

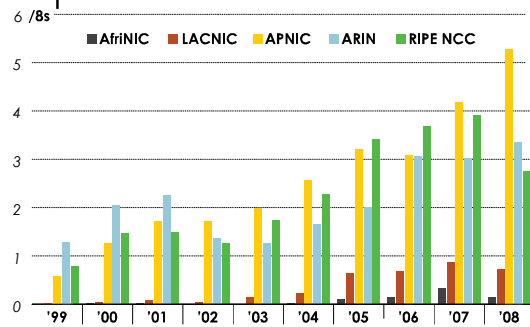
IPv6 : Internet Protocol Version 6

History

- o Internet growth was faster than anticipated
- o In early 1990's, it was realized that we may run out of IPv4 addresses somewhere between 2000 and 2010
- o Also, experiences with IPv4 showed its flaws
- o Work on IPv6 started

- o RFC 1883 in Dec. 1995 gives the basic specs (obsoleted by RFC 2460 later)
- o Experimental deployment started in 1995
- o Designed to work alongside IPv4 (Dual stack)
 - Early years of transition will have much more IPv6 machines than IPv4
 - Transition expected to take some time
- o Do we have time?
 - NAT, DHCP, and burst of dotcom bubble slowed down the rate of IPv4 address usage
 - Current estimate varies

IPv4 Year-Wise Allocation (/8s)





Types of IPv6 Addresses

- Unicast
 - An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
- Multicast
 - An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.
- Anycast:
 - An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance).



Broadcast

- There is no broadcast in IPv6.
- This functionality is taken over by multicast.
- Scope field in multicast address defines the scope of the multicast



IPv6 Addresses

- 128 bit
- Write as sequence of eight sets of four hex digits (16 bits each) separated by colons
 - Leading zeros in group may be omitted
 - Contiguous all-zero groups may be replaced by "::"
 - Only one such group can be replaced



Example

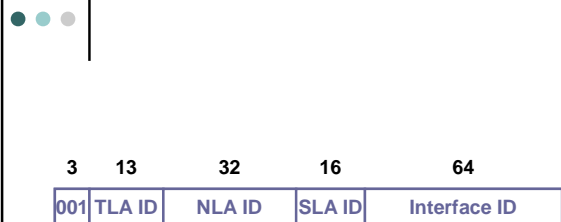
- Consider
 - 3ffe:3700:0200:00ff:0000:0000:0000:0001
- This can be written as
 - 3ffe:3700:200:ff:0:0:0:1 or
 - 3ffe:3700:200:ff::1

Types of Unicast Addresses

- Unspecified address
 - All zeros (::)
 - Used as source address during initialization
 - Also used in representing default
- Loopback address
 - Low-order one bit (::1)
 - Same as 127.0.0.1 in IPv4

Types of Unicast Addresses

- Aggregatable global unicast address
 - Public addresses on the internet
 - From range 2000::/3 (top 3 bits 010)
 - Three fields in /64 prefix
 - 13-bit Top Level Aggregator (TLA)
 - 8-bit reserved
 - 24-bit Next Level Aggregator (NLA)
 - 16-bit Site Level Aggregator (SLA)



- TLA
 - Allocated by RIRs to top level providers
- NLA
 - Allocated by TLAs to customers
 - Similar to subnetting by providers to give to customers
 - Can have multiple levels
- SLA
 - Similar to subnetting by customers within their organization
 - Can have multiple levels
- Interface ID
 - Usually derived from physical address
 - Identifies host



- In practice, terms like TLA, NLA are not used much anymore, but the intent of provider level aggregation is still there
 - Reduce routing entries in core routers
- RIRs allocate global routing prefixes, sizes may vary from the /16 specified by TLA (usually larger sized prefix, /32 etc.)



Types of Unicast Addresses

- Link-local address
 - Unique on a subnet
 - Auto configured
 - High-order: FE80::/64
 - Low-order: interface identifier
 - Routers must not forward any packets with link-local source or destination addresses.



Types of Unicast Addresses

- Site-local address
 - Unique to a "site"
 - High-order: FEC0::/48
 - Low-order: interface identifier
 - Used when a network is isolated and no global address is available.
 - Similar to private IP addresses in IPv4



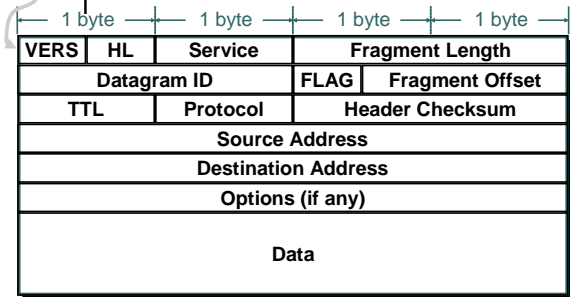
Types of Unicast Addresses

- Mapped IPv4 addresses/IPv4-mapped IPv6 addresses
 - Of form ::FFFF:a.b.c.d
 - Ex: ::FFFF:144.16.192.57
 - Used by dual-stack machines to communicate over IPv4 using IPv6 addressing
- Compatible IPv4 addresses/IPv4-compatible IPv6 addresses
 - Of form ::a.b.c.d
 - Used by IPv6 hosts to communicate over automatic tunnels

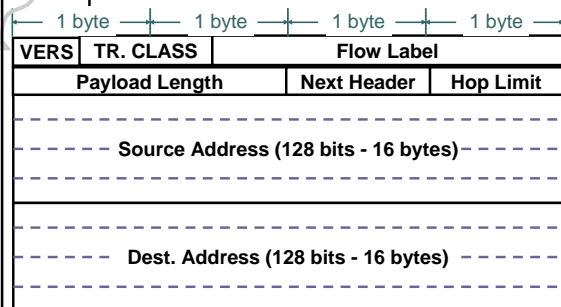
IPv6 Headers

- Simpler header - faster processing by routers.
 - Smaller no. of fields in base header
 - Fixed size for base header (40 bytes)
 - No checksum
 - No fragmentation information in base header
- Support for multiple headers
 - more flexible than simple "protocol" field
 - Each header linked to next header
 - Can be IPv6 extension headers or other higher layer protocol header (ex. TCP)

IPv4 Header



IPv6 Header



IPv6 Header Fields

- VERS (4 bits): 6 (IP version number)
- Traffic Class (8 bits): distinguish between different traffic class
- Flow Label (20 bits): experimental - sender can label a sequence of packets as being in the same flow.
- Payload Length (16 bits): number of bytes in everything following the 40 byte header, or 0 for a *Jumbogram*
 - If 0, jumbo payload option in hop-by-hop extension header will contain the actual length



IPv6 Header Fields

- Next Header (8 bits) is similar to the IPv4 “protocol” field - indicates what type of header follows the IPv6 header.
 - Can be IPv6 extension headers also
- Hop Limit (8 bits) is similar to the IPv4 TTL field (but now it really means hops, not time).



Extension Headers

- Hop-by-Hop option
 - Processed by routers
- Routing Header - source routing
- Fragmentation Header - supports fragmentation of IPv6 datagrams
 - Fragmentation supported only at source, reassembled only at destination
 - No fragmentation at intermediate routers
- Authentication Header
- Encapsulating Security Payload Header