IPv6

Internet Protocols

CSC / ECE 573

Fall, 2005

N. C. State University

Announcements

I. *HW6 is online*

Today's Lecture

I. Motivation
II. IPv6 Addressing
III. IPv6 Base Header
IV. IPv6 Extension Headers (Options)
V. Transition from IPv4 to IPv6

MOTIVATION

Goals

- Fix IPv4's addressing problems → need larger address space
  - support tens or hundreds of billion hosts
  - every light bulb, toaster, fire alarm, ..., with IP address
- Simplify IP protocol
  - better header format
  - eliminate seldomly-used or unused functions
- Improved options, greater extensibility

Goals (cont'd)

- Support for resource allocation and QoS
- Provide built-in security (encryption and authentication)
- More levels of address hierarchy → better address aggregation
- Improved autoconfiguration
- Aid multicasting

History

- IETF call for white papers on IPng (RFC 1550, 1993)
  - 21 proposals, 7 serious ones
- Full spec: RFC2460 (1998)

IPv6 ADDRESSING

Address Space (RFC 3513)

- IPv6 addresses = 128 bits
- **340,282,366,920,938,463,463,374,607,431,768,211,456** addresses
  - **665,570,793,348,866,943,898,599** hosts per square meter of the earth's surface!
- Hierarchical assignment
  - somewhat inefficient use of bits, but very helpful for administration and routing
- 15% of address space currently allocated, 85% reserved for future use

Hexadecimal Colon Notation

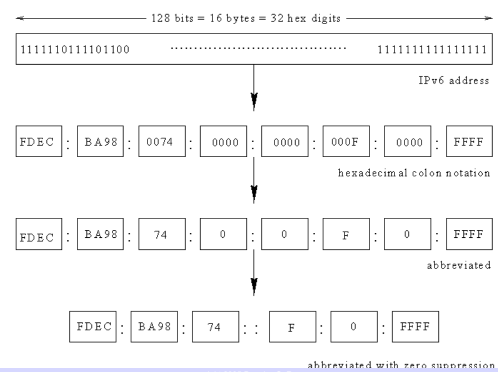

Hexadecimal Colon Notation (cont'd)

- Can also use dotted-decimal style
  - e.g., 253.236.186.152.0.116.0.0.0.0.0.15.0.0.255.255, or 253.236.186.152.0.116::15.0.0.255.255
  - e.g., 0:0:0:0:0:0:128.10.2.1, or ::128.10.2.1
- Can use "/" notation to indicate the length of the address
  - e.g., 12AB::CD30:0:0:0:0/60 = 12AB00000000CD3

Categories of Addresses

- *Unicast*: defines a single entity (host, etc.)
- *Multicast*: defines a group of entities
  - may or may not share the same address prefix
  - may or may not be connected to same physical network
  - packet must be delivered to each member of the group

Categories of Addresses (cont'd)

- *Anycast*: defines group of entities having same address prefix
  - packet should be delivered to just one member of the group ("nearest," "most easily accessible," …)
  - e.g., a group of servers offering the same service
  - allocated from same space as unicast addresses

Address Structure (RFC 3513)

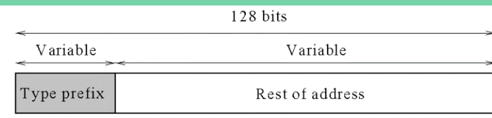

Prefix	Address Type
0…0 (128 bits)	Unspecified
0…01 (128 bits)	Loopback address
1111 1110 10	Link local address
1111 1110 11	Site local address
1111 1111	Multicast address
(everything else)	Global Unicast

Unspecified and Loopback Addresses

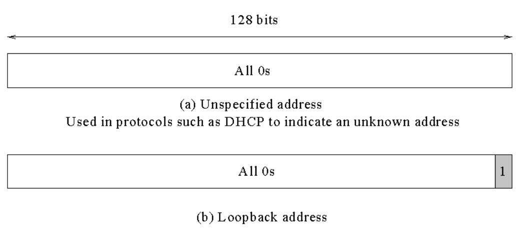

Multicast Addresses

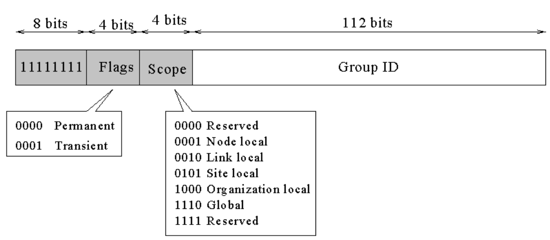

Local Addresses

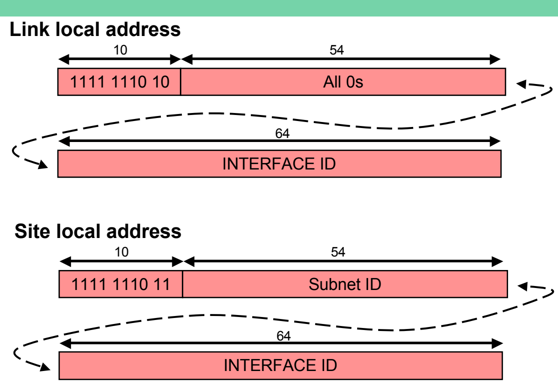

Local Addresses (cont'd)

- These addresses can be reused at each organization, i.e., represent a private address space
- Packets with such addresses can only be routed locally
  - Link-local cannot be propagated outside the same physical network
    - used for autoconfiguration, …
  - Site-local cannot be propagated outside organizational boundaries

IPv4 Addresses (inside IPv6 Addresses)

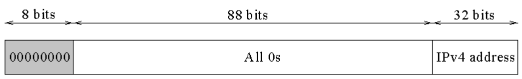

(a) Compatible address

Used by a node that can speak both IPv6 and IPv4

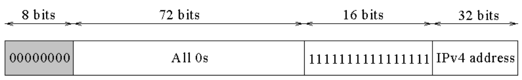

(b) Mapped address

Indicates a node speaking only IPv4, but packet must pass through IPv6 net

Mapping Ethernet MAC Addresses to Interface IDs

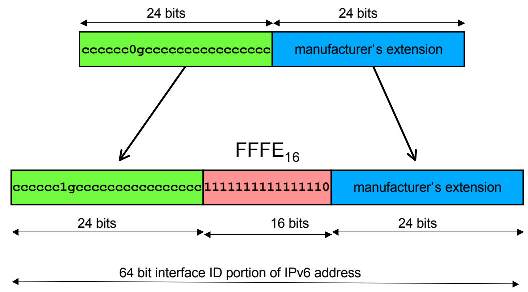

IPv6 BASE HEADER

IPv6 Datagram Format

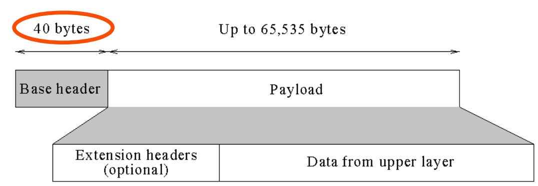

IPv6 Base Header

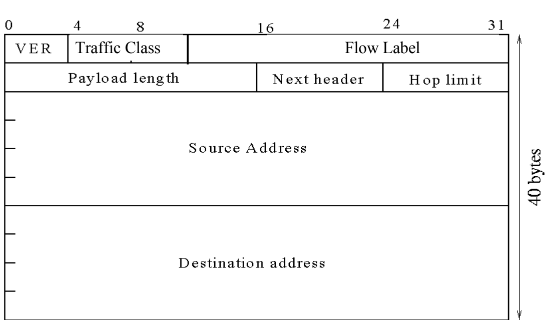

Header Fields

- **`Payload Length`**
  - 40 header bytes no longer counted as part of length
  - 16 bits: payload length < 64K (includes extensions)
- **`Next Header`**: type of the next header
  - optional headers may follow
  - last header points to higher-layer protocol (TCP, UDP)

Header Fields (cont'd)

- **`Hop Limit`**
  - name reflects the way the **`TTL`** field in IPv4 is used
- **`Version`**: always 6
  - useful during transition from IPv4
- Checksum gone!
  - rely on lower layers for header protection, and higher layers (transport) for payload protection
- **`Traffic Class`**
  - Same interpretation as DiffServ DSCP field

Flow Label

- Allows source and destination(s) to set up a "pseudoconnection"
  - still experimental
  - an attempt to have it both ways: the flexibility of datagram network, and the guarantees of a virtual circuit network
- Flow: uniquely identified by source, destination, flow label
  - multiple flows (audio, video, graphic windows, etc.)
  - a flow may comprise a single or multiple TCP connections

Flow Label (cont'd)

- Router's point of view: packets with non-zero flow label require special treatment
  - router tables specify treatment
  - requirements and reservations must be negotiated ahead of time

IPv4 ↔ IPv6 Header Comparison

IPv4	IPv6
20 bytes fixed header	40 bytes
12 fields	8 fields
4-byte addresses	16-byte addresses
final destination	*intermediate* destination
precedence, TOS field	flow id, priority
header length, total length	payload length
$\leq$ 40 bytes of options	"unlimited"
options	header extensions
fragmentation fields	fragmentation header
header checksum	–
higher layer protocol	next header type
TTL time based	TTL hops only

IPv6 EXTENSION HEADERS (OPTIONS)

Extension Headers

- Fixed Base Header followed by optional Extension Headers
  - intermediate routers seldom need to process all the extension headers
- IPv6 Extension Headers similar to IPv4 Options

Extension Headers (cont'd)

- Compromise between…
  - generality: must include mechanisms for source routing, etc.
  - efficiency: most datagrams do not use all mechanisms
- **`Next Header`** field helps in parsing the datagram
- 6 extension headers defined

Extension Header Types

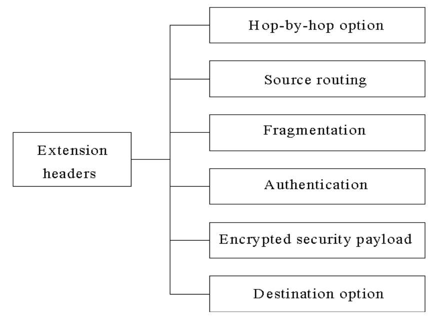

Extension Header Format

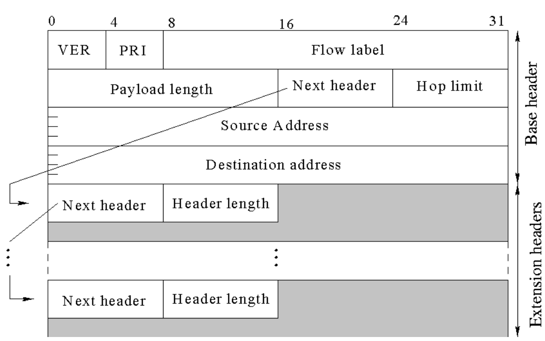

Fragmentation and Reassembly

- Fragmentation header
  - Fragment Offset
  - Datagram Identification
  - MF flag
- Destination performs reassembly (as in IPv4)

Fragmentation and Reassembly (cont'd)

- Only source may do fragmentation
  - end-to-end → no fragmentation at intermediate routers
  - source must perform path MTU discovery or use guaranteed minimum MTU of 1280 bytes
- If a route change requires smaller fragments…
  - new type of ICMPv6 error message generated
  - source does new MTU discovery

IPv6 Fragmentation Example

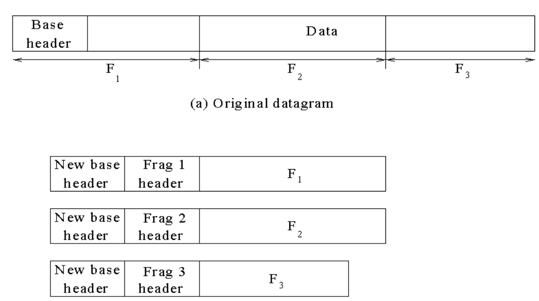

(a) Original datagram

(b) Resulting fragments

Source Routing

- Routing Extension Header allows up to 24 intermediate addresses
  - loose routing initially defined
- Ultimate destination address not in Base Header; instead, it's the last address in the Routing Header
  - Base Header destination address is address of first router in path
- Destination node required to reverse routes in a packet containing a routing header when replying to sender

Source Routing Header

Next Header	Header Length	Routing Type	Addresses Left
Reserved			
First Address			
Second Address			
Last Address			

Source Routing Example

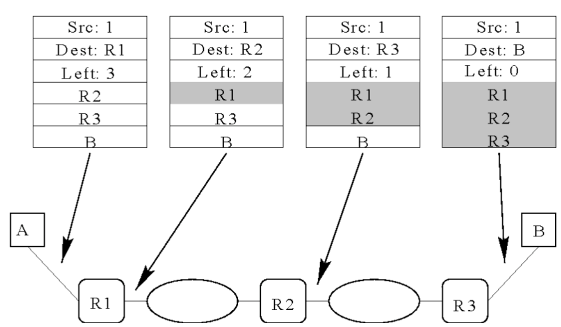

IPv6 Security

- Basically IPSec as an extension header

Hop by Hop Extension Headers

- Source passes info to routers (like "router alert")
- If router does not support the option: first 2 bits of Type field indicate whether to…
  1. skip option
  2. discard packet but do not send ICMP error message
  3. discard and send ICMP error message to source
- Third bit specifies whether field is mutable (i.e., replace with zeros for authentication purposes)

Hop-by-Hop Extension (cont'd)

- *Jumbograms*: datagrams 64KB – 4GB in size
  - extension header specifies length, Payload Length in base header not used

Other Protocols Affected by IPv6

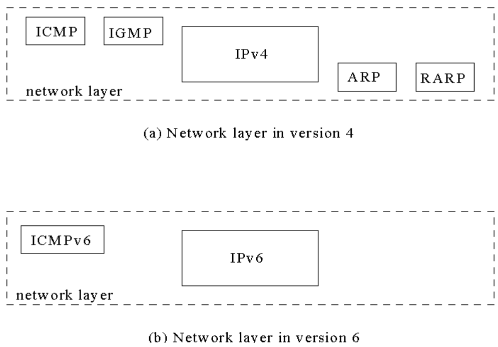

Other Protocols Affected by IPv6 (cont'd)

- ARP/RARP eliminated
  - MAC address mapped directly into 64-bit Interface ID field
  - e.g., Ethernet 48-bit address maps as: 24 bits, 16 bit padding (0xFFFE), 24 bits
- IGMP eliminated
  - Group membership Query, Report, and Termination messages are added to ICMPv6
- DNS support for IPv6 addresses
  - new AAAA resource record

Serverless Auto Configuration

- Host generates link-local address, sends Router Solicitation
- Router responds with Router Advertisement
  - what default router to use
  - whether to use DHCP
  - prefixes to use for site-local and global addresses
  - may include a "lifetime" for prefixes; allows renumbering of networks

TRANSITION FROM IPv4 TO IPv6

IPv4 → IPv6 Transition

- How much has IPv6 been deployed?
- Some reasons why IPv4 has lasted longer than expected
  - CIDR
  - NAT (network address translator)
  - DHCP: improved ability to configure IPv4 addresses

Transition Strategies (RFC 2893)

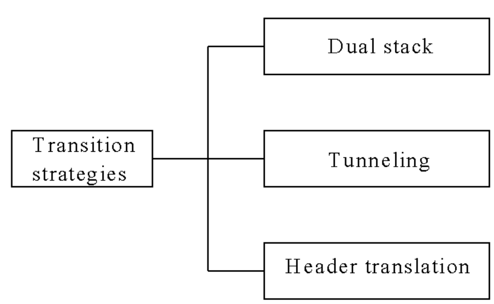

Dual Stack

- Host has complete support for both IPv4 and IPv6

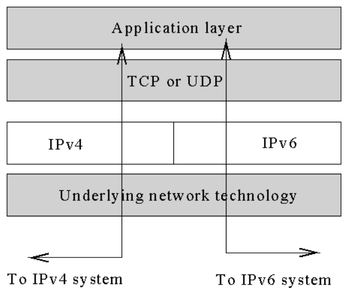

Tunneling

- Allows packets between IPv6 hosts to pass through an IPv4 region
  - encapsulate IPv6 packets within IPv4 headers
- *Automatic tunneling*: IPv4 tunnel endpoint determined from IPv4 address embedded in IPv4-compatible destination address of IPv6 packet
  - only end-hosts involved

Tunneling (cont'd)

- *Configured tunneling*: IPv4 tunnel endpoint determined by configuration info at the encapsulating node
  - routers translate headers

Configured Tunneling (IPv6 over IPv4)

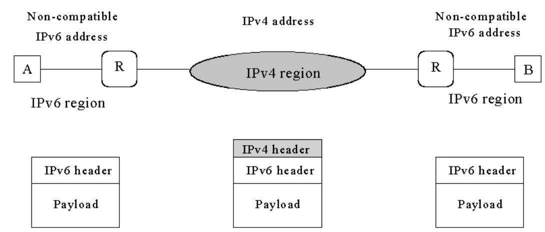

Header Translation

- Sender wants to use IPv6, but destination only understands IPv4
- Translation needed: special translator nodes
- IPv4 nodes that do not support IPv6: ::FFFF:128.10.2.1

Header Translation (cont'd)

- *Problem*: TCP/UDP layer at destination verifies address with checksum of pseudo header
  - *solution*: 1's complement checksum of IPv4 address and IPv6 encoding identical
  - i.e., changing the address has no effect

Summary

- IPv6 provides a number of new capabilities and improvements
  - most visible/important: larger addresses
  - Other: simplify IPv4, more extensible, more capabilities
- IPv6: needed, but when?
- Transition taking a lot longer than expected
  - but available in most desktop OSes now
  - required by DoD
  - basis for mobile IP (next gen cell phones)

Next Lecture

- Mobile IPv6