IP Next Generation (IPv6)

Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

is677-98/

- Limitations of current Internet Protocol (IP)
- How many addresses do we need?
- IPv6 Addressing
- IPv6 header format

IP Addresses

- **Example**: 164.107.134.5
  = 1010 0100 : 0110 1011 : 1000 0110 : 0000 0101
  = A4:6B:86:05 (32 bits)
- Maximum number of address = 2^{32} = 4 Billion
- Class A Networks: 15 Million nodes
- Class B Networks: 64,000 nodes or less
- Class C Networks: 254 nodes or less

IP Address Format

- Three all-zero network numbers are reserved
- 127 Class A + 16,381 Class B + 2,097,151 Class C networks = 2,113,659 networks total
- Class B is most popular.
- 20% of Class B were assigned by 7/90 and doubling every 14 months ⇒ Will exhaust by 3/94
- Question: Estimate how big will you become?
  Answer: More than 256!
  Class C is too small. Class B is just right.

How Many Addresses?

- 10 Billion people by 2020
- Each person will be served by more than one computer
- Assuming 100 computers per person $\Rightarrow 10^{12}$ computers
- More addresses may be required since
  - Multiple interfaces per node
  - Multiple addresses per interface
- Some believe 2^6 to 2^8 addresses per host
- Safety margin $\Rightarrow 10^{15}$ addresses
- IPng Requirements $\Rightarrow 10^{12}$ end systems and 10^9 networks. Desirable 10^{12} to 10^{15} networks

Address Size

- H Ratio = $\log_{10}$(number of objects)/available bits
- 2^n objects with n bits: H-Ratio = $\log_{10}2 = 0.30103$
- French telephone moved from 8 to 9 digits at 10^7 households $\Rightarrow$ H = 0.26 (assuming 3.3 bits/digit)
- US telephone expanded area codes with 10^8 subscribers $\Rightarrow$ H = 0.24
- SITA expanded 7-character address at 64k nodes $\Rightarrow$ H = 0.14 (assuming 5 bits/char)
- Physics/space science net stopped at 15000 nodes using 16-bit addresses $\Rightarrow$ H = 0.26
- 3 Million Internet hosts currently using 32-bit addresses $\Rightarrow$ H = 0.20 $\Rightarrow$ A few more years to go

IPv6 Addresses

- 128-bit long. Fixed size
- $2^{128} = 3.4\times10^{38}$ addresses
  $\Rightarrow 665\times10^{21}$ addresses per sq. m of earth surface
- If assigned at the rate of $10^6/\mu s$, it would take 20 years
- Expected to support 8×10^{17} to 2×10^{33} addresses
  $8\times10^{17} \Rightarrow$ 1,564 address per sq. m
- Allows multiple interfaces per host.
- Allows multiple addresses per interface
- Allows unicast, multicast, anycast
- Allows provider based, site-local, link-local

Colon-Hex Notation

- **Dot-Decimal**: 127.23.45.88
- **Colon-Hex:**
  FEDC:0000:0000:0000:3243:0000:0000:ABCD
  - Can skip leading zeros of each word
  - Can skip one sequence of zero words, e.g.,
    FEDC::3243:0000:0000:ABCD
    ::3243:0000:0000:ABCD
  - Can leave the last 32 bits in dot-decimal, e.g.,
    ::127.23.45.88
  - Can specify a prefix by /length, e.g.,
    2345:BA23:7::/40

Initial IPv6 Prefix Allocation

Allocation	Prefix	Allocation	Prefix
Reserved	0000 0000	Unassigned	101
Unassigned	0000 0001	Unassigned	110
NSAP	0000 001	Unassigned	1110
IPX	0000 010	Unassigned	1111 0
Unassigned	0000 011	Unassigned	1111 10
Unassigned	0000 1	Unassigned	1111 110
Unassigned	0001	Unassigned	1111 1110
Unassigned	001	Unassigned	1111 1110 0
Provider-based	010	Link-Local	1111 1110 10
Unassigned	011	Site-Local	1111 1110 11
Geographic	100	Multicast	1111 1111

Local-Use Addresses

- Link Local: Not forwarded outside the link, FE:80::xxx

10 bits	n bits	118-n
1111 1110 10	0	Interface ID

- Site Local: Not forwarded outside the site, FE:C0::xxx

10 bits	n bits	m bits	118-n-m bits
1111 1110 11	0	Subnet ID	Interface ID

- Provides plug and play

Multicast Addresses

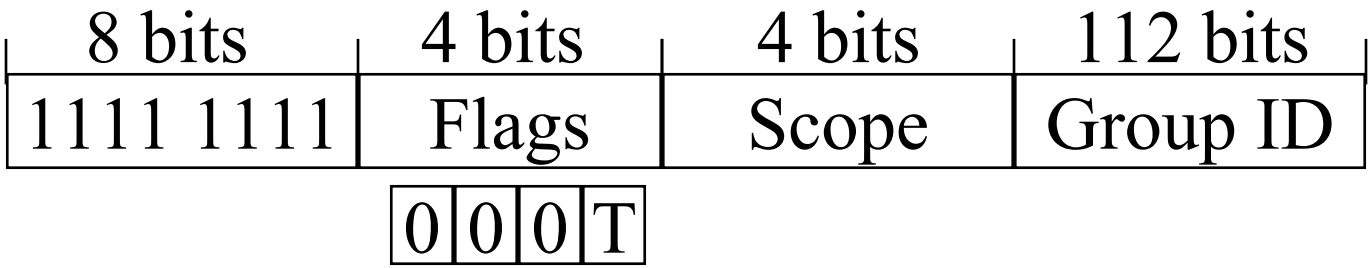

- T = 0 ⇒ Permanent (well-known) multicast address, 1 ⇒ Transient
- Scope: 1 Node-local, 2 Link-local, 5 Site-local, 8 Organization-local, E Global
- Predefined: 1 ⇒ All nodes, 2 ⇒ Routers, 1:0 ⇒ DHCP servers

Multicast Addresses (Cont)

- Example: 43 ⇒ Network Time Protocol Servers
  - FF01::43 ⇒ All NTP servers on this node
  - FF02::43 ⇒ All NTP servers on this link
  - FF05::43 ⇒ All NTP servers in this site
  - FF08::43 ⇒ All NTP servers in this organization
  - FF0F::43 ⇒ All NTP servers in the Internet

Header

- IPv6:

Version	Priority	Flow Label	
Payload Length		Next Header	Hop Limit
Source Address			
Destination Address			

- IPv4:

Version	IHL	Type of Service	Total Length
Identification		Flags	Fragment Offset
Time to Live	Protocol	Header Checksum	
Source Address			
Destination Address			
Options			Padding

Protocol and Header Types

Decimal	Keyword	Header Type
	HBH	Hop-by-hop (IPv6)
1	ICMP	Internet Control Message (IPv4)
2	IGMP	Internet Group Management (IPv4)
2	ICMP	Internet Control Message (IPv6)
3	GGP	Gateway-to-Gateway
4	IP	IP in IP (IPv4 Encaptulation)
5	ST	Stream
6	TCP	
17	UDP	
29	ISO-TP4	
43	RH	Routing Header (IPv6)
44	FS	Fragmentation Header (IPv6)
45	IDRP	Interdomain Routing
51	AH	Authentication header (IPv6)
52	ESP	Encrypted Security Payload
59	Null	No next header
60	ISO-IP	CLNP
88	IGRP	
89	OSPF	Open Shortest Path First

IPv6 vs IPv4

- 1995 vs 1975
- IPv6 only twice the size of IPv4 header
- Only version number has the same position and meaning as in IPv4
- Removed: header length, type of service, identification, flags, fragment offset, header checksum
- Datagram length replaced by payload length
- Protocol type replaced by next header
- Time to live replaced by hop limit
- Added: Priority and flow label
- All fixed size fields.

- No optional fields. Replaced by extension headers.
- 8-bit hop limit = 255 hops max (Limits looping)
- Next Header = 6 (TCP), 17 (UDP),

Extension Headers

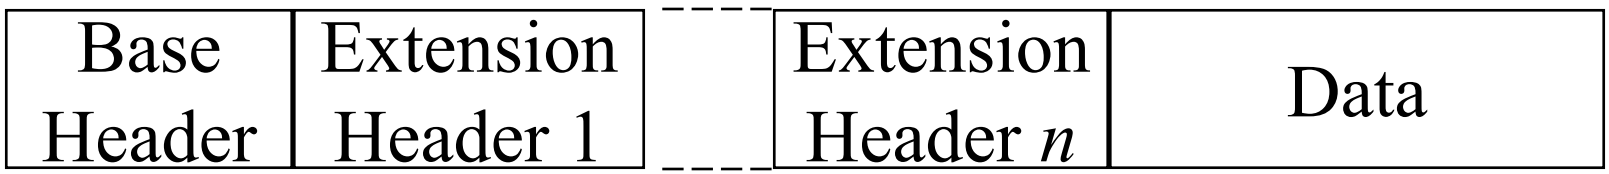

- Most extension headers are examined only at destination
- Routing: Loose or tight source routing
- Fragmentation: All IPv6 routers can carry 536 Byte payload
- Authentication
- Security Encaptulation: Confidentiality
- Hop-by-Hop Option
- Destination Options:

Extension Header (Cont)

- Only Base Header:

Base Header Next = TCP	TCP Segment

- Only Base Header and One Extension Header:

Base Header Next = TCP	Route Header Next = TCP	TCP Segment

- Only Base Header and Two Extension Headers:

Base Header Next = TCP	Route Header Next = Auth	Auth Header Next = TCP	TCP Segment

Routing Header

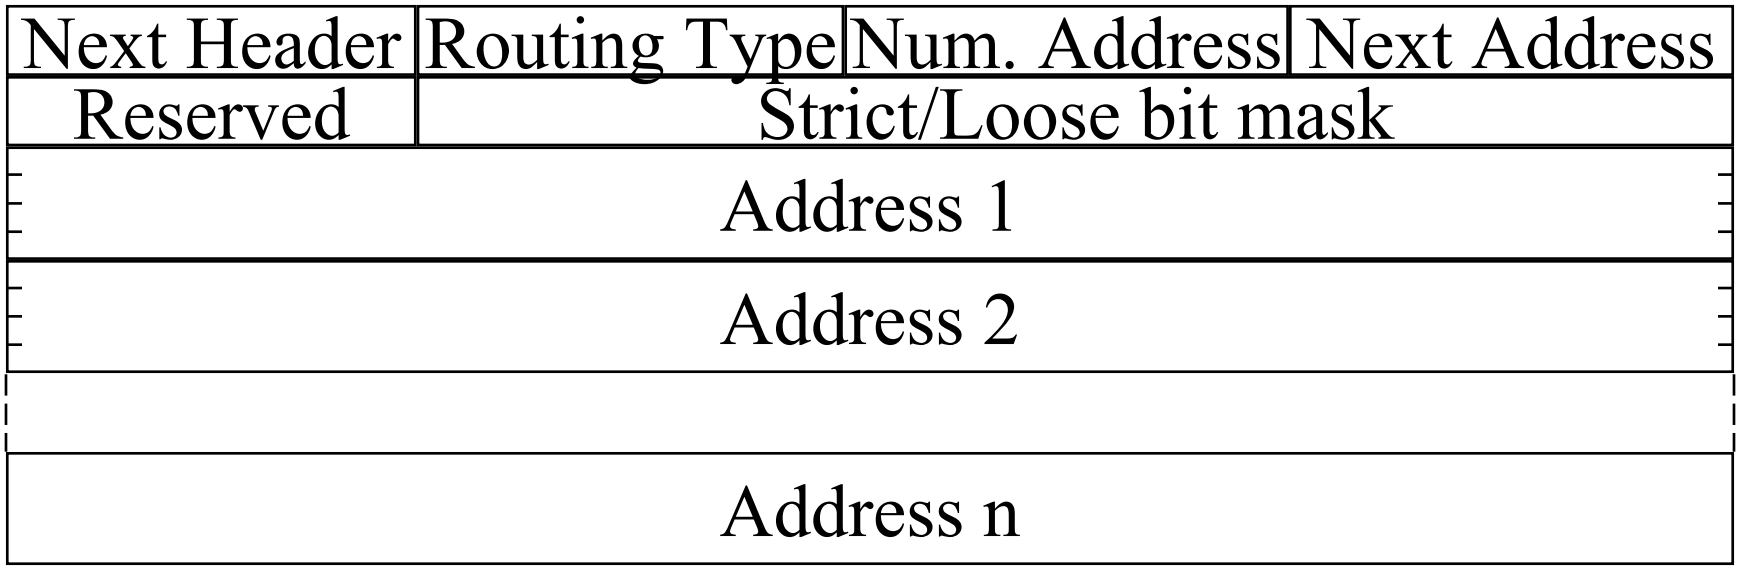

- Strict ⇒ Discard if Address[Next-Address] ≠ neighbor
- Type = 0 ⇒ Current source routing
- Type > 0 ⇒ Policy based routing (later)
- New Functionality: Provider selection, Host mobility, Auto-readdressing (route to new address)

Provider Selection

- Possible using routing extension header
- Source specified intermediate systems
- No preference: H1, H2
- P1 Preferred: H1, P1, H2
- H1 becomes Mobile: H1, PR, P1, H2

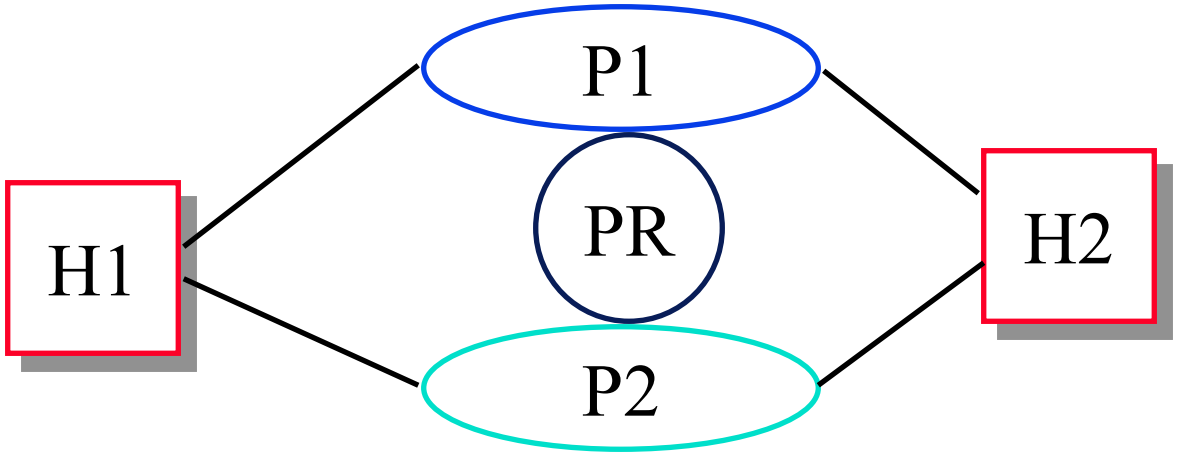

Summary

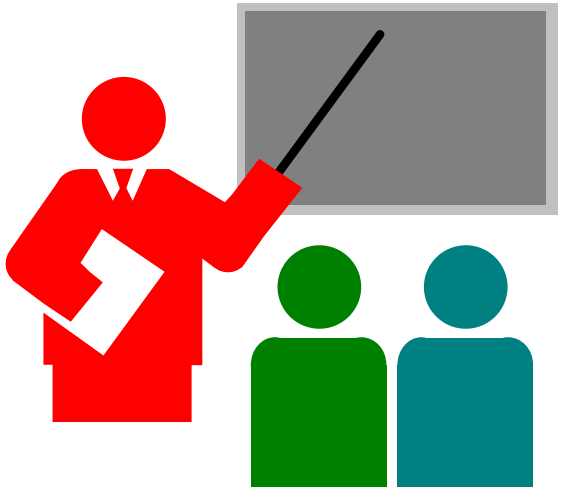

- IPv6 uses 128-bit addresses
- Allows provider-based, site-local, link-local, multicast, anycast addresses
- Fixed header size. Extension headers instead of options. Extension headers for provider selection, security

Homework

- Read Section 16.5 of Stallings
- Submit answer to Exercise 16.19
- Due Date: Next Class