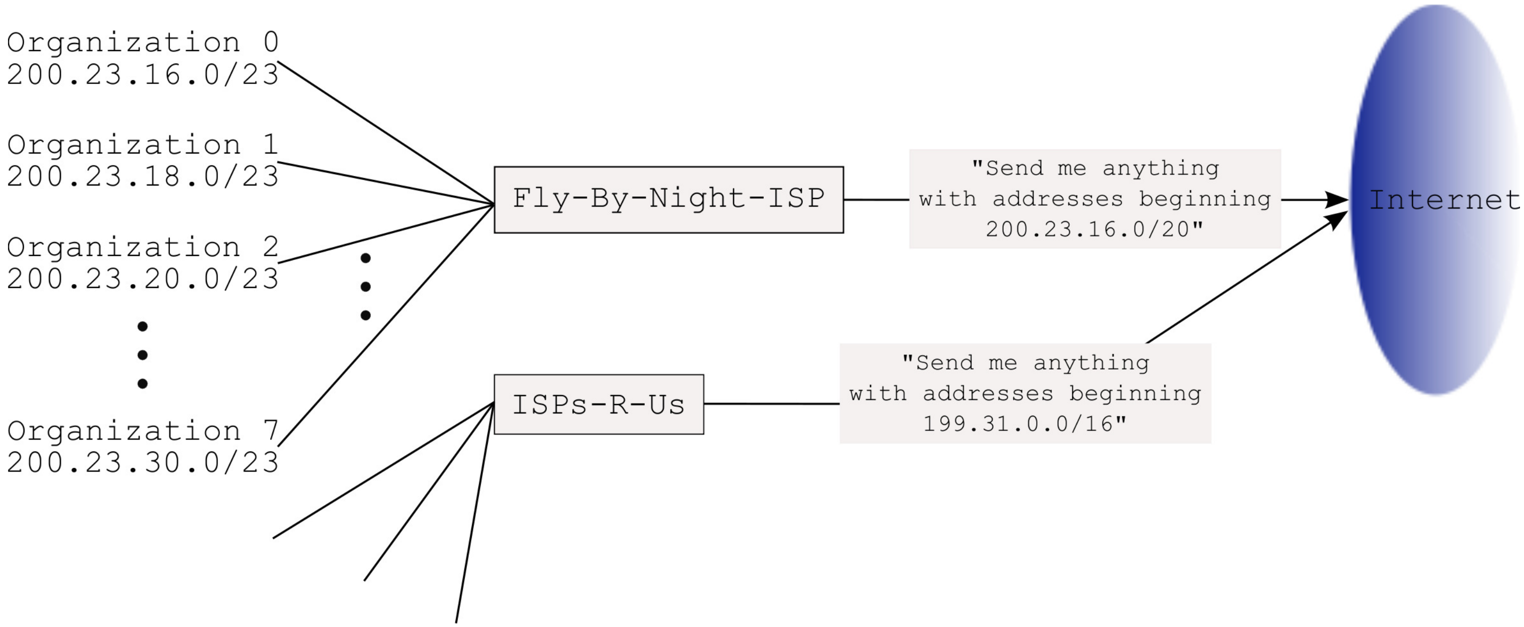

Route aggregation

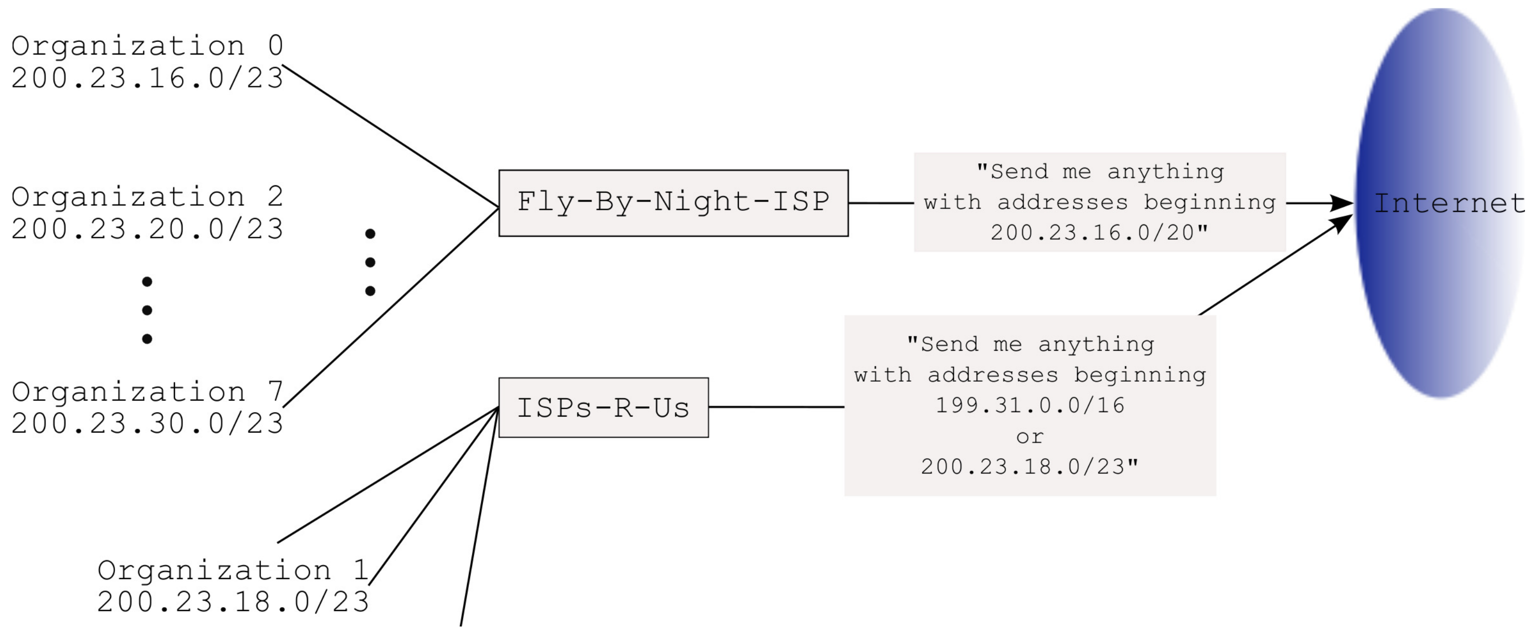

The IP Header

Offset	0-3	4-7	8-15	16-18	19-31
0	Version	Hdr Len	TOS	Total Length	
32	Identification			Flags	Frag. Offset
64	Time to Live		Protocol	Hdr. Checksum	
96	Source Address				
128	Destination Address				
160	Options (optional)				
160+	Data				

Where is the IP Header?

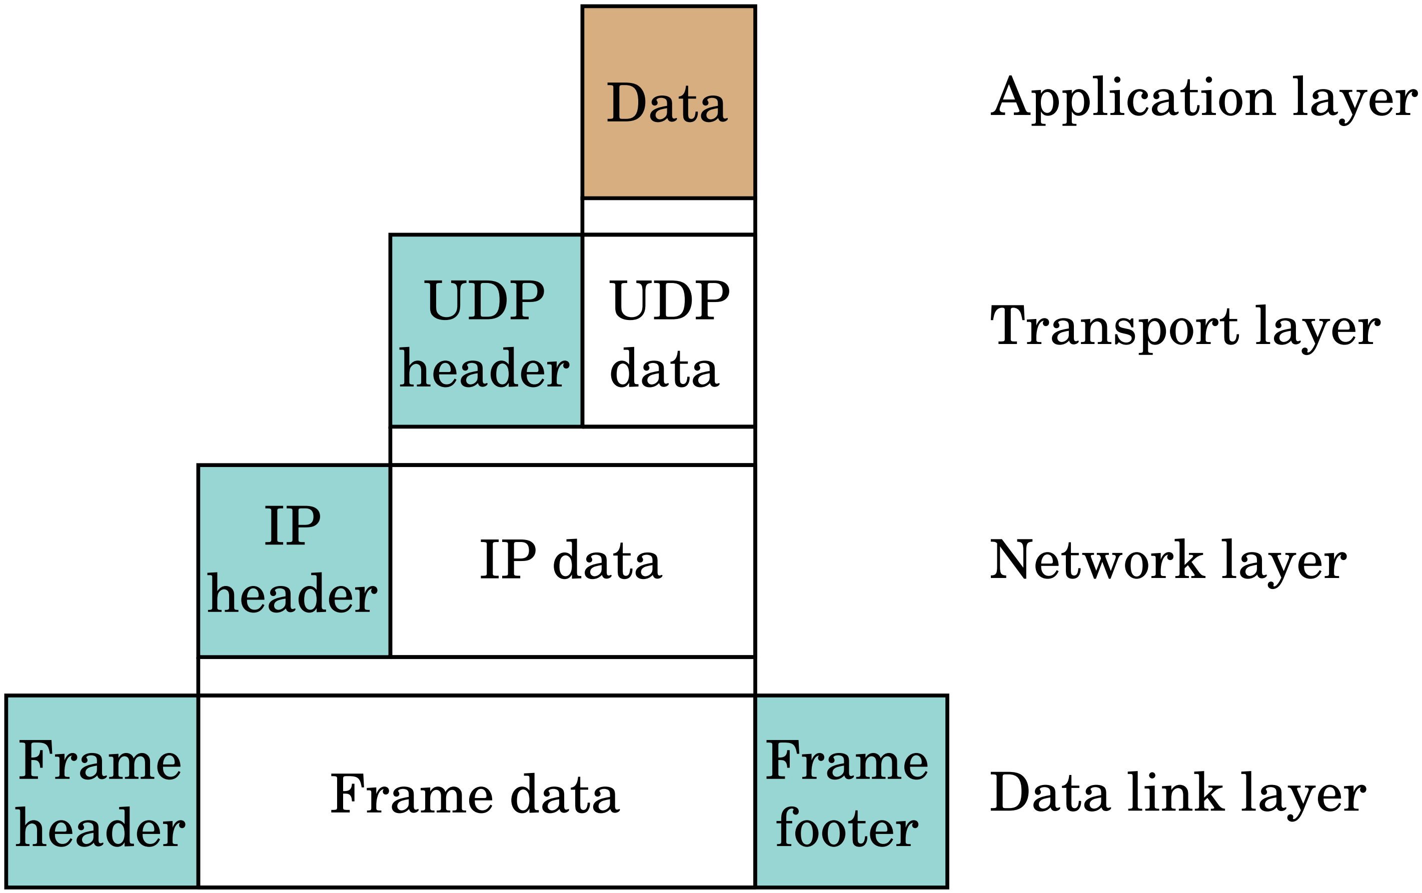

IP Fragmentation & Reassembly

- Network links have a Maximum Transfer Unit (MTU)
- IP datagrams that exceed the MTU are divided on the network
- Reassembly only happens at the final destination
- IP header bits are used to identify, order related fragments

Example

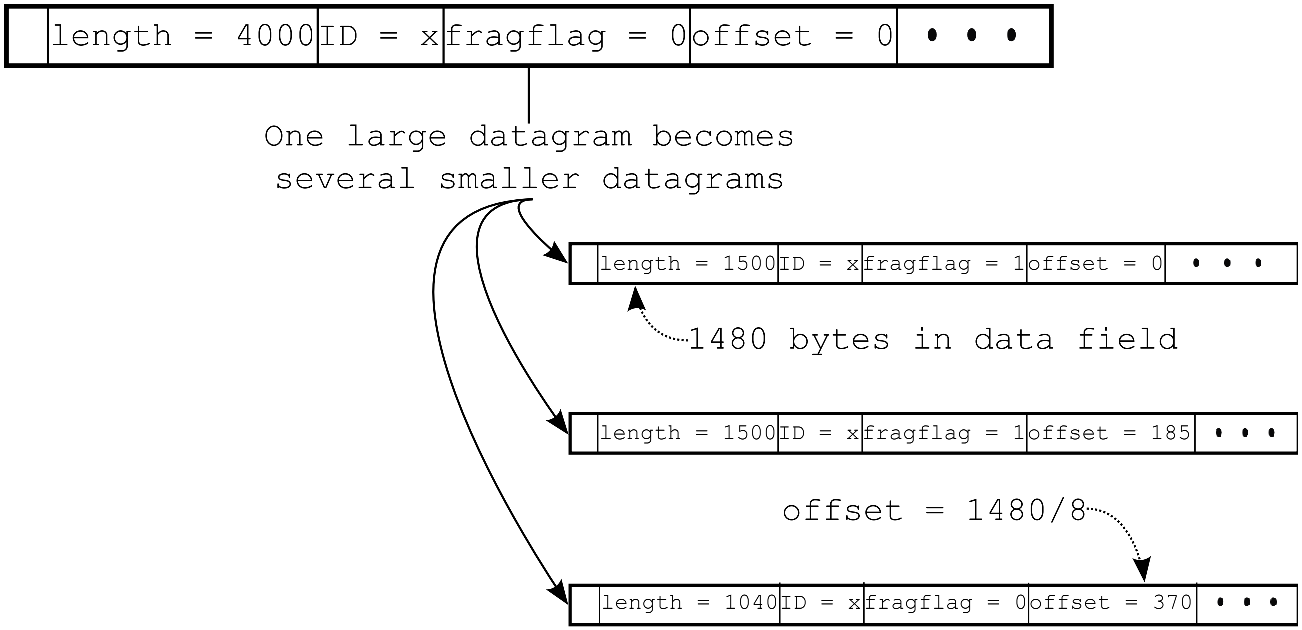
length = 4000
ID = x
fragflag = 0
offset = 0
One large datagram becomes several smaller datagrams
length = 1500
ID = x
fragflag = 1
offset = 0
1480 bytes in data field
length = 1500
ID = x
fragflag = 1
offset = 185
offset = 1480/8
length = 1040
ID = x
fragflag = 0
offset = 370

How does a host get an IP address?

- Hard-coded by system administrator in configuration
- Determined by network service: DHCP
- other options ($\Rightarrow$ IPv6)

The Dynamic Host Configuration Protocol

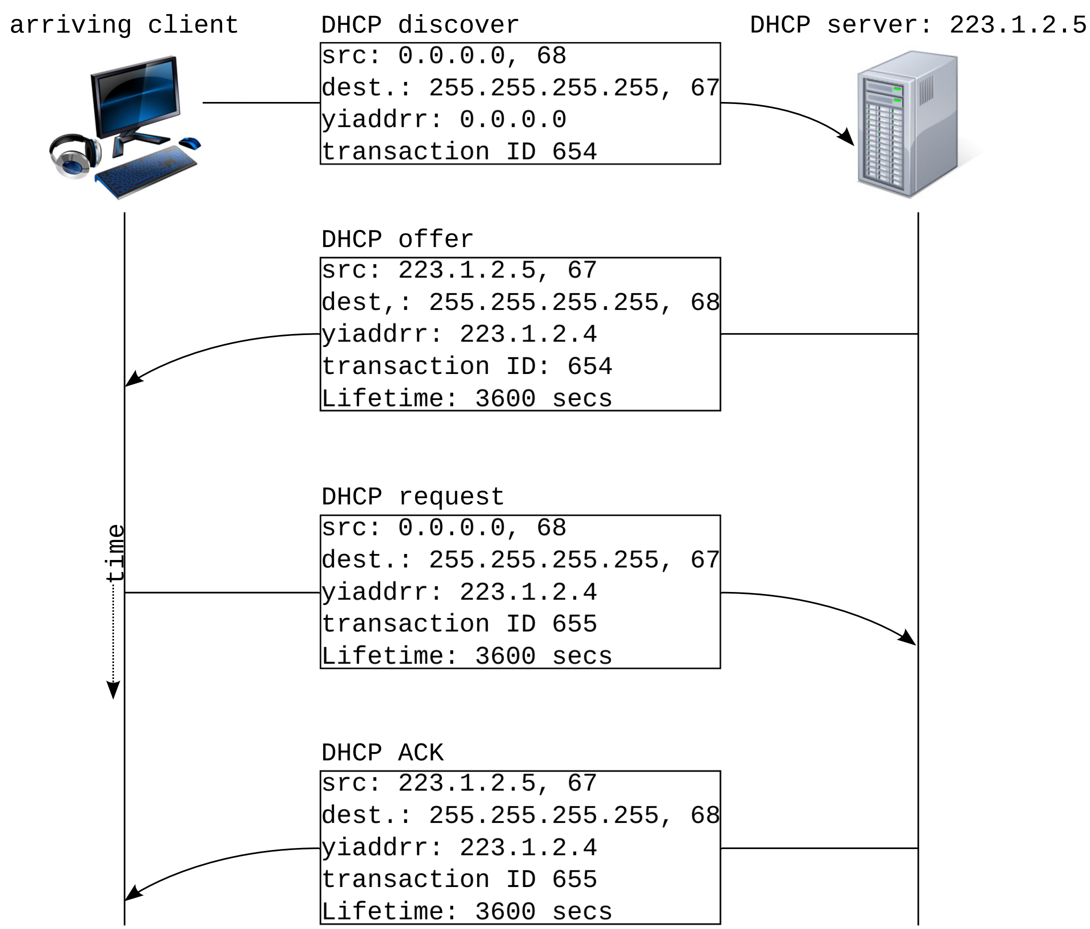

The Internet Control Message Protocol

Communicate network-level information:

- Error reporting
- Diagnostics

Type	Code	Function
3	0	dest net unreachable
3	1	dest host unreachable
3	2	dest prot unreachable
3	3	dest port unreachable
3	6	dest net unknown
3	7	dest host unknown
0	0	Echo reply
8	0	Echo request
4	0	source quench
9	0	router advertisement
10	0	router discovery
11	0	TTL expired
12	0	bad IP header

Traceroute

- Source send series of UDP segments to destination
- First has TTL=1, second TTL=2, etc.
- n-th router discards and sends ICMP TTL expired
- Ultimate destination returns ICMP PORT UNREACHABLE
- `traceroute` measures average RTTs for each step

How does an ISP get a block of addresses?

- Internet Corporation for Assigned Names and Numbers (ICANN)

How does an ISP get a block of addresses?

- ICANN (not really)
- Internet Assigned Numbers Authority (IANA)

How does an ISP get a block of addresses?

- ICANN (not really)
- IANA (not really)
- Regional Internet Registry (RIR)

How does an ISP get a block of addresses?

- ICANN (not really)
- IANA (not really)
- RIR (IPv4 only until 2011-2012)

"Their Internet usage is growing very rapidly, and even they can do the math: If everyone in China needed an IPv4 address — just one — this country would use up one third of the entire public IP address space." – Vinton Cerf

Questions

?

Copyright