IP: Next Generation—A Tutorial

Bruce A. Mah
bmah@{CS.Berkeley.EDU,research.ATT.COM}

University of California at Berkeley
International Computer Science Institute
AT&T Bell Laboratories

29 August 1994

Disclaimer

I am not an IPng researcher.

I don't even play one on TV.

Outline

- Introduction
- IPng features
  - Better header
  - 128-bit address space
  - Better options support
  - Routing
  - QOS support (sort of)
  - IPv4 interoperability (Simple SIPP Transition)
- Current hot topics

Problems with the Internet Protocol

IP address space exhaustion

Routing table explosion

Inefficient headers for high-speed networks

Some features (e.g. source routing) not well supported

New features needed (e.g. security)

No QOS support

IP: Next Generation

Replace IP with a new internetwork layer

Retain the same basic philosophy, but try to solve IP's problems

Avoid changing other protocols (i.e. TCP, UDP) or applications (i.e. `telnet`) where possible

Need a transition plan

Proposals

- SIPP (Simple Internet Protocol Plus)
- TUBA (TCP and UDP with Bigger Addresses)

Terminology

IPv4	Internet Protocol, Version 4
SIP	Simple Internet Protocol
SIPP	Simple Internet Protocol Plus
SIPP-8	Original version, 64-bit addresses
SIPP-16	Revised version, 128-bit addresses
IPng	IP: Next Generation
IPv6	Internet Protocol, Version 6

The IPv6 Header

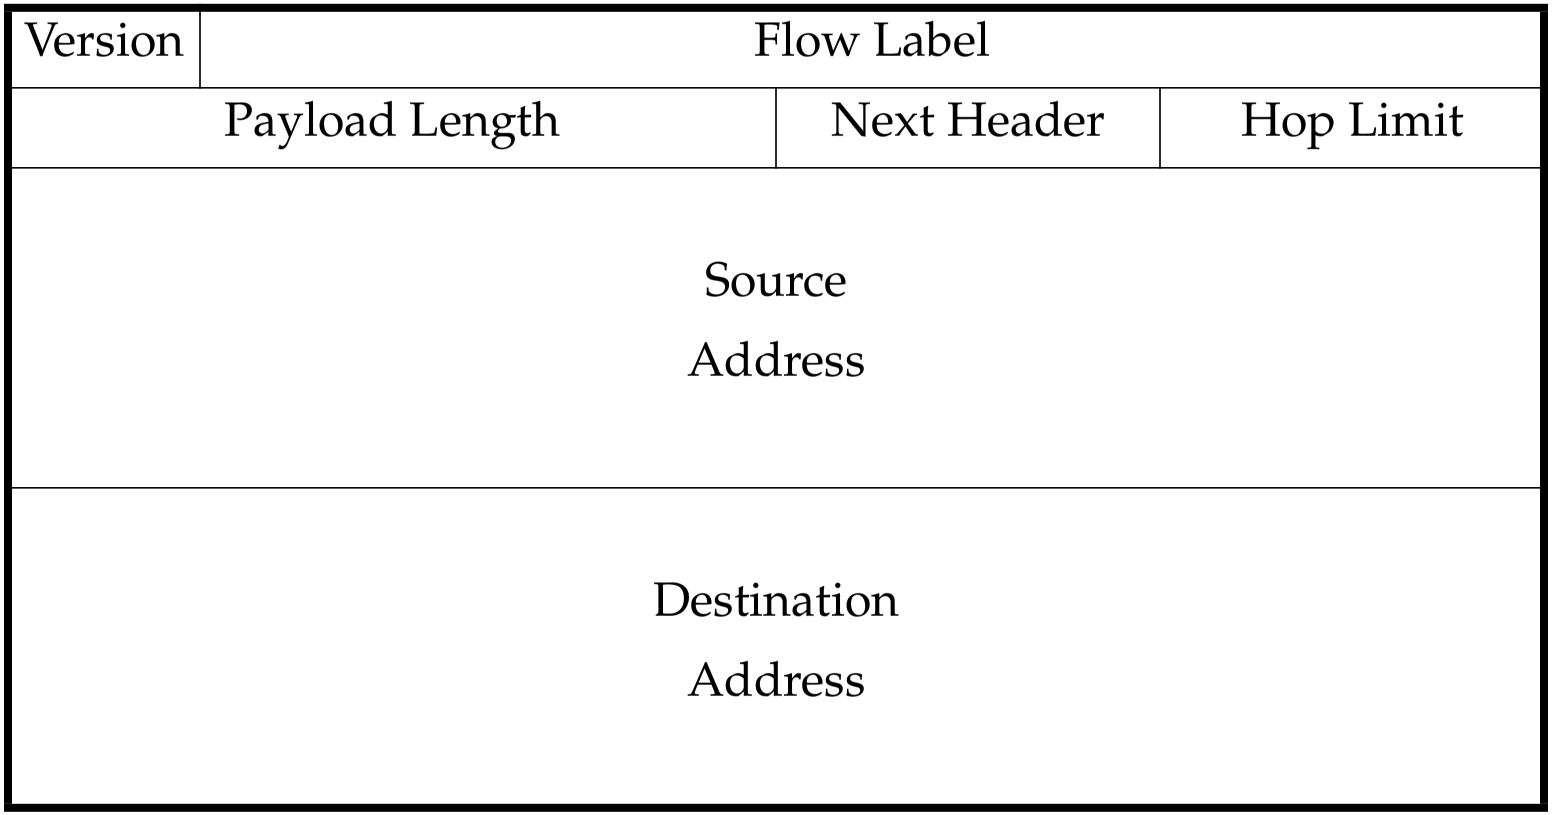

Changes from IPv4 Header

Version	Hdr Ln	Prec	TOS	Total Length	
ID				Flags	Fragment Offset
TTL		Protocol		Header Checksum	
Source Address					
Destination Address					

Fixed size header (no need for header length)

Precedence/TOS handled with QOS and Flow Label

TTL now a Hop Limit

Fragmentation now an option

Header protected by transport layer pseudo-header and checksums

Addressing

Each address identifies an interface, possibly multiple addresses per interface

Types of addresses distinguished by prefix

Unicast Addresses

- IPv4
- Provider-Based
- Geographic-Based (sort of, space reserved)
- NSAP
- IPX
- Local Use

Cluster Addresses

Multicast Addresses

IPv4 Compatable Addresses

IPv4

0000.....0000	IP Address

IPv4 with IPv6 Support

0000.....FFFF	IP Address

Examples

0:0:0:0:0:FFFF:8020:C941

0:0:0:0:0:FFFF:128.32.201.65

::FFFF:8020:C941

::FFFF:128.32.201.65

Provider-Based Unicast Addresses

- Globally unique
- No fixed field boundaries
- Open issues
  - How to easily reconfigure when switching providers?
  - How to handle multi-homed hosts with multiple providers?

Local-Use Addresses

11111110	000.....00	Subnet ID	Node ID

Address scope limited to a single subscriber site

Subnet ID used for routing

Node ID *can be* an IEEE 802 address (for example)

Applications

- Private internetworks (i.e. not attached to the Internet)
- Autoconfiguration and bootstrapping

Cluster Addresses

Cluster Prefix	0000.....0000

"Nearest" boundary router in a cluster of nodes

Intended for use in source routing

Multicast Addresses

Flags: Transient/Permanent

Scope: Control extent of propagation

- Analagous to use of TTL for IPv4 multicast

Group ID: Identifies multicast group

- Similar to IPv4 multicast groups

No broadcast addresses, pre-defined groups used

- All Nodes
- All Hosts
- All Routers

DNS Modifications

New `ASEQ` records contain IPv6 addresses (or sequences of addresses for source routed addreses)

New `sipp-addr.arpa` domain for reverse name translation

`A` records continue to hold IP addresses for IPv4-compatable hosts

IPv6 Options

IPv6 Header (Next Header= Opt1)
Opt1 Header (Next Header= Opt2)
Opt2 Header (Next Header= TCP)
TCP Header
TCP Payload

Options (usually) only examined at destination specified in IPv6 header

IPv6 Options

Hop-by-hop options (TLV format)

Routing

Fragmentation

Authentication

End-to-end (TLV format)

Routing Option

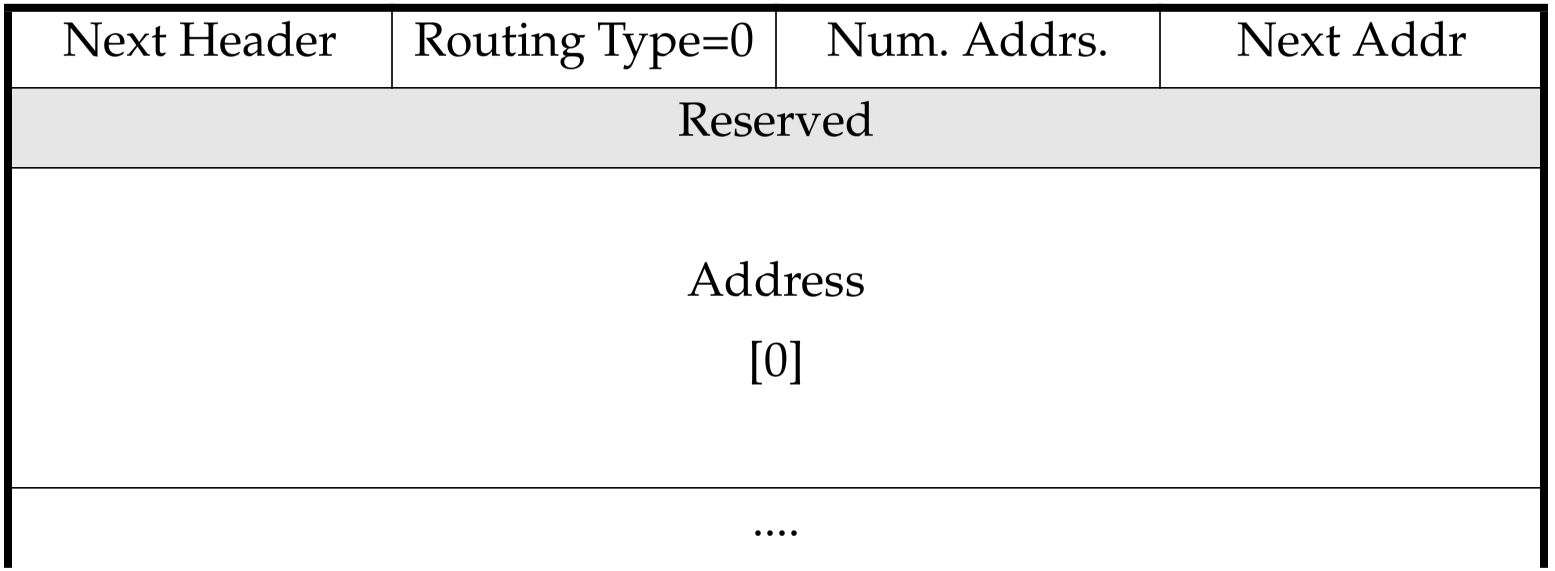

Routing option processing off fast path

Cluster addresses can be used to force routing through a given service provider or network

SRDP has a similar routing option with Routing Type = 1

Fragmentation Option

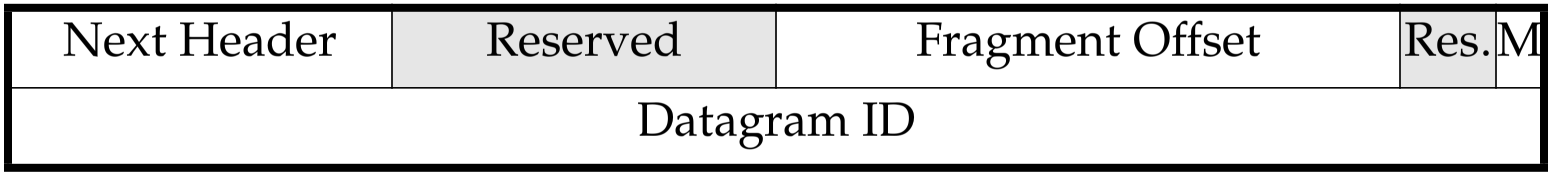

M bit: 1 = More Fragments

Same functionality as IPv4 (datagram ID is bigger)

Not a part of common-case processing but easy to detect at receiver

Path MTU discovery algorithms mandatory

- ICMP messages now return next-hop MTU

Security Options

SIPP Authentication Header

- Authentication and Integrity Assurance
- MD5 (128-bit key) recommended, other algorithms optional
- Want an exportable (outside USA) algorithm

SIPP Security Encapsulation Protocol

- Authentication, Integrity, and Confidentiality
- DES CBC proposed, other algorithms optional

Open Issues

- Key Management?

Routing

Very similar to IPv4 CIDR

Routing lookup based on longest prefix matches

Relies on reasonable assignment of addresses for routing aggregation

Common-case routing code only examines destination address in IPv6 header, regardless of routing headers

Source Routing

Version	Flow Label		
Payload Length		Next = Routing	Hop Limit
Source Address = `my:provider:my.net.subnet.host`			
Destination Address = `her:provider:0:0`			
Next = TCP	Type = 0	Num Addrs = 2	Next Addr = 0
Reserved			
Address[0] = `his:provider:0:0`			
Address[1] = `your:provider:your.net.subnet.host`			

QOS Support

Flow is defined by Flow ID (24 bits) and source IPv6 address

TClass = Traffic Class

- Flow controlled vs. non-flow controlled
- "Priority" within flow controlled or non-flow controlled traffic types

Open issues:

- Resource model?
- Signalling to set up flows?

Simple SIPP Transition Goals

Support IPv6 over IPv4 infrastructure

Interoperability between IPv4 hosts and IPv6 hosts, where possible

Operational requirements

- No "flag days"
- Gradual transition
- Uneven rates of IPv6 deployment
- Some hosts may never transition

SST Addressing

Special IPv6 addresses for IPv4 compatability

- 0:0:0:0:0:0:128.32.201.65 (IPv4 only)
- 0:0:0:0:0:FFFF:128.32.201.65 (IPv4 and IPv6)

A and ASEQ records in DNS servers

Some interfaces may have multiple addresses

- IPv4 compatable, local to an IPv4 area
- IPv4 incompatable, presumed globally unique

SST Mechanisms

Dual protocol stacks in hosts and routers

Tunnelling via IPv4 encapsulation

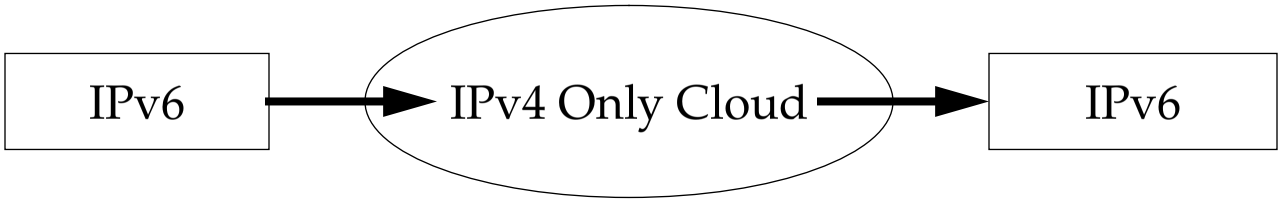

Header Translation

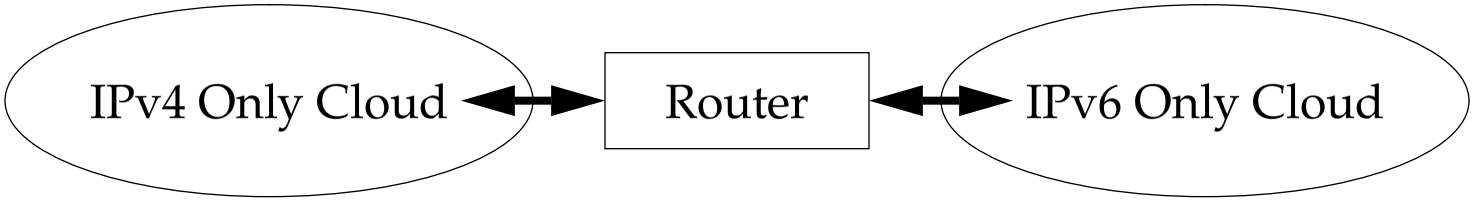

SST Routing

IPv4 to anything

Route using normal IPv4 routing

IPv6 to IPv6

If possible to send directly (same subnet) do so

Else if there is an on-subnet IPv6 router, route via it

Else if there is an off-subnet IPv6 router, tunnel to it

Else tunnel to the destination

IPv6 to IPv4

If a dual-stack machine, send as IPv4

Else, compute IPv6 address and send as to an IPv6 host

Topics Flambé

Should the IPv6 address space be expanded to 20+ bytes to accomodate OSI NSAP addresses?

Provider based addressing?

Autoreconfiguration

- Is it necessary?
- How to make it work?

Authentication and source routing?

For More Information

`ipng` mailing list

- Send mail to `Majordomo@sunroof.eng.sun.com`
- Body text `subscribe ipng`

`ftp://ds.internic.net/internet-drafts`

- `draft-ietf-sipp-spec-01.txt`
- `draft-ietf-sipp-routing-addr-02.txt`
- `draft-ietf-sipp-sst-overview-00.txt`

`ftp://parcftp.xerox.com/pub/sipp`

- SIPP archive

`http://town.hall.org`

- Information and pointers to SIP/SIPP implementations