Internet Engineering

241-461

Robert Elz

kre@munnari.OZ.AU
kre@fivedots.coe.psu.ac.th

http://fivedots.coe.psu.ac.th/~kre

IPv6 - Why?

- IPv4 running out of addresses
  - 2^32 ==> 4 Billion addresses max
  - But many "wasted" because of allocation policy
  - 2 million connected organisations max
- IPv4 routing tables becoming unmanageable
  - ~50K routes
    - now ~200K routes
  - needs recalculating frequently

IPv6 - Answers

- IPv6 addresses 128 bits
  - 2^128 ==> too big for words to describe
  - 340282366920938463463374607431768211456
  - (39 digits)
    - Very Very many will be wasted

IPv6 Header

Header Comparison

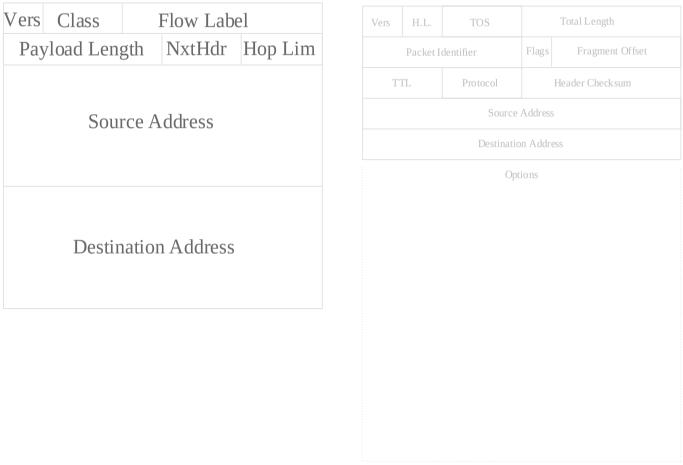

Equivalent fields

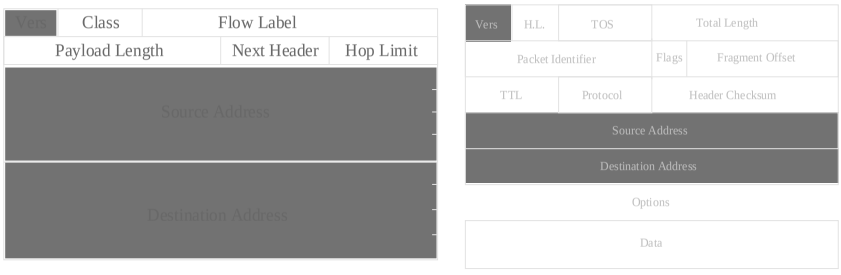

- Version
  - 4 for IPv4
  - 6 for IPv6
    - Hence the names...
- Version identification is one way to allow upgrades

Modified Fields

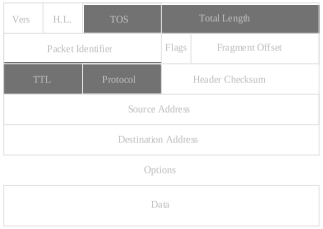

The Modified Fields

- Time to Live
  - Hop Limit
- Total Length
  - Payload length
- Protocol
  - Next Header
- TOS
  - Traffic Class (DSCP) & Flow Label
    - Currently not well defined

Missing Fields

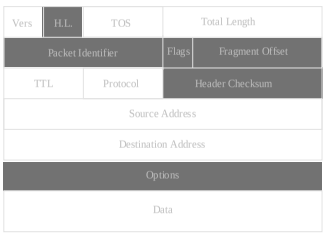

The Missing Fields

- Header Length
  - IPv6 header is fixed length ==> not needed
- Options
  - Exist in a different form
- Fragmentation fields (ID / Offset / Flags)
  - Exist in a different form
- Header Checksum
  - The controversial change - gone completely

An IPv4 Packet

```
0:10:a4:f:41:cf 0:1:3:40:8a:e5 0800 98:

    172.30.0.77 > 172.30.0.161:
                icmp: echo request

     4500 0054 7e24 0000 ff01 e459 ac1e 004d
     ac1e 00a1 0800 a334 0d82 0000 b613 3a3b
     5ff7 0c00 0809 0a0b 0c0d 0e0f 1011 1213
     1415 1617 1819 1a1b 1c1d 1e1f 2021 2223
     2425 2627 2829 2a2b 2c2d 2e2f 3031 3233
     3435 3637
```

An IPv6 Packet

```
0:10:a4:f:41:cf 0:1:3:40:8a:e5 86dd 70:

    3ffe:8001:2:181:210:a4ff:fe0f:41cf >
    3ffe:8001:2:181:201:3ff:fe40:8ae5:
                 icmp6: echo request

     6000 0000 0010 3a40 3ffe 8001 0002 0181
     0210 a4ff fe0f 41cf 3ffe 8001 0002 0181
     0201 03ff fe40 8ae5 8000 e8a3 0d9c 0000
     7616 3a3b d908 0700
```

IPv6 Header Chains

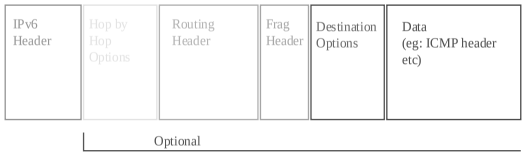

- Each Header contains a "Next Header" field
  - Until the transport header
  - Specifies the type of the header that follows
  - The first header is always an IPv6 header
- Each Header is of fixed length,
  - or provides a mechanism to allow its length to be determined
  - Length field in most headers
- All Headers even multiple of 8 bytes
  - Until the transport header
  - Minimum of 8 bytes
  - Length field usually counts multiples of 8
    - beyond first 8

Header Processing

- Headers processed "left to right" through the packet
  - Unrecognised header is an error,
    - ICMP error report
    - Drop packet
- Routers look only at IPv6 header, and Hop-by-Hop Options
  - Easy to tell if HBH is present
    - Next Header field in IPv6 Header (0)

Header Processing

- Headers processed "left to right" through the packet
  - Unrecognised header is an error,

IPv6 Addresses

- 128 bits
- Like IPv4, divided into
  - network identifier
  - host identification on that network

  Eg:

  172.30.2.60 /24

  netmask 255.255.255.0

  172.30.2.0 is the network identifier

  0.0.0.60 is the host identifier on that network
- Netmask can divide anywhere,
  - not just at one of the dots

IPv6 Prefix

- The "network part" of an IPv6 is the prefix
- Prefix identified by a length, rather than a mask
  - written /nn (as in /32 or /57)
- Prefix is never longer than 64
  - On a standard network
- That is:
  - 64 bits are always available
  - to identify a host on the network

Generating IPv6 Addresses

```
00:10:a4      :0f:41:cf
02 10:a4 ff:fe 0f:41 cf
```

- FFFE is the defined way
  - to transform a 48 bit MAC address
  - into a 64 bit MAC address (EUI-64)
- FFFE is inserted
  - between the 3rd and 4th octets of 48 bit address
- 02 in the first address octet is the "locally defined" bit
  - In MAC address bit set indicates a locally defined address
  - In IPv6 address, bit is inverted

When and How to Autoconfigure

- Not all sites want any random node to be able
- to acquire an address and use it without authorisation
  - Auto-config does not enable that
    - it does make it easy
- DHCP used to assign addresses in IPv4,
    - can implement policy
  - Auto-config has no policy
- Some users concerned about privacy
  - Every IPv6 address they use contains their MAC address
- Hence need alternative methods

IPv6 Address Types

- Local Loopback
- Link Local Address
- Site Local Address
- Global Address
- IPv4 compatible Address
- Multicast Address

Loopback Address

- ::1
  - :: (0::0) is the "unspecified address"
- IPv6 Address Notation
  - nnnn:nnnn::nnnn
    - The :: indicates as many 0's as are needed
    - Only one :: in any address
    - 16 bits in each other numeric block (between :'s)
- Loopback address means "this node"
  - 127.0.0.1 in IPv4

Link Local Addresses

- Defined Prefix
  - FE80::/10
- Low 64 bits contain host identifier
  - fe80::210:a4ff:fe0f:41cf
  - fe80::1

Link Local Addresses

- These addresses can be used to communicate with other nodes
- on the same link
  - Routers do not route packets containing link local addresses
  - Useful for all on-link communications
    - eg: router advertisement
  - Used where off link addressing is incorrect
    - eg: redirect
    - eg: Neighbour Discovery

Site Local Addresses

- Defined Prefix
  - FEC0::/10
- Used for communications with a site
  - "site" can mean whatever is appropriate
  - Often a company/university/...
- Packets using these addresses
  - not forwarded beyond the boundaries of the site
  - FECr:rrrr:rrrr:SSSS:EUI-64
- Now deprecated (obsolete)
    - (Before IPv6 was really ever used)

Local Addresses

- Defined Prefix
  - FC00::/7
- Unroutable address
  - Not useful to reach random destination
- Possibly unique
  - some would say probably
  - some would hope certainly
  - nothing enforces uniqueness
    - no way to test either
- `FCaa:aaaa:aaaa:SSSS:EUI-64`
  - Assigned by number authority
- `FDrr:rrrr:rrrr:SSSS:EUI-64`
  - Generated by random number generator

Global Addresses

- 48 bit site prefix
- 16 bit subnet number
- 64 bit EUI-64

- FP: Format Prefix (3 bits: 001)
  - 001 010 011 100 101 110
    - All except 000 and 111
- Prefix: Site Identification (45 bits ... 48 bits with FP included)
  - Internal aggregation boundaries exist
- Subnet: Network within site (16 bits)

IPv4 Compatible Addresses

- ::a.b.c.d
  - 96 bits of all 0
  - followed by an IPv4 address (32 bits)
- ::FFFF:a.b.c.d
  - 80 bits of all 0
  - followed by FFFF (16 bits)
  - then IPv4 address (32 bits)
- First form
  - used to number any node that has IPv6
  - and also has an IPv4 address
    - IPv4 compatibility mode
- Second form
  - used to number all IPv4 nodes that have no IPv6

IPv4 Compatible Addresses

- ::a.b.c.d
- ::FFFF:a.b.c.d
- Packet translation is possible
  - These addresses generate the same
  - transport protocol checksums
  - as the IPv4 addresses they represent
- Used for API purposes only
  - Not routed on IPv6 network

IPv4 Mapped Addresses

- 2002:a.b.c.d:SLA:EUI-64
  - Cannot be written this way
  - Must use
    - 2001:AABB:CCDD:SLA:EUI-64
  - a.b.c.d must be a global IPv4 address
- Any site with an IPv4 address
  - can use this as an IPv6 prefix
    - IPv4 internet is the IPv6 backbone

IPv6 Fragment Header

- Fragment Offset & MF
  - Identical to IPv4
    - Except now flags are where they belong
- Identification
  - Identical purpose, but 32 bits
    - Less chance of accidental collision
- DF
  - Not needed

IPv6

- Requires use of Fragment header
  - Only source nodes add headers
  - Only source node can fragment packets
    - No router complexity
- No overheads when no fragmentation
- PMTUD is required
  - Or packets must remain smaller
  - than guaranteed PMTU
    - 1280 for IPv6