IPv6 Protocols & Standards

ISP/IXP Workshops

So what has really changed?

- Expanded address space
  - Address length quadrupled to 16 bytes
- Header Format Simplification
  - Fixed length, optional headers are daisy-chained
  - IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- No checksum at the IP network layer
- No hop-by-hop segmentation
  - Path MTU discovery
- 64 bits aligned
- Authentication and Privacy Capabilities
  - IPsec is mandated
- No more broadcast

IPv4 and IPv6 Header Comparison

IPv4 Header **IPv6 Header**

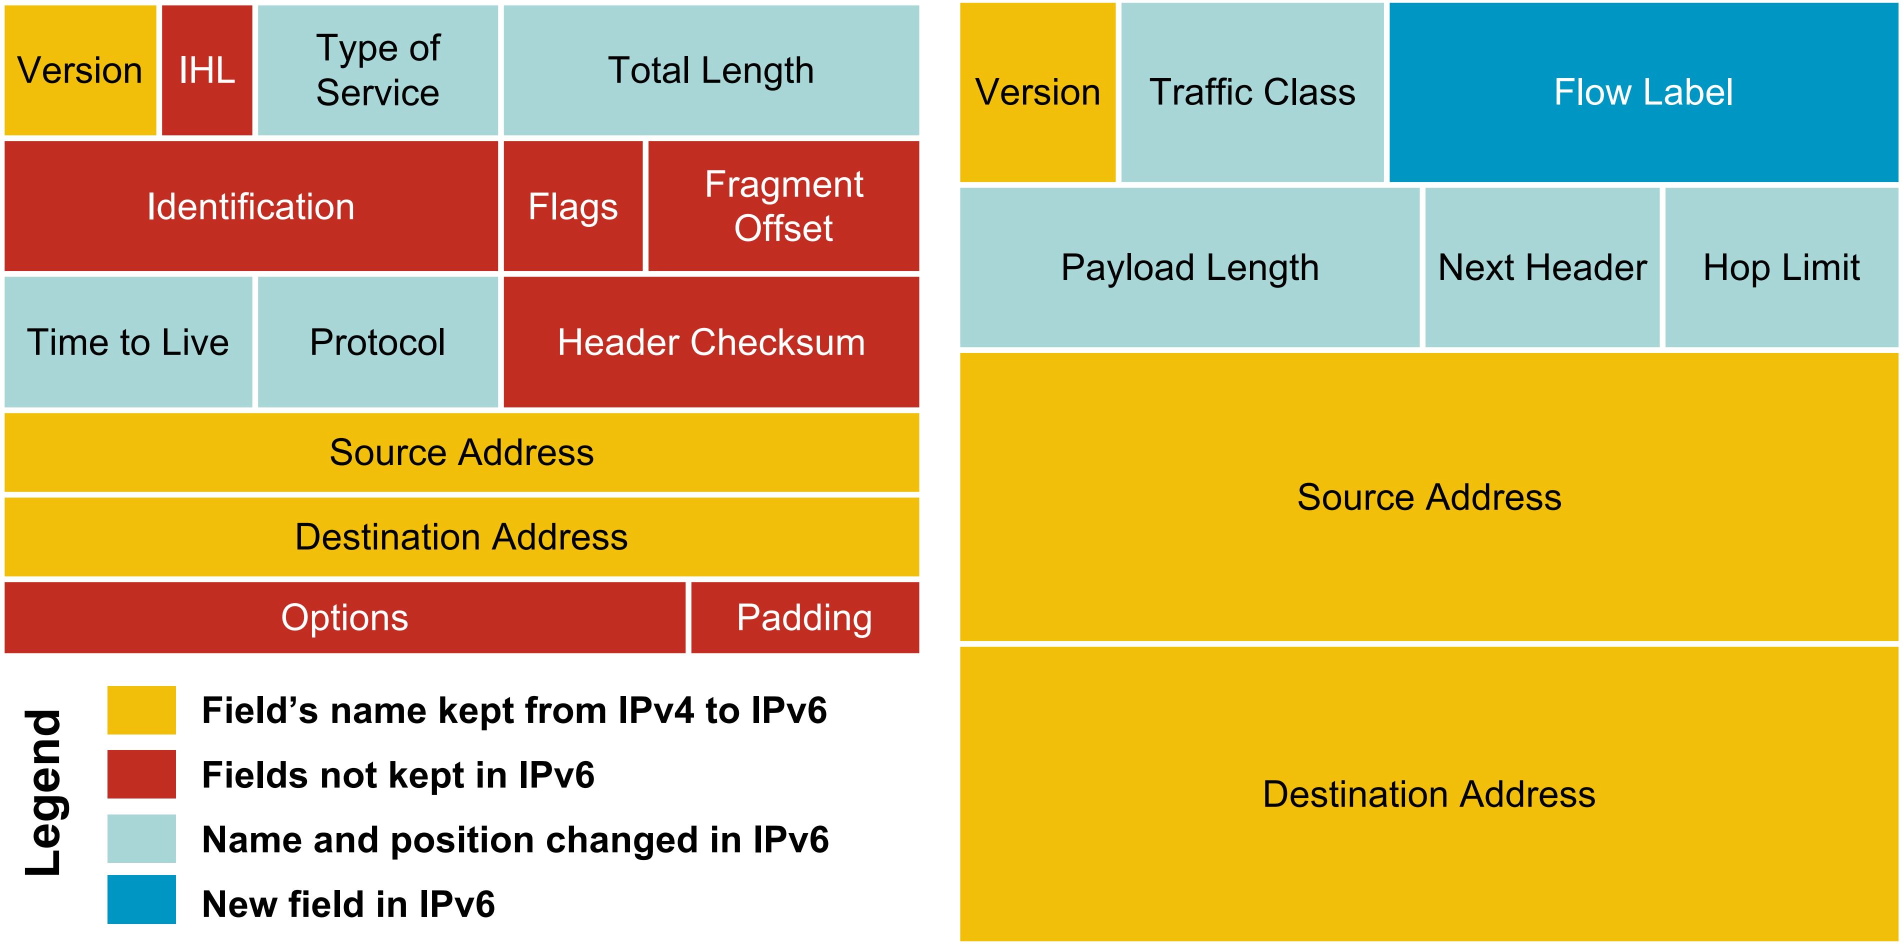

Larger Address Space

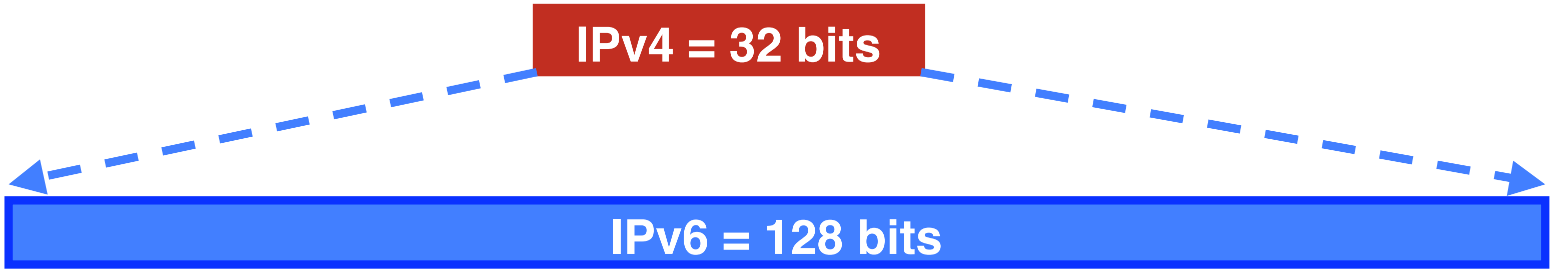

IPv4

- 32 bits
- = 4,294,967,296 possible addressable devices

IPv6

- 128 bits: 4 times the size in bits
- = 3.4×10^{38} possible addressable devices
- = 340,282,366,920,938,463,463,374,607,431,768,211,456
- ~ 5×10^{28} addresses per person on the planet

How was the IPv6 Address Size Chosen?

- Some wanted fixed-length, 64-bit addresses
  - Easily good for 1012 sites, 1015 nodes, at .0001 allocation efficiency (3 orders of magnitude more than IPv6 requirement)
  - Minimizes growth of per-packet header overhead
  - Efficient for software processing
- Some wanted variable-length, up to 160 bits
  - Compatible with OSI NSAP addressing plans
  - Big enough for auto-configuration using IEEE 802 addresses
  - Could start with addresses shorter than 64 bits & grow later
- Settled on fixed-length, 128-bit addresses

IPv6 Address Representation

- 16 bit fields in case insensitive colon hexadecimal representation

  2031:0000:130F:0000:0000:09C0:876A:130B

- Leading zeros in a field are optional:

  2031:0:130F:0:0:9C0:876A:130B

- Successive fields of 0 represented as ::, but only once in an address:

  2031:0:130F::9C0:876A:130B is ok

  2031::130F::9C0:876A:130B is NOT ok

  0:0:0:0:0:0:0:1 → ::1 (loopback address)

  0:0:0:0:0:0:0:0 → :: (unspecified address)

IPv6 Address Representation

- IPv4-compatible (not used any more)

  0:0:0:0:0:0:192.168.30.1

  = ::192.168.30.1

  = ::C0A8:1E01

- In a URL, it is enclosed in brackets (RFC2732)

  http://[2001:db8:4f3a::206:ae14]:8080/index.html

  Cumbersome for users

  Mostly for diagnostic purposes

  Use fully qualified domain names (FQDN)

- ⇒ The DNS has to work!!

IPv6 Address Representation

- Prefix Representation

  Representation of prefix is just like IPv4 CIDR

  In this representation you attach the prefix length

  Like IPv4 address:

    198.10.0.0/16

  IPv6 address is represented in the same way:

    2001:db8:12::/40

IPv6 Addressing

- IPv6 Addressing rules are covered by multiples RFCs
  - Architecture defined by RFC 4291
- Address Types are :
  - Unicast : One to One (Global, Unique Local, Link local)
  - Anycast : One to Nearest (Allocated from Unicast)
  - Multicast : One to Many
- A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, multicast)
  - No Broadcast Address → Use Multicast

IPv6 Addressing

Type	Binary	Hex
Unspecified	000…0	::/128
Loopback	000…1	::1/128
Global Unicast Address	0010	2000::/3
Link Local Unicast Address	1111 1110 10	FE80::/10
Unique Local Unicast Address	1111 1100 1111 1101	FC00::/7
Multicast Address	1111 1111	FF00::/8

IPv6 Global Unicast Addresses

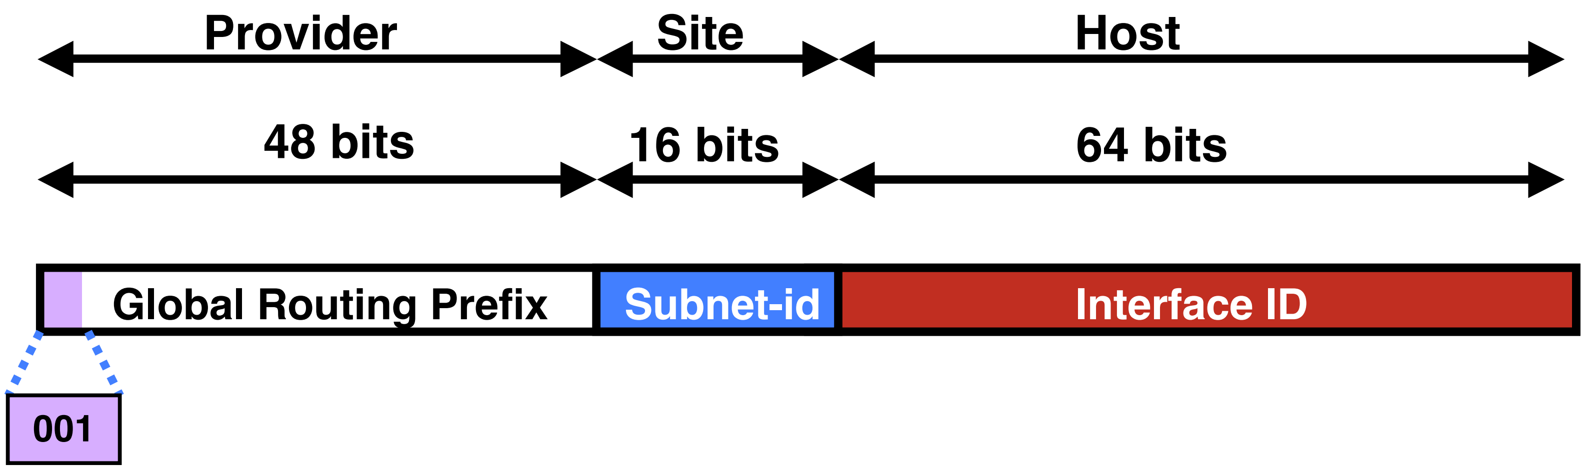

- IPv6 Global Unicast addresses are:
  - Addresses for generic use of IPv6
  - Hierarchical structure to simplify aggregation

IPv6 Address Allocation

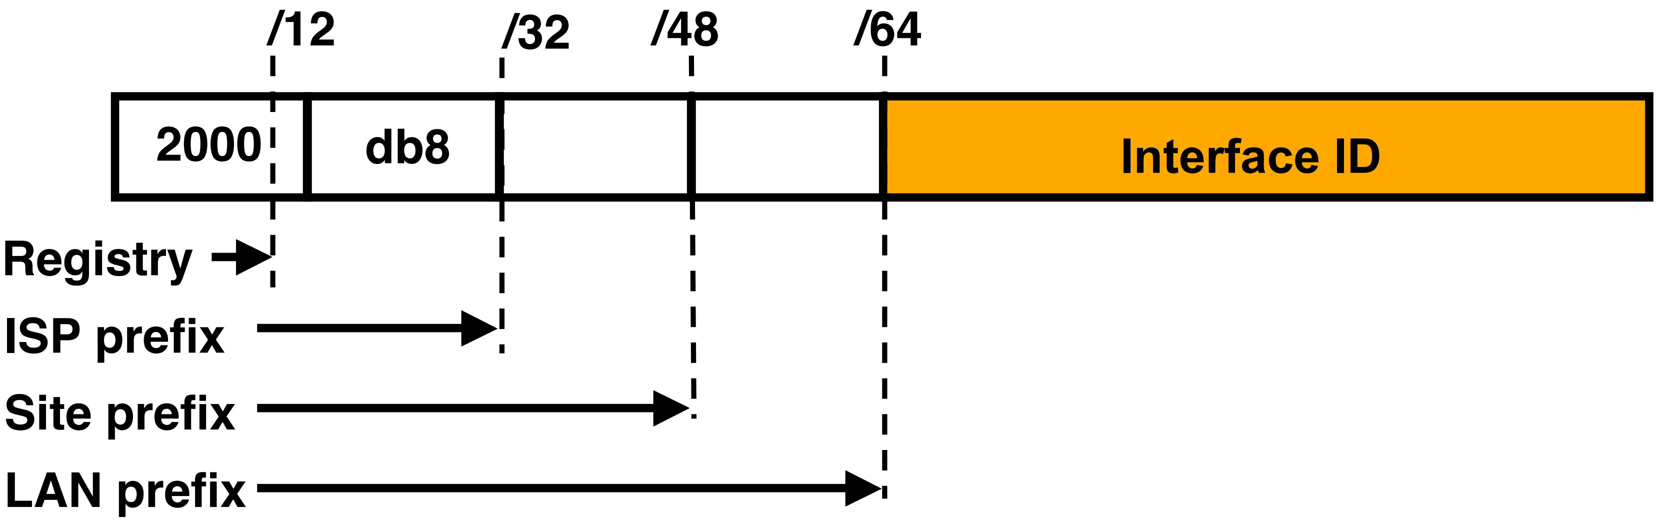

- The allocation process is:

    The IANA is allocating out of 2000::/3 for initial IPv6 unicast use

    Each registry gets a /12 prefix from the IANA

    Registry allocates a /32 prefix (or larger) to an IPv6 ISP

    Policy is that an ISP allocates a /48 prefix to each end customer

IPv6 Addressing Scope

- 64 bits reserved for the interface ID
  - Possibility of 2^{64} hosts on one network LAN
  - Arrangement to accommodate MAC addresses within the IPv6 address
- 16 bits reserved for the end site
  - Possibility of 2^{16} networks at each end-site
  - 65536 subnets equivalent to a /12 in IPv4 (assuming 16 hosts per IPv4 subnet)

IPv6 Addressing Scope

- 16 bits reserved for the service provider

  Possibility of 2^{16} end-sites per service provider

  65536 possible customers: equivalent to each service provider receiving a /8 in IPv4 (assuming a /24 address block per customer)

- 32 bits reserved for service providers

  Possibility of 2^{32} service providers

  i.e. 4 billion discrete service provider networks

    Although some service providers already are justifying more than a /32

  Equivalent to the size of the entire IPv4 address space

How to get an IPv6 Address?

- IPv6 address space is allocated by the 5 RIRs:

  AfriNIC, APNIC, ARIN, LACNIC, RIPE NCC

  ISPs get address space from the RIRs

  Enterprises get their IPv6 address space from their ISP

- 6to4 tunnels 2002::/16
- (6Bone)

  Was the IPv6 experimental network since the mid 90s

  Now retired, end of service was 6th June 2006 (RFC3701)

Aggregation hopes

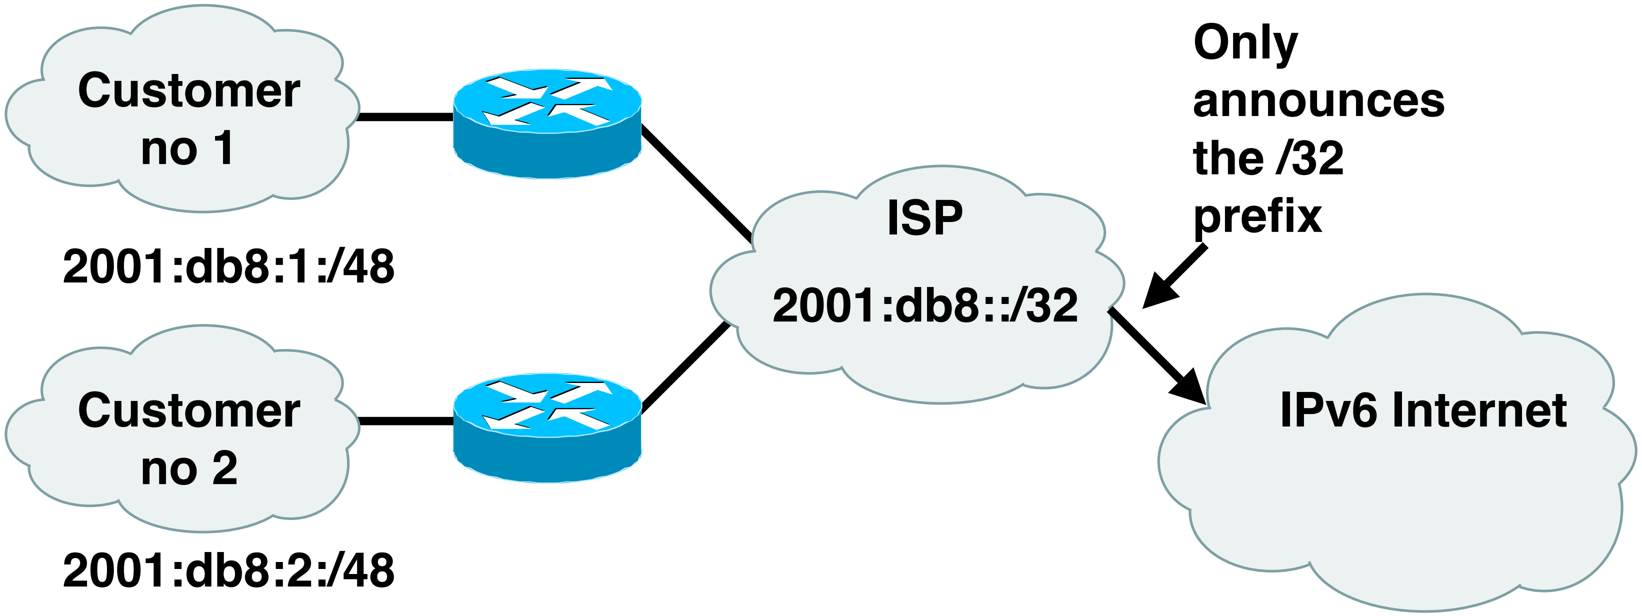

- Larger address space enables aggregation of prefixes announced in the global routing table
- Idea was to allow efficient and scalable routing
- **But current Internet multihoming solution breaks this model**

Interface IDs

- Lowest order 64-bit field of unicast address may be assigned in several different ways:
    - Auto-configured from a 64-bit EUI-64, or expanded from a 48-bit MAC address (e.g., Ethernet address)
    - Auto-generated pseudo-random number (to address privacy concerns)
    - Assigned via DHCP
    - Manually configured

EUI-64

Ethernet MAC address (48 bits)

64 bits version

Uniqueness of the MAC

Eui-64 address

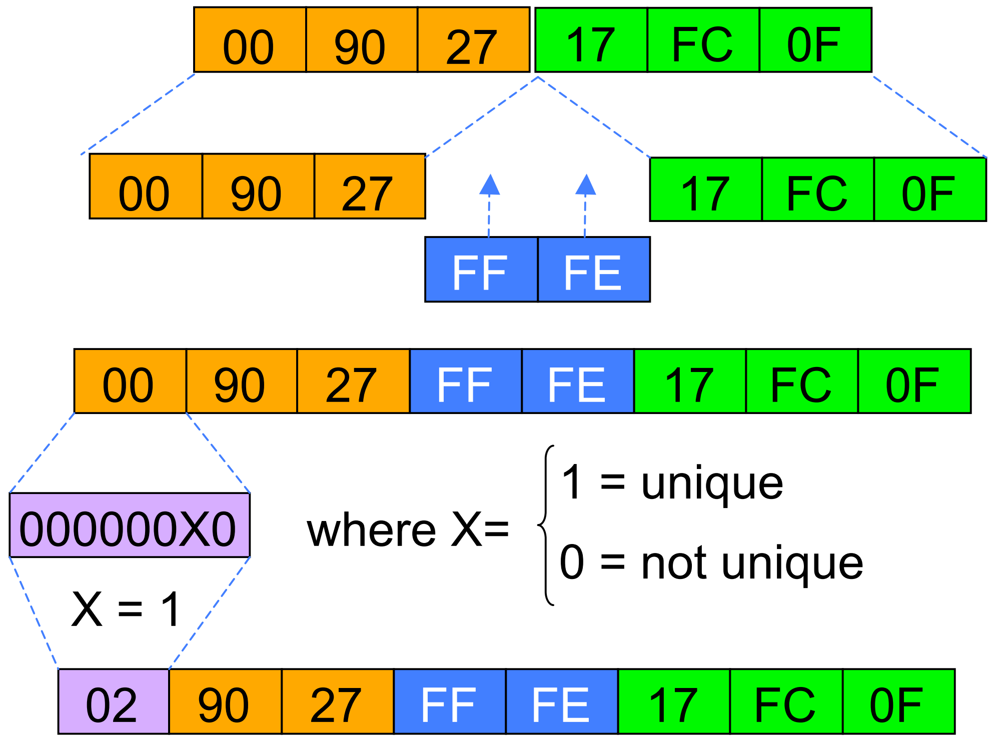

- EUI-64 address is formed by inserting FFFE and OR’ing a bit identifying the uniqueness of the MAC address

IPv6 Addressing Examples

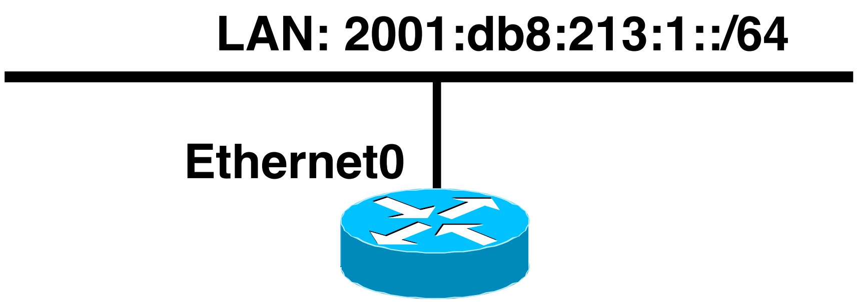


```
interface Ethernet0
 ipv6 address 2001:db8:213:1::/64 eui-64
```

MAC address: 0060.3e47.1530

```
router# show ipv6 interface Ethernet0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::260:3EFF:FE47:1530
Global unicast address(es):
    2001:db8:213:1:260:3EFF:FE47:1530, subnet is 2001:db8:213:1::/64
  Joined group address(es):
    FF02::1:FF47:1530
    FF02::1
    FF02::2
  MTU is 1500 bytes
```

IPv6 Address Privacy (RFC 3041)

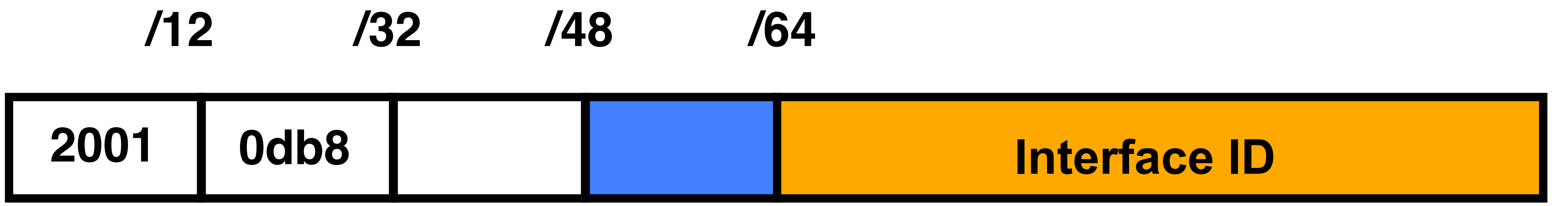

- Temporary addresses for IPv6 host client application, e.g. Web browser
- Intended to inhibit device/user tracking but is also a potential issue
  - More difficult to scan all IP addresses on a subnet
  - But port scan is identical when an address is known
- Random 64 bit interface ID, run DAD before using it
- Rate of change based on local policy
- **Implemented on Microsoft Windows XP only**

IPv6 Auto-Configuration

- Stateless (RFC2462)

  Host autonomously configures its own Link-Local address

  Router solicitation are sent by booting nodes to request RAs for configuring the interfaces.

- Stateful

  DHCPv6 – required by most enterprises

- Renumbering

  Hosts renumbering is done by modifying the RA to announce the old prefix with a short lifetime and the new prefix

  Router renumbering protocol (RFC 2894), to allow domain-interior routers to learn of prefix introduction / withdrawal

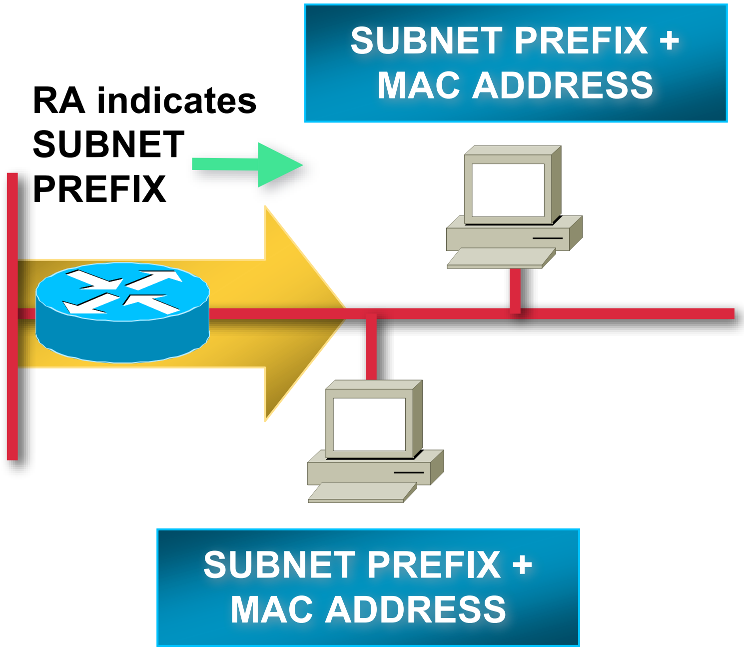

At boot time, an IPv6 host build a Link-Local address, then its global IPv6 address(es) from RA

Auto-configuration

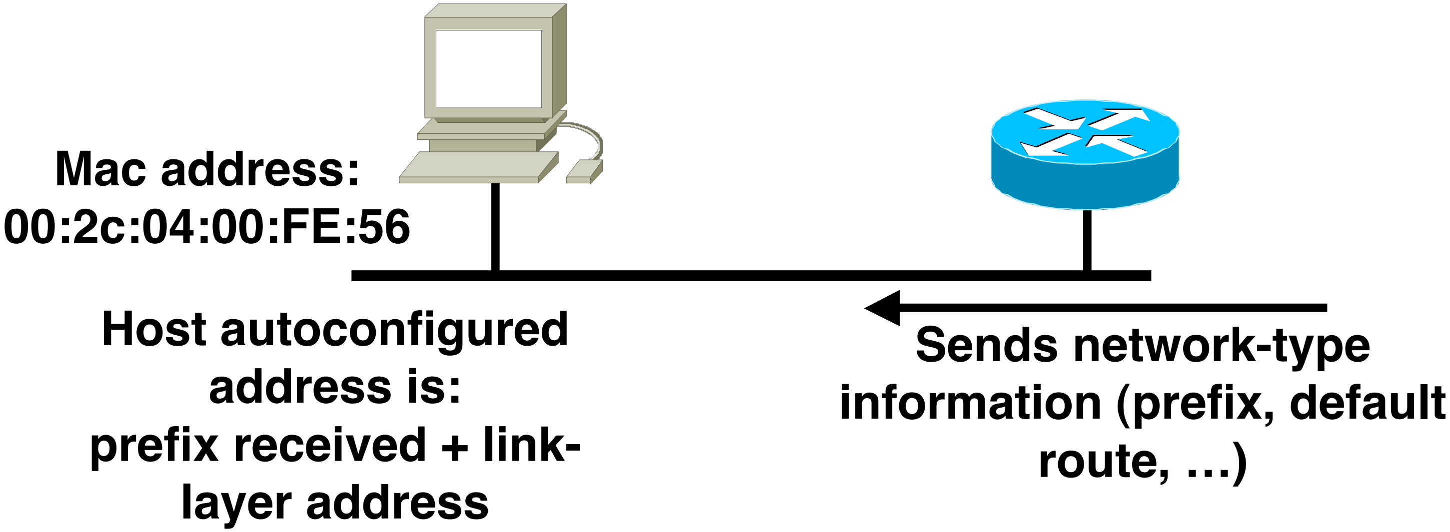

- Larger address space enables:
  - The use of link-layer addresses inside the address space
  - Auto-configuration with "no collisions“
  - Offers "Plug and play"

Renumbering

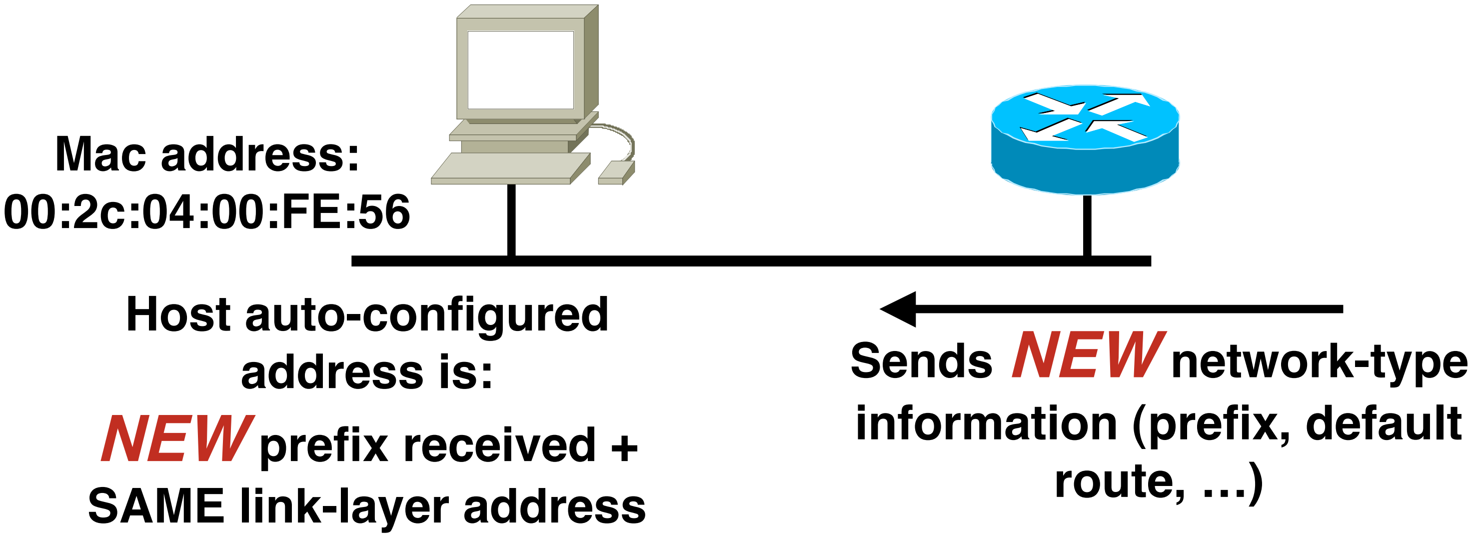

- Larger address space enables:
  - Renumbering, using auto-configuration and multiple addresses

Unique-Local

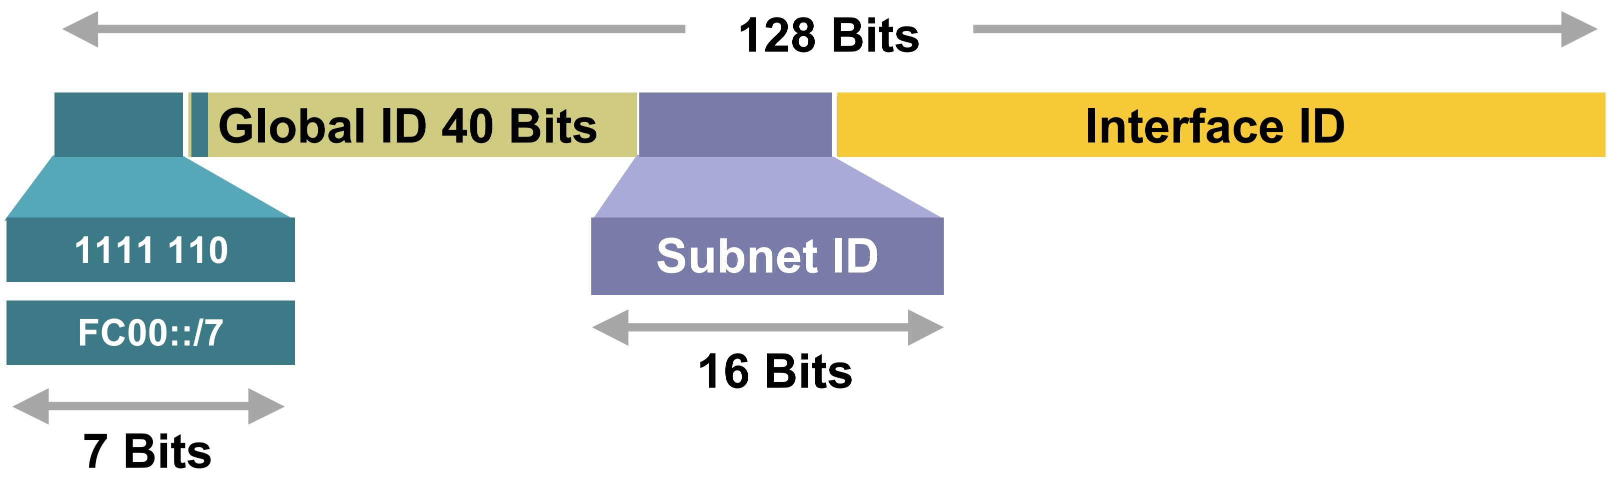

- Unique-Local Addresses Used For:
    - Local communications
    - Inter-site VPNs
- **Not** routable on the Internet
- Reinvention of the deprecated site-local?

Link-Local

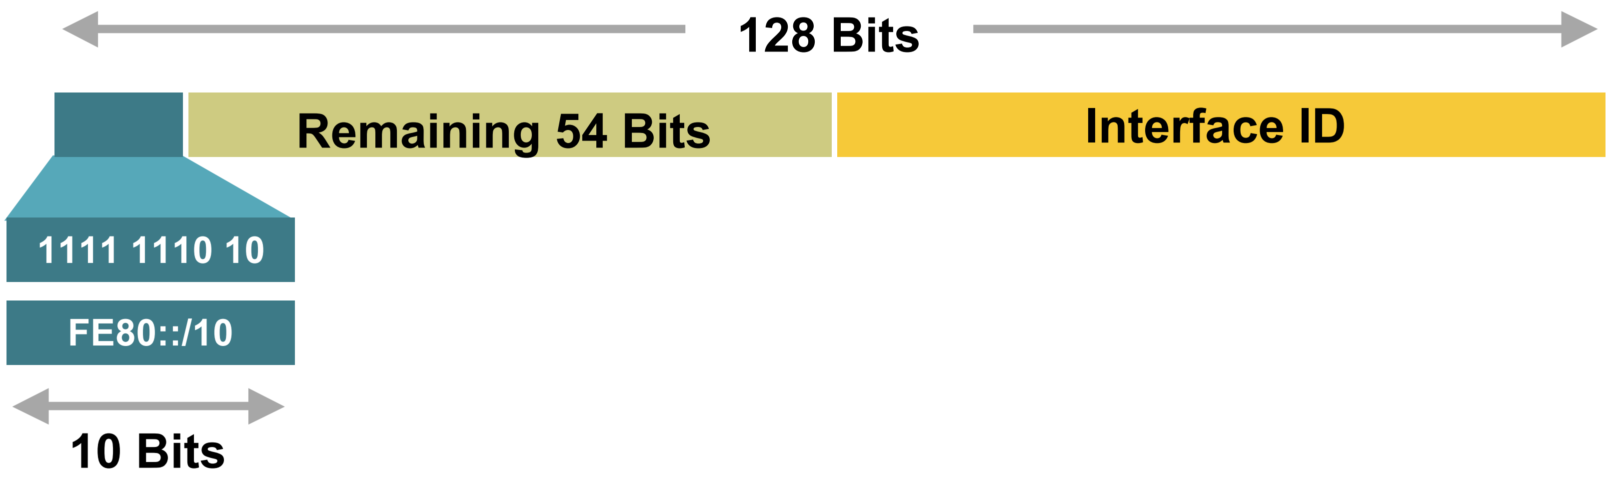

- Link-Local Addresses Used For:
  - Communication between two IPv6 device (like ARP but at Layer 3)
  - Next-Hop calculation in Routing Protocols
- Automatically assigned by Router as soon as IPv6 is enabled
  - Mandatory Address
- Only Link Specific scope
- Remaining 54 bits could be Zero or any manual configured value

Multicast use

- Broadcasts in IPv4
  - Interrupts all devices on the LAN even if the intent of the request was for a subset
  - Can completely swamp the network (“broadcast storm”)
- Broadcasts in IPv6
  - Are not used and replaced by multicast
- Multicast
  - Enables the efficient use of the network
  - Multicast address range is much larger

IPv6 Multicast Address

- IP multicast address has a prefix FF00::/8
- The second octet defines the lifetime and scope of the multicast address.

8-bit	4-bit	4-bit	112-bit
1111 1111	Lifetime	Scope	Group-ID

Lifetime	
0	If Permanent
1	If Temporary

Scope	
1	Node
2	Link
5	Site
8	Organization
E	Global

IPv6 Multicast Address Examples

- RIPng

  The multicast address AllRIPRouters is FF02::9

    Note that 02 means that this is a permanent address and has link scope

- OSPFv3

  The multicast address AllSPFRouters is FF02::5

  The multicast address AllDRouters is FF02::6

- EIGRP

  The multicast address AllEIGRPRouters is FF02::A

Solicited-Node Multicast

- Solicited-Node Multicast is used for Duplicate Address Detection as part of Neighbour Discovery
  - Replaces ARP
  - Duplicate IPv6 Addresses are rare, but still have to be tested for
- For each unicast and anycast address configured there is a corresponding solicited-node multicast address
  - This address is only significant for the local link

Solicited-Node Multicast Address

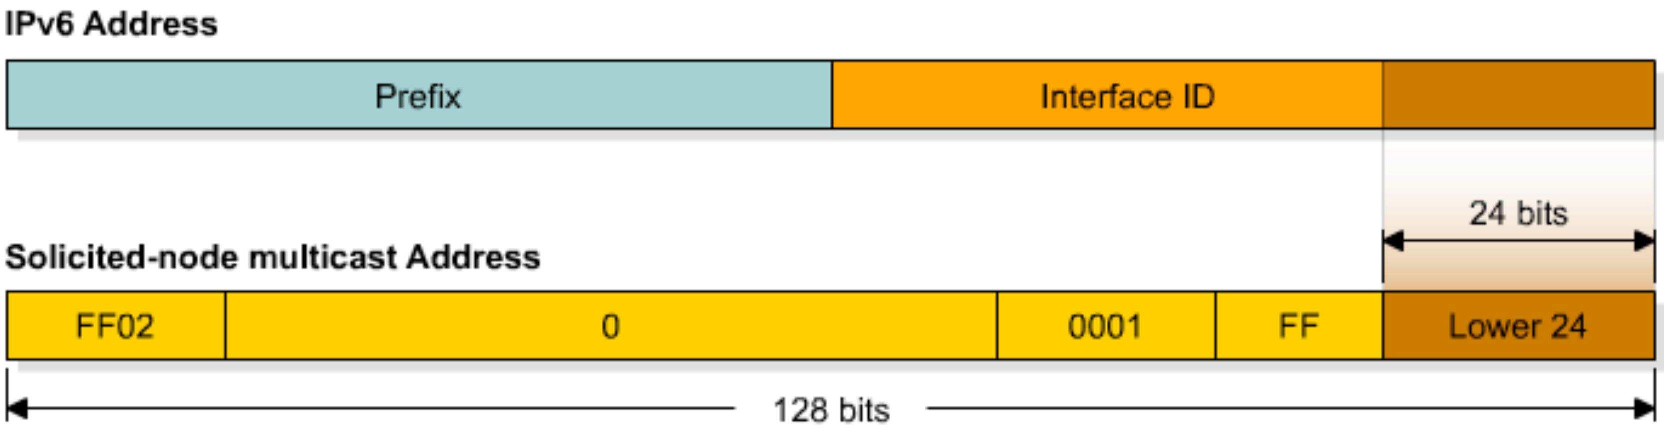

- Solicited-node multicast address consists of FF02:0:0:0:0:1:FF::/104 prefix joined with the lower 24 bits from the unicast or anycast IPv6 address

Solicited-Node Multicast

```
R1#sh ipv6 int e0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF3A:8B18
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
R1#
```

Solicited-Node Multicast Address

IPv6 Anycast

- An IPv6 anycast address is 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 protocol’s measure of distance).

  RFC4291 describes IPv6 Anycast in more detail

- In reality there is no known implementation of IPv6 Anycast as per the RFC

  Most operators have chosen to use IPv4 style anycast instead

Anycast on the Internet

- A global unicast address is assigned to all nodes which need to respond to a service being offered
  - This address is routed as part of its parent address block
- The responding node is the one which is closest to the requesting node according to the routing protocol
  - Each anycast node looks identical to the other
- Applicable within an ASN, or globally across the Internet
- Typical (IPv4) examples today include:
  - Root DNS and ccTLD/gTLD nameservers
  - SMTP relays within ISP autonomous systems

MTU Issues

- Minimum link MTU for IPv6 is 1280 octets (versus 68 octets for IPv4)
  - ⇒ on links with MTU < 1280, link-specific fragmentation and reassembly must be used
- Implementations are expected to perform path MTU discovery to send packets bigger than 1280
- Minimal implementation can omit PMTU discovery as long as all packets kept ≥ 1280 octets
- A Hop-by-Hop Option supports transmission of "jumbograms" with up to 2^{32} octets of payload

Neighbour Discovery (RFCs 2461 & 4311)

- Protocol built on top of ICMPv6 (RFC 4443)
  - combination of IPv4 protocols (ARP, ICMP, IGMP,…)
- Fully dynamic, interactive between Hosts & Routers
  - defines 5 ICMPv6 packet types:
    - Router Solicitation / Router Advertisements
    - Neighbour Solicitation / Neighbour Advertisements
    - Redirect

IPv6 and DNS

	IPv4	IPv6
Hostname to IP address	A record: www.abc.test. A 192.168.30.1	AAAA record: www.abc.test AAAA 3FFE:B00:C18:1::2
IP address to hostname	PTR record: 1.30.168.192.in-addr.arpa. PTR www.abc.test.	PTR record: 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.0.0.b.0.e.f.f.3.ip6.arpa PTR www.abc.test.

IPv6 Technology Scope

IP Service	*IPv4 Solution*	*IPv6 Solution*
Addressing Range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	Serverless, Reconfiguration, DHCP
Security	IPSec	IPSec Mandated, works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality-of-Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
IP Multicast	IGMP/PIM/Multicast BGP	MLD/PIM/Multicast BGP,Scope Identifier

What does IPv6 do for:

- Security
  - Nothing IPv4 doesn't do – IPSec runs in both
  - But IPv6 mandates IPSec
- QoS
  - Nothing IPv4 doesn't do –
    - Differentiated and Integrated Services run in both
    - So far, Flow label has no real use

IPv6 Security

- IPsec standards apply to both IPv4 and IPv6
- All implementations required to support authentication and encryption headers ("IPsec")
- Authentication separate from encryption for use in situations where encryption is prohibited or prohibitively expensive
- Key distribution protocols are not yet defined (independent of IP v4/v6)
- Support for manual key configuration required

IP Quality of Service Reminder

- Two basic approaches developed by IETF:
- "Integrated Service" (int-serv)

  Fine-grain (per-flow), quantitative promises (e.g., x bits per second), uses RSVP signaling

- "Differentiated Service" (diff-serv)

  Coarse-grain (per-class), qualitative promises (e.g., higher priority), no explicit signaling

- Signaled diff-serv (RFC 2998)

  Uses RSVP for signaling with course-grained qualitative aggregate markings

  Allows for policy control without requiring per-router state overhead

IPv6 Support for Int-Serv

- 20-bit Flow Label field to identify specific flows needing special QoS

  Each source chooses its own Flow Label values; routers use Source Addr + Flow Label to identify distinct flows

  Flow Label value of 0 used when no special QoS requested (the common case today)

- This part of IPv6 is standardised as RFC 3697

IPv6 Support for Diff-Serv

- 8-bit Traffic Class field to identify specific classes of packets needing special QoS
  - Same as new definition of IPv4 Type-of-Service byte
  - May be initialized by source or by router enroute; may be rewritten by routers enroute
  - Traffic Class value of 0 used when no special QoS requested (the common case today)

IPv6 Standards

- Core IPv6 specifications are IETF Draft Standards → well-tested & stable

  IPv6 base spec, ICMPv6, Neighbor Discovery, PMTU Discovery,...

- Other important specs are further behind on the standards track, but in good shape

  Mobile IPv6, header compression,...

  For up-to-date status: www.ipv6tf.org

- 3GPP UMTS Rel. 5 cellular wireless standards mandate IPv6; also being considered by 3GPP2

IPv6 Status – Standardisation

- Several key components on standards track…
  - Specification (RFC2460)
  - ICMPv6 (RFC4443)
  - RIP (RFC2080)
  - IGMPv6 (RFC2710)
  - Router Alert (RFC2711)
  - Autoconfiguration (RFC4862)
  - DHCPv6 (RFC3315 & 4361)
  - IPv6 Mobility (RFC3775)
  - GRE Tunnelling (RFC2473)
  - DAD for IPv6 (RFC4429)
  - Neighbour Discovery (RFC4861 & 4311)
  - IPv6 Addresses (RFC4291 & 3587)
  - BGP (RFC2545)
  - OSPF (RFC2740)
  - Jumbograms (RFC2675)
  - Radius (RFC3162)
  - Flow Label (RFC3697)
  - Mobile IPv6 MIB (RFC4295)
  - Unique Local IPv6 Addresses (RFC4193)
  - Teredo (RFC4380)
- IPv6 available over:
  - PPP (RFC5072)
  - FDDI (RFC2467)
  - NBMA (RFC2491)
  - Frame Relay (RFC2590)
  - IEEE1394 (RFC3146)
  - Ethernet (RFC2464)
  - Token Ring (RFC2470)
  - ATM (RFC2492)
  - ARCnet (RFC2497)
  - FibreChannel (RFC4338)

Recent IPv6 "Hot Topics" in the IETF

- **Multi-homing**
- Address selection
- Address allocation
- DNS discovery
- 3GPP usage of IPv6
- Anycast addressing
- Scoped address architecture
- Flow-label semantics
- API issues

  (flow label, traffic class, PMTU discovery, scoping,…)

- Enhanced router-to-host info
- Site renumbering procedures
- Inter-domain multicast routing
- Address propagation and AAA issues of different access scenarios
- End-to-end security vs. firewalls
- **And, of course, transition/ co-existence/interoperability with IPv4** (a bewildering array of transition tools and techniques)

Conclusion

- Protocol is “ready to go”
- The core components have already seen several years field experience

IPv6 Protocols & Standards

ISP/IXP Workshops