Introduction to IPv6

Thomas Narten

narten@us.ibm.com

April 6, 2003

ARIN XI

Memphis, Tennessee

Overview

- Background/History
- Benefits
- Technical Overview
- Coexistence with IPv4
- State of Deployment

Background and History

- 1991: IAB observed that more routing flexibility needed and that address exhaustion would happen
- 1991: ROAD WG begins studying issue
- 1993: CIDR (RFC 1519) produced
- 1994: IETF settles on IPv6 as basis for IPng
  - Evolution, *not* revolution
- 1996: Base IPv6 protocols as Proposed Standard
- 1998: Base protocols to Draft Standard
- Today: Many products, experimentation, initial deployment

Benefits

- Provides almost unlimited addresses
  - More addresses *cannot* be retrofitted into IPv4
- Plus
  - Improved autoconfiguration
  - Improved support for site renumbering
  - Mobility with route optimization (important for wireless)
  - Miscellaneous minor improvements

The Real Costs of NAT

- IPv4 address shortage has led to extensive deployments of network address translators (NAT)
- NATs delay, but do not obviate need for IPv6
  - Pain of NATs depends on one's perspective
- NATs are barrier to continued Internet scaling
  - Assume simple client-server programming and deployment model
  - NATs break protocols that rely on globally unique addresses, e.g., IPsec security, some audio/video
  - NATs have operational and administrative scaling problems
  - Always-on devices need permanent, global addresses (NATs prevent this)
  - Barrier to deployment of new types of applications (e.g., peer-to-peer, IP telephony, multi-party applications)
- IPv6 alleviates these problems and removes barriers to continued Internet expansion

Important IPv6 Technical Features

- IPv6 header and extension headers
  - Stream-lined IPv6 header
  - Optional extension headers for fragmentation, security, etc
- Routers no longer fragment forwarded datagrams
- Extended IP Address
  - 32 bits => 128 bits (but only 64 bits for routing)
- Neighbor Discovery & Stateless Auto-Configuration
  - Router Discovery and Neighbor Unreachability Detection (NUD)
  - Address configuration with no manual or server-based configuration
- IPv4/IPv6 Coexistence & Transition Mechanisms
  - Co-existence of IPv4 & IPv6
  - Tunneling and translation mechanisms

IPv6 Headers

- 40-byte IPv6 header (vs. 20 bytes for IPv4)
  - 16-byte IPv6 vs. 4-byte IPv4 address
- No IPv6 header checksum
  - End-to-end (e.g., TCP, UDP) checksum more appropriate
- "Next header" facility for chained extension headers
  - Extension headers used for routing, security, options
  - Fragmentation requires an extension header
- Flow label field (no IPv4 counterpart)
  - Minimizes need to parse through extension headers for upper layer ports
  - Potential long-term benefit, no proposed usage today

IPv6 Header Format

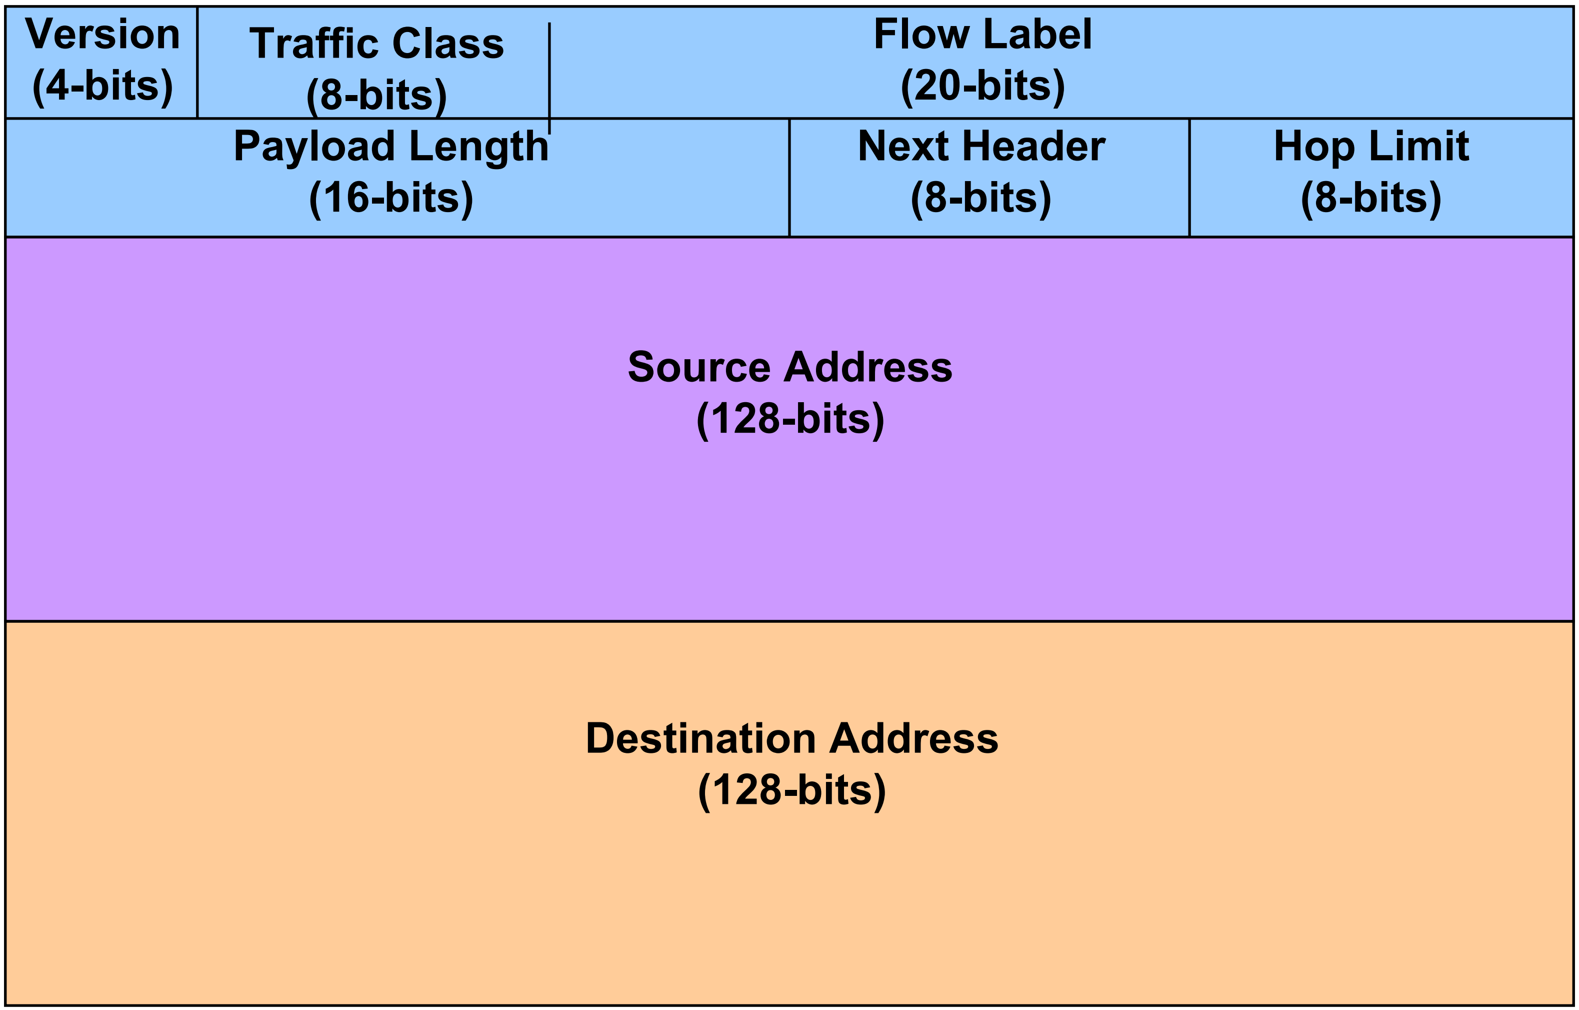

IPv6 Extension Headers

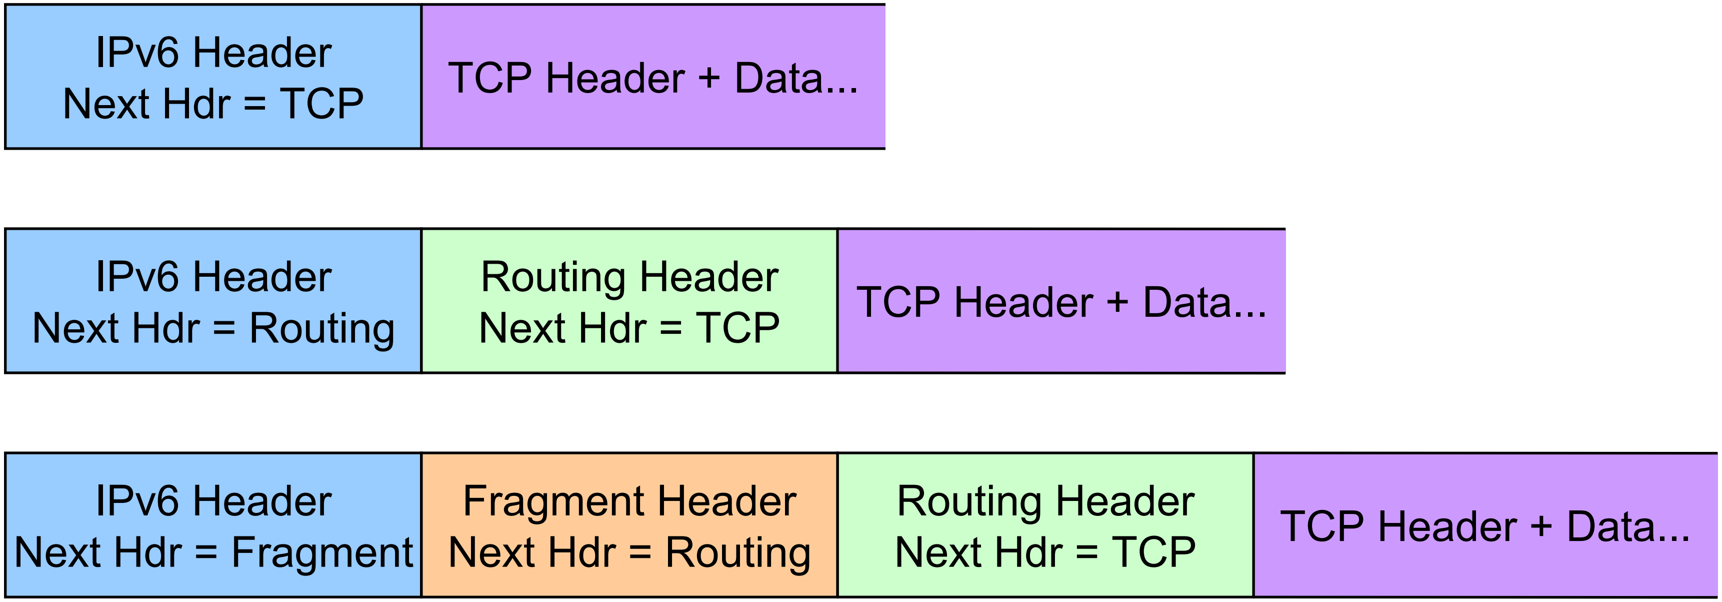

IPv6 Addressing

- IPv6 addresses are 128-bits long
  - A 64-bit subnet prefix identifies the link of node
  - followed by a 64-bit Interface Identifier (IID)
- IID derived from IEEE identifier (I.e., MAC address)
- Only left-most 64 bits available for routing and “network addressing”

Global Routing Prefix (n bits)	Subnet ID (64-n bits)	Interface Identifier (IID) (64-bits)

IPv6 Address Text Representation

- Addresses are represented as 8 sets of 4 hex digits (16-bits), separated by colons

2001:183E:0:0:240:2BFF:FE3D:71AD

- Two colons in a row can be used to denote one or more sets of zeroes, usually used between the prefix and the interface ID

2001:183E::240:2BFF:FE3D:71AD

- The prefix length can be indicated after a slash at the end

2001:183E::240:2BFF:FE3D:71AD/64

- A prefix alone is represented as if the interface ID bits are all zero

2001:183E::/64

IPv6 Address Allocation Model

- Architecture:
  - End nodes addresses numbered with 64-bit IID
  - Assumed by stateless address autoconfig (RFC 2461)
- RIR policy (with input from IETF, e.g., RFC 3177):
  - Simple sites get a /64
  - More complex sites get /48
  - /128 assignments possible, but:
    - Inconsistent with privacy extensions (RFC 3041)
    - No need from address conservation perspective
- Host-Density ratio (RFC 3194), indicates:
  - 100's of millions of /48 prefixes can be assigned
  - Ability to address is not the limit; aggregation and route scaling is
- Explicit engineering tradeoff to simplify aspects of end sites at expense of "wasting" address space

Neighbor Discovery

- Provides three different functions
- Router Discovery
  - Router Solicitations and Router Advertisements used to find and keep track of neighboring routers
  - Additional information for IP stack configuration
- Neighbor Discovery
  - Neighbor Solicitations and Neighbor Advertisements perform address resolution (i.e., ARP functions)
  - Uses ICMP rather than running over link layer
- Neighbor Unreachability Detection (NUD)
  - Keep track of reachability of neighbors
  - If path to router fails, switch to another router before TCP timeouts

Stateless Address Autoconfiguration

- Address Configuration without separate DHCP server
  - Router *is* the server, advertising key address configuration info
- Addresses formed by combining routing prefix with IID
- Link-local address configured when an interface is enabled
  - Allows immediate communication with devices on the local link
  - Primarily used for bootstrapping and discovery
  - Well-known prefix combined with locally-generated 64-bit IID
- Other addresses configured via Routing Advertisements
  - RA advertises 64-bit prefixes (e.g, on-link, form an address)
  - Public (e.g, server) addresses formed from interface IID
  - Randomly-generated IIDs support privacy addresses (RFC 3041)
  - Prefixes lifetimes enable graceful prefix changes/renumbering

Prefix Delegation

- Need way to delegate address prefixes from ISP to customer
- Customer edge router (e.g, DSL) requests address prefix from ISP (after providing appropriate credentials)
- Edge router redistributes route internally, e.g.:
  - Routers advertise individual prefix in RAs
  - Hosts generate addresses via stateless address autoconfiguration
- See draft-ietf-dhc-dhcpv6-opt-prefix-delegation (work nearly complete)

Site Renumbering Support

- Nodes can have multiple addresses
  - One from each ISP
  - One from “old” ISP, one from “new” ISP
- Addresses have associated lifetimes
  - Valid Lifetime: how long the address can be used (e.g. is routed and works)
  - Preferred Lifetime: At what point the address should stop being used (gracefully)
- To renumber a site:
  - Introduce new prefix (e.g, from new ISP)
  - Use both during transition
  - Phase out old address when new addresses working satisfactorily

IPv6 Scoped Unicast Addressing

- Concept of scoped unicast addresses part of architecture
- Link-local addresses for use on a single link
  - Primarily used for bootstrapping and infrastructure protocols such as Neighbor Discovery
  - Address = well-known link-local prefix plus node-generated IID
- Site-local addresses for use within a site
  - Like net 10
  - Full (negative) implications only recently understood
    - Application complexity
    - Nodes in multiple sites simultaneously
  - IETF may deprecate
- Global address prefixes are provided by ISPs

Elimination of IP Broadcast

- IPv6 eliminates broadcast addresses
  - All-nodes multicast address used only when all nodes targeted (relatively rare)
  - Targeted multicast address groups are used by many protocols
    - All routers multicast address
    - Solicited-nodes multicast address for Neighbor Discovery (based on interface identifier)
  - Eliminates unnecessary interrupts to handle broadcast traffic

IPv4/IPv6 Coexistence

- IPv6 was designed for co-existence with IPv4
  - IPv4-only, IPv6-only and IPv4/IPv6 nodes on a single network
- Dual Stack implementations include IPv4 & IPv6 in a single node
  - IPv4 is used to reach IPv4-only nodes and services
  - IPv6 is used to reach IPv6-only nodes and services
  - DNS lookups return A or AAAA Resource Records
  - Either version can be used to reach other dual stack implementations

IPv4/IPv6 Coexistence (cont.)

- Many IPv6 deployments will not require cutover transition mechanisms
  - Side-by-side deployment of native IPv4 and native IPv6
  - Notion of building IPv6-only site seems distant
- Transition mechanisms exist for special cases
  - Tunnels for IPv6 service over IPv4-only transit networks
    - Configured (e.g, by ISPs)
    - Automatic (e.g, by home users via 6to4)
  - NAT-PT translation to allow IPv6-only nodes to communicate to the IPv4 Internet
    - But with all the features of NAT (and maybe more)

6to4 Implicit Tunnels (RFC 3056)

- End sites want globally unique IPv6 /48 prefix, even when ISP doesn't support IPv6
- Form IPv6 addresses from IPv4 address:
  - Globally unique IPv4 address forms IPv6 /48 prefix
  - Within site, 100% standard IPv6
- Three types of routers
  - Normal router (exactly what you expect)
  - 6to4 router (e.g, sits between edge customer and ISP)
  - 6to4 relay router (glues the IPv6 and IPv4 worlds together)

6to4 Implicit Tunnels (cont.)

- 6to4 Router:
  - Advertises default route for IPv6 internally via IGP
  - Tunnels 6to4 packets for 6to4 destinations directly to IPv4 address (IPv6 in IPv4 tunneling)
  - Forwards packets for native IPv6 destinations (if it has a route)
  - Tunnels other IPv6 packets to Relay Router
- Relay Router
  - Advertises 2002::/16 prefix into IPv6 network (e.g, IGP/BGP)
  - Advertises 192.88.99.0/24 prefix into IPv4 network (IGP/BGP)
  - Relays transit traffic between native IPv6 and 6to4 sites

NAT-PT (Protocol Translation)

- Analogous to IPv4 NAT, but more to translate
  - Same limitations as NAT for IPv4 (plus more?)
- Some feeling that one is better off using IPv4 NAT and dual stack
  - IPv4 NAT is more of a known quantity
  - New corner cases and implementation features will arise with NAT-PT
- Does IPv6-only site make sense, in the short term?

IETF IPv6 Activities

- The entire IETF is assuming responsibility for IPv6
- Internet Protocol version 6 (IPv6) Working Group
  - Standardization of IPv6 protocols
- IPv6 Operations (v6ops) Working Group
  - Identifying IPv6 operational and security issues
  - Planning for the transition from IPv4 to a shared IPv4/IPv6 Internet
- IPv6 Site Multihoming (multi6) Working Group
  - Finding an scalable approach to site multihoming in IPv6
  - Very hard problem, no clear approach yet…
- Much activity in other IETF areas
  - IPv6 addresses in MIBs, IPv6 routing protocols, IPv6 applications...

IETF V6OPS Working Group

- Ngtrans -> v6ops; focus on operational issues
- Analyze common deployment scenarios that mirror how one would deploy IPv6
- Address problem:
  - Too many individual tools, unclear which ones are key
  - Ensure that key ones are robust and fully understood
  - Fully consider operational impacts/limitations
  - Some proposed tools are not on Standards Track yet (e.g., DSTM, Teredo, ISATAP)

IPv6 Deployment Status

- Good News
  - Extensive research and test network deployment
    - 6Net, 6Bone, WIDE, UNH, TAHI, etc.
  - Commercial IPv6 deployment has begun
    - Backbone networks in Europe, U.S. and Japan
    - Commercial DSL service in Japan
- Bad News
  - Minimal commercial availability
  - Very little IPv6 traffic
  - Robust, fully-functional products still not universal
  - IPv6 connectivity is currently slow and unreliable for many users
    - Often achieved via tunneling to public relays

Drivers for IPv6 Deployment

- Government support
  - Wide-scale IPv6 promotion underway in Japan, Korea & Taiwan
    - Continue Internet growth and establish improved competitive position, roll roll-out by 2005
  - European Commission (EC) encourages IPv6 research, education, and adoption in member countries
    - Continue growth of the Internet, becoming most dynamic and competitive knowledge based economy by 2011
  - Discussions within US DoD on need for IPv6
- Continuing rapid growth of the Internet
  - China plans to roll out ~1 billion Internet nodes
    - Starting with a 320 million student educational network
  - Real Internet-connected cell phones are coming (slowly)
  - Billions of Internet-connected consumer devices are on the way
    - Including many low cost, minimally configurable devices

Obstacles to IPv6 Deployment

- Economic slowdown has slowed growth and spending
  - Network infrastructure vendors are not introducing new products quickly
  - Service providers are not upgrading and expanding networks
- IPv6 upgrades to network infrastructure are expensive
  - IPv6 routing performance requires hardware upgrades
  - New technology requires staff training
  - New code/additional complexity will cause added support burdens
  - No current revenue stream to justify the costs
- Major technology markets are comfortable with IPv4
  - U.S., Europe have (relatively) many IPv4 addresses
  - Address shortages have been mitigated by use of NAT
- Benefits of IPv6 are not widely understood or not compelling
  - Desire that it solves more problems (e.g., multihoming)

Preparing for a Long Transition

- The transition from IPv4 to a shared IPv4/IPv6 network will be long and difficult
  - IPv6 deployment will begin at different times in different parts of the network
  - IPv6 may be deployed for specialized products and services before native IPv6 is available from many providers
  - There may be operational or security issues with layering a flat IPv6 architecture over a NAT-based IPv4 network
  - Extensive use of tunneling is expected, to allow IPv6 connectivity to IPv4 networks, and vice versa
- The transition must be carefully planned and managed to avoid serious operational and security problems

Summary

- IPv6 is a stable base
- IPv6 will be introduced gradually
- IPv6/IPv4 coexistence for foreseeable future

Questions?

Acknowledgments

- This presentation derived from an earlier presentation by Margaret Wasserman