IP Next Generation (IPv6)

Raj Jain
The Ohio State University
Columbus, OH 43210
Jain@CIS.Ohio-State.Edu
http://www.cis.ohio-state.edu/~jain/

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

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: 250 nodes or less

IP Address

Class A:

0	Network	Local	
1	7	24	bits

Class B:

10	Network	Local	
2	14	16	bits

Class C:

110	Network	Local	
3	21	8	bits

Class D:

1110	Host Group (Multicast)	
4	28	bits

Local = Subnet + Host (Variable length)

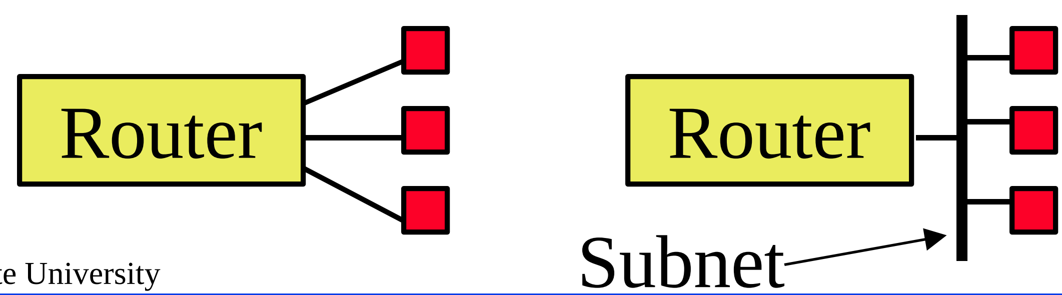

IP Address Format

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?

0 Billion people by 2020

ach person will be served by more than one comput

ssuming 100 computers per person $\Rightarrow 10^{12}$ compu

ore addresses may be required since

- Multiple interfaces per node
- Multiple addresses per interface

ome believe 2^6 to 2^8 addresses per host

afety margin $\Rightarrow 10^{15}$ addresses

Png Requirements $\Rightarrow 10^{12}$ end systems and 10^9 etworks. 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 ⇒ H = 0.26 (assuming 3.3 bits/digit)
- US telephone expanded area codes with 10^8 subscribers ⇒ H = 0.24
- SITA expanded 7-character address at 64k nodes ⇒ H = 0.14 (assuming 5 bits/char)
- Physics/space science net stopped at 15000 nodes using 16-bit addresses ⇒ H = 0.26
- 3 Million Internet hosts currently using 32-bit addresses ⇒ H = 0.20 ⇒ A few more years to go

IPv6 Addresses

28-bit long. Fixed size

8 = 3.4×10^{38} addresses

⇒ 665×10^{21} addresses per sq. m of earth surface

f assigned at the rate of $10^{6}/\mu s$, it would take 20 yea

xpected to support 8×10^{17} to 2×10^{33} addresses

$\times 10^{17} \Rightarrow$ 1,564 address per sq. m

llows multiple interfaces per host.

llows multiple addresses per interface

llows unicast, multicast, anycast

llows provider based, site-local, link-local

5% of the space is unassigned

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 <u>one</u> 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

location	Prefix	Allocation	Prefix
served	0000 0000	Unassigned	101
assigned	0000 0001	Unassigned	110
AP	0000 001	Unassigned	1110
X	0000 010	Unassigned	1111 0
assigned	0000 011	Unassigned	1111 10
assigned	0000 1	Unassigned	1111 110
assigned	0001	Unassigned	1111 1110
assigned	001	Unassigned	1111 1110 0
ovider-based	010	Link-Local	1111 1110 10
assigned	011	Site-Local	1111 1110 11
ographic	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

8 bits	4 bits	4 bits	112 bits
1111 1111	Flags	Scope	Group ID

0	0	0	T

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

Pv6:

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

IPv4:

'ersion	IHL	Type of Service		Total Length	
Identification			Flags	Fragment Offset	
'ime 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

Base Header	Extension Header 1	-----	Extension Header *n*	Data

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

Extension Header (Cont)

Only Base Header:

ase Header ext = TCP	TCP Segment

Only Base Header and One Extension Header:

ase Header ext = TCP	Route Header Next = TCP	TCP Segment

Only Base Header and Two Extension Headers:

ase Header ext = TCP	Route Header Next = Auth	Auth Header Next = TCP	TCP Segmen

Routing Header

ext Header	Routing Type	Num. Address	Next Addre
Reserved	Strict/Loose bit mask		
Address 1			
Address 2			
Address n			

- Strict $\Rightarrow$ Discard if Address[Next-Address] $\neq$ neighb
- Type = 0 $\Rightarrow$ Current source routing
- Type > 0 $\Rightarrow$ 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

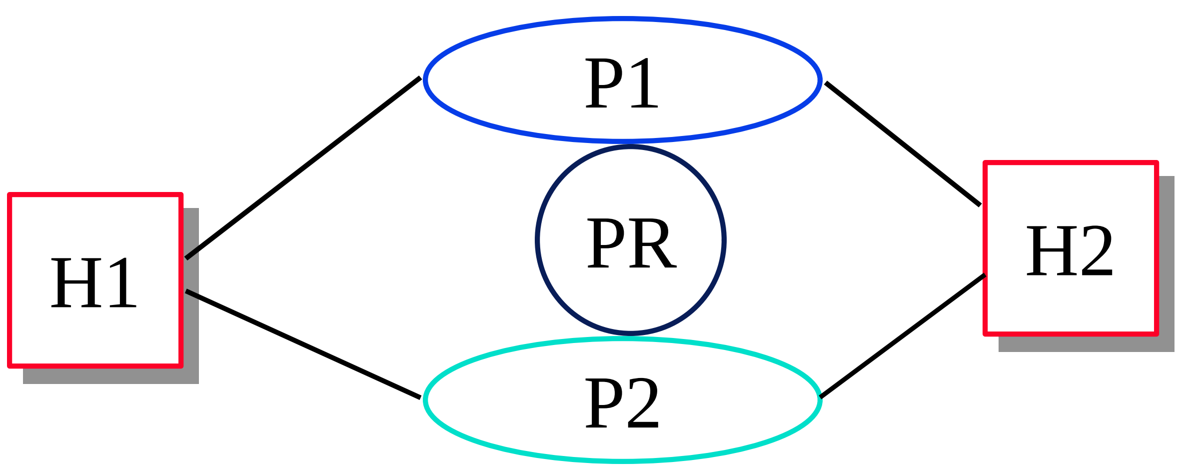

IPv6 Features

- Larger addresses
- Flexible header format
- Improved options
- Support for resource allocation
- Provision for protocol extension
- Built-in Security:
  Both authentication and confidentiality

Address Autoconfiguration

- Allows plug and play
- BOOTP and DHCP are used in IPv4
- DHCPng will be used with IPv6
- Two Methods: Stateless and Stateful
- Stateless:
  - A system uses link-local address as source and multicasts to "All routers on this link"
  - Router replies and provides all the needed prefix info

Address Autoconfiguration (Cont)

- All prefixes have a associated lifetime
- System can use link-local address permanently if no router

Stateful:

- Problem w stateless: Anyone can connect
- Routers ask the new system to go DHCP server (by setting managed configuration bit)
- System multicasts to "All DHCP servers"
- DHCP server assigns an address

Transition Mechanisms

- Dual-IP Hosts, Routers, Name servers
- Tunneling IPv6 over IPv4
- Hosts and Routers can be gradually upgraded to IPv6
- It is better (though not required) to upgrade routers before upgrading hosts

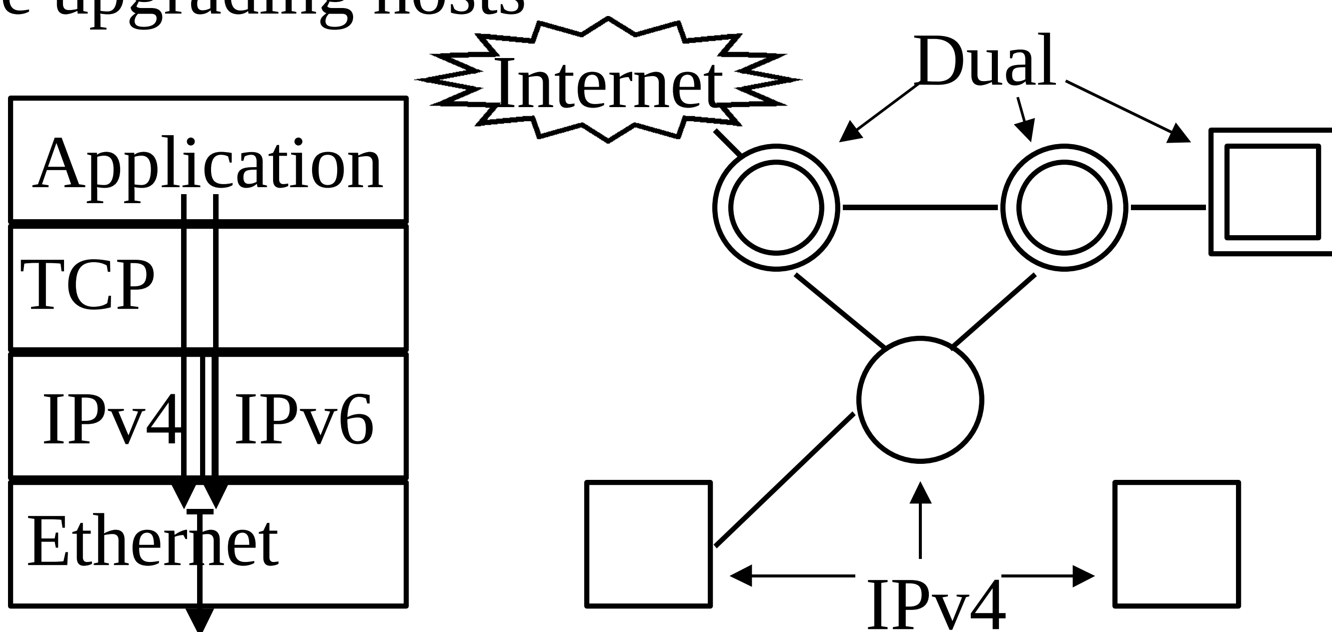

Application Issues

- Most application protocols will have to be upgraded FTP, SMTP, Telnet, Rlogin
- 27 of 51 Full Internet standards, 6 of 20 draft standards, 25 of 130 proposed standards will be revised for IPv6
- No checksum ⇒ checksum at upper layer is mandatory, even in UDP
- non-IETF standards: X-Open, Kerberos, ... will be updated
- Should be able to request and receive new DNS records

Implementations

- 4.4-lite BSD by US Naval Research Laboratory
- UNIX, OPEN-VMS by DEC
- DOS/WINDOWS by FTP Software
- HP-UX SICS (Swedish Institute of Comp. Science)
- Linux
- NetBSD by INRIA Rocquencourt
- Solaris 2 by Sun
- Streams by Mentat

Summary

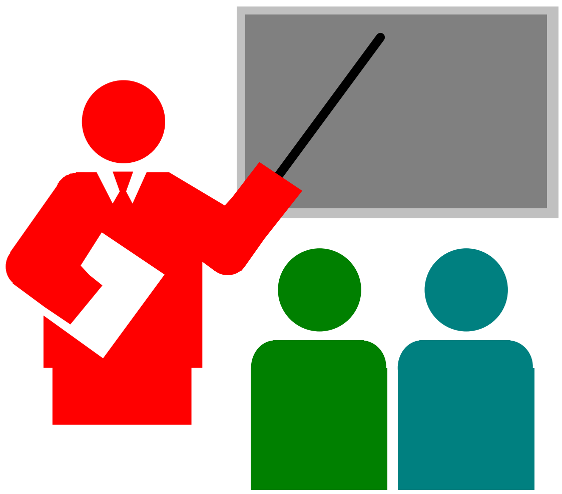

Pv6 uses 128-bit addresses
llows provider-based, site-local, link-local, multica
nycast addresses
ixed header size. Extension headers instead of optic
xtension headers for provider selection, security
llows auto-configuration
ual-IP router and host implementations for transitic