



The Internet Protocol Part 3: Addressing

Jeffrey Shafer

IP Addresses

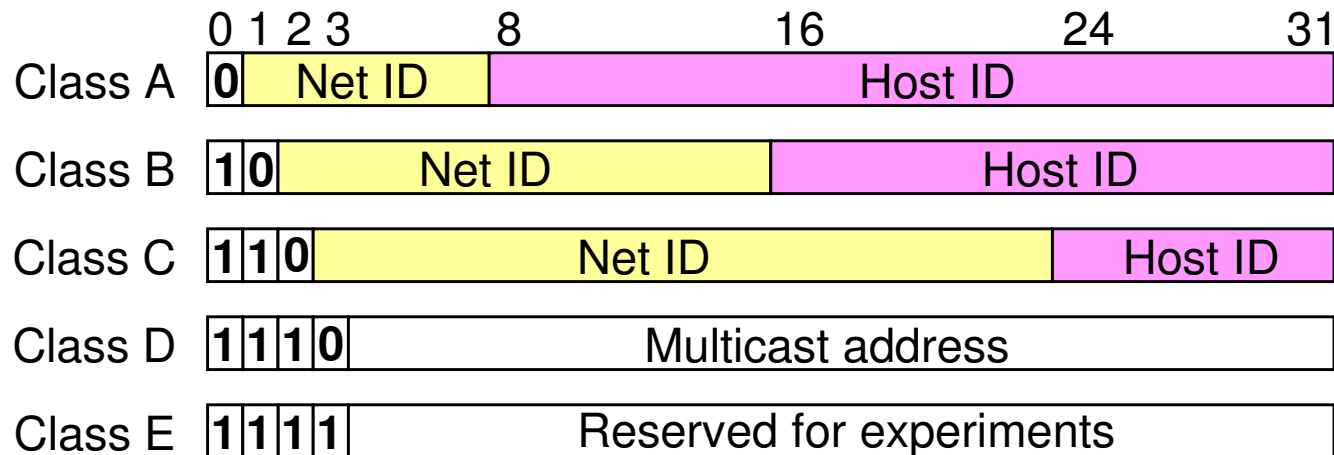
- IP version 4 addresses are 32-bits
 - Version 6 address are 128 bits
- Every network interface has at least one IP address
 - A computer might have 2 or more IP addresses
 - A router has many IP addresses
 - These addresses can be assigned statically or dynamically
- IP addresses are always in big-endian byte order (network byte order)
 - True in general for any integer transferred in a packet header from one machine to another
 - E.g., the port number used to identify a TCP connection

IP Address Format

- IPv4 addresses are usually displayed in dotted decimal notation
 - Each byte represented by decimal value
 - Bytes are separated by a period
 - IP address $0x8002C2F2 = 128.2.194.242$
- IP addresses are hierarchical
 - Address is composed of a network ID and a host ID
 - `www.rice.edu: 128.42.206.11`

IP Address Structure

- IPv4 Address space divided into classes:



- Special IP addresses

- Loop-back address: 127.0.0.1
- Unrouted (private) IP addresses:
 - 10.0.0.0 – 10.255.255.255
 - 172.16.0.0 – 172.31.255.255
 - 192.168.0.0 – 192.168.255.255

IP Address Examples

- www.apple.com: 17.251.200.32
 - Class A: $0 < 17 < 128$
- www.mit.edu: 18.7.22.83
 - Class A: $0 < 18 < 128$
- www.rice.edu: 128.42.206.11
 - Class B: $128 \leq 128 < 192$ (128+64)



Routing

- Routers in the Internet should only need to know about network IDs
- Routers inside an organization can route based on host IDs
- Problems?

Problems

- Address classes were too “rigid”
 - Class C is too small and Class B is too big in many situations
 - Inefficient use of address space
 - Leads to a shortage of addresses
- Small organizations wanted Class B networks
 - In case they grew to more than 255 hosts
 - But there are only about 16,000 Class B network IDs
- Larger organizations wanted many Class C networks
 - Separate network ID for each router link
- Every router in the Internet had to know about every network ID in every organization
 - Leads to large address tables in every router




Subnetting


- Divide the network within an organization
 - Basically consider one Class B network to be a collection of many smaller networks
 - Size of smaller networks can be selected by the organization (don't have to be Class C sized networks)
- Internet routers don't need to know about subnetting within an organization
 - Just route their traffic to the organization

Subnetting


Company's
Class B Network




Building 1 Network



Department 4 network in Building 1



Floor 13 network in Building 6



Subnetting

- Can recursively subnet addresses down to as fine a granularity as you want
 - Almost...
 - Minimum-sized subnet is 4 addresses (0 names the subnet, 3 is broadcast, 1-2 are hosts)
- Subnet sizes don't have to be the same
 - One building divided by department, one by floor
 - Department/floor subnets not the same size


Subnet Notation

- A.B.C.D/X
 - IP address of the subnet (with 0's in all host ID bits)
 - X = number of bits in the subnet network address
- Examples:
 - 17.0.0.0/8 – Apple's entire class A address space
 - 17.2.3.0/24 – A class C sized subnet in Apple's network
- Alternatively represented by subnet IP and a bit mask (netmask)
 - 17.0.0.0/255.0.0.0
 - 17.2.3.0/255.255.255.0

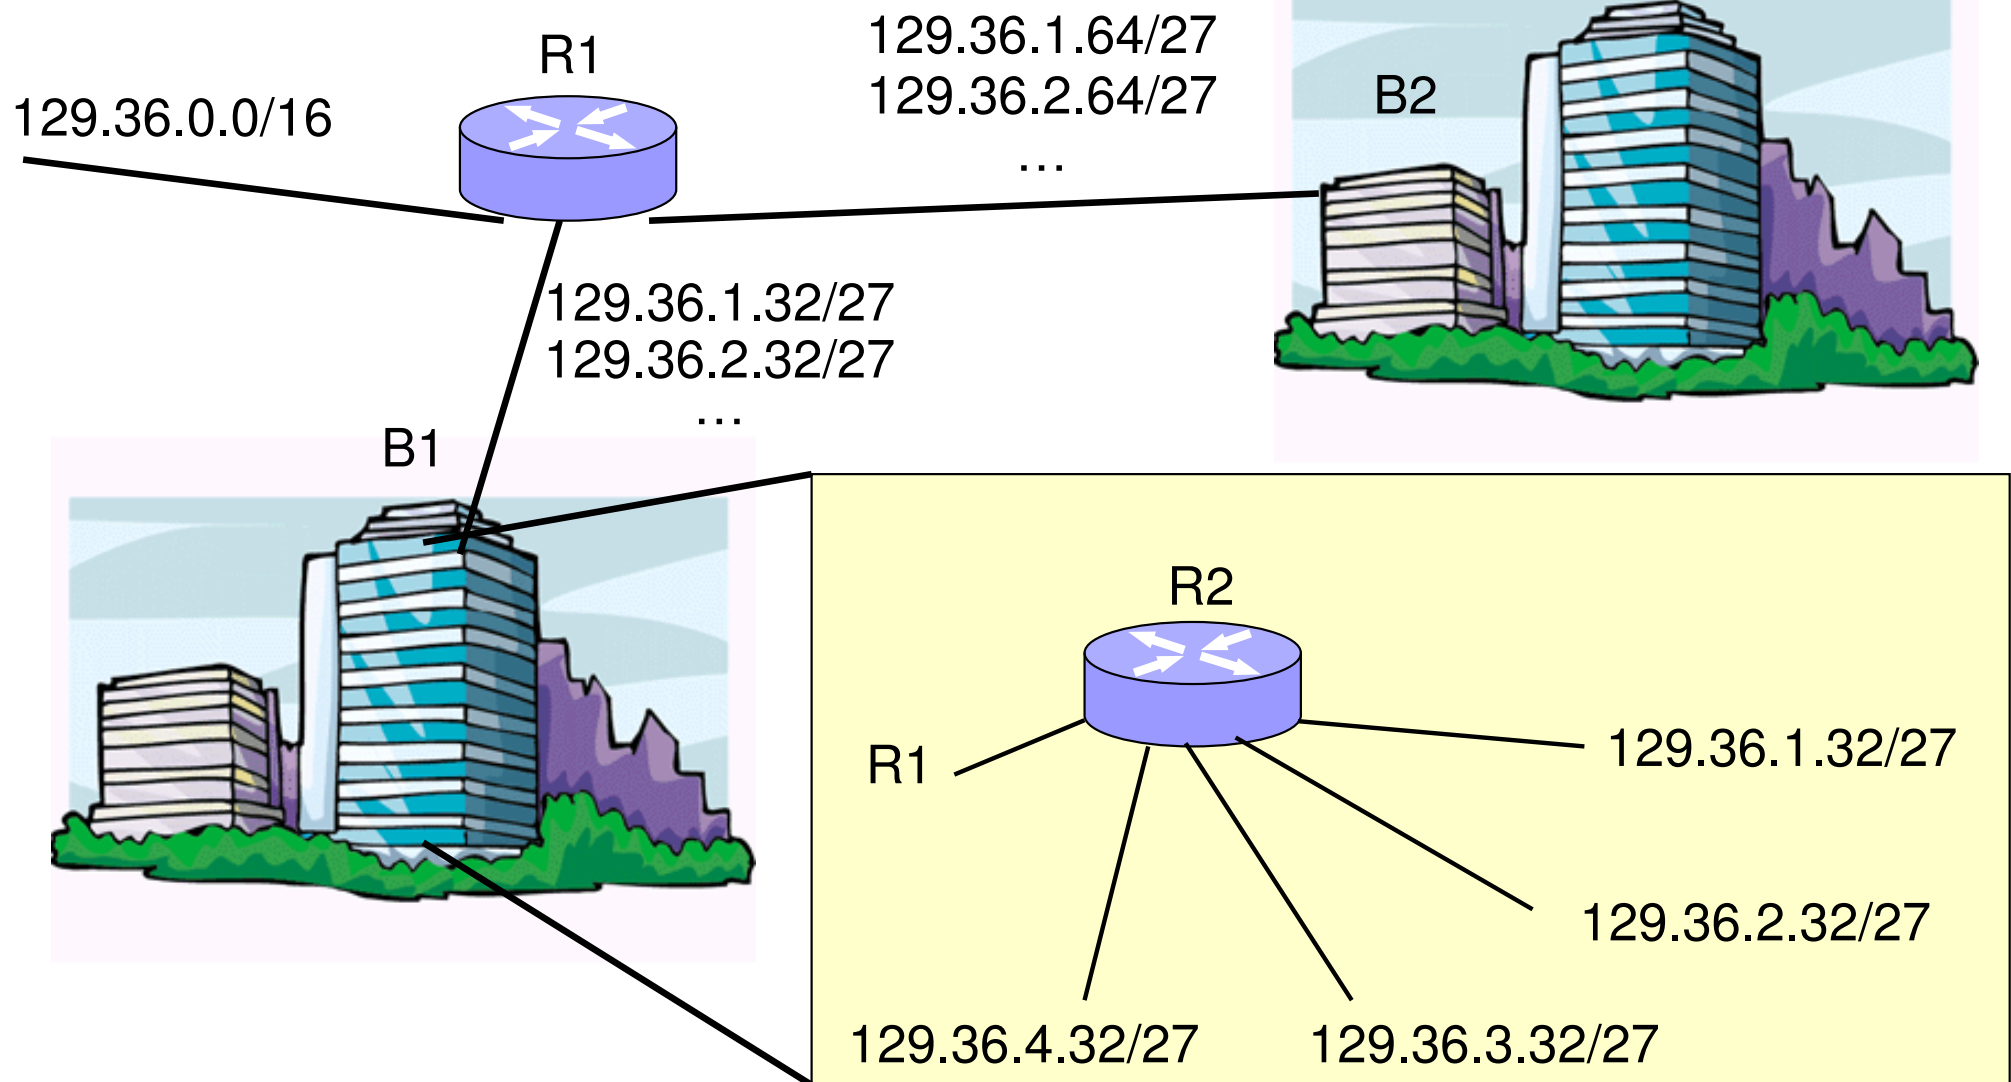
Subnet Meaning

- Subnets don't have to have physical meaning
 - Although easier to keep track of if they do...
- Good subnet assignment simplifies routing for internal routers
 - All traffic for “building 1” goes through this port
 - All traffic for “department 3” goes through that port
 - ...

Route Aggregation



Little/No Route Aggregation



Subnetting at Rice

- Rice uses private IP addresses
 - 10.0.0.0/8
- Affinity group is encoded in the second octet:
 - 10.A.0.0/16
- Building is encoded in the third octet:
 - 10.A.B.0/24
- Routers can aggregate routes for affinity groups, then buildings

Classless InterDomain Routing

- CIDR introduced in 1993
 - Meant to provide more flexible routing
 - Eliminate dependences on “class” networks in routing
- “Supernetting”
 - Combine multiple contiguous networks into one larger network
 - Effectively reduces the number of entries needed in each routing table
 - Inverse of subnetting which takes one larger network and breaks it into multiple contiguous smaller networks

CIDR Idea

- Break up IP address space into *prefixes*
 - Same idea as subnets (128.42/16)
- Each prefix has its own routing entry
 - All traffic to Rice (128.42/16) within the Internet should be routed the same way, regardless of how Rice subnets its address space


Route Aggregation

- Fundamental idea behind supernetting
 - Same type of aggregation as with subnets, but on a larger scale
- E.g., one ISP handles traffic for two corporate networks (129.32/16 and 129.33/16)
 - Aggregate route to 129.32/15 for both networks
 - Only break them apart when necessary for the last (few) hop(s)
- E.g., one ISP handles all 64/8 networks
 - All other routers know to route 64/8 traffic to that ISP
 - External routers don't care how the ISP breaks up the network addresses internally!
 - Internal ISP routers must route at a finer grain

What if there are holes?

- Rice builds a second campus
 - 128.42/16 needs to be routed to Houston
 - 128.42.8/24 needs to be routed to our satellite campus in Italy...
- Do we need to break routes up?
 - 128.42.0/20 (.0-.7)
 - **128.42.8/24 (.8)**
 - 128.42.9/24 (.9), 128.42.10/24 (.10), 128.42.11/24 (.11)
 - 128.42.12/22 (.12-.15)
 - 128.42.16/20 (.16-.31)
 - 128.42.32/19 (.32-.63)
 - 128.42.64/18 (.64-.127)
 - 128.42.128/17 (.128-.255)

IP Prefixes



- IP address space can be viewed as a number line
 - Each segment represents an aggregated route
 - Segments can overlap
- Look for smallest segment that matches the destination address

Longest Prefix Match

- Allow more specific entries to supersede more general ones
 - 128.42.8/24
 - Route this traffic to Italy
 - 128.42/16
 - Route this traffic to Houston
 - Except for addresses that match a route with a longer prefix (i.e., 128.42.8/24)
- Allows significantly more route aggregation
- Simplifies things if companies move (physically or to another ISP) their block of IP addresses

IP Address Classes

- CIDR makes address classes less important
- With CIDR, routing is based on arbitrary subdivisions of the address space
 - Aggregate routes into largest possible group
 - Use longer prefixes to deal with exceptions
- Routing
 - Routers use longest prefix matching to determine routes
 - No longer deal with exact matches on class network IDs

Implementing Software LPM

- Simple software solution
 - Keep list of prefixes and masks
 - Search for all prefix matches
 - Choose result with longest prefix
- Trick
 - Sort list based on decreasing prefix length
(stop search after first match)
- This is extremely inefficient for large forwarding tables!!!

Implementing Hardware LPM

■ Ternary CAM

- Use “don’t care” for bits outside the prefix
- All matching prefixes will yield a hit
- Pick the one with the longest prefix

■ Trick

- Sort CAM entries based on decreasing prefix length (use lowest matching entry)

■ This does not necessarily scale well to large forwarding tables!!!

ARIN WHOIS Database Search

Relevant Links: [ARIN Home Page](#) [ARIN Site Map](#) Training: [Querying ARIN's WHOIS](#)

Search ARIN WHOIS for: 128.42.0.0

```
OrgName: Rice University
OrgID: RICEUN
Address: Networking MS 119
Address: 6100 Main Street
City: Houston
StateProv: TX
PostalCode: 77005
Country: US

NetRange: 128.42.0.0 - 128.42.255.255
CIDR: 128.42.0.0/16
OriginAS: AS8
NetName: RICE-NET
NetHandle: NET-128-42-0-0-1
Parent: NET-128-0-0-0-0
NetType: Direct Assignment
NameServer: NS1.RICE.EDU
NameServer: NS2.RICE.EDU
NameServer: NS.PURDUE.EDU
Comment:
RegDate: 1983-12-02
Updated: 2007-11-07
```

```
RTechHandle: RUH-ORG-ARIN
RTechName: Rice University Networking
RTechPhone: +1-713-348-4989
RTechEmail: ipadmin@rice.edu
```

```
OrgTechHandle: RUH-ORG-ARIN
OrgTechName: Rice University Networking
OrgTechPhone: +1-713-348-4989
OrgTechEmail: ipadmin@rice.edu
```

```
# ARIN WHOIS database, last updated 2009-02-01 19:10
# Enter ? for additional hints on searching ARIN's WHOIS database.
```

Other WHOIS Servers: [Afrinic](#) [APNIC](#) [LACNIC](#) [RIPE](#) [InterNIC](#)

Request Bulk Copies of ARIN WHOIS Data

Copyright © 1997-2007 American Registry for Internet Numbers. All Rights Reserved.



Next Time

- ARP – Address Resolution Protocol