CIS 551 / TCOM 401

Computer and Network Security

Spring 2008
Lecture 12

Announcements

- Reminder:
  – Project 2 is due Friday, March 7th at 11:59 pm

Internet Protocol Interoperability

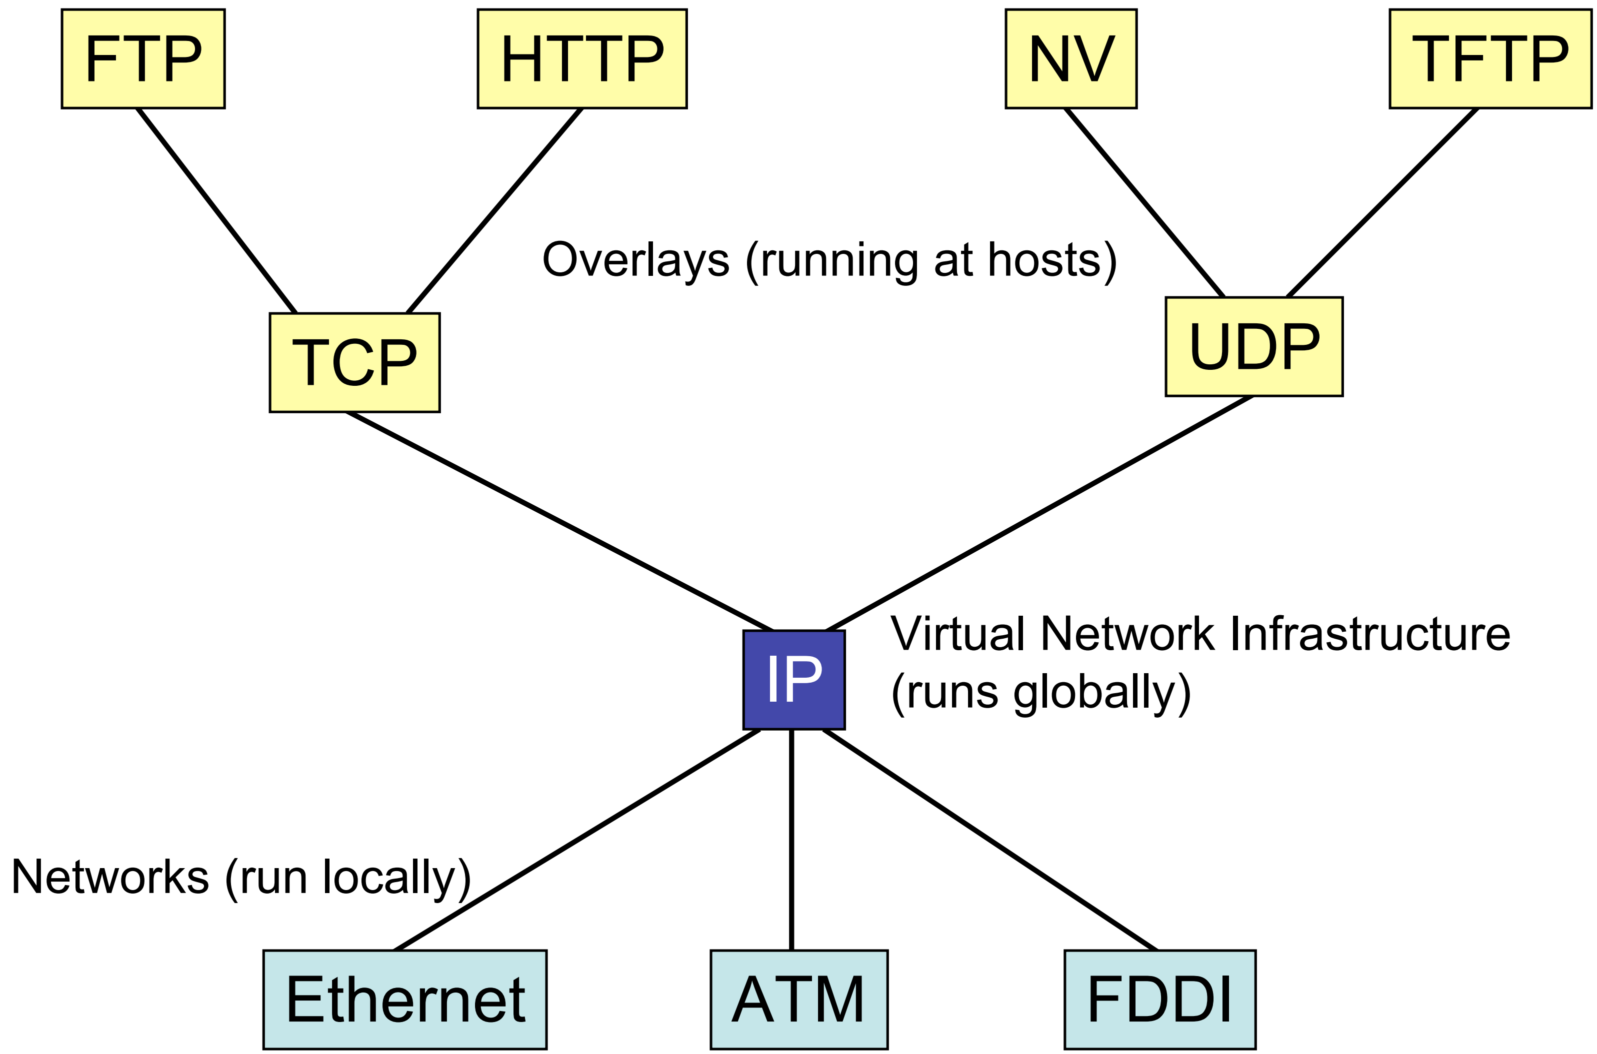

Internetworks

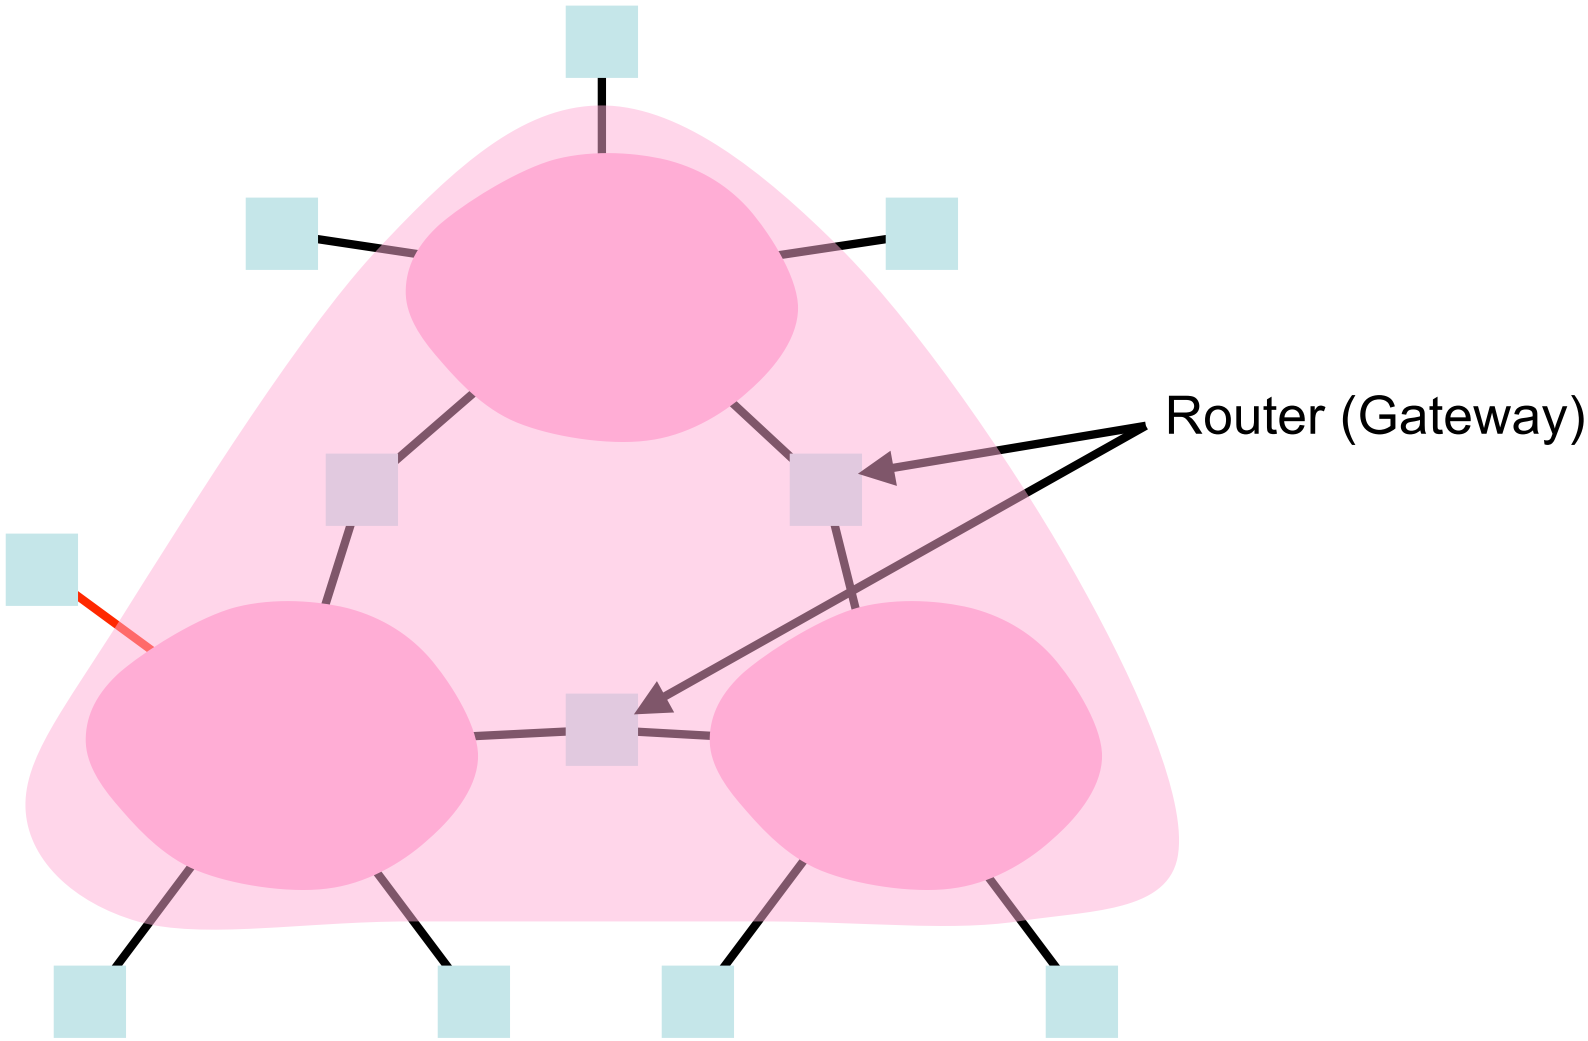

Internetworks

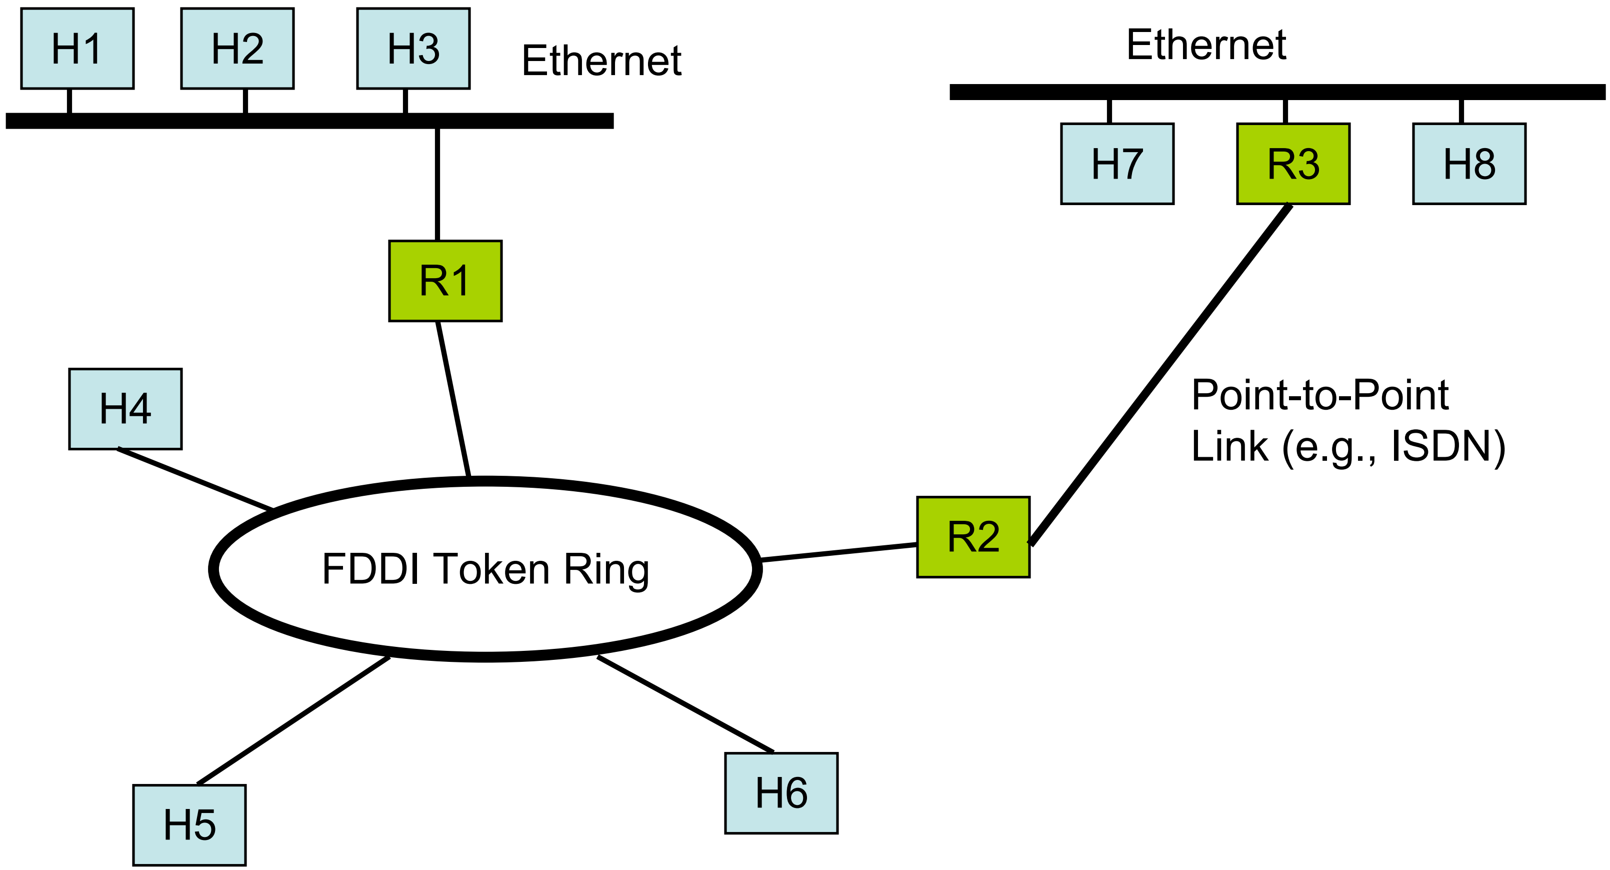

IP Encapsulation

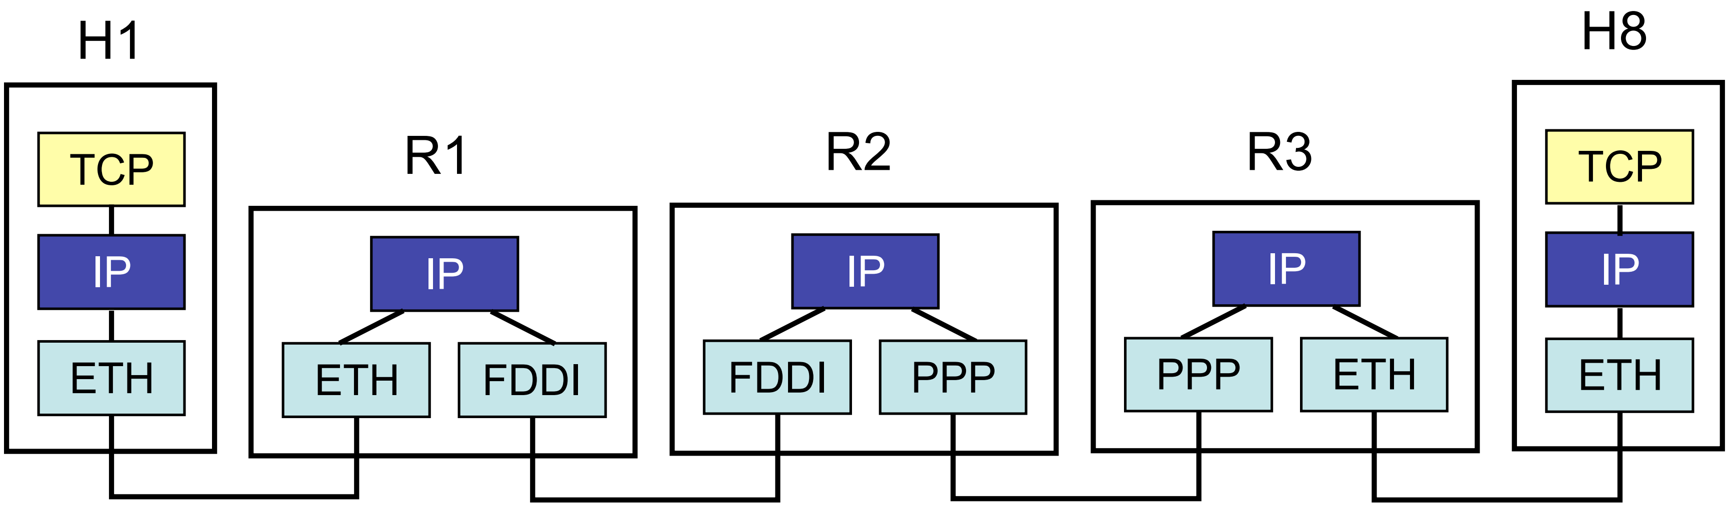

Example of protocol layers used to transmit from H1 to H8 in network shown on previous slide.

IP Service Model

- Choose minimal service model
  - All nets can implement
  - “Tin cans and a string” extremum
- Features:
  - Best-effort datagram delivery
  - Reliability, etc. as *overlays* (as in TCP/IP)
  - Packet format standardized

IPv4 Packet Format

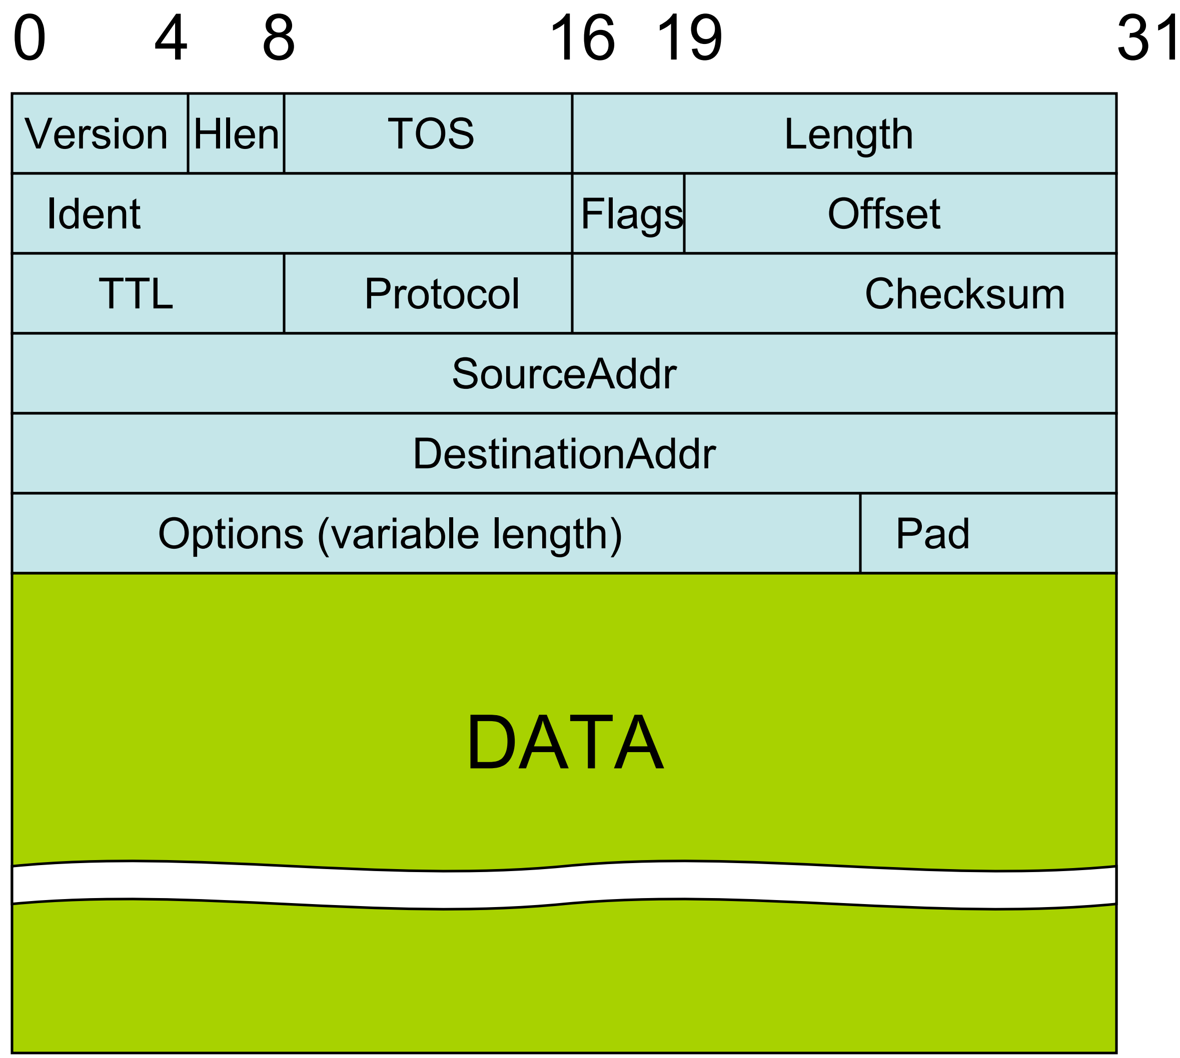

Fields of IPv4 Header

- Version
  - Version of IP, example header is IPv4
  - First field so easy to implement case statement
- Hlen
  - Header length, in 32-bit *words*
- TOS
  - Type of Service (rarely used)
  - Priorities, delay, throughput, reliability
- Length
  - Length of datagram, in *bytes*
  - 16 bits, hence max. of 65,536 bytes
- Fields for *fragmentation* and *reassembly*
  - Identifier
  - Flags
  - Offset

Header fields, continued

- TTL
  - Time to live (in reality, hop count)
  - 64 is the current default (128 also used)
- Protocol
  - e.g., TCP (6), UDP(17), etc.
- Checksum
  - Checksum of header (not CRC)
  - If header fails checksum, discard the whole packet
- SourceAddr, DestinationAddr
  - 32 bit IP addresses - global, IP-defined
- Options
  - length can be computed using Hlen

IP Datagram Delivery

- Every IP packet (datagram) contains the destination IP address
- The network part of the address uniquely identifies a single network that is part of the larger Internet.
- All hosts and routers that share the same network part of their address are connected to the same physical network.
- Routers can exchange packets on any network they're attached to.

IP addresses

- Hierarchical, not flat as in Ethernet

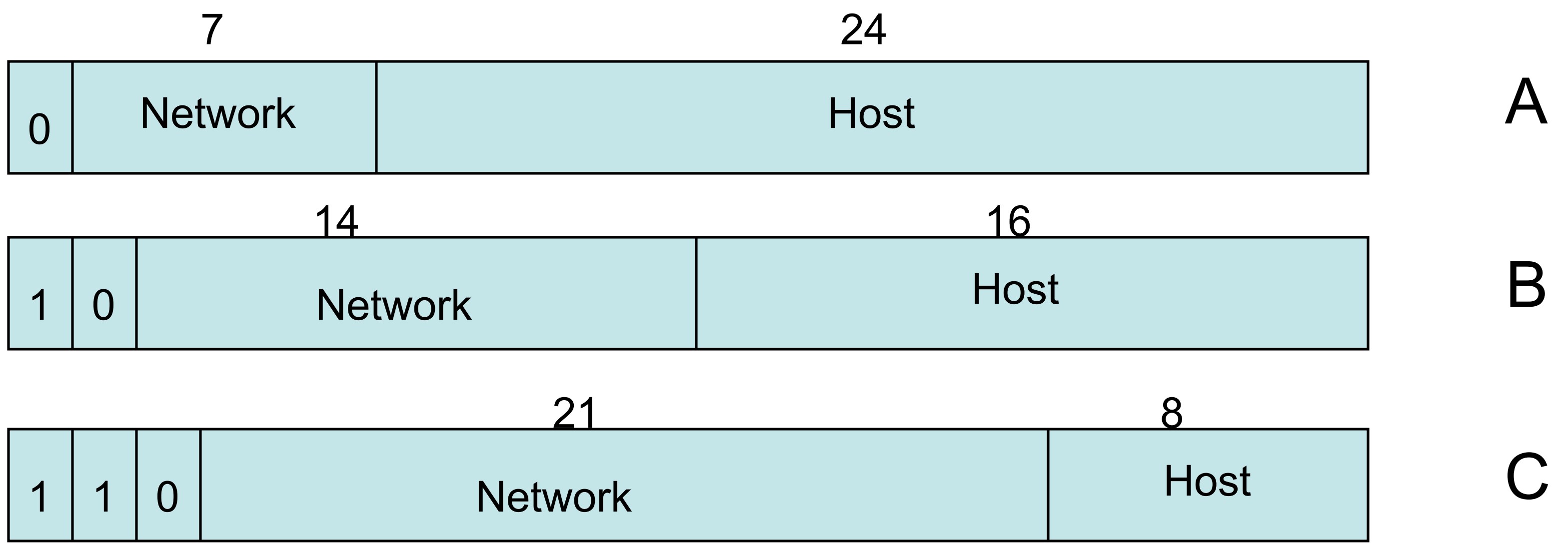

- Written as four decimal numbers separated by dots: 158.130.14.2

Network Classes

Class	*# of nets*	*# of hosts per net*
A	126	~16 million
B	8192	65534
C	~2 million	254

IP Forwarding algorithm

- If (Network # dest == Network # interface) then deliver to destination over interface
- else if (Network # dest in forwarding table) deliver packet to NextHop router
- else deliver packet to default router

- Forwarding tables
  - Contain (Network #, NextHop) pairs
  - Additional information
  - Built by routing protocol that learns the network topology, adapts to changes

Subnetting

- Problem: IP addressing scheme leads to fragmentation
  - A class B network with only 300 machines on it wastes > 65,000 addresses
  - Need a way to divide up a single network address space into multiple smaller subnetworks.

- Idea: One IP network number allocated to several physical networks.
  - The multiple physical networks are called *subnets*
  - Should be close together (why?)
  - Useful when a large company (or university!) has many physical networks.

Subnet Numbers

- Solution: *Subnetting*
  - All nodes are configured with *subnet mask*
  - Allows definition of a *subnet number*
    - All hosts on a physical subnetwork share the same *subnet number*

Subnet Mask (255.255.255.0)

111111111111111111111111	00000000

Subnetted Address:

Network number	Subnet ID	Host ID

Example of Subnetting

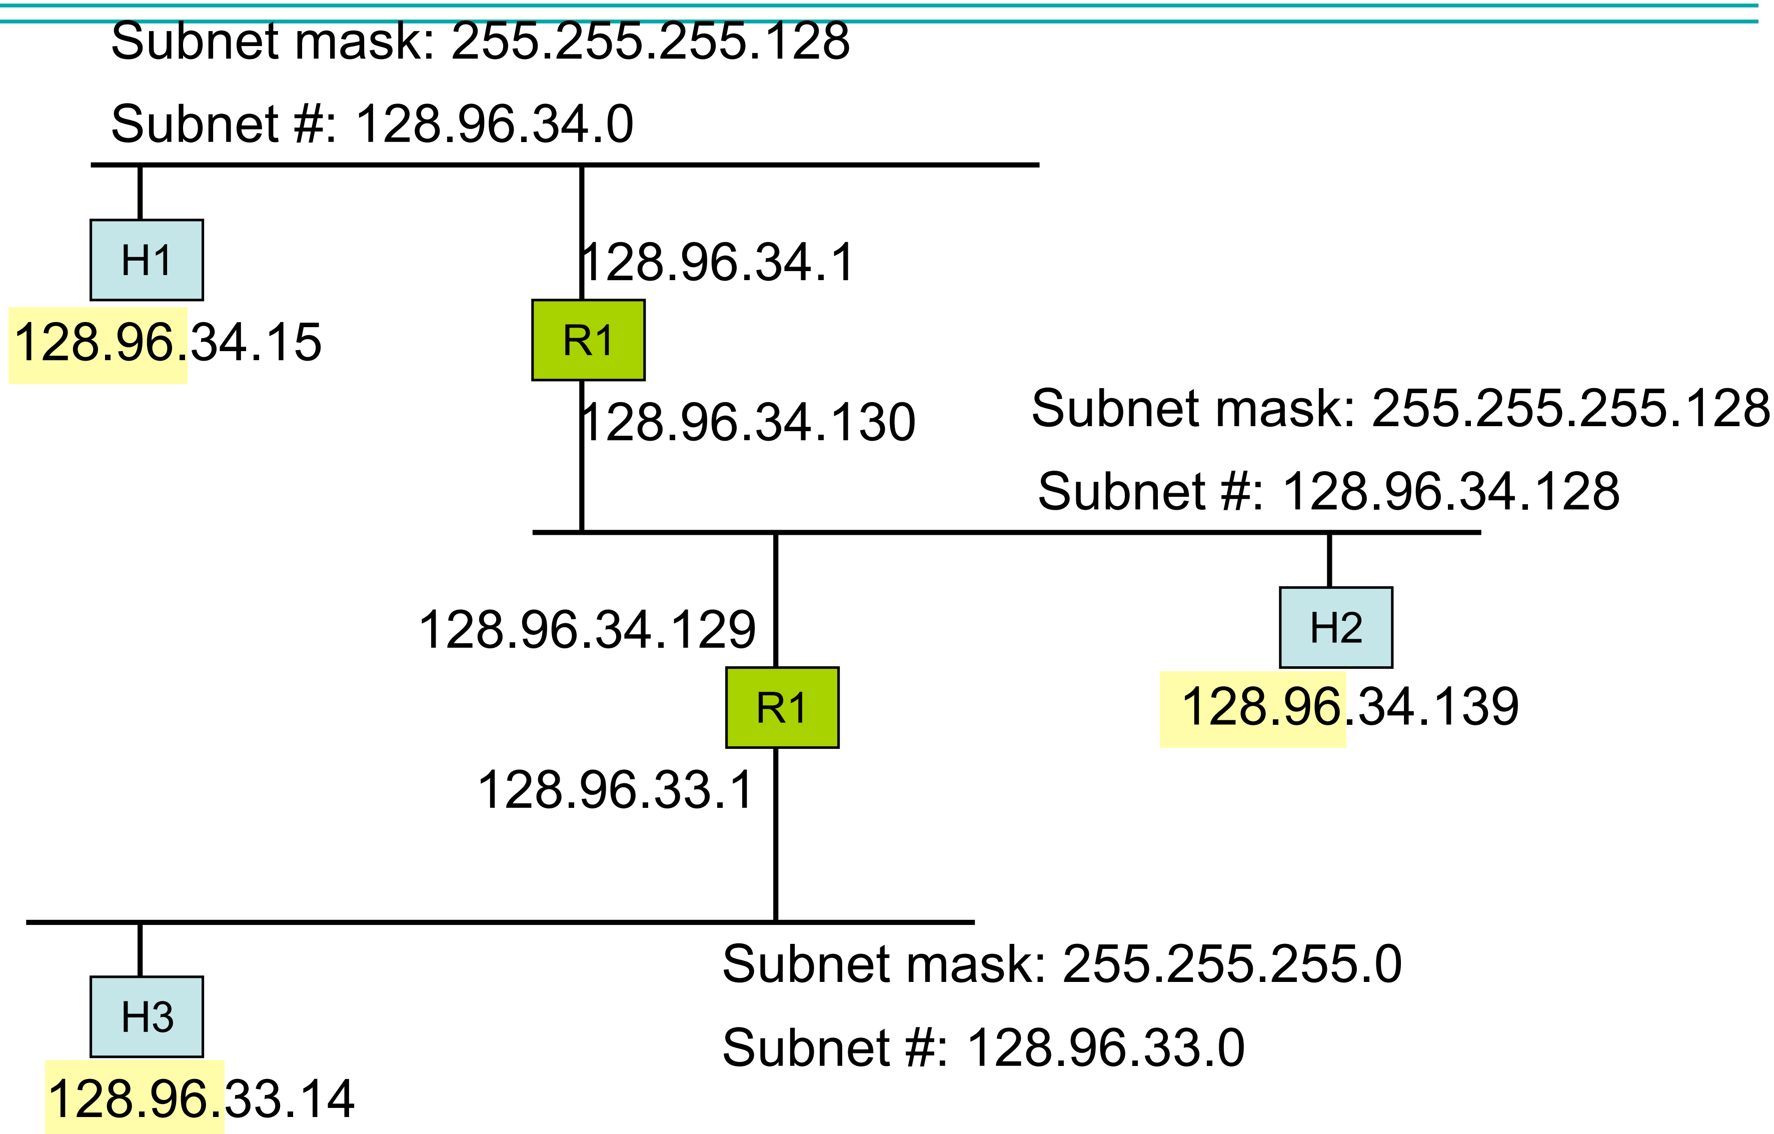

Subnets, continued

- Mask is bitwise-ANDed with address
- This is done at routers
- Router tables in this model:
  - <Subnet #, Subnet Mask, NextHop>
- Subnetting allows a set of physical networks to look like a single logical network from elsewhere

Forwarding Algorithm

```
D = destination IP address
for each forwarding table entry
(SubnetNumber, SubnetMask, NextHop)
    D1 = SubnetMask & D
    if  D1 = SubnetNumber
        if NextHop is an interface
            deliver datagram directly to destination
        else
            deliver datagram to NextHop (router)

Deliver datagram to default router (if above fails)
```

ARP - Address Resolution Protocol

- Problem:
  - Need mapping between IP and link layer addresses.

- Solution: ARP
  - Every host maintains IP–Link layer mapping table (cache)
  - Timeout associated with cached info (15 min.)
- Sender
  - Broadcasts “Who is IP addr X?”
  - Broadcast message includes sender’s IP & Link Layer address
- Receivers
  - Any host with sender in cache “refreshes” time-out
  - Host with IP address X replies “IP X is Link Layer Y”
  - Target host adds sender (if not already in cache)

ICMP: Internet Control Message Protocol

- Collection of error & control messages
- Sent back to the source when Router or Host cannot process packet correctly
- Error Examples:
  - Destination host unreachable
  - Reassembly process failed
  - TTL reached 0
  - IP Header Checksum failed
- Control Example:
  - Redirect – tells source about a better route

Domain Name System

- System for mapping mnemonic names for computers into IP addresses.

  zeta.cis.upenn.edu → 158.130.12.244
- Domain Hierarchy
- Name Servers
  – 13 Root servers map top-level domains such as ".com" or ".net"
  – (Why 13? Early UDP protocol supported only 512 bytes…)
- Name Resolution
  – Protocol for looking up hierarchical domain names to determine the IP address
  – Protocol runs on UDP port 53

Domain Name Hierarchy

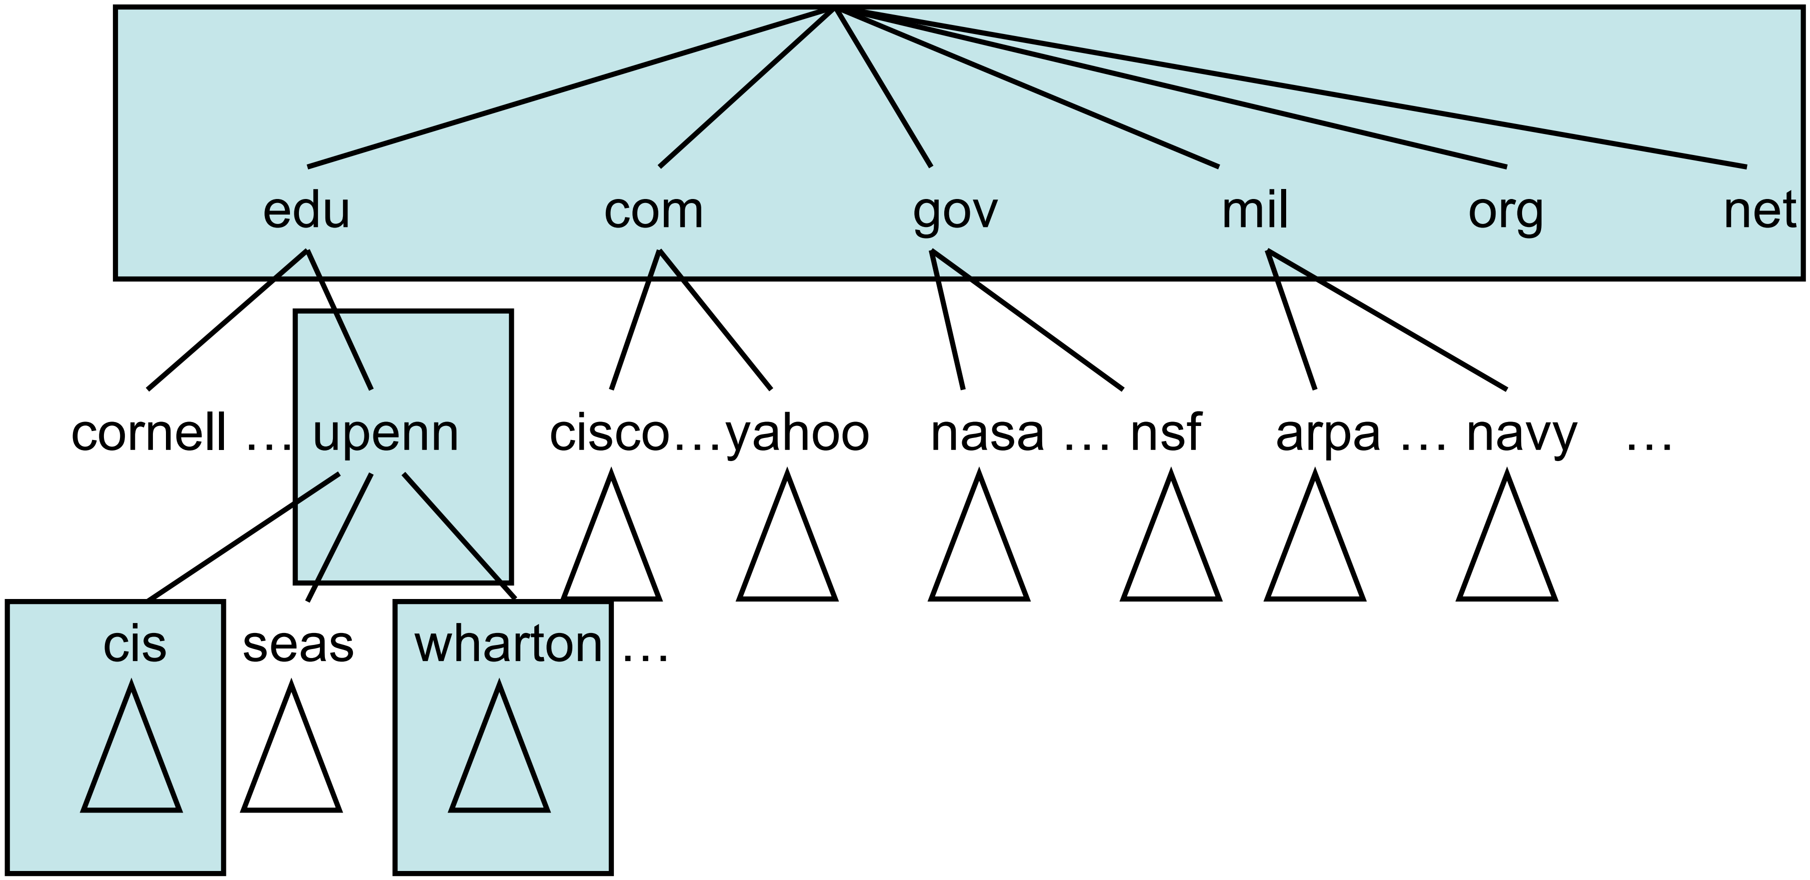

Hierarchy of Name Servers

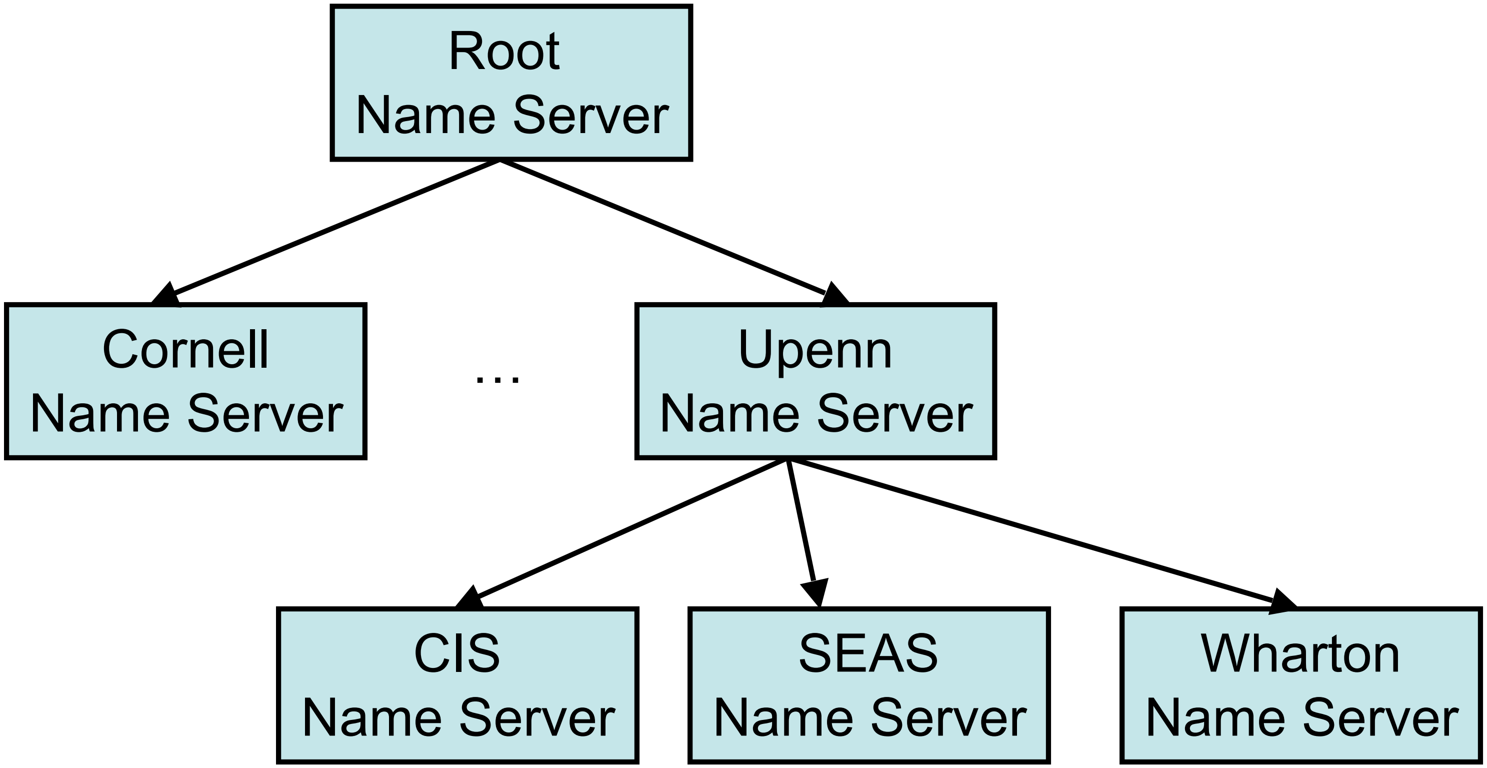

Records on Name Servers

- < Name, Type, Class, TTL, RDLength, RDATA >
- Name of the node
- Types:
  - A Host to address mappings
  - NS Name server address mappings
  - CNAME Aliases
  - MX Mail exchange server mappings
  - … others
- Class IN for IP addresses

Name resolution

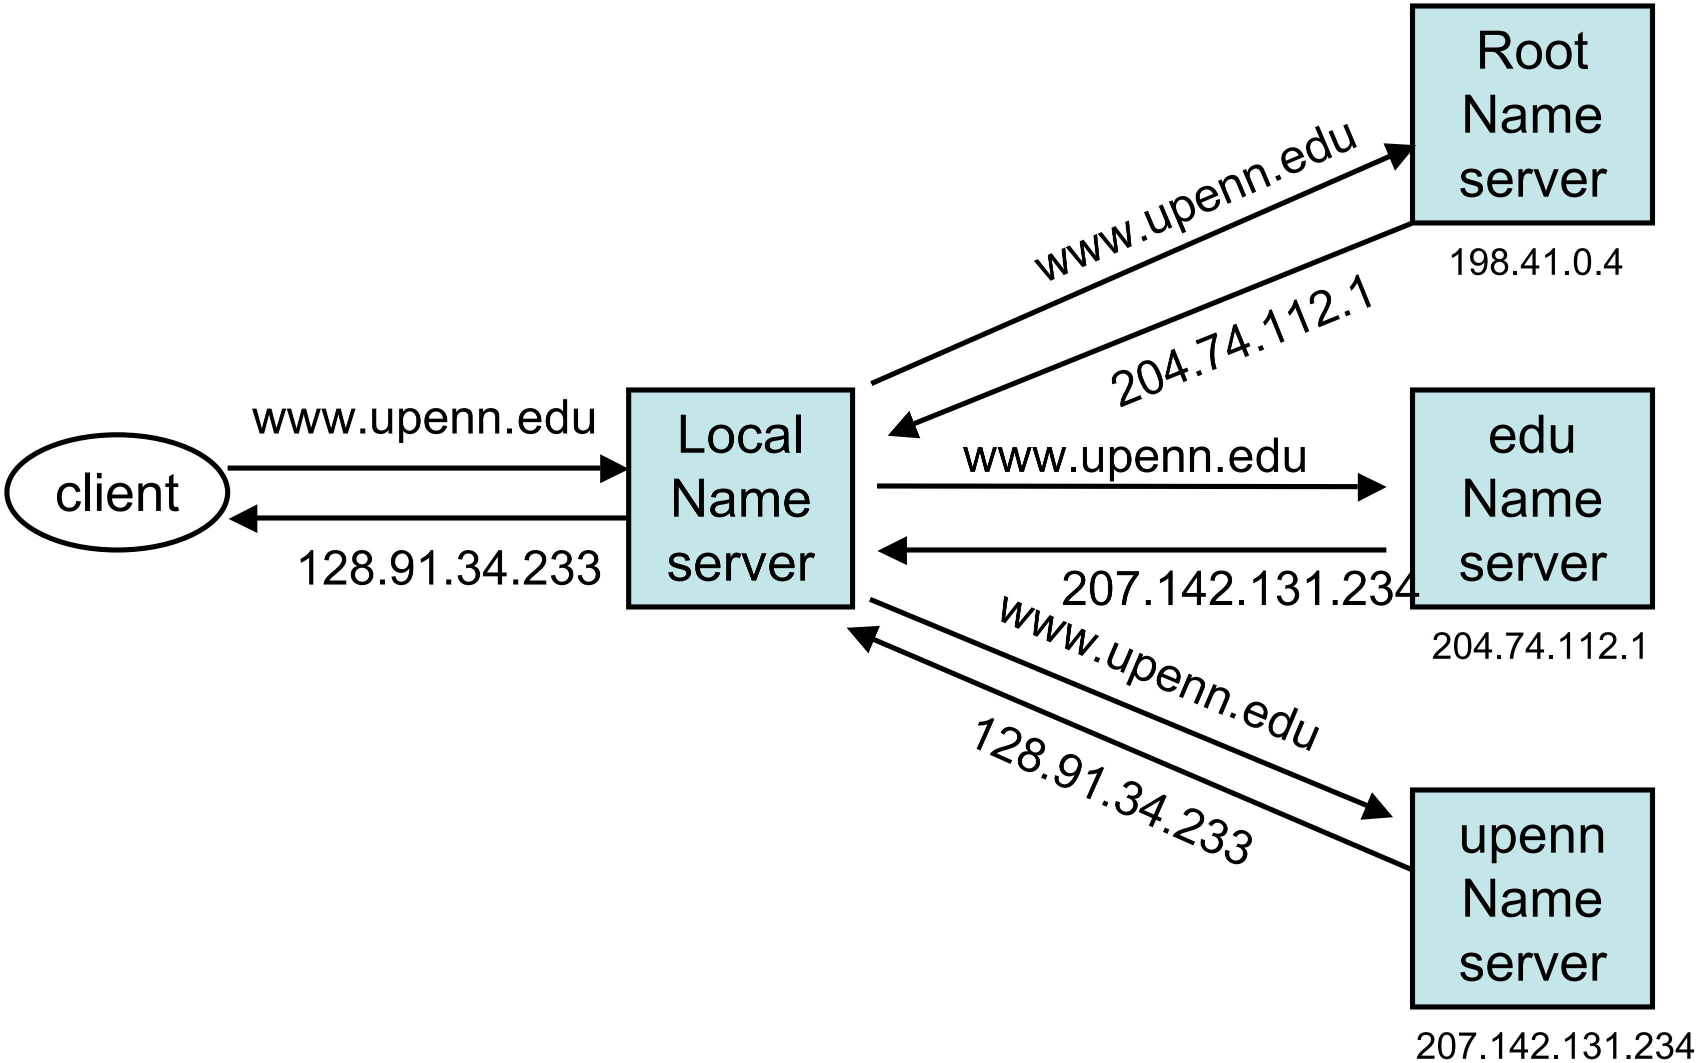

DNS Vulnerabilities

- See *"Corrupted DNS Resolution Paths: The rise of a malicious resolution authority"* by Dagon et al.
- Rogue DNS Servers
  - Compromised DNS servers that answer incorrectly
- DNS Cache Poisoning
  - Request: `subdomain.example.com IN A`
  - Reply:

```
Answer:
(no response)

Authority section:
example.com. 3600 IN ns.wikipedia.org.

Additional section:
ns.wikipedia.org IN A w.x.y.z
```

Reflected denial of service

- ICMP message with an "echo request" is called 'ping'
- Broadcast a ping request
  - For sender's address put target's address
  - All hosts reply to ping, flooding the target with responses
- Hard to trace
- Hard to prevent
  - Turn off ping? (Makes legitimate use impossible)
  - Limit with network configuration by restricting scope of broadcast messages
- Sometimes called a "smurf attack"

(Distributed) Denial of Service

- Coordinate multiple subverted machines to attack
- Flood a server with bogus requests
  - TCP SYN packet flood
  - > 600,000 packets per second
- Detection & Assessment?
  - 12,800 attacks at 5000 hosts! (in 3 week period during 2001)
  - IP Spoofing (forged source IP address)
  - http://www.cs.ucsd.edu/users/savage/papers/UsenixSec01.pdf

- Feb. 6 2007: 6 of 13 root servers suffered DDoS attack
- Oct. 21 2002: 9 of 13 root servers were swampted
  - Prompted changes in the architecture

- Prevention?
  - Filtering?
  - Decentralized file storage?