Transition to IPv6

Hui Chen, Ph.D.
Dept. of Engineering & Computer Science
Virginia State University
Petersburg, VA 23806

Acknowledgements

- Some pictures used in this presentation were obtained from the Internet
- The instructor used the following references
  - Larry L. Peterson and Bruce S. Davie, Computer Networks: A Systems Approach, 5th Edition, Elsevier, 2011
  - Andrew S. Tanenbaum, Computer Networks, 5th Edition, Prentice-Hall, 2010
  - James F. Kurose and Keith W. Ross, Computer Networking: A Top-Down Approach, 5th Ed., Addison Wesley, 2009
  - Larry L. Peterson's (http://www.cs.princeton.edu/~llp/) Computer Networks class web site

Outline

- IPv6
  - Addressing
- IP Multicast
- Mobile IP

IPv6: Major Features

- 128-bit addresses
- Multicast
- Real-time service
- Authentication and security
- Auto-configuration
- End-to-end fragmentation
- Enhanced routing functionality, including support for mobile hosts

IPv6 Addresses

- Classless addressing/routing
  - Similar to CIDR
- 128 bits/16 bytes in length

IPv6 Address: Notation

- IPv6 Notation: a human friendly text representation
- x:x:x:x:x:x:x:x where x is a 16-bit (or 2-byte) hexadecimal number, e.g.,
  - `47CD:1234:4422:ACO2:0022:0022:1234:A456`
- Contiguous 0s can be compressed, e.g.,
  - `47CD:0000:0000:0000:0000:0000:A456:0124`
  - can be written as
  - `47CD::A456:0124`

IPv6 Addresses with Embedded IPv4 Addresses

- IPv4-mapped IPv6 address
  - Prefixing with 2 bytes of all 1's and then zero-extending to 128 bits, e.g, :
  - `::FFFF:150.174.2.101`
- IPv4-compatible IPv6 address
  - *Deprecated*
  - Zero-extending IPv4 addresses to 128 bits, e.g.,
  - `::150.174.2.101`
- See https://tools.ietf.org/html/rfc4291

IPv6 Address Types

Address Type	Binary Prefix	IPv6 Notation
Unspecified	00…0 (128 bits)	::/128
Loopback	00…1 (128 bits)	::1/128
Multicast	1111 1111	FF00::/8
Link-local Unicast	1111 1110 10	FE80::/10
Global Unicast	Everything else	

IPv6 Address Assignment

- Provider-based
  - Assign an address prefix to a direct provider
  - The direct provider assigns longer prefixes to its subscribers
  - The scheme allows the provider to aggregate prefixes and advertise a single prefix for all of its subscribers
- Geographic
  - Provider-based scheme creates a challenge for a subscriber to change its provider
  - Ongoing research using other schemes, such as, geographic addressing

IPv6 Autoconfiguration

- IPv4 autoconfiguration: via DHCP
- IPv6 autoconfiguration: can be done via “stateless” autoconfiguration, no server needed
  - How?
    - Obtain an interface ID that is unique on the link to which the host is attached
    - Obtain the correct address prefix for this subnet
    - Prefix + a number of 0’s + interface ID

IPv6 Autoconfiguration: Example

- Creating a link-local IPv6 address from Ethernet address
  - Obtain an interface ID that is unique on the link to which the host is attached
    - Ethernet MAC address → IEEE EUI-64

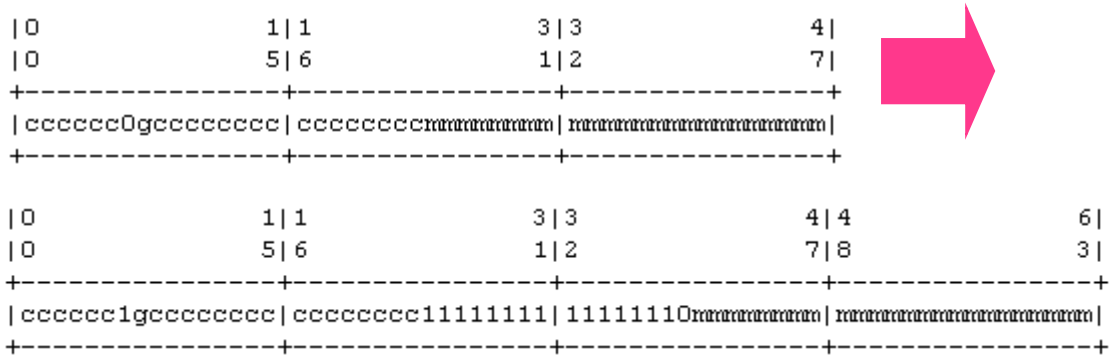

Example: an Ethernet adapter has hardware address: 00:1E:C9:2E:F4:6D → 02:1E:C9:FF:FE:2E:F4:6D

  - Obtain the correct address prefix for this subject
    - e.g., Link-local prefix = FE80::/64
  - Prefix + a number of 0's + interface ID

    FE80::021E:C9FF:FE2E:F46D/64

IPv6 Packet Header

- Header
  - 40-byte “base” header
- Extension headers
  - fixed order, mostly fixed length
    - hop-by-hop options
    - destination options
    - routing header
    - fragment header
    - authentication header
    - encapsulating security payload header
    - other options

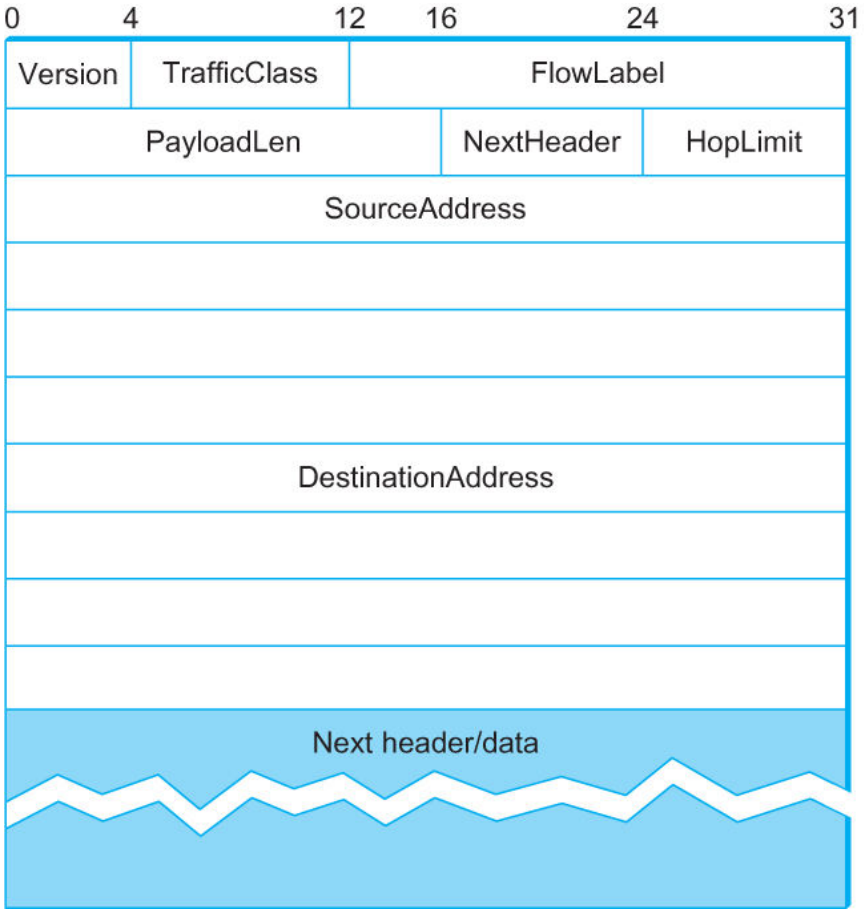

IPv6 "Next Header" Example

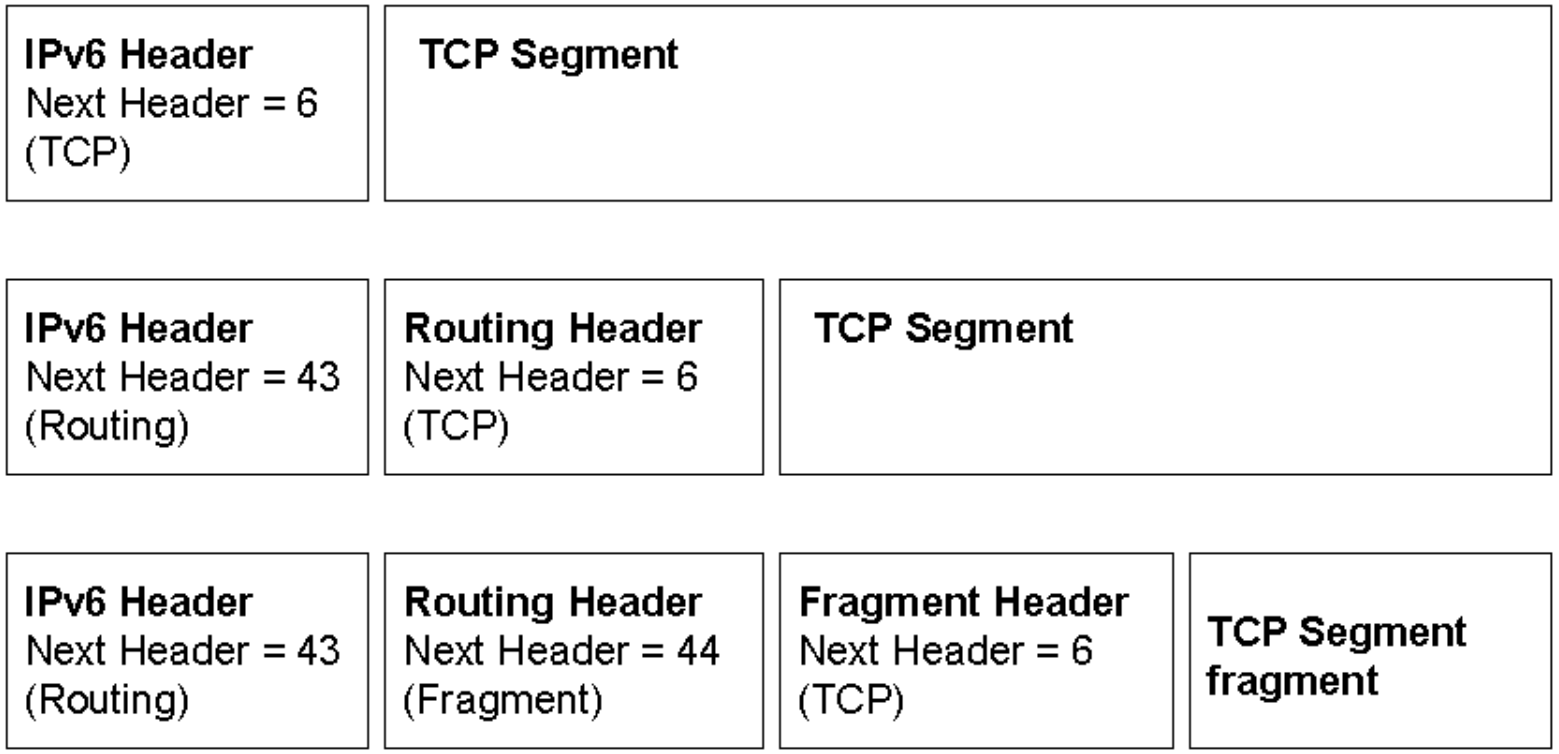

(From: http://www.gdt.id.au/~gdt/presentations/2007-10-31-aarnet-ipv6/)

IPv6 Advanced Routing

- IPv6 Routing header
  - Without the header, IPv6 differs little from IPv4 CIDR
  - It defines nodes or topological areas that the packet should visit en route to its definition
    - A topological area could be a back-bone provider's network
    - Supports provider selection on a packet-by-packet basis
      - e.g., select cheap one, select one that meets certain security requirement
    - the packet sent to the anycast address will go to the "nearest" one of the interfaces
      - Support also mobile IP

Anycast Address

- Assigned to a set of interfaces
- Allocated from the unicast address space, syntactically indistinguishable from unicast addresses
  - Assigning a unicast address to more than one interface makes a unicast address an anycast address
- Nodes to which the anycast address is assigned must be explicitly configured to recognize that the address is an anycast address.
- Anycast addresses can be used only by a device, not a host,
- Anycast addresses must not be used as the source address of an IPv6 packet

Internet Multicast

- IPv4
  - class D addresses
  - demonstrated with MBone
  - uses tunneling
- Integral part of IPv6
  - problem is making it scale

Multicast: Two Types

- One-to-many
  - Radio station broadcast
  - Transmitting news, stock-price
  - Software updates to multiple hosts
- Many-to-many
  - Multimedia teleconferencing
  - Online multi-player games
  - Distributed simulations

What If No Multicast?

- Without support for multicast for many-to-many and one-to-many communications
  - A source needs to send a separate packet with the identical data to each member of the group
    - This redundancy consumes more bandwidth
    - Redundant traffic is not evenly distributed, concentrated near the sending host
  - Source needs to keep track of the IP address of each member in the group
    - Group may be dynamic

Introducing IP Multicast

- To support *efficient* many-to-many and one-to-many communications, IP provides an IP-level multicast
- Basic IP multicast model is many-to-many based on multicast groups
  - Each group has its own IP multicast address
  - Hosts that are members of a group receive copies of any packets sent to that group's multicast address
  - A host can be in multiple groups
  - A host can join and leave groups

IP Multicast: Benefit

- Using IP multicast to send the identical packet to each member of the group
  - A host sends a single copy of the packet addressed to the group's multicast address
  - The sending host does not need to know the individual unicast IP address of each member
  - Sending host does not send multiple copies of the packet

One-to-Many Multicast

- Many-to-many multicast
  - Any source multicast (ASM)
  - Any hosts can send to a multicast group identified by its multicast address
  - The hosts do not have to be in the multicast group
  - There may be any number of such senders to a given group
- One-to-many multicast
  - Source specific multicast (SSM)
  - A receiving host specifies both a multicast group and a specific sending host

Forming Multicast Group

- A host signals its desire to join or leave a multicast group by communicating with its local router using a special protocol
  - In IPv4, the protocol is Internet Group Management Protocol (IGMP)
  - In IPv6, the protocol is Multicast Listener Discovery (MLD)
- The router has the responsibility for making multicast behave correctly with regard to the host
  - Router periodically polls the LAN

Multicast Addresses

IPv4 Address Type	Binary Prefix	IPv4 CIDR Notation
Multicast	1110	E0.00.00.00/4

IPv6 Address Type	Binary Prefix	IPv6 Notation
Multicast	1111 1111	FF00::/8

Hardware Multicasting

- Internet authorities allocates Ethernet and Fiber Distributed Data Interface (FDDI) media access control (MAC) addresses
  - 01-00-5E-00-00-00 to 01-00-5E-7F-FF-FF
  - 23 bits
- For IPv4
  - 28 bits of possible multicast address
  - Problematic to work with hardware multicasting (e.g., Ethernet or FDDI)
    - Mapping 28 bits to 23 bits is many on one (2^5 = 32 IP address per Ethernet address)
    - Hosts or switches must examine destination IPv4 address in a packet to see if it is really for them.

Unicast vs. Multicast Forwarding

- Unicast forwarding tables
  - For any IP address, which link to use to forward the unicast packet
  - Collectively specify a set of paths
- How about multicast forwarding tables?

Multicast Routing

- Multicast forwarding tables
  - Based on multicast address, which links to use to forward the multicast packet
  - Collectively specify a set of trees
    - Multicast distribution trees
- To support source specific multicast
  - Multicast forwarding tables must indicate which links to use based on
    - the combination of multicast address
    - the unicast IP address of the source

Multicast Routing

- Multicast routing is the process by which multicast distribution trees are determined
- Many approaches to multicast routing
  - Distance Vector Multicast Routing Protocol (DVMRP)
  - Protocol Independent Multicast and variants (PIM-SM)
  - Multicast Source Discovery Protocol (MSDP)
  - Source-Specific Multicast (PIM-SSM)
- Difficult problem space in which to find optimal solutions
  - Consider bandwidth usage, router state, path length, etc.

Distance Vector Multicast (1)

- Distance Vector Multicast Routing Protocol
  - Extension to Distance Vector Routing
- From unicast Distance Vector Routing table
  - Each router already knows that shortest path to source *S* goes through router *N*
  - When receive multicast packet from S, forward on all outgoing links (except the one on which the packet arrived), if and only if packet arrived from N.

Distance Vector Multicast (2)

- Eliminate duplicate broadcast packets by only letting
  - "parent" for LAN (relative to S) forward
    - shortest path to S (learn via distance vector)
    - smallest address to break ties

Distance Vector Multicast (3)

- Reverse Path Broadcast (RPB)
  - Goal: Prune networks that have no hosts in group G
  - Step 1: Determine of LAN is a *leaf* with no members in G
    - leaf if parent is only router on the LAN
    - determine if any hosts are members of G using IGMP
  - Step 2: Propagate "no members of G here" information
    - augment **<Destination, Cost>** update sent to neighbors with set of groups for which this network is interested in receiving multicast packets.
    - only happens when multicast address becomes active.

Protocol Independent Multicast (PIM)

Shared Tree

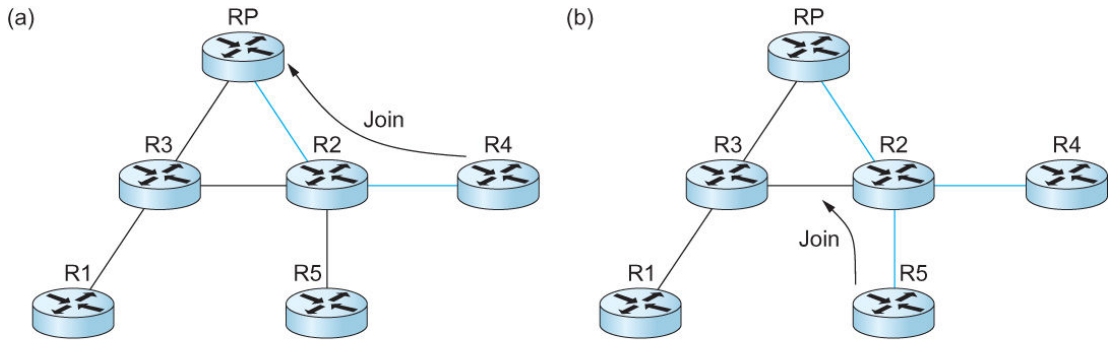

Source specific tree

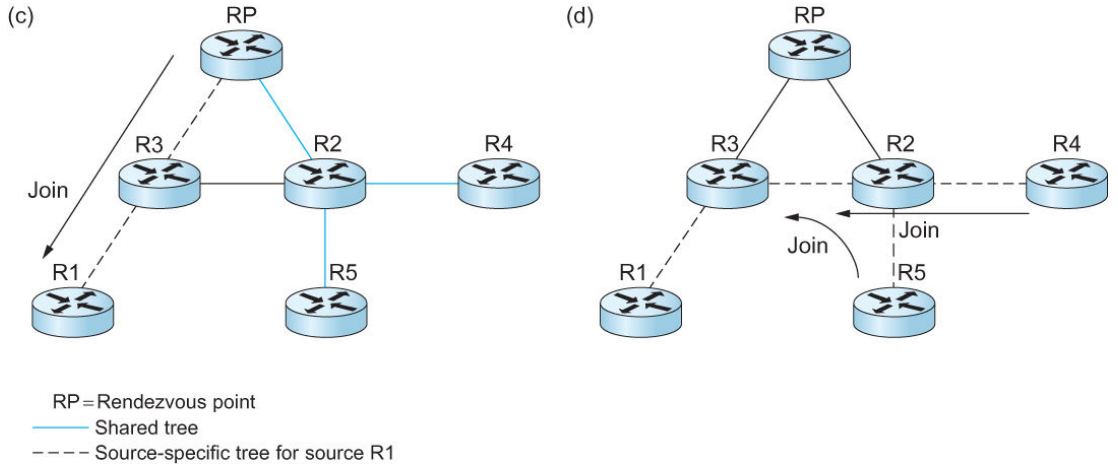

Protocol Independent Multicast (PIM)

- Delivery of a packet along a shared tree. R1 tunnels the packet to the RP, which forwards it along the shared tree to R4 and R5.

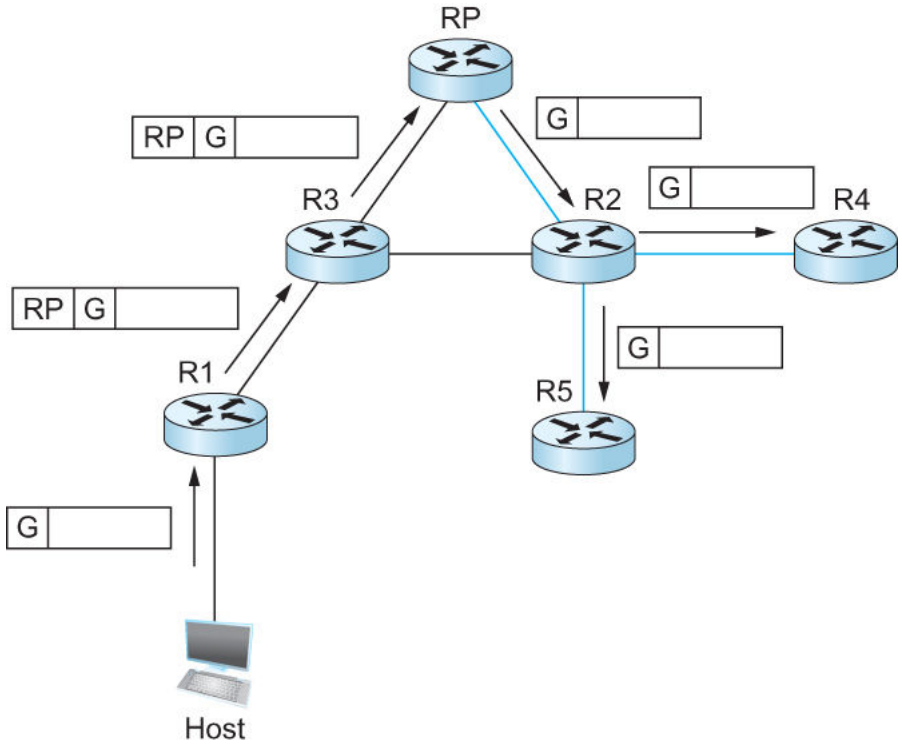

Inter-domain Multicast

- Multicast Source Discovery Protocol (MSDP)

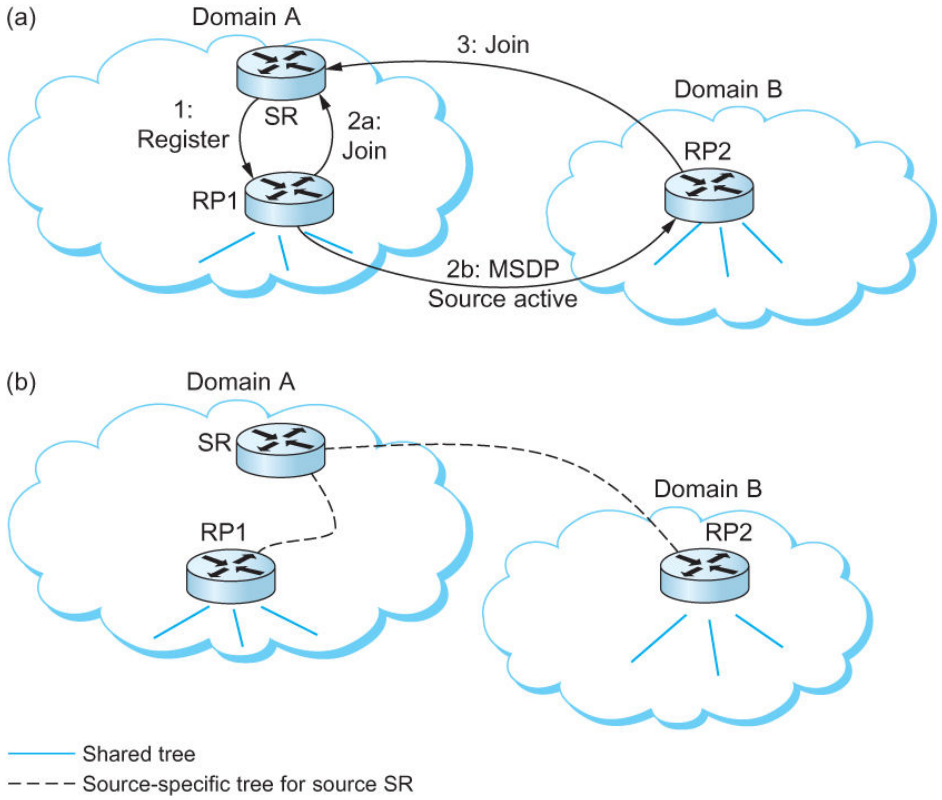

Routing for Mobile Hosts

- Mobile IP
  - *home agent*
    - Router located on the home network of the mobile hosts
  - *home address*
    - The permanent IP address of the mobile host.
    - Has a network number equal to that of the home network and thus of the home agent
  - *foreign agent*
    - Router located on a network to which the mobile node attaches itself when it is away from its home network

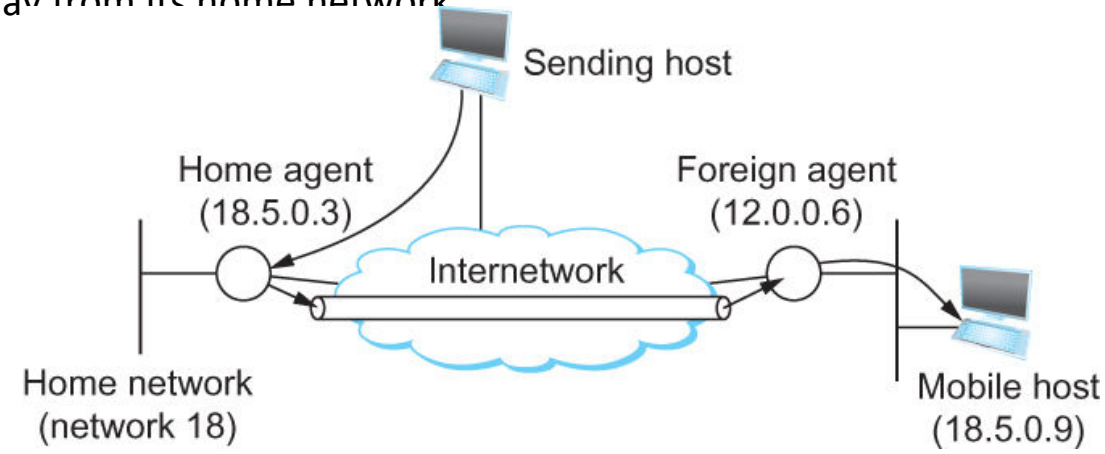

Routing for Mobile Hosts

- Problem of delivering a packet to the mobile node
  - How does the home agent intercept a packet that is destined for the mobile node?
    - Proxy ARP
  - How does the home agent then deliver the packet to the foreign agent?
    - IP tunnel
    - Care-of-address
  - How does the foreign agent deliver the packet to the mobile node?

Route optimization in Mobile IP

- The route from the sending node to mobile node can be significantly sub-optimal
- One extreme example
  - The mobile node and the sending node are on the same network, but the home network for the mobile node is on the far side of the Internet
    - Triangle Routing Problem
- Solution
  - Let the sending node know the care-of-address of the mobile node. The sending node can create its own tunnel to the foreign agent
  - Home agent sends binding update message
  - The sending node creates an entry in the binding cache
  - The binding cache may become out-of-date
    - The mobile node moved to a different network
    - Foreign agent sends a binding warning message

Mobility in IPv6

- Mobility support built into IPv6 from beginning
- Include necessary capabilities to act as a foreign agent in every IPv6 node
  - IPv6 does aware with the foreign agent
- Use extension headers to optimize routing
  - IPv6 includes a flexible set of extension headers
  - IPv6 node can send an IP packet to the care-of address with the home address contained in a routing header
    - No tunneling is necessary. More efficient

Development in Mobile IP

- Many challenging problems
  - Power consumption of mobile devices
  - Ad hoc mobile networks
  - Sensor networks
  - Wireless communications and mobility

Transition to IPv6

- Completely impossible to have a “flag day” to switch to IPv6
- IPv4 and IPv6 will coexist from many years to come
- Type of Nodes
  - IPv4-only node
  - IPv6-only node
  - IPv6/IPv4 node
  - IPv6 nodes include both IPv6-only and IPv6/IPv4 nodes
  - IPv4 nodes include both IPv4-only and IPv6/IPv4 nodes

Transition Mechanisms

- Dual IP Layer Operation
- Tunneling Mechanisms
- Further reading
  - http://www.rfc-editor.org/rfc/rfc4213.txt

Dual IP Layer Operation

- Providing a complete IPv4 implementation on IPv6 nodes
  - IPv6/IPv4 nodes, have the ability to send and receive both IPv4 and IPv6 packets.
  - Directly interoperate with IPv4 nodes using IPv4 packets
  - Directly interoperate with IPv6 nodes using IPv6 packets.
- Address configuration
  - IPv4: e.g., DHCP
  - IPv6: autoconfiguration or DHCPv6

Domain Name Resolution

- Introduced “AAAA” type resource records
- Extending DNS resolver libraries to handle BOTH “AAAA” and “A” types of resource records
- Applications SHOULD be able to specify whether they want IPv4, IPv6, or both records

Example of "AAAA" Records

```
$ dig aaaa www.google.com  +nocmd +noques

; <<>> DiG 9.9.5-3ubuntu0.5-Ubuntu <<>> aaaa www.google.com +nocmd +noques
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 57830
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; MBZ: 0005 , udp: 4000
;; ANSWER SECTION:
www.google.com.         5       IN      AAAA    2607:f8b0:4004:807::1011

;; Query time: 4 msec
;; SERVER: 192.168.101.2#53(192.168.101.2)
;; WHEN: Wed Nov 18 14:06:49 EST 2015
;; MSG SIZE  rcvd: 71
```

IPv6 to IPv4 Tunnels

- Configure tunnels to use existing IPv4 routing infrastructure to carry IPv6 traffic
- Tunneling can be used in a variety of ways:
  - Router-to-Router
  - Host-to-Router
  - Host-to-Host
  - Router-to-Host

Tunnel Mechanisms

- Entry node of the tunnel (the encapsulator)
  - creates an encapsulating IPv4 header
  - transmits the encapsulated packet.
  - may need to maintain soft-state information
    - e.g., MTU of the tunnel
- Exit node of the tunnel (the decapsulator)
  - receives the encapsulated packet
  - reassembles the packet if needed
  - removes the IPv4 header, and processes the received IPv6 packet.

Encapsulation

- Two Issues
  - Determine when to fragment and when to report an ICMPv6 "packet too big" error back to the source.
  - How to reflect ICMPv4 errors from routers along the tunnel path back to the source as ICMPv6 errors

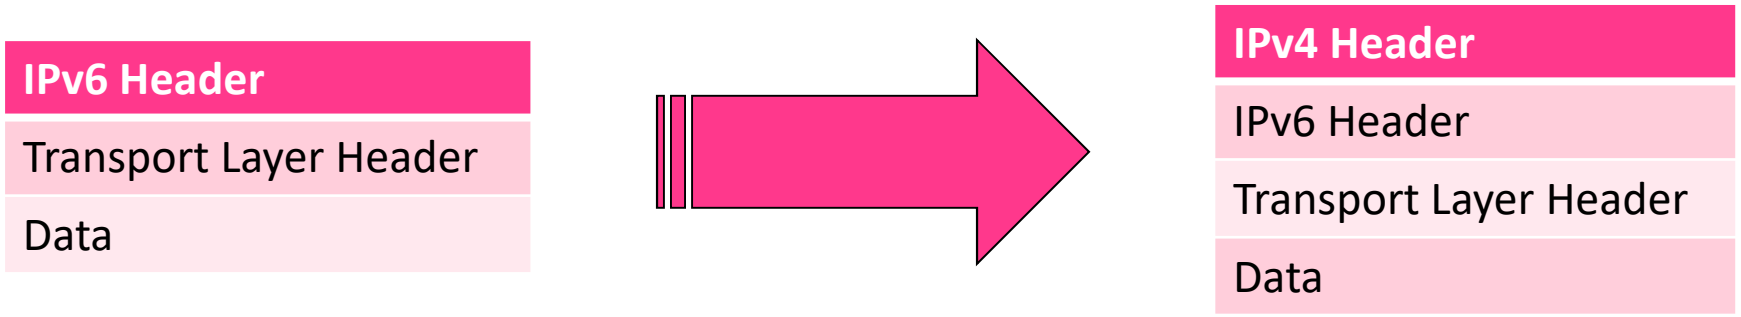

Tunnel MTU and Fragmentation

- Use the fixed static MTU or OPTIONAL dynamic MTU determination based on the IPv4 path MTU to the tunnel endpoint
  - IPv4 layer fragmentation should be avoided due to performance problems
  - Any IPv4 fragmentation occurring inside the tunnel should also be avoided due to performance and memory problems.
  - Encapsulator has no way of knowing that the decapsulator is able to defragment IPv4 packets

Static Tunnel MTU

- By default, the MTU MUST be between 1280 and 1480 bytes (inclusive), but it SHOULD be 1280 bytes
- When using the static tunnel MTU, the Don't Fragment bit MUST NOT be set in the encapsulating IPv4 header

Dynamic Tunnel MTU

- Using IPv4 Path MTU Discovery Protocol

```
if (IPv4 path MTU - 20) is less than 1280
   if packet is larger than 1280 bytes
        Send ICMPv6 "packet too big" with MTU = 1280. Drop packet.
    else
        Encapsulate but do not set the Don't Fragment flag in the IPv4
        header. The resulting IPv4 packet might be fragmented by the IPv4
        layer on the encapsulator or by some router along the IPv4 path.
    endif
else
     if packet is larger than (IPv4 path MTU - 20)
        Send ICMPv6 "packet too big" with MTU = (IPv4 path MTU - 20). Drop
packet.
     else
        Encapsulate and set the Don't Fragment flag in the IPv4 header.
     endif
endif
```

Other Issues

- Hop limits
- Handling ICMPv4 errors
- Decapsulation
- Link-local addresses
- Neighbor discovery over tunnels
- Threat related to source address spoofing and security considerations

Tunneling Implementation

- Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)
  - Windows, Linux, Cisco IOS
- IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)
  - Linux, Cisco IOS
- Connection of IPv6 Domains via IPv4 Clouds (6to4)
  - Linux, Cisco IOS
  - Recommended to be deprecated.
- Teredo Extensions
  - Windows

Summary

- IPv6 Addressing
- IPv6 Routing
- Multicasting
- Mobile IP
- IPv6 and IPv4 coexistence