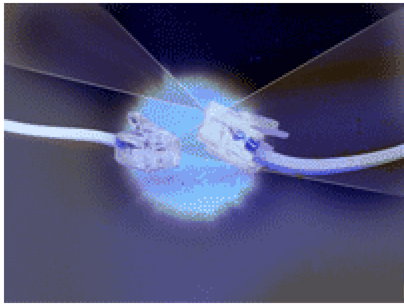

IPv6 - The Next Generation Internet

- Subnetting and Classless Inter-domain Routing (CIDR) improve utilization of IP address space and slow growth of routing information, but at some point, they will not be sufficient – more than 32 bits of IP address will be required
- Problem first examined by IETF in 1991
  - Expansion of IP address means change to header of every IP packet, and new software in every host and router running IP – big deal!

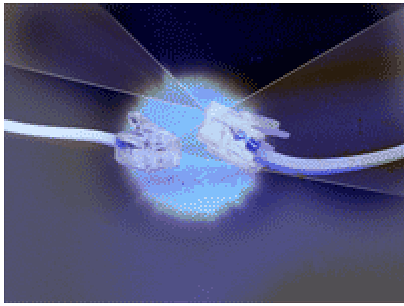

History (cont.)

- First known as IP Next Generation (IPng)
  - Changed to IPv6 when version number officially assigned
- Since a big change was being made, everyone wanted to fix other problems with IP at the same time
  - Support for real-time services
  - Security support
  - Autoconfiguration
  - Enhanced routing, support for mobile hosts

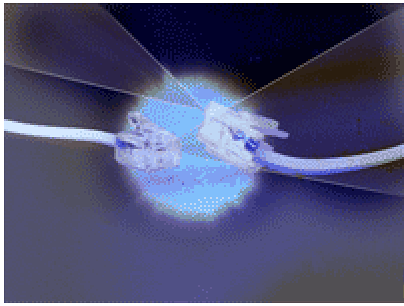

History (cont.)

- IPv6 took so long to materialize that several of these features have been incorporated into IPv4 in the meantime
- Big requirement was IPv4 → IPv6 transition plan, allowing for gradual switch-over of Internet

Addressing and Routing

- Address space increased from 32 bits to 128 bits
  - Even based on typical efficiency of address space usage, will provide over 1,500 addresses per square foot of Earth’s surface
- Addresses do not have classes like IPv4, but they are still subdivided based on leading bits
  - Bits specify use of address

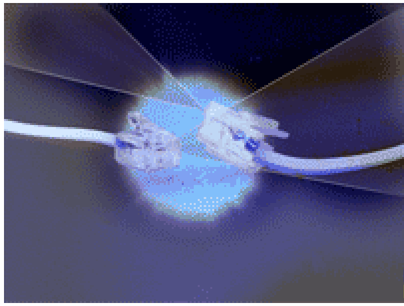

Addressing (cont.)

- All current addresses fit in the “Aggregatable Global Unicast Addresses” group, with 001 prefix. (1/8 of address space)
- Two sections reserved for encoding non-IP addresses – NSAP and IPX
- One section reserved for “Link Local Use” – addresses that will work on LAN and may not be globally unique
- One section for “Site Local Use” – similar to link local, for a local internetwork

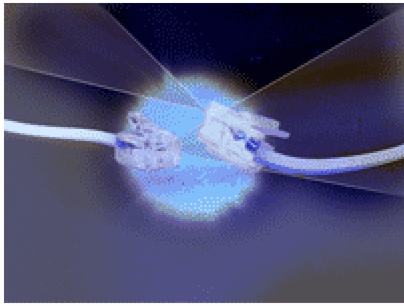

Addressing (cont.)

- One section for Multicast
- IPv4 addresses can be converted to IPv6 by zero-filling start of address (IPv4-compatible) or adding two bytes of 1s then zero-filling (IPv4-mapped)

• Notation: 1234:2346:1122:3344:5566:7788:99AA:BBCC

- One set of contiguous zeroes can be omitted
  - 1234::AABB:CCDD
- IPv4-mapped - ::FFFF:64.25.129.148

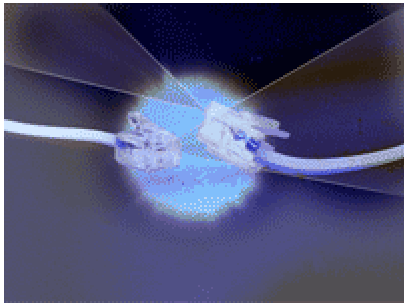

Unicast Addressing

- Subdivide *autonomous systems* (ASes)
  - – Subscriber – endpoint network
  - – Provider – transit network (i.e. intermediate)
    - Direct – connect to subscribers
    - Indirect – connect direct providers (backbones)
- Want to aggregate multiple networks to reduce routing information required
  - – Assign address prefix to direct provider
  - – Provider extends prefix for each subscriber

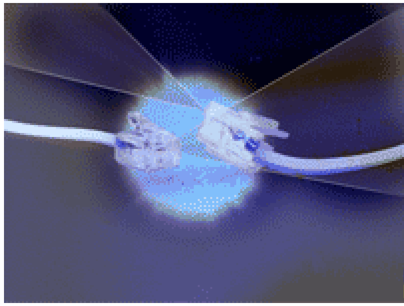

Unicast Addressing (cont.)

- Allows all subscribers to be routed through direct provider using single prefix
- Drawback: if site changes providers, needs to renumber every node on network
  - Ongoing area of research to find alternative
- Might aggregate at a higher level – i.e. by continent

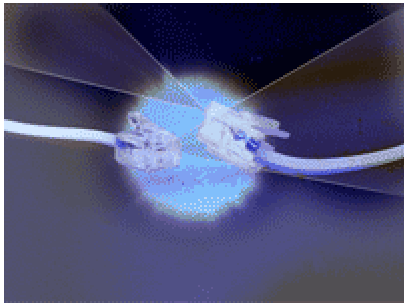

Packet Format

- 40-byte IPv6 header

0 4 8 12 16 19 24 31

Version	Traffic Class	FlowLabel
Payload Length	NextHeader	HopLimit
Source address (4 32-bit words)		
Destination address (4 32-bit words)		

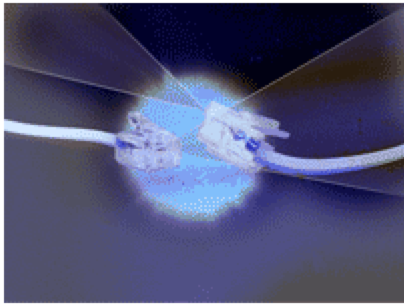

Packet Format (cont.)

- Header fields
  - Version set to 6 – in same place as IPv4
  - Traffic Class, Flow Label for QoS
  - Payload length in bytes
  - NextHeader combines options and protocol fields from IPv4
    - Fragmentation handled as optional header
  - Hop limit is like TTL field
  - Addresses 4 times length of IPv4

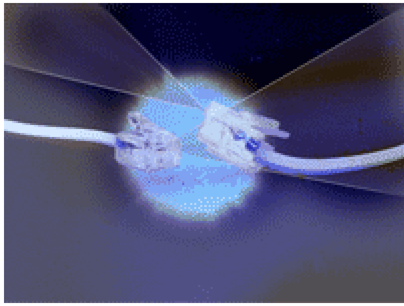

Optional Headers

- Extension headers, if present, must appear in predetermined order
  - NextHeader indicates type of following header – field must be included in every header
  - Can be of arbitrary length
  - Last extension header followed by transport header, which will contain protocol field
  - Example – fragmentation header

0	8	16	29	31
NextHeader	Reserved	Offset	Res	M
Ident				

Autoconfiguration

- Similar to DHCP, but *stateless* (doesn't require server)
- Two step process
  - Obtain interface ID unique on LAN to which interface is attached
  - Obtain correct address prefix for subnet
- First step – use MAC address, extend into link local address
  - If device doesn't need to communicate outside network, this step is sufficient

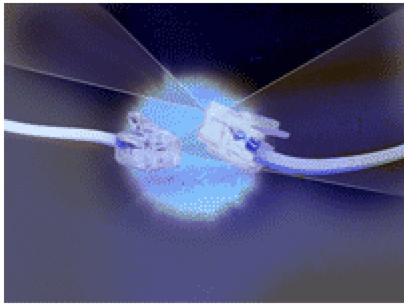

Autoconfig (cont.)

- Second step – routers periodically advertise correct prefix over LAN
  - Node replaces link-local prefix with correct subnet prefix
  - Requires prefix to leave at least 48 bits at end to insert MAC address

Enhanced Routing

- Another extension header for routing
  - – Allows *source-directed* routing
  - – Contains list of IPv6 addresses of nodes or areas that packet should visit en route to destination
    - Can route packets through specific providers; i.e. the cheap one vs. the expensive one that provides good QoS
  - – This will be used to provide mobile routing support; details are still being defined

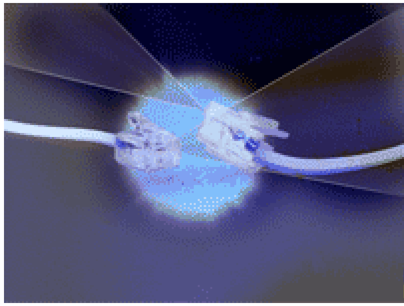

Additional Features

- QoS, security also incorporated
- Most features have already been fit into IPv4, so main reason for IPv6 is larger address space

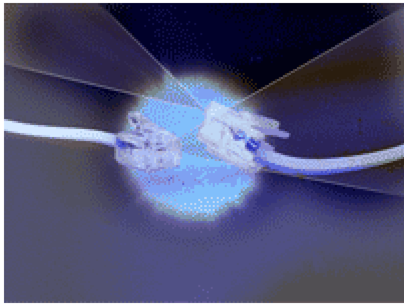

Network Address Translation (NAT)

- Big reason IPv6 hasn't been necessary sooner
- Assign addresses to hosts that are not globally unique, only unique within local context (i.e. within company network)
  - Network numbers 10 and 192.168.0 reserved for this use

NAT (cont.)

- If node needs to communicate outside local internetwork, a NAT device (possibly the router connected to the Internet) maps the local address to some globally unique address (i.e. the router’s outside IP)
  - NAT device may have small pool of IP addresses and assign them on a first-come, first-serve basis

NAT (cont.)

- The NAT device may also need to translate IP addresses carried inside application protocol – makes process complicated, limits introduction of new applications
- One advantage of understanding underlying protocol – NAT device can use information like TCP/UDP port to help do mapping between addresses