IPv6: Features, Functions, Prospects and Schedule

Scott Bradner

Harvard University

sob@harvard.edu

Session Overview

- Background - Why do we need an IPng?
- History - How did we go about selecting one?
- Overview - What is IPv6?
- Issues - Unresolved conflicts.
- Schedule - When will what happen?
- Effects - What will change?
- Prospects - Will it play in Peoria?

Background

- August 1990

  projected exhaustion of Class B space by March 1994
- 32 bit address space can identify 4 billion hosts

  assignment inefficiency reduces utilization

  use of classfull addresses reduces efficiency
- routing table bloat

  table space increasing faster than memory technology

  thus can not just use multiple Class C addresses

History

- November 1991
  - Routing and Addressing (ROAD) group formed
- March 1992
  - ROAD report
    - do CIDR
    - issue call for IPng proposals
- July 1992
  - IAB issues "IP version 7"
  - IETF issues call for IPng proposals

History, contd.

- July 1993
  - ipdecide BOF & IESG plenary
    - IESG take on responsibility making IPng recommendation
    - (not let the market decide)
- August 1993
  - temporary IETF area formed to consolidate IPng activity
  - Allison Mankin & Scott Bradner Area co-Directors
- December 1993
  - RFC 1550 call for IPng White Papers

History, contd.

- July 1993

  IPng recommendation

- October 1994

  IESG approved recommendation

- April 1995

  base documents ready for Proposed Standard

IPng Proposals

- multiple working groups
- different approaches to solve addressing and routing problems
- different views on problems
- optimize different aspects of problems
- not right or wrong
- learned from all efforts

Available Timeframe

- Address Lifetime Expectations (ALE) working group
  - Frank Solensky, FTP Software <solensky@ftp.com>
  - Tony Li, Cisco Systems <tli@cisco.com>
- made prediction at Seattle,Toronto & San Diego IETF meetings
  - 2005 - 2011
- mixed view of confidence level
  - questions on base data & assumes no paradigm shifts
  - routing tables are still going to be a problem
- projection at Danvers meeting pushes out time

Scope of IPng

- development, testing & deployment will take time
- still we seem to have adequate time in IPv4 address space but not excessive (excluding paradigm shifts)
- can do more than 'just' fix addresses
- use requirements process to determine actual scope of IPng effort

IPng Technical Requirements

- IPng requirements process
  - Frank Kastenholz, FTP Software <kasten@ftp.com>
  - Jon Crowcroft, UCL <J.Crowcroft@cs.ucl.ac.uk>
- RFC1550 request for white papers
- requirements document
  - based on Frank Kastenholz/Craig Partridge draft
  - criteria, discussion & time frame
- RFC 1726

IPng Criteria

- at least 10^9 networks, 10^{12} end-systems
  safer goal 10^{12} nets, 10^{15} end-systems
- conservative routing schemes
- topologically flexible
- high performance
- straightforward transition plan from IPv4
- robust service
- media independent

IPng Criteria, contd.

- datagram service
- autoconfiguration
- secure operation
- globally unique names
- access to standards
- support multicasting
- extensible
- support service classes

IPng Criteria, contd.

- support mobility
- include control protocol (ping etc.)
- support for private networks (tunneling)

Result of Proposal Reviews

- significant flaws seen in all offered proposals
- revised proposal offered by SIPP WG after reviews
- answers most of the perceived problems
  - routing and addressing
  - transition
  - autoconfiguration
  - source routing support
- synthesis of multiple efforts

Address Length

- hotly discussed issue
- four basic views
  - 8 bytes is enough, more is inefficient
  - 16 bytes is about right, 8 is not enough
  - use 20 byte NSAPs, provide global harmonization
  - variable length gives best safety and efficiency
- many detailed arguments
- consensus is that 16 bytes is enough

IPv6 Overview

- expanded from IPv4 addressing capability (16 byte addresses)
- simple header
- support for extension headers and options
- support for authentication and privacy
- support for autoconfiguration
- support for source routes
- simple and flexible transition from IPv4
- flow ID

IPng Addresses

- propose mapping algorithms from and to other environments
- where addresses are globally unique and assigned with regard to network topology
- IETF should work with other organizations for development of such mappings
- common addresses facilitate transition to IPng
- goal to provide a 1:1 mapping between address types (e.g. IPX, NSAP, E164)

IPv6 Header

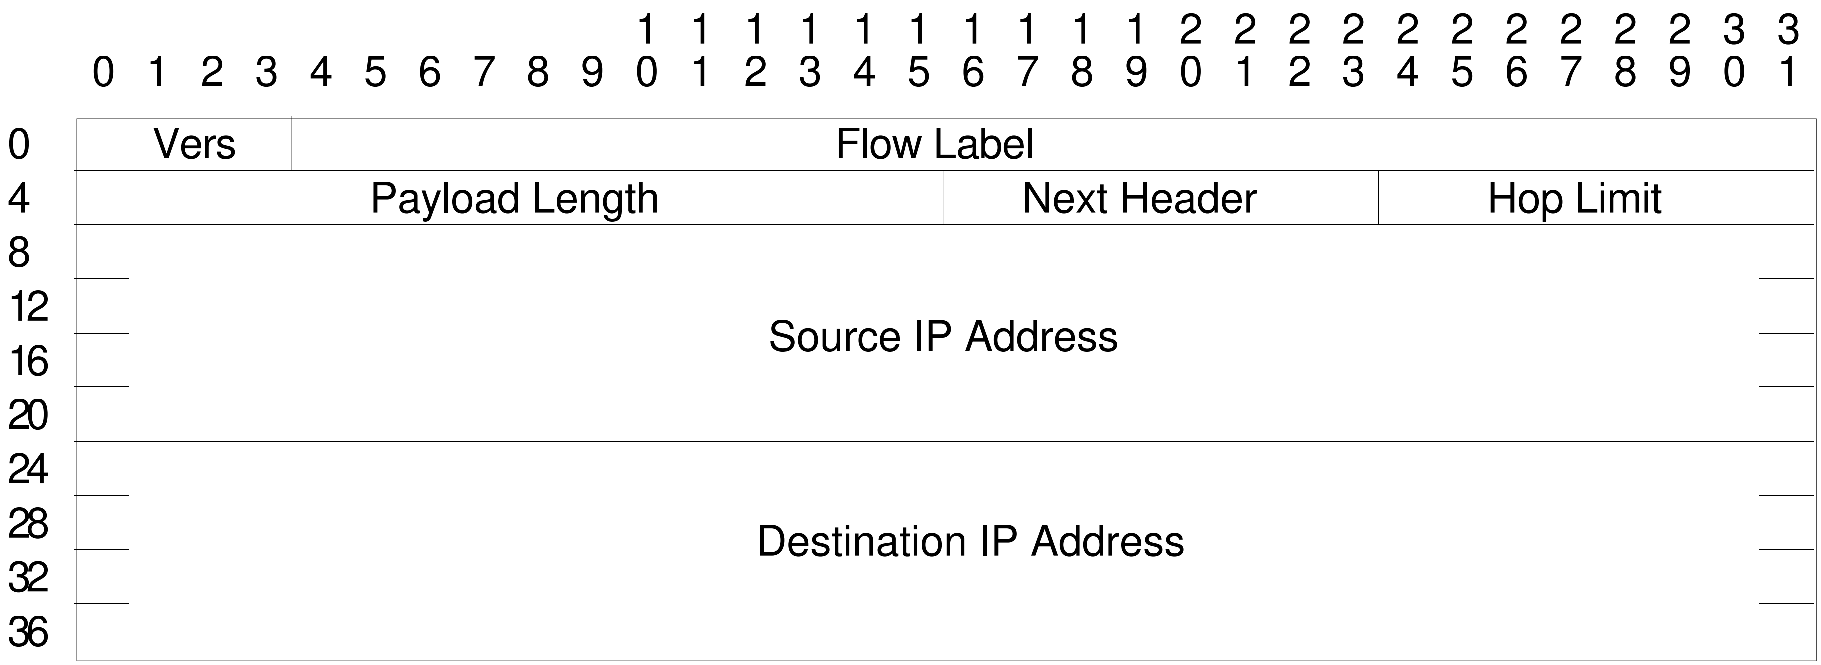
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
0 4 8 12 16 20 24 28 32 36
Vers
Flow Label
Payload Length
Next Header
Hop Limit
Source IP Address
Destination IP Address

IPv4 Header

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

0	Vers	Hlen	Pre	D	T	R	C		Total Length
4	Identifier		N	M	Fragment Offset				
8	TTL	Protocol	Header Checksum						
12	Source IP Address								
16	Destination IP Address								
20	Options	Padding							

IPv6 Header Changes from IPv4

- longer address - 32 bits -> 128 bits
- fragmentation fields moved to separate header
- header checksum field eliminated
- header length field eliminated (fixed length header)
- length field excludes IPv6 header
- "Time to Live" -> "Hop Limit"
- "Protocol" -> "Next Header"
- 64 bit field alignment
- added Flow Label
- removed TOS bits

IPv6 Extension Headers

- less used functions moved to Extension Headers
- only present when needed
- extensible
  - Hop-by-Hop
  - Routing
  - Fragment
  - Authentication
  - Privacy
  - End-to-End

IPv6 Header Option Handling

- highest-order two bits of each option specify the action to be taken if unknown option
- 00 - skip this option
- 01 - discard the packet
- 10 - discard the packet & send ICMP message
- 11 - undefined

eases introduction of new options

- third-highest-order bit in Hop-by-Hop options

include option in the Authentication integrity assurance computation (option data that changes en route must be excluded)

Flow Label

1	3	24
T	TC	flow ID

- T - time sensitivity

  0 = yes

  1 = no

- TC - Traffic Class

  type of flow

- Flow ID - random, unique-to-source value

  combined with source address to identify traffic flow

IPv6 Address Types

Allocation	Prefix (binary)	Fraction
reserved	0000 0000	1/256
reserved	0000 0001	1/256
NSAP Allocation	0000 001	1/128
IPX Allocation	0000 010	1/128
reserved	0000 011	1/128
reserved	0000 1	1/32
reserved	0001	1/16
reserved	001	1/8
provider-based unicast	010	1/8
reserved	011	1/8
reserved for geographic	100	1/8
reserved	101	1/8
reserved	110	1/8
reserved	1110	1/16
reserved	1111 0	1/32
reserved	1111 10	1/64
reserved	1111 110	1/128
local use address	1111 1110	1/256
multicast address	1111 1111	1/256

IPv6 Transition Goals

- allow incremental upgrade from IPv4 hosts to IPv6
- few sequence dependencies
- support what vendors will do
- allow IPv4-only hosts to talk to IPv6-only hosts
- finish before IPv4 addresses run out

IPv6 Transition Techniques

- IPv4 compatible addresses

  IPv4 address embedded in IPv6 address
- IPv6 in IPv4 encapsulation

  tunnel IPv6 across IPv4 topology
- IPv4 <-> IPv6 header translation

  optional

Address Autoconfiguration

- two types of autoconfiguration
  - server-less
  - state-full server
- DHCPng deals with state-full server
- security policy an issue
- trying for *plug & play* in dentist's office

IPv6 Routing

- hierarchical addresses used in IPv6
- 1st version ‘provider based’ hierarchy
- working on geographic or interconnect-point based
- address assignment a concern from the start
- easy renumbering may be important in maintaining efficient use of routing table space

Routing Paradigm

- longest-match routing will be used
- also source routing - ERP header
  - provider selection
- existing routing protocols will be modified for IPv6
  - RIPv2
  - OSPF
  - IS-IS
  - IDRP

Issues - Security

- network security seen as a continuing problem
  actually - lack of
- much bad press - often justified
  front page - *New York Times* , etc.
- network security was a key concern in the IPng selection process

From the IPng Recommendation

"We feel that an improvement in the basic level of security in the Internet is vital to its continued success. Users must be able to assume that their exchanges are safe from tampering, diversion and exposure. Organizations that wish to use the Internet to conduct business must be able to have a high level of confidence in the identity of their correspondents and in the security of their communications. **The goal is to provide strong protection as a matter of course throughout the Internet**."

IPng Security Recommendations

- support for the Authentication Header be required in all compliant IPv6 implementations
- support for a specific authentication algorithm be required
- support for the Privacy Header be required in all compliant IPv6 implementations
- support for a privacy authentication algorithm be required

IPv6 Security Design Goals

- provide authentication and integrity separate from encryption to minimize export control issues
- provide standard encryption mechanism
- ensure that encryption can deal with whole datagrams or portions
- keep mechanisms algorithm-independent
- decouple key management
- retain the ability to use firewalls
- same system for both IPv4 and IPv6

IPv6 Security Recommendation Issues

- performance impact
- key management
- what if kernel lies about having security
- “wrong layer” - should be in application
- export and use restrictions on encryption

Security Export Issue

- strong encryption faces export controls in some countries
- strong encryption faces usage restrictions in some countries
- if IPv6 requires the implementation of strong encryption in order to be compliant to the standard
  - U.S. vendors have to ship non-compliant versions for export
  - all vendors must ship non-compliant versions for some countries
- some vendors do not like to be non-compliant
- big discussion during Danvers IETF meeting

Issues - Addressing Architecture

- routing table efficiency requires that addressing hierarchy follow network hierarchy
- otherwise tend toward flat routing
  - require entry for all non hierarchical routes
- addressing that follows network hierarchy is provider based addressing
  - tends to bind customers to vendors
- geographic addressing requires richly connected network
  - must get vendor cooperation at many places

Addressing Issues, contd.

- in future may have charges for addresses that can not be aggregated
- IPv6 auto configuration may help
  - make renumbering easier
  - support for multi-homed domains
    - multiple addresses per interface
  - can migrate from one provider to another

Schedule

- most documents are ready for standards process
- 2 main ones not ready yet
  - neighbor discovery - mostly editing required
  - address autoconfiguration - open DNS related issues
- some test implementations working
- beta code expected this summer
  - hosts and routers
- release code Q4 ‘95
- test networks (using tunnels) Q3 ‘95
- network provider support early ‘96

Effects

- confusion
  - when should IPv6 be used?
  - conformance issues
  - new technology
- feature support
  - some features must wait native network support
    - flow ID, ERP routing
  - some features can be supported in tunnels
    - authentication, encryption, NSAP & IPX addresses

Effects, contd.

- migration from OSI CLNP
  - larger addresses supported
  - can map U.S. gosip NSAPs
  - support for OSI apps over IPv6
- set top box autoconfiguration
  - will help to get router autoconfiguration

Prospects

- carrots may be hard to see
  - plug & play, security, flow ID, advanced routing
- NAT box option
  - network address translation box
  - use private IPv4 address and map to subset for external
- IPv6 too little, too early?
- a few key decisions will determine IPv6 fate
- high level of assumption that IPv6 is the right thing
- disagreement over size of step taken