IPv6

Introduction

- Christian Huitema (www.huitema.net)
  - IPv6, IP telephony, SIP
- Reasons
  - RT-explosion
  - Class B is empty !
  - Address space exhaustion until 2009
- IPv4 address space: 2^32 = 4,294,967,296
  - Minus Class D and E (536,870,912)
  - Minus Net 0 and 127 (33,554,432)
  - Minus RFC 1918 (17,891,328)
  - Result: 3,706,650,624 usable addresses

Address assignment efficiency vs. H-Ratio (RFC 1715)

- Assignment efficiency = Effective number of hosts / theoret. Maximum
- H = log(No of addresses in use) / No of bits available
  - H <= 0.30103 = log(2^32) / 32
- Other numbering plans use H={0.14 .. 0.26} which is *considered* a typical utilization range
- "IPng should be capable of numbering 10^15 systems"

Net-ID Len of IPv6	H=0.14 (bad util.)	H=0.26 (good util.)
32 bits	3*10^4	2*10^8
64 bits	9*10^8	4*10^16
80 bits	1.6*10^11	2.6*10^27
128 bits	8*10^17	2*10^33

NAT Drawbacks

- No end-to-end IP but client/server model
- Network needs to maintain state of connection
- Fast rerouting difficult if NAT router fails
- Makes development of new application difficult ("NAT friendly?")
- Security (end-to-end)
- Manageability

IPv6 Main Features (1)

- Large address space
  - ~10^30 pro person
- Global reachability
  - With end-to-end security
- Hierarchical addressing
  - Plus scoping: site local vs. link local
- Aggregation
  - Single large prefix per organization
- Multiple addresses per host or network
  - Simple multi-homing through large address space => high reliability
- Autoconfiguration
  - 64+64, gateway tells prefix
- Simple renumbering
  - Gateway tells prefix

IPv6 Main Features (2)

- No broadcasts
  - To reduce DoS and bc-storms
- Simpler header
  - 64-bit aligned, fewer fields, no checksum (UDP checksum now required)
- 20-bit flow label
  - Still not specified in detail (?)
- Extension headers
- Mobile-IP included
  - With direct routing
- IPsec included
  - Even mandatory!

IPv6 Main Features (3)

- "Transition richness"
  - Lots of migration possibilities available
- Serverless configuration
  - No DHCP required
- Multicast
  - IGMP => MLD, Scope Identifier

Header Details (RFC 2460)

- 40 bytes
- 8-bit traffic class (for DiffServ)
- 20 bit flow label
- Payload length (not total length as in IPv4—next headers are part of payload)
- Next header (instead of protocol field)
- Hop Limit (same as TTL)

Extension Headers (in correct order)

- Hop-by-hop Options
  - Used for Router Alert and Jumbogram
  - Only header which is processed by all hops along path
- Destination Options
  - For Mobile IP
  - Only if Routing header is used (see below)
- Routing
  - To specify a Loose Source Route
- Fragment
- AH
- ESP
- Destination options

The Routing Header

- Contains “segments” (= next hops) and “counter” (= segments left)
- Next-hop list is decremented
- DA = next hop (“segment”)

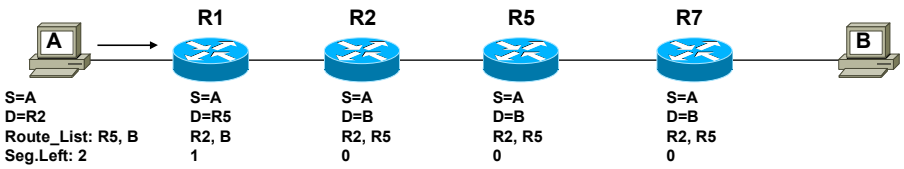

The Fragment Header

- Used when host MUST send packets with length > MTU of path
- Similar usage as with IPv4
  - Fragment offset
  - Identification number

Packet Lengths / MTU

- IPv4
  - Minimum = 68 Bytes
  - Recommended min. = 576 Bytes
- IPv6
  - Minimum = 1280 Bytes
  - Recommended min. = 1500 Bytes
  - Path MTU discovery strongly recommended
  - Maximum = 64 KB or dedicated Jumbograms

Addresses Overview

- All unicast have 64-bit interface ID (EUI-64): X:X:X:X:I:I:I:I (each X or I stands for 16 bit)
  - Except those starting with 000 (binary) which are special purpose addresses
- Global unicast: Glob_rt_prefix|Subnet-ID|If-ID
  - Consumes biggest part of address space
  - IANA allocates 2000::/3 block (start with binary 001)
- Site local: prefix=FEC0::/10
  - 1111 1110 11<38 bits><16 bits subnet> <64 bits If-ID>
  - NOT automatically configured
  - Like RFC 1918 addresses
  - When site connects to Internet, only renumber the prefix part (/48)
- Link local: prefix=FE80::/10
  - 1111 1110 10<54 zeroes><64-bit If-ID>
  - Automatically configured
  - Used for neighbor and router discovery
  - Used by many routing protocols

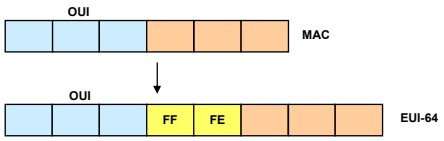

Special Addresses

- Unspecified address
  - 0:0:0:0:0:0:0:0 or simply "::"
  - Same as IPv4 0.0.0.0
  - Usually a placeholder (DHCP request, etc.)
- Loopback address
  - 0:0:0:0:0:0:0:1 or "::1"
- IPv4 compatible address
  - <96 bits "0"><32 bit IPv4 address>
  - E.g. "::192.168.30.1" or "::C0A8:1E01"
  - Assigned to dual-stack nodes and used to allow automatic tunneling
    - Scalability problems!
    - Deprecated
- IPv4-mapped address
  - <80 bits "0">FFFF<32 bit IPv4 address>
  - On dual-stack node, THIS node sends IPv4 packets
  - E.g. "::FFFF:192.168.30.1" or "::FFFF:C0A8:1E01"

Multicast Addresses

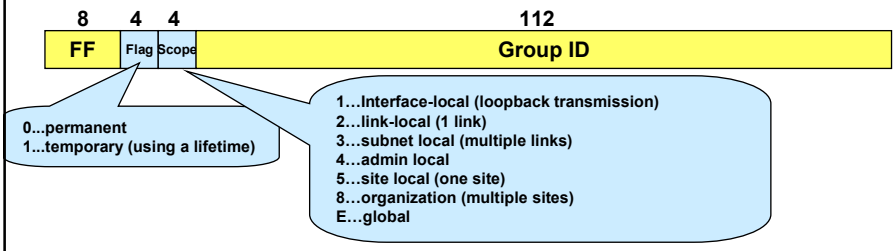

- Use 1/256 of total range
- Assigned addresses:
  - FF02::1 => all nodes on this link
  - FF02::2 => all routers on this link
  - FF02::9 => all RIP-routers on this link
  - FF02::1:FFxx:xxxx => “Solicited node”, used for duplicate address detection and neighbor discovery (Note: only 24 bits of MAC address are used)

Anycast (as proposed 1993)

- Assigned to multiple interfaces
- Packet routed to *nearest* interface
- Part of unicast range
- Little experience
  - Router-subnet anycast
  - Mobile-IPv6 home agent anycast
- Only for DA, never SA
- Only assigned to routers

IPv6 Hosts

- Always have multiple IPv6 addresses
  - Link-local (REQUIRED)
  - Unicast, anycast
  - Loopback
  - All-nodes multicast
  - Solicited multicast FOR EACH unicast/anycast
  - Multicast addresses
- On a router, ADDITIONALLY:
  - Subnet-router anycast addresses for all interfaces
  - Other anycast
  - All-routers multicast

IPv6 over Ethernet

- Type = 0x86DD
- Multicast mapping L3->L2:
  - MAC=3333<last 32 bits of IPv6 address>
- EUI-64:

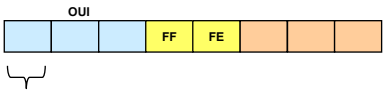

xxxx xx1x if MAC is unique (for future use)

Configuring IPv6

- (config)# ipv6 unicast-routing
- (config-if)# ipv6 address <addr/len> [link-local]
  - Used instead of automatically configured link-local address
  - Without "link-local" keyword: site-local or global address (requires specification of entire address)
- (config-if)# ipv6 address <prefix/len> eui-64
  - Short form, configures site-local and global addresses
  - Also link-local address automatically configured
- (config-if)# ipv6 unnumbered <interface>
  - Use global address of specified interface as SA
- (config-if)# ipv6 enable
  - Automatically configures link-local address (only)

Router Advertisements

- The previous commands (red and blue) automatically enables sending of Router Advertisements (RAs)
  - Autoconfiguration only works when advertised prefix-length == 64 bits !!!
- (config-if)# ipv6 nd prefix <prefix|default> [<valid-lifetime> <preferred lifetime>] [off-link] [no-autoconfig]
  - Only used to override default values
  - Default…all configured prefixes used

ICMPv6

- Optionally can use IPsec authentication & encryption
  - If SA between peers exists
- Used for Path MTU Discovery
  - Routers along path might reply with ICMP “packet too big, use MTU=…”
- Used for neighbor discovery (ND)
  - Like ARP: find MAC addresses of neighbors on that link
    - ICMP-type 135, SA=A, DA=Solicited-node mcast of B, Data=MAC A
    - Reply: ICMP-type 136, SA=B, DA=A, Data=MAC B
  - Also used to verify reachability of neighbor (with unicast DA)
  - Neighbor Advertisements can be sent when MAC changes (with all-nodes mcast DA)

ICMPv6 – RA

- Router Advertisements (RA)
  - ICMP type=134
  - Sent periodically (each 200 secs) and on request
  - DA=all-nodes mcast
- Contains
  - Prefixes used (=> autoconf hosts)
  - Lifetimes of prefixes (default: inf)
  - Flags (for autoconf)
  - Def. GW lifetime
  - Other (Note: no DNS information!)

ICMPv6 – RS

- Router Solicitation
  - ICMP type = 133
  - Sent by hosts at boot time (max 3 times) to trigger RAs
    - To immediately start autoconf
- Also used for Duplicate Address Detection (DAD)
  - RS with SA=:: and DA=solicited-node mcast of OWN address
  - Used during autoconf for safety

DNS

- AAAA records
  - www.xyz.at AAAA 3fee:b000:…
  - …0.0.b.e.e.f.3.ip6.arpa PTR www.xyz.at

DHCPv6

- More control than stateless autoconf with RS and RA
- Also usable for renumbering
- Also for dynamic DNS (automatic name registration)
  - 1) client detects router
  - 2) client receives RA
  - 3) client decides whether DHCPv6 can be used
  - 4) DHCP solicit if no router found
- Clients: DA=FF02::1:2 … “all-DHCP agents” (servers or relays)
  - Link-local scope
- Relays or clients: DA=FF05::1:3 …”all-DHCP servers
  - Site-local scope
  - SA must not be link-local

Mobile-IPv6 (TODO)

- Without loosing connection
- SA remains the home address
- IPv4 used tunneling which led to “triangle routing”