

Internet Engineering

241-461

Robert Elz

kre@munnari.OZ.AU

kre@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
 - (that "now" was > 10 years ago!)
 - Needs recalculating frequently
 - ▷ Every time any link changes status
 - ▷ Calculation scales $O(n \log(n))$

© 2004

▪ Double number of routes
Also the double computation time

IPv6 - Answers

- ◊ IPv6 addresses 128 bits
 - 2^{128} ==> too big for words to describe
 - **34028236692093846374607431768211456**
 - (39 digits)
 - ▷ Very Very many will be wasted
 - ▷ But there will still be plenty to use
- ◊ Routing
 - Later ...
 - No Real Change
 - ▷ Situation improves because we can better assign addresses

IPv6 Header

| | | |
|---------------------|-------|-------------|
| Vers | Class | Flow Label |
| Payload Length | | Next Header |
| Source Address | | Hop Limit |
| Destination Address | | |

Header Comparison

| | | |
|---------------------|-------|------------|
| Vers | Class | Flow Label |
| Payload Length | | Next Hdr |
| Source Address | | Hop Lim |
| Destination Address | | |

| | | | |
|-------------------|----------|---------------------|-----------------|
| Vers | HL | TOS | Total Length |
| Packet Identifier | | Flags | Fragment Offset |
| TTL | Protocol | | Header Checksum |
| Source Address | | Destination Address | |
| Options | | | |


Equivalent fields

| | | |
|---------------------|-------|-------------|
| Vers | Class | Flow Label |
| Payload Length | | Next Header |
| Source Address | | Hop Limit |
| Destination Address | | |

| | | | |
|-------------------|----------|---------------------|-----------------|
| Vers | HL | TOS | Total Length |
| Packet Identifier | | Flags | Fragment Offset |
| TTL | Protocol | | Header Checksum |
| Source Address | | Destination Address | |
| Options | | | |
| Data | | | |

- ◊ **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 (Type Of Service)
 - Traffic Class (TC) (or DSCP)
 - DSCP now used instead of TOS in IPv4
 - Flow Label
 - > Currently not well defined
 - > But improving as time passes
 - More ideas
 - More experience

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)

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