COMP 790-088
Networked and Distributed Systems

IP Addressing

Jasleen Kaur

September 14, 2009

IP Addressing
How to Deal with Heterogeneity and Scale?

- Requirements:
  - » Should be globally unique
  - » Should facilitate easy mapping to link-layer addresses
  - » Should facilitate scalable assignment
  - » Should facilitate scalable routing
- Basic Idea:
  - » Assign addresses to interfaces (and not to hosts)
    - ❖ Facilitates meaningful mapping to MAC addresses
  - » Use hierarchical addressing
    - ❖ Aids both address assignment and routing

IP Addressing

Hierarchical Addressing

- Network+Host based address classes:
  - » Class A: *10.3.2.4*
    - ❖ 126 Class A networks
  - » Class B: *128.96.33.81*
    - ❖ 2^{14} Class B networks
  - » Class C: *192.12.69.77*
    - ❖ Each network can have 254 hosts
  - » Class D: multicast
  - » Class E: unused

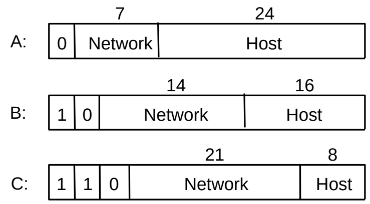

Datagram Forwarding

How Does Hierarchy Help Scale Packet Forwarding?

- Route to networks, not destinations!

If connected to destination (same network-number as an interface)
    Directly deliver over physical network on that interface
Else find next-hop for this network-number in forwarding table
    Send to next-hop over physical network
Else   Send to default router

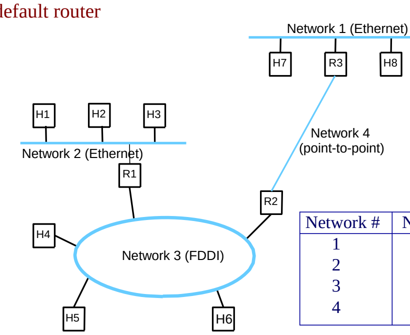

Network #	Next Hop
1	R3
2	R1
3	interface 1
4	interface 0

Address Resolution

How to Deal with Heterogeneity?

- Different layer-2 protocols may use different addressing schemes
  - » How does the "deliver to a host on same physical network" happen?
  - » IP knows only the layer-3 address of the immediate destination
    - ❖ How does link-layer know which layer-2 address to deliver the packet to?
- Solution: translate/map IP addresses to MAC addresses
- Address Resolution Protocol (ARP):
  - » Exploits the fact that many link-layer technologies support broadcast
  - » Broadcasts a query "who has this IP address?"
  - » Target node responds back with link-layer address
    - ❖ Enough bits reserved in ARP messages to accommodate maximum possible link-layer address size (48 bits)

Address Assignment

How Does Hierarchical Addressing Help Scalability?

- Class A addresses typically assigned to large ISPs
  - » Customers assigned addresses from this larger pool
- Class B addresses assigned to medium-sized ISPs/organizations
  - » Customers/departments assigned addresses from this larger pool
- Class C addresses assigned to smaller organizations
  - » Individuals assigned from this pool
- DHCP helps recycle a smaller pool among larger number of hosts
  - » Machine new on the network can obtain an IP address from the DHCP server
  - » Self-configuring: can discover DHCP server using link-layer broadcast

Hierarchical addressing and DHCP help achieve scalability of network management !

IP Addressing
Summary

- How does IP accommodate heterogeneity?
  - By offering a simple service model
    - That any link-layer can help support
  - By having a common packet format
    - Adding a layer that everyone understands
  - By having a global address space
    - That can be mapped to layer-2 addresses
- How does it achieve scalability?
  - Routing:
    - By making forwarding decisions based on hierarchical addressing
  - Manageability:
    - By delegating the job of address assignment to administrative entities

Scaling of Address Space
Exhaustion of IPv4 Addresses

- IP addresses likely to run out
  - Can address at most 4 billion hosts
  - Allows only 16,000 class B addresses
  - Inefficient address assignment due to class-based addressing
    - A 2-node network will use a class C address (0.78% efficient)
    - A 257-node network will use a class B address (0.39% efficient)
- Solutions:
  - DHCP, NAT
  - Sub-netting
  - Super-netting
  - IPv6

Subnetting

Introducing Another Level of Hierarchy

- How about assigning many class C addresses per organization?
  - » ☹: increases forwarding table size in routers!

- Subnets add another level of hierarchy in IP addresses

  Subnet number = network part + subnet part

  IP address = subnet number + host part

  - » Nodes on same physical network have same subnet number
    - ❖ Subnet masks applied to derive subnet number
      - ◆ Define variable portion of host part
  - » Different physical networks may share same network part
    - ❖ Subnets visible only within a site

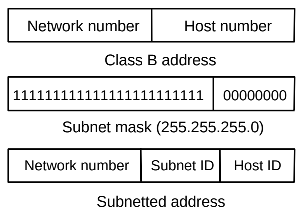

Subnet Masks

An Example

- Each host is configured with: IP address, Subnet mask

  Subnet number = (IP address) & (Subnet mask)

  - » Routers store subnet mask, subnet number for each attached subnet

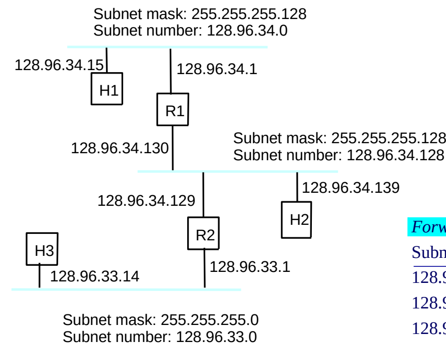

Forwarding table at router R1:

Subnet Number	Subnet Mask	Next Hop
128.96.34.0	255.255.255.128	interface 0
128.96.34.128	255.255.255.128	interface 1
128.96.33.0	255.255.255.0	R2

Packet Forwarding with Subnets

How are Subnet Masks Used?

- Nodes:

```
if destinationIP & selfSubnetMask = selfSubnetNumber,
  deliver packet directly over physical network
else
  send to router
```

- Routers:

```
D = destination IP address
for each entry (SubnetNum, SubnetMask, NextHop)
   S = SubnetMask & D
   if S = SubnetNum
      if NextHop is an interface
         deliver datagram directly to D
      else
         deliver datagram to NextHop
```


Subnetting

Summary

- What table entries do routers outside the organization keep?
  - » Only one for the entire network!
    - ❖ Only routers inside see subnets
    - ❖ Not everyone has same routing view / information
- How does subnetting help address scalability?
  - » Improves address assignment
    - ❖ Don't have to assign a Class B/C for every new physical network
  - » Improves forwarding efficiency for outside routers
    - ❖ By aggregating routing information
- Multiple subnets can also be defined on same physical network
  - » Forced to talk through router
  - » Helps provide isolation between different departments sharing a LAN

Supernetting
Why Do We Need More?

- Subnetting does not solve:
  - » Exhaustion problem of Class B addresses
    - Any organization with more than 255 nodes would still need a Class B address
  - » Growth in routing table sizes
    - State in forwarding tables still grows in proportion to the number of networks
- Supernetting:
  - » Goal:
    - Help scale assignment efficiency as well as routing table sizes
  - » Approach:
    - Aggregate routes !

Supernetting
Basic Idea

- Breaks the rigid boundaries between address classes
  - » Address space efficiency:
    - Prefixes can be of arbitrary length (and not just 8, 16, or 24)
  - » Forwarding table control:
    - Allow routing state aggregation at several levels
      - And not just at subnets-within-a-network level
    - Assign contiguous network numbers (Class C addresses) to nearby networks
      - eg: consider 16 "nearby" networks: 192.2.16 – 192.4.31
        - Top 20 bits are common in these prefixes
        - Create a 20-bit network prefix to represent all of thse
        - Use a single routing table entry at routers "far away"
      - How?
        - Assign regional address space to service providers, which distribute among regional customers

Supernetting also known as Classless Interdomain Routing (CIDR)

Supernetting vs. Subnetting

Philosophical Difference

- Subnetting
  - » Shares one address among multiple physical networks
  - » Prevents wastage of Class C addresses
    - ❖ By re-using same network address among several smaller networks
- CIDR:
  - » Collapses multiple addresses that would be assigned to a single AS onto one address
  - » Prevents wastage of Class B addresses
    - ❖ By assigning multiple Class C addresses instead

Forwarding Algo

Use of Network Prefixes

- Same as before, but prefix-length included in order to identify network number
  - » Network number is not class-based, but an aggregated prefix
  - » eg: 192.4.16/20
- Caveat:
  - » Some networks can interfere with the process of aggregation
    - ❖ eg: if a customer changes service providers (but retains IP addresses)
      - Has to be routed through new service provider
      - Has different routing treatment than "nearby" networks
  - » Therefore, need to store "exceptions" in addition to aggregated routes
    - ❖ Eg: 2 routing table entries: 208.12.16/20 and 208.12.21/24
  - » Router needs to find the ***longest matching prefix*** when forwarding packets

Finding the longest match makes forwarding a complicated task !

Exceptions: A Headache

Impact on Routing Tables

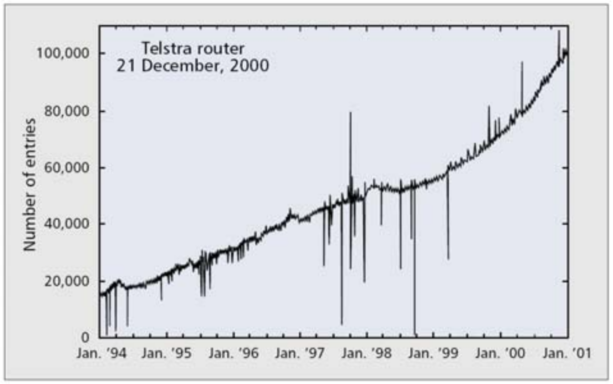

- Exponential growth before CIDR
- Linear growth till exceptions take over

Exceptions: A Headache

Impact on Routing Tables

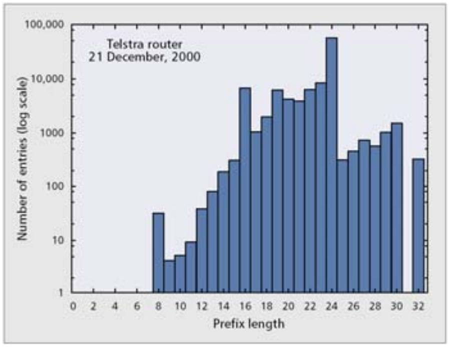

- Smaller prefixes would imply greater aggregation
- Class C addresses still dominate routing tables!
- Can geographic addressing help reduce exceptions?

IPv6
128-bit Address Space

- IP version 6 uses 128-bit address spaces
  - » Enough to assign 1500 addresses per square foot of earth's surface
- Hosts and routers need to understand new packet format
  - » A feasible transition plan a must!
- Facilitating transitions:
  - » Dual-stack operation:
    - ❖ IPv6 nodes that can process both IPv4 and IPv6 packets
  - » Tunneling:
    - ❖ Tunnels used to send an IPv6 packet over an IPv4-only network