Overview: Internet vs Data Center Networks

Hakim Weatherspoon
Assistant Professor, Dept of Computer Science

CS 5413: High Performance Systems and Networking

August 29, 2014

Goals for Today

- Overview
  - What is the Internet?
    - What is it and how did we get here?
  - What is about Data Centers?
- Network Structure
  - Internet
  - Data center

What is the Internet?

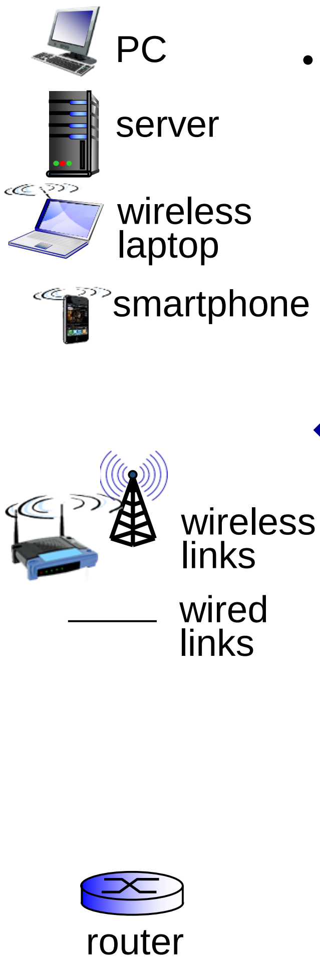

- millions of connected computing devices:
  - *hosts = end systems*
  - running *network apps*

❖ *communication links*
  - fiber, copper, radio, satellite
  - transmission rate: *bandwidth*

❖ *Packet switches:* forward packets (chunks of data)
  - *routers* and *switches*

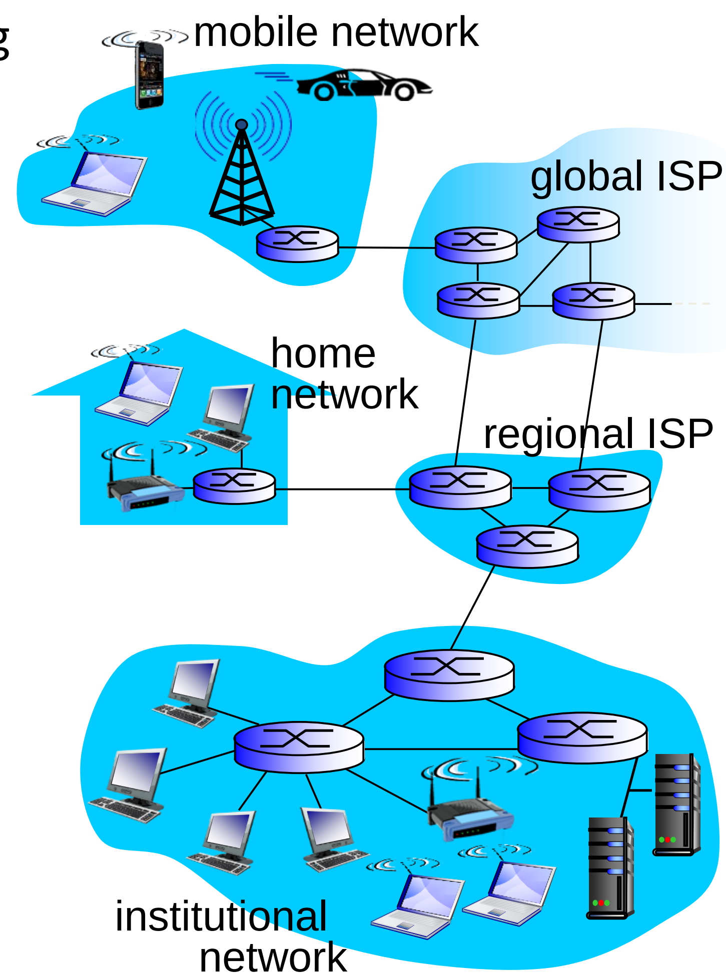

What is the Internet?

IP picture frame
http://www.ceiva.com/

Web-enabled toaster +
weather forecaster

Tweet-a-watt:
monitor energy use

Internet
refrigerator

Slingbox: watch,
control cable TV remotely

Internet phones

What is the Internet?

- *Internet:* "network of networks"
  - Interconnected ISPs
- *protocols* control sending, receiving of msgs
  - e.g., TCP, IP, HTTP, Skype, 802.11
- *Internet standards*
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

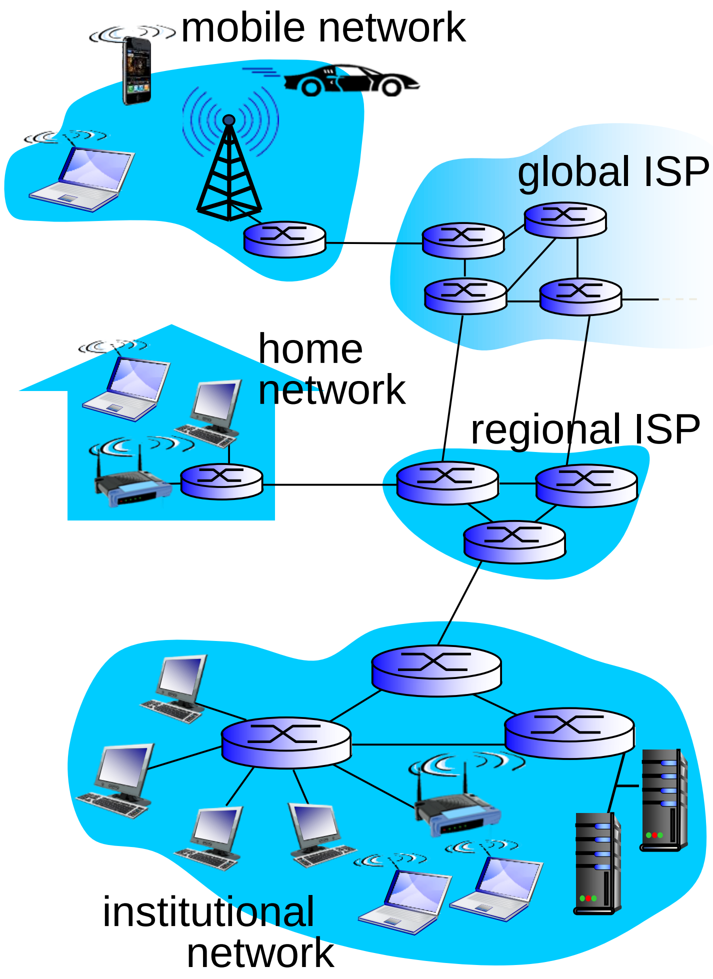

What is the Internet?

- *Infrastructure that provides services to applications:*
  - Web, VoIP, email, games, e-commerce, social nets, …
- *provides programming interface to apps*
  - hooks that allow sending and receiving app programs to "connect" to Internet
  - provides service options, analogous to postal service

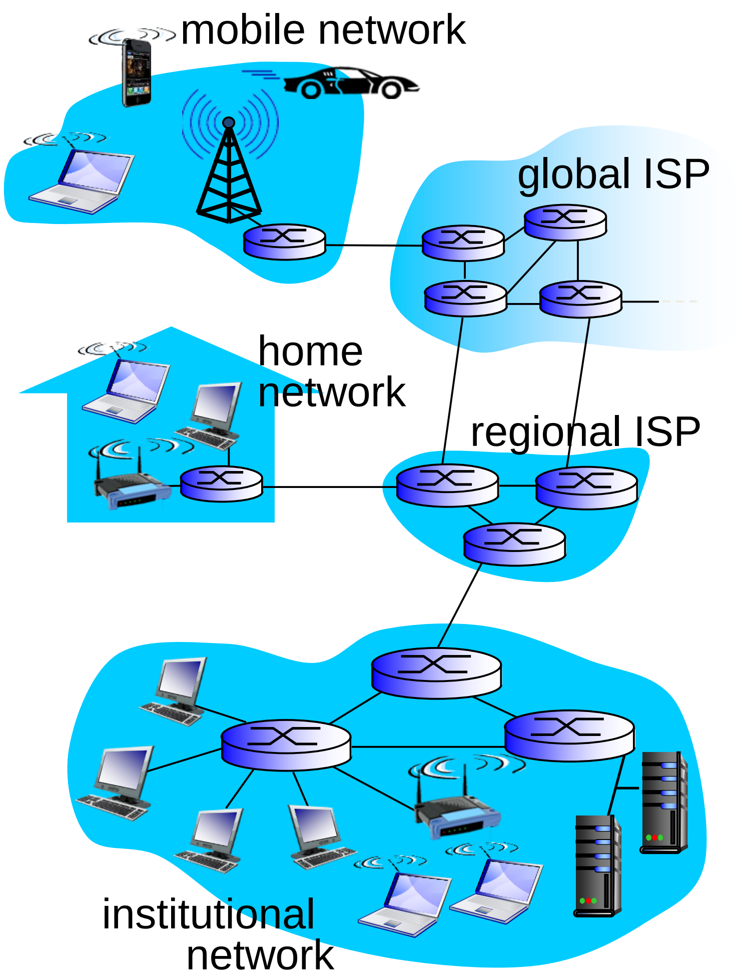

How did we get here: Internet History

1961-1972: Early packet-switching principles

- 1961: Kleinrock - queueing theory shows effectiveness of packet-switching
- 1964: Baran - packet-switching in military nets
- 1967: ARPAnet conceived by Advanced Research Projects Agency
- 1969: first ARPAnet node operational

- 1972:
  - ARPAnet public demo
  - NCP (Network Control Protocol) first host-host protocol
  - first e-mail program
  - ARPAnet has 15 nodes

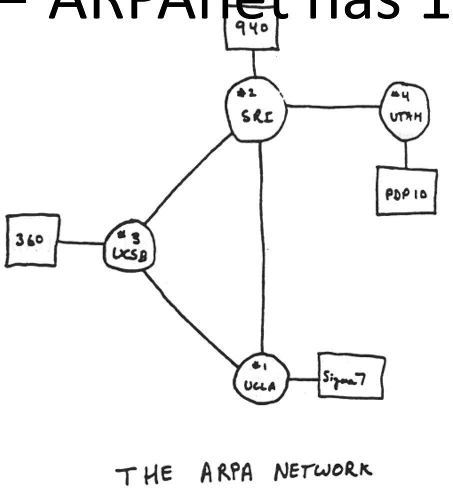

How did we get here: Internet History

1972-1980: Internetworking, new and proprietary nets

- 1970: ALOHAnet satellite network in Hawaii
- 1974: Cerf and Kahn - architecture for interconnecting networks
- 1976: Ethernet at Xerox PARC
- late70's: proprietary architectures: DECnet, SNA, XNA
- late 70's: switching fixed length packets (ATM precursor)
- 1979: ARPAnet has ***200 nodes***

Cerf and Kahn's internetworking principles:
- minimalism, autonomy - no internal changes required to interconnect networks
- best effort service model
- stateless routers
- decentralized control

define today's Internet architecture

Goals for Today

- Overview
  - What is the Internet?
    - What is it and how did we get here?
  - What is about Data Centers?
- Network Structure
  - Internet
  - Data center

How did we get here: Internet History

1980-1990: new protocols, a proliferation of networks

- 1983: deployment of TCP/IP
- 1982: smtp e-mail protocol defined
- 1983: DNS defined for name-to-IP-address translation
- 1985: ftp protocol defined
- 1988: TCP congestion control

- new national networks: Csnet, BITnet, NSFnet, Minitel
- *100,000 hosts connected to confederation of networks*

How did we get here: Internet History

1990, 2000's: commercialization, the Web, new apps

- early 1990's: ARPAnet decommissioned
- 1991: NSF lifts restrictions on commercial use of NSFnet (decommissioned, 1995)
- early 1990s: Web
  - hypertext [Bush 1945, Nelson 1960's]
  - HTML, HTTP: Berners-Lee
  - 1994: Mosaic, later Netscape
  - late 1990's: commercialization of the Web

late 1990's – 2000's:

- more killer apps: instant messaging, P2P file sharing
- network security to forefront
- ***est. 50 million host, 100 million+ users***
- backbone links running at Gbps

How did we get here: Internet History

2005-present

- ***~750 million hosts***
  - Smartphones and tablets
- Aggressive deployment of broadband access
- Increasing ubiquity of high-speed wireless access
- Emergence of online social networks:
  - Facebook: soon one billion users
- Service providers (Google, Microsoft) create their own networks
  - Bypass Internet, providing "instantaneous" access to search, email, etc.
- E-commerce, universities, enterprises running their services in "cloud" (eg, Amazon EC2)

What is different about Data Centers?

- Cost
  - "It is the economics stupid"
    - James Hamilton, VP & Distinguished engineer, Amazon Web Services
  - In 2008, data center staff to servers was 1:1000

    Today, closer to 1:10,000
- Scale
  - Millions of servers, billions of users, trillions of objects
  - Scale out instead of scale up
- Efficient
  - Massive scale in the same location eases design and lowers costs
- Global scale data centers
  - Data Centers strategically placed where power is cheap and close to consumers

Where do the costs go?

- Breakdown
  - 45% Servers - CPU, memory, storage subsystems
  - 25% Infrastructure - Power distribution and cooling
  - 15% Power draw - Electrical utility costs
  - 15% Network - Links, transit, equipment

Where do the costs go?

- Breakdown
  - 45% Servers - CPU, memory, storage subsystems
  - 25% Infrastructure - Power distribution and cooling
  - 15% Power draw - Electrical utility costs
  - 15% Network - Links, transit, equipment
- How to reduced costs
  - Servers and Infrastructure
    - Let servers fail and infrastructure fail
    - Software, Replication and network efficiency can help
  - Power and Network
    - High utilization (better on than off)
    - Agility (ability to run applications anywhere in data center)

Networking in Data Centers

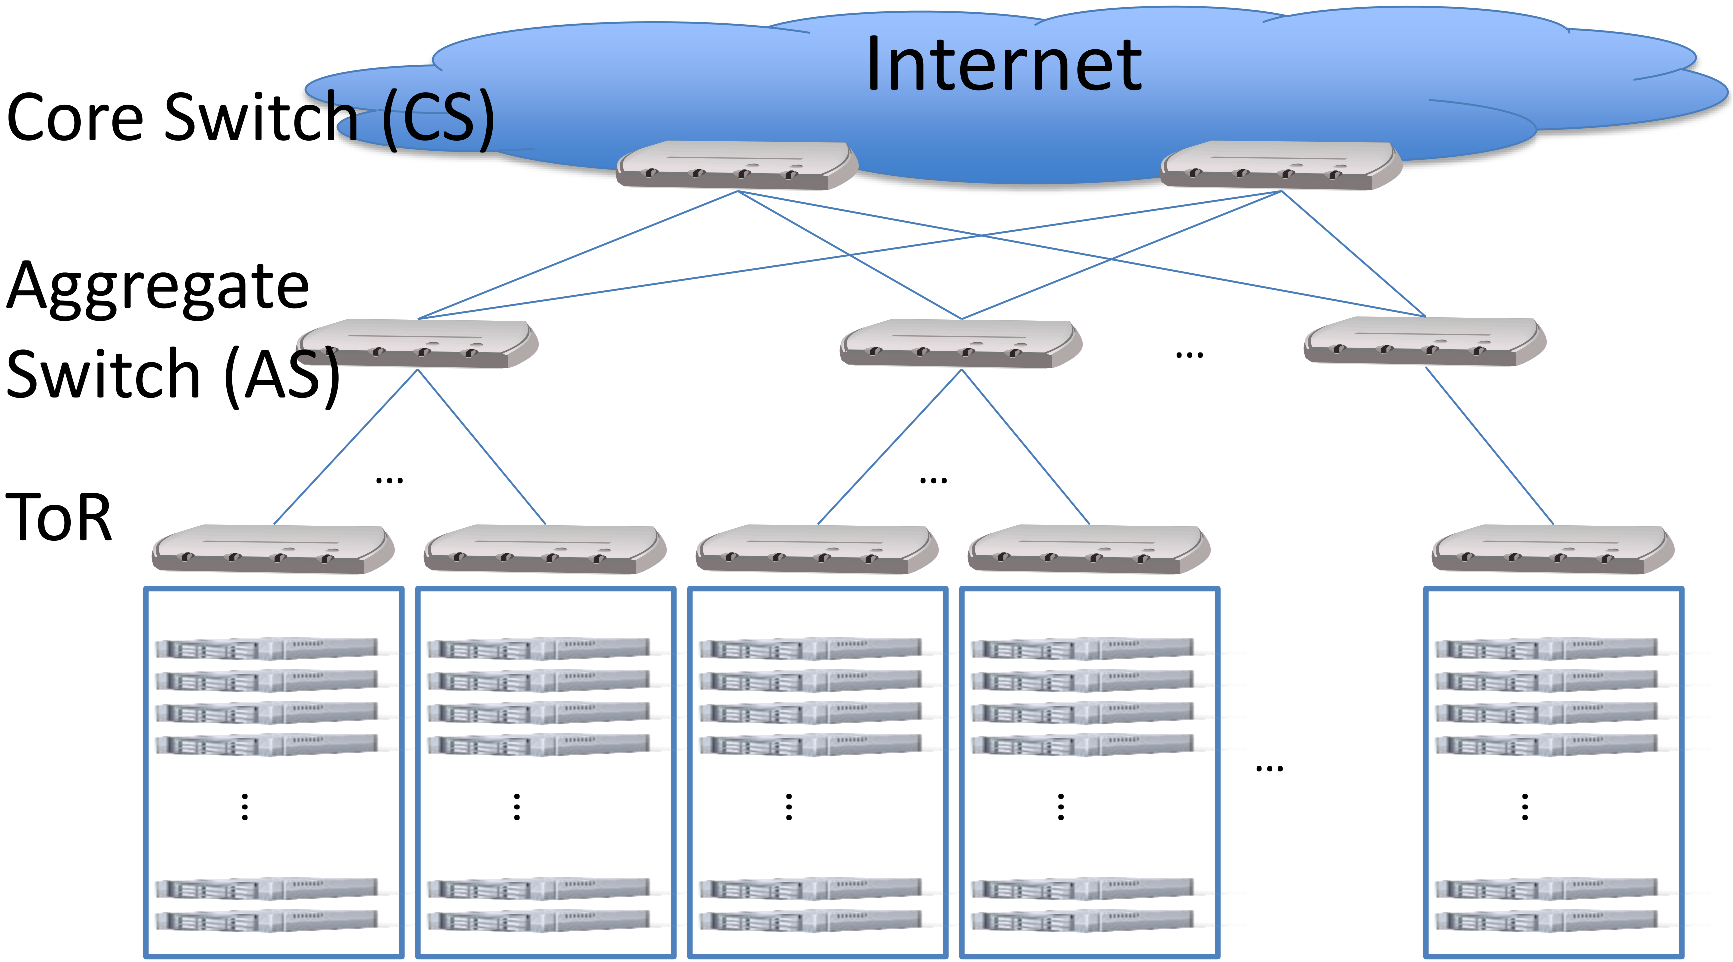

Networking in Data Centers

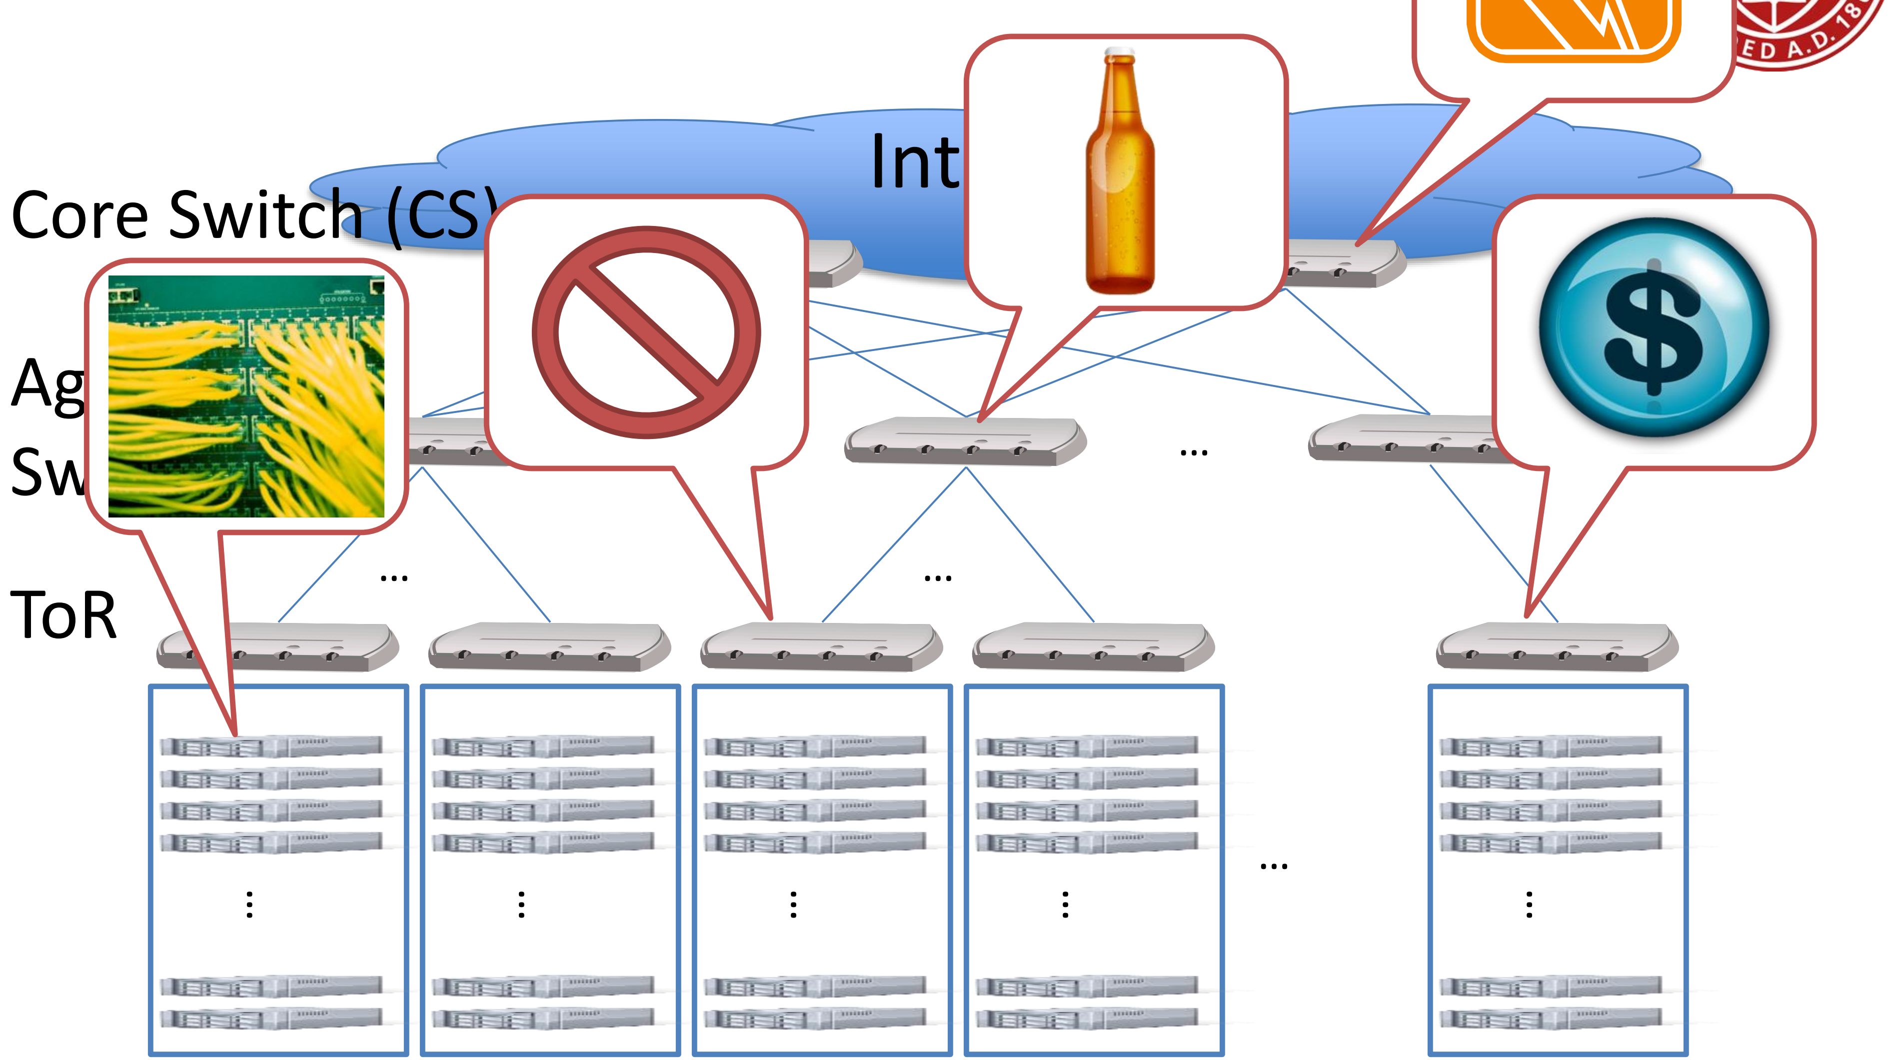

Geo-distributed Data Centers

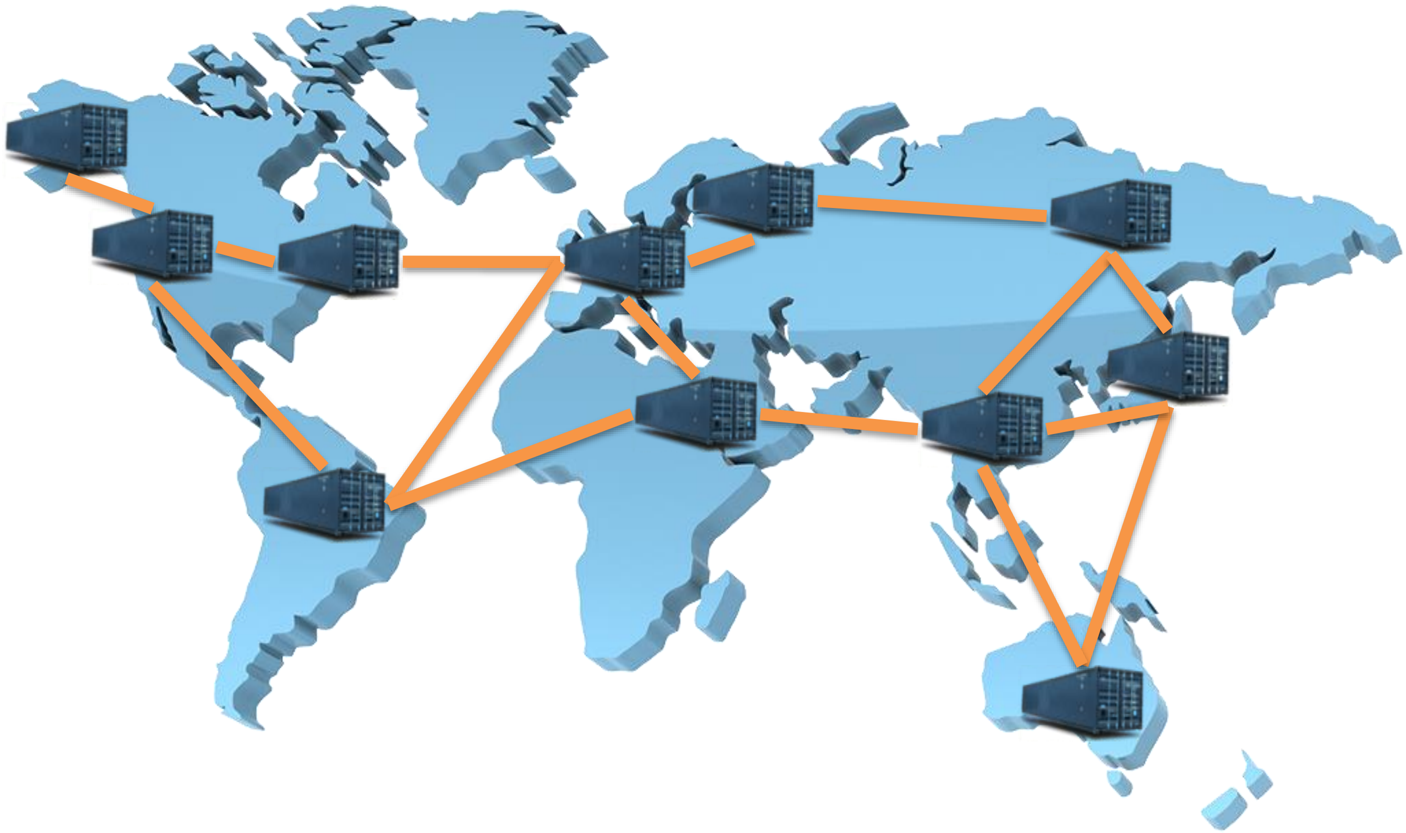

Goals for Today

- Overview
  - What is the Internet?
    - What is it and how did we get here?
  - What is about Data Centers?
- Structure
  - Internet
  - Data Centers

Internet Structure

- Network Protocols
- Edge Network
- Core Networks

Network Protocols

human protocols:

- "what's the time?"
- "I have a question"
- introductions

... specific msgs sent

... specific actions taken when msgs received, or other events

network protocols:

- machines rather than humans
- all communication activity in Internet governed by protocols

protocols define format, order of msgs sent and received among network entities, and actions taken on msg transmission, receipt

Network Protocols

a human protocol and a computer network protocol:

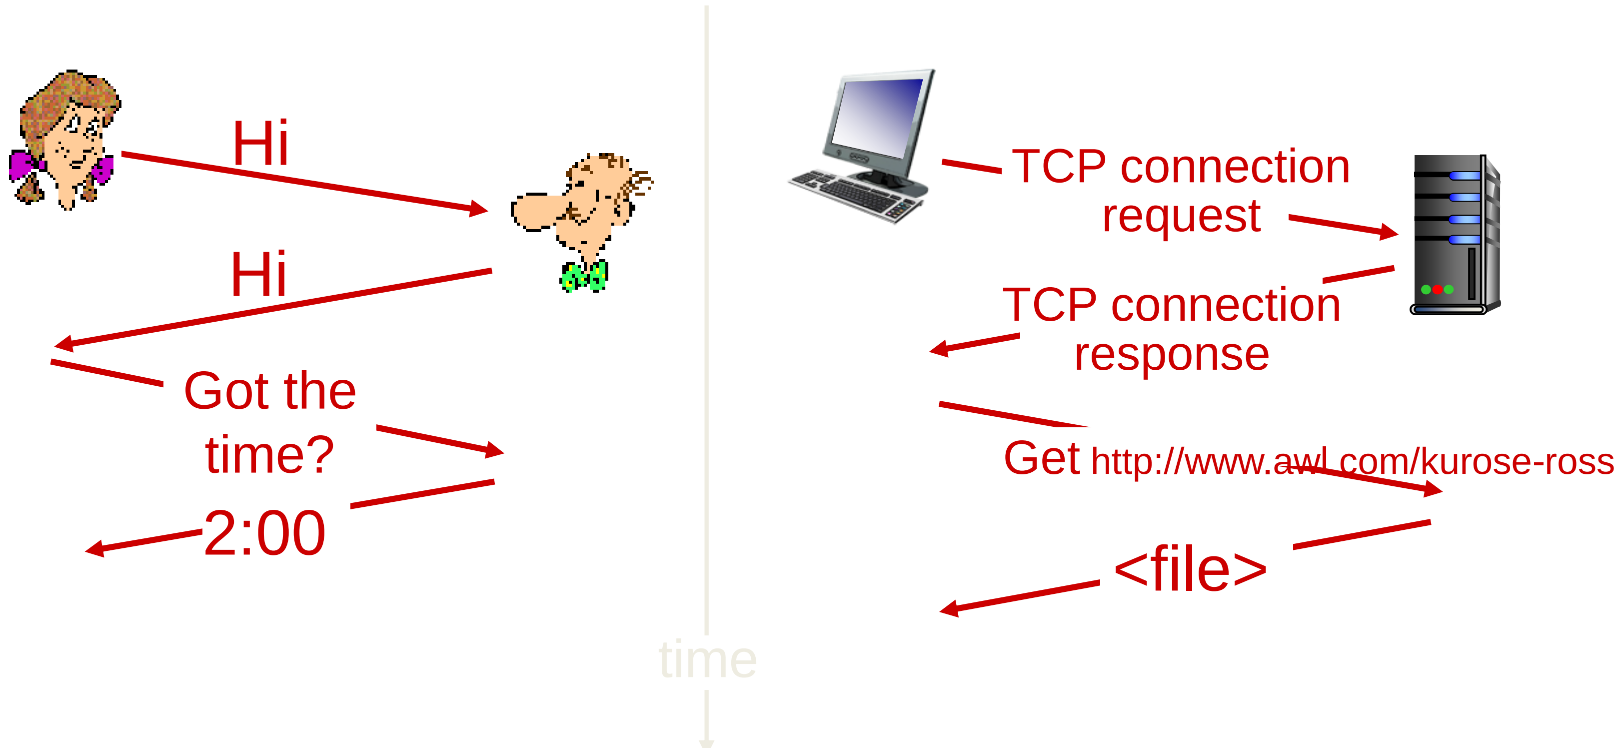

Network Protocol "Layers"

Networks are complex,
with many "pieces":

- hosts
- routers
- links of various media
- applications
- protocols
- hardware, software

Question:

is there any hope of *organizing* structure of network?

Network Protocol “Layers”

Similar to Traveling protocol

ticket (purchase)	ticket (complain)
baggage (check)	baggage (claim)
gates (load)	gates (unload)
runway takeoff	runway landing
airplane routing	airplane routing

airplane routing

- a series of steps

Network Protocol “Layers”

Similar to Traveling protocol

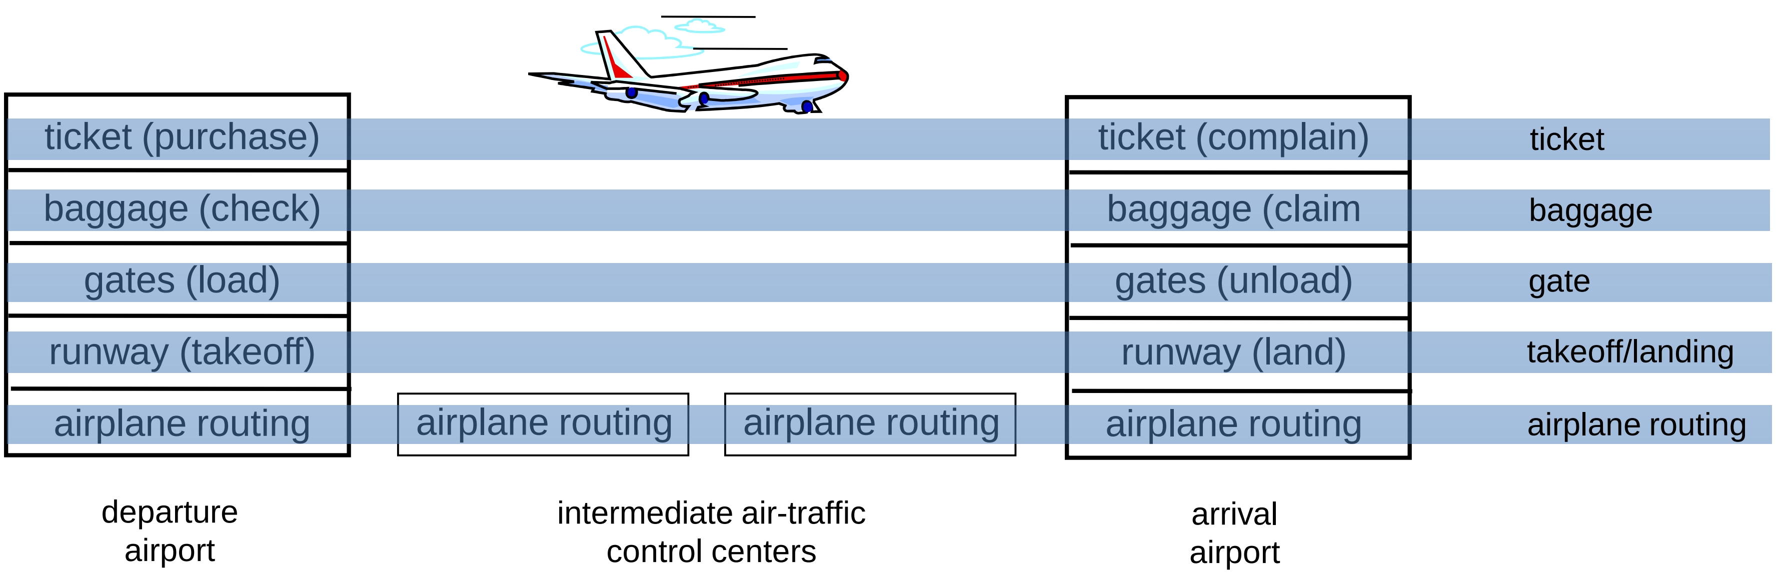

layers: each layer implements a service

- via its own internal-layer actions
- relying on services provided by layer below

Why layering?

dealing with complex systems:

- explicit structure allows identification, relationship of complex system's pieces
  - layered *reference model* for discussion
- modularization eases maintenance, updating of system
  - change of implementation of layer's service transparent to rest of system
  - e.g., change in gate procedure doesn't affect rest of system
- layering considered harmful?

Internet Protocol Stack

- *application:* supporting network applications
  - FTP, SMTP, HTTP
- *transport:* process-process data transfer
  - TCP, UDP
- *network:* routing of datagrams from source to destination
  - IP, routing protocols
- *link:* data transfer between neighboring network elements
  - Ethernet, 802.111 (WiFi), PPP
- *physical:* bits "on the wire"

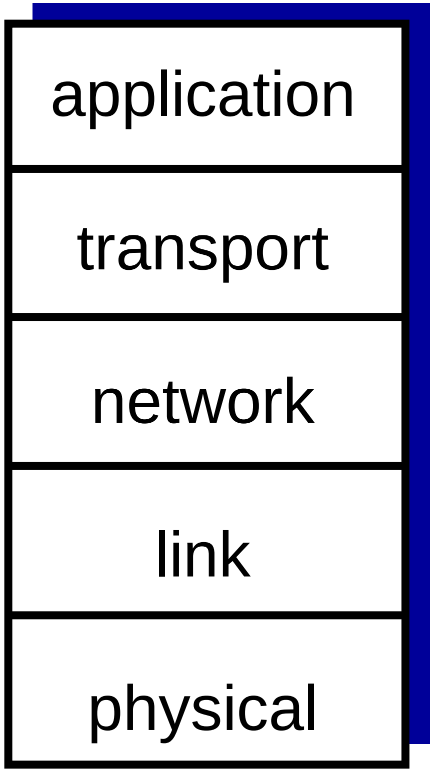

ISO/OSI Reference Model

Not used with due to Internet Protocol Stack

- *presentation:* allow applications to interpret meaning of data, e.g., encryption, compression, machine-specific conventions
- *session:* synchronization, checkpointing, recovery of data exchange
- Internet stack "missing" these layers!
  - these services, *if needed,* must be implemented in application
  - needed?

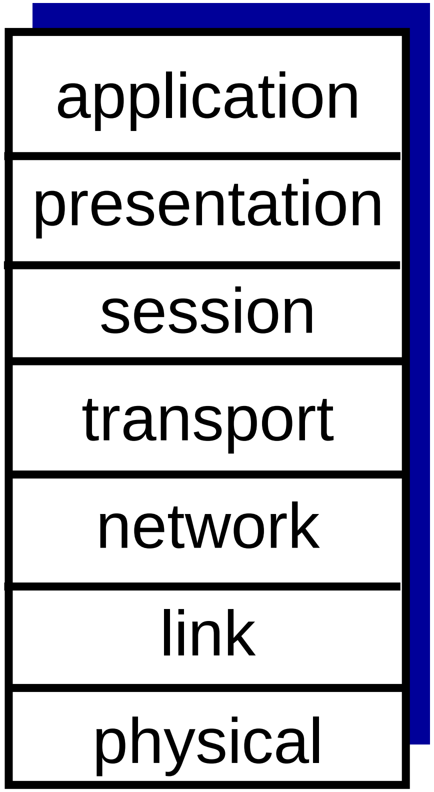

Encapsulation

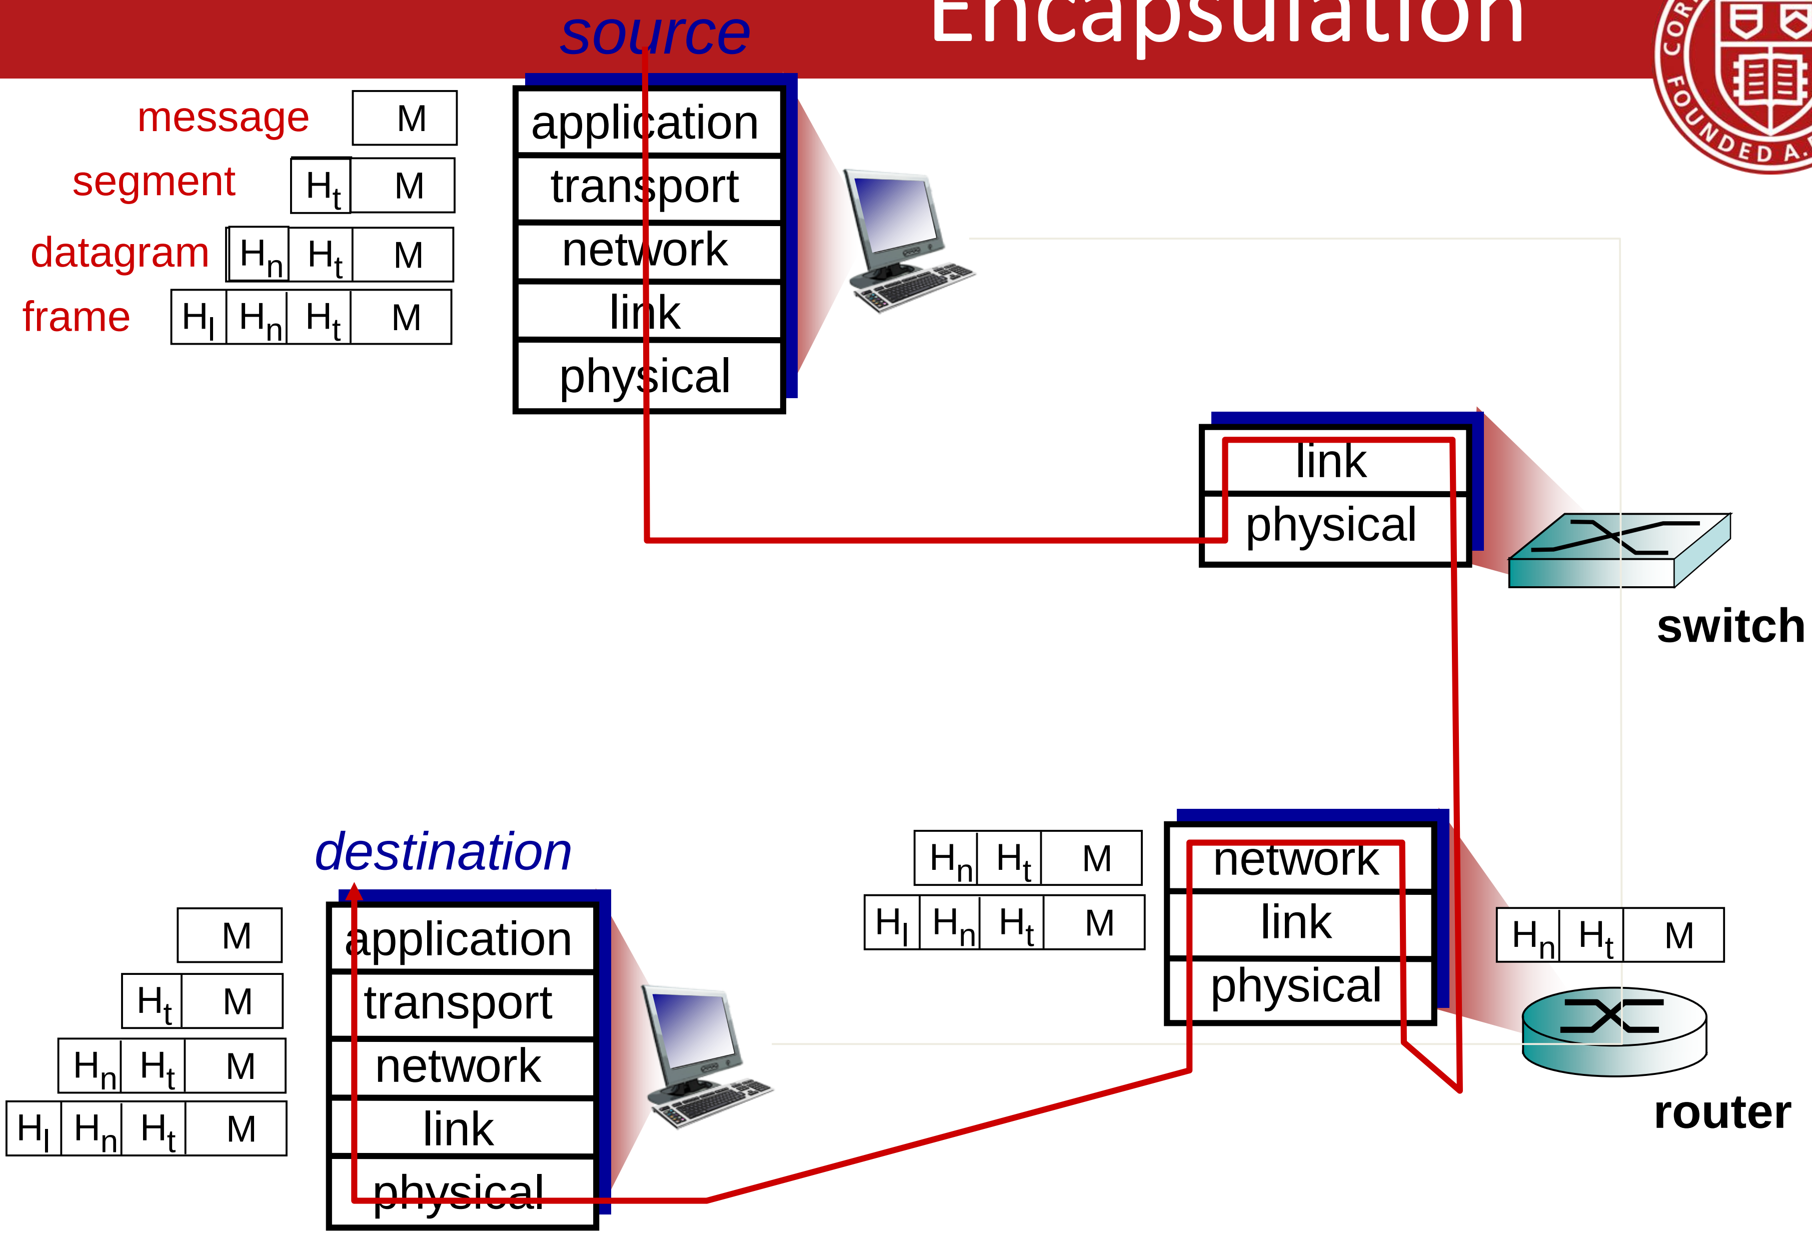

What About Data Centers?

- Data Centers use the same network protocol stack
- But, is this a good thing?

What About Data Centers?

- Data Centers use the same network protocol stack
- But, is this a good thing?
  - Pro
    - Standard for all applications and services
  - Con
    - Efficiency
    - Can prevent full utilization of data center resourcces

Internet Structure

- Network Protocols
- Edge Network
- Core Networks

Internet Structure

- *network edge:*
  - hosts: clients and servers
  - servers often in data centers

- *access networks, physical media:* wired, wireless communication links

- *network core:*
  - interconnected routers
  - network of networks

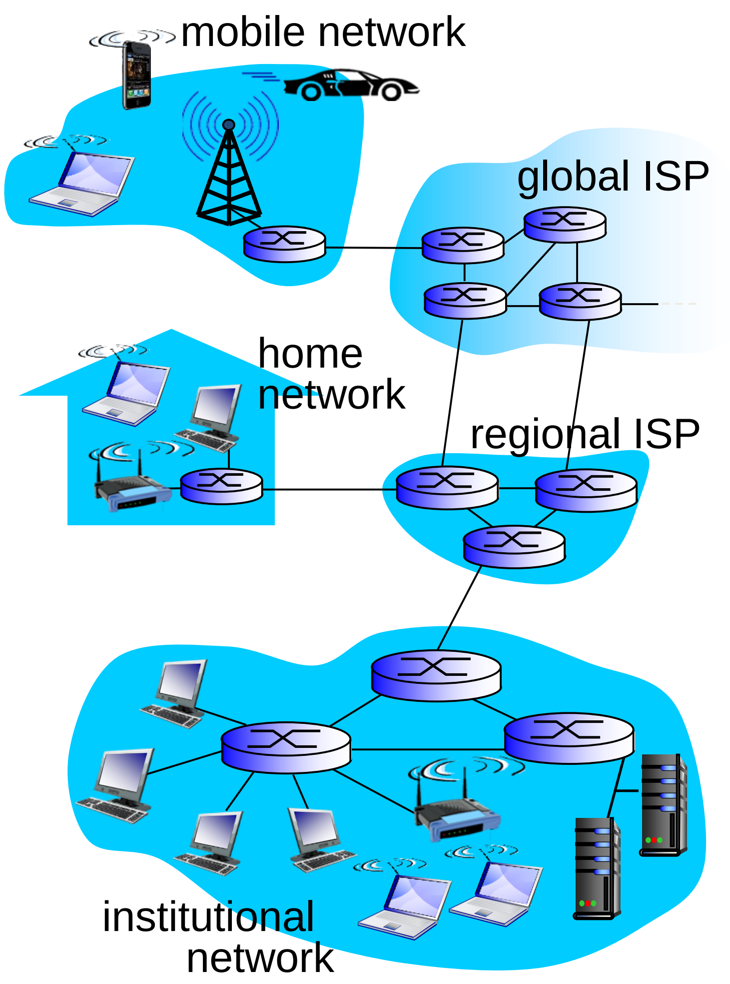

Internet Structure

Q: How to connect end systems to edge router?

- residential access nets
- institutional access networks (school, company)
- mobile access networks

keep in mind:

- bandwidth (bits per second) of access network?
- shared or dedicated?

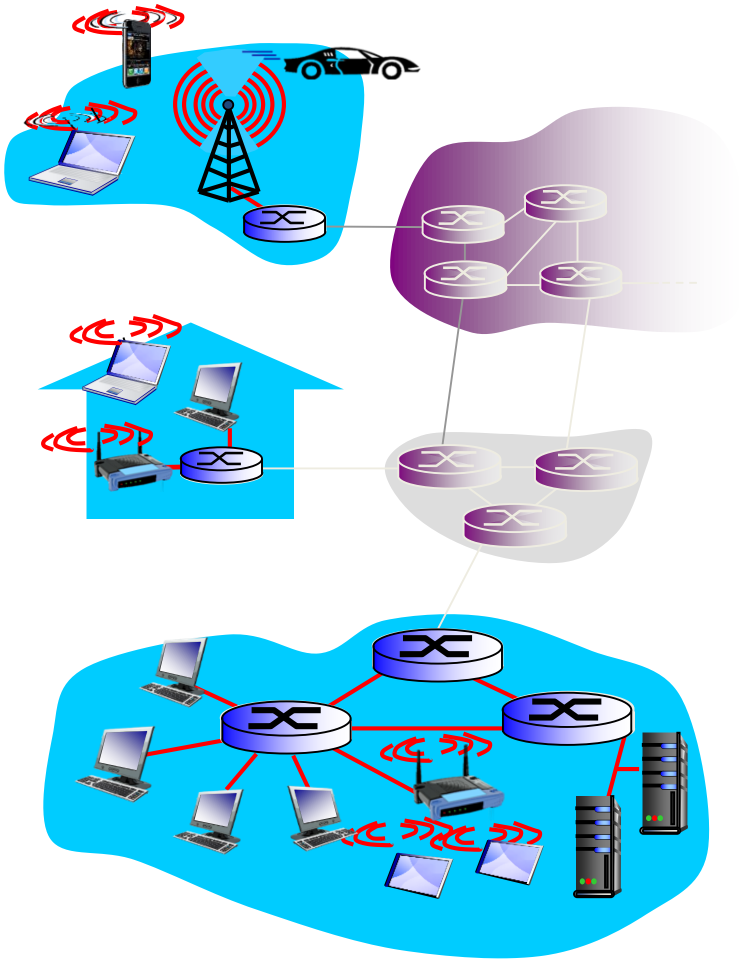

The network Core

- mesh of interconnected routers
- packet-switching: hosts break application-layer messages into *packets*
  - forward packets from one router to the next, across links on path from source to destination
  - each packet transmitted at full link capacity

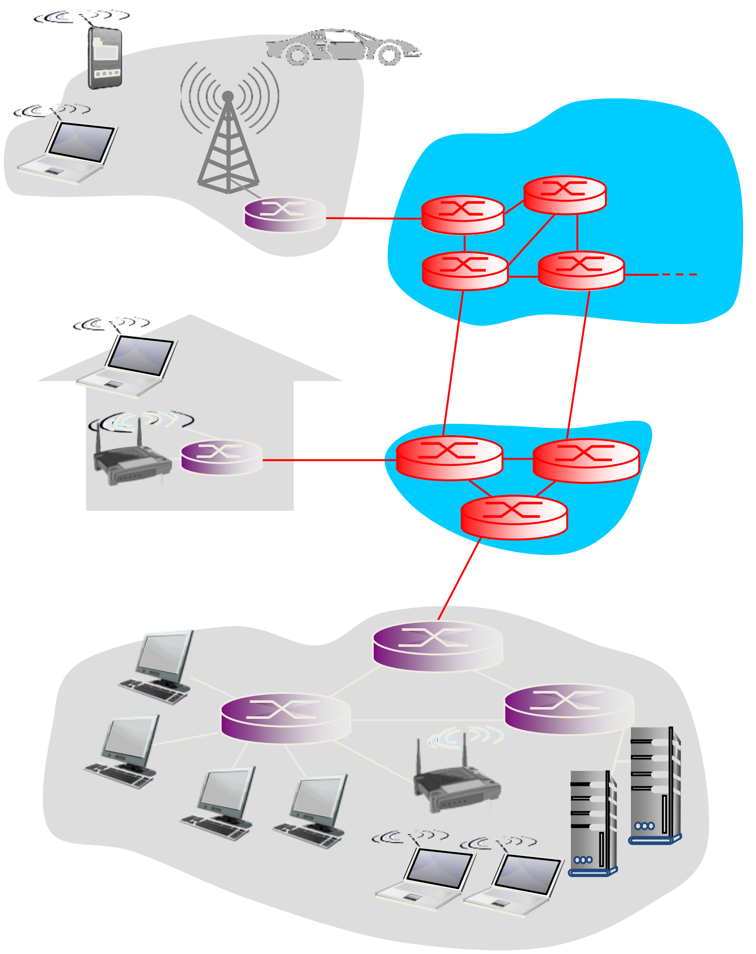

Two key network core functions

routing: determines source-destination route taken by packets

- *routing algorithms*

forwarding: move packets from router's input to appropriate router output

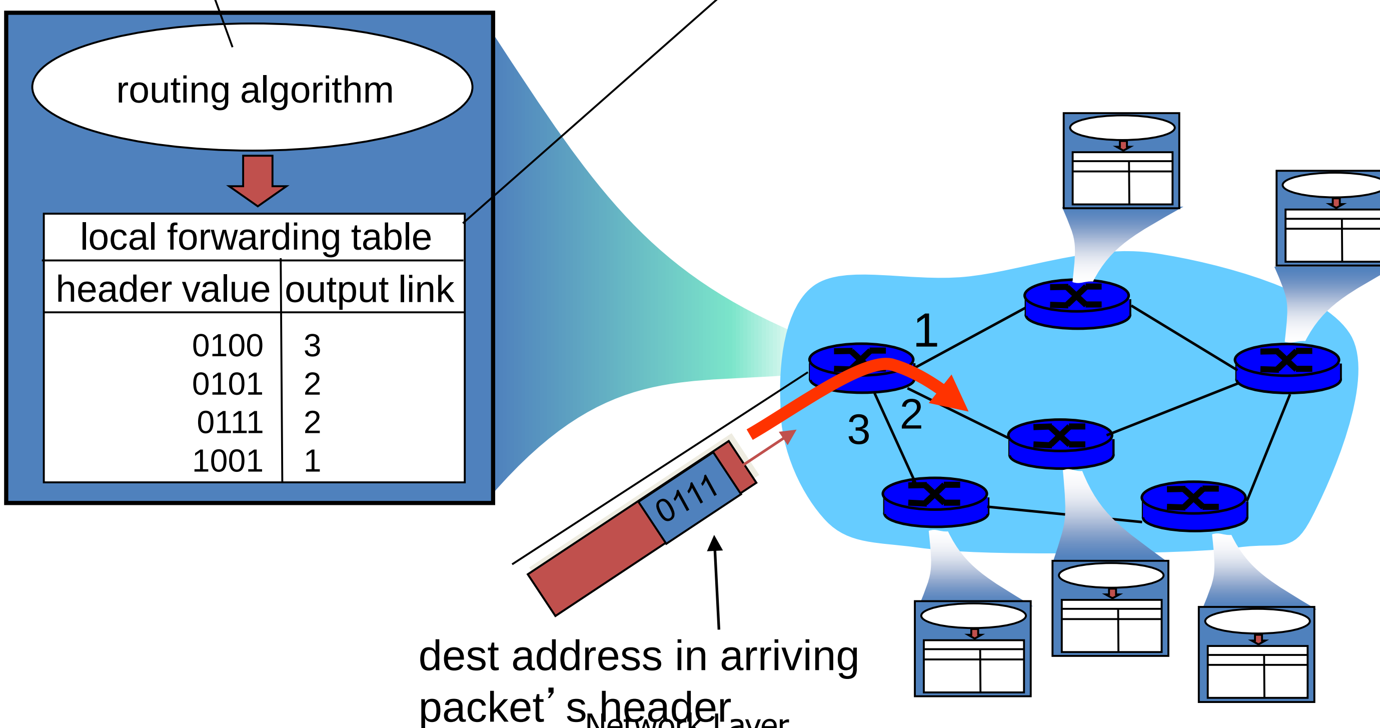

Alternative Core: Circuit Switching

end-end resources allocated to, reserved for "call" between source & dest:

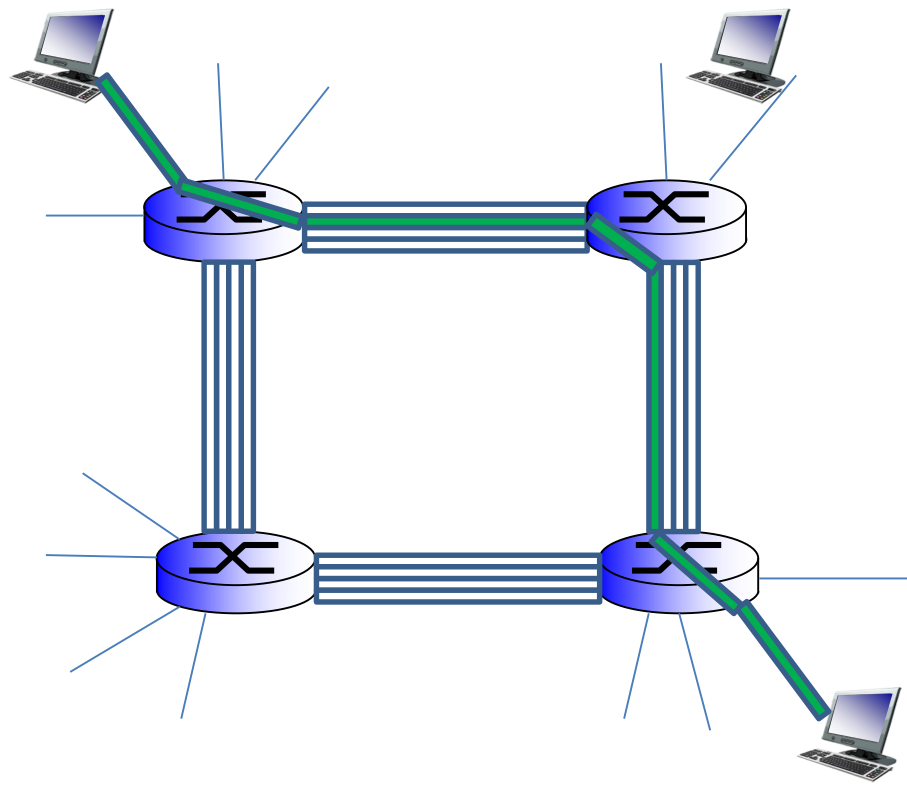

- In diagram, each link has four circuits.
  - call gets 2^{nd} circuit in top link and 1^{st} circuit in right link.
- dedicated resources: no sharing
  - circuit-like (guaranteed) performance
- circuit segment idle if not used by call *(no sharing)*
- Commonly used in traditional telephone networks

Circuit Switching: FDM vs TDM

Example:

4 users

FDM

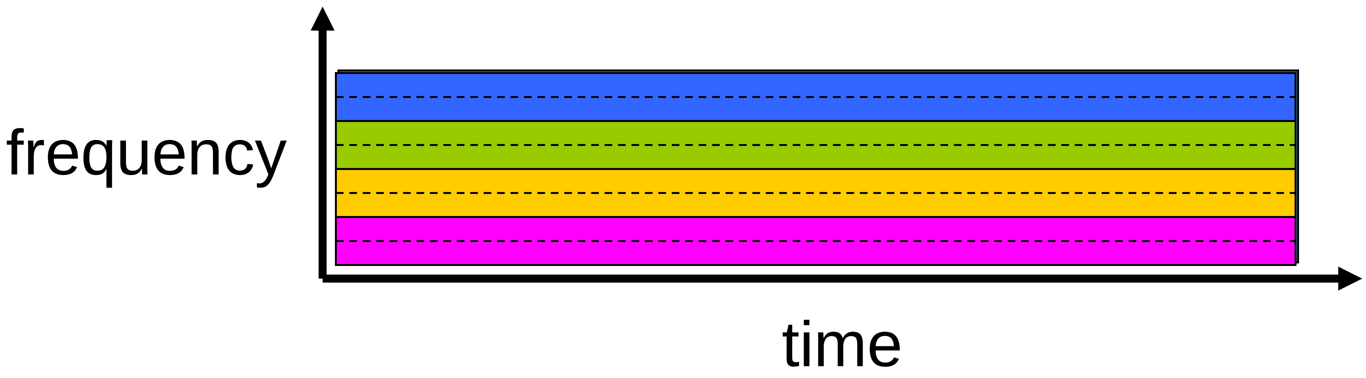

TDM

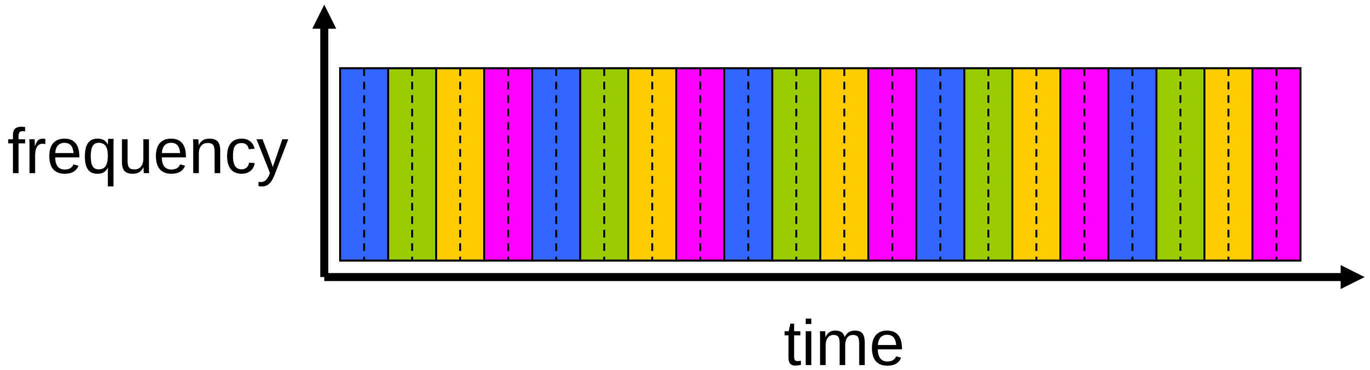

Packet switching versus circuit switching

packet switching allows more users to use network!

example:

- 1 Mb/s link
- each user:
  - 100 kb/s when "active"
  - active 10% of time

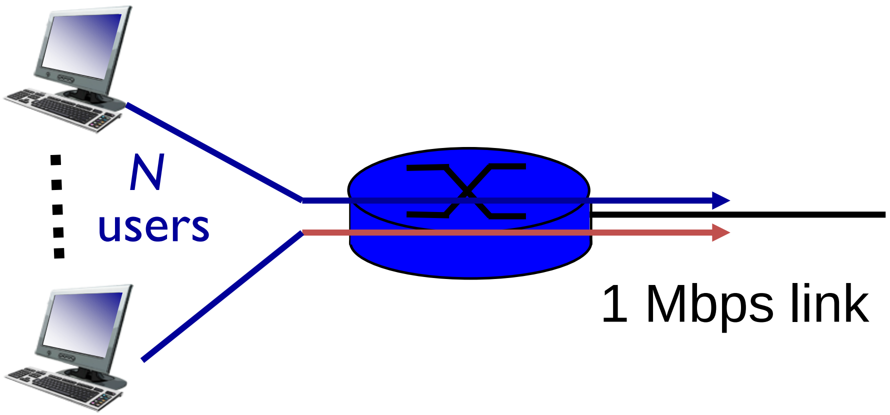

- *circuit-switching:*
  - 10 users
- *packet switching:*
  - with 35 users, probability > 10 active at same time is less than .0004 *

Q: how did we get value 0.0004?

Q: what happens if > 35 users ?

* Check out the online interactive exercises for more examples

Internet Structure

is packet switching a "slam dunk winner?"

- great for bursty data
  - resource sharing
  - simpler, no call setup
- excessive congestion possible: packet delay and loss
  - protocols needed for reliable data transfer, congestion control
- *Q:* How to provide circuit-like behavior?
  - bandwidth guarantees needed for audio/video apps
  - still an unsolved problem (chapter 7)

Q: human analogies of reserved resources (circuit switching) versus on-demand allocation (packet-switching)?

Internet Structure: Network of Networks

- End systems connect to Internet via access ISPs (Internet Service Providers)
  - Residential, company and university ISPs
- Access ISPs in turn must be interconnected.
  - So that any two hosts can send packets to each other
- Resulting network of networks is very complex
  - Evolution was driven by economics and national policies
- Let's take a stepwise approach to describe current Internet structure

Internet Structure: Network of Networks

Question: given *millions* of access ISPs, how to connect them together?

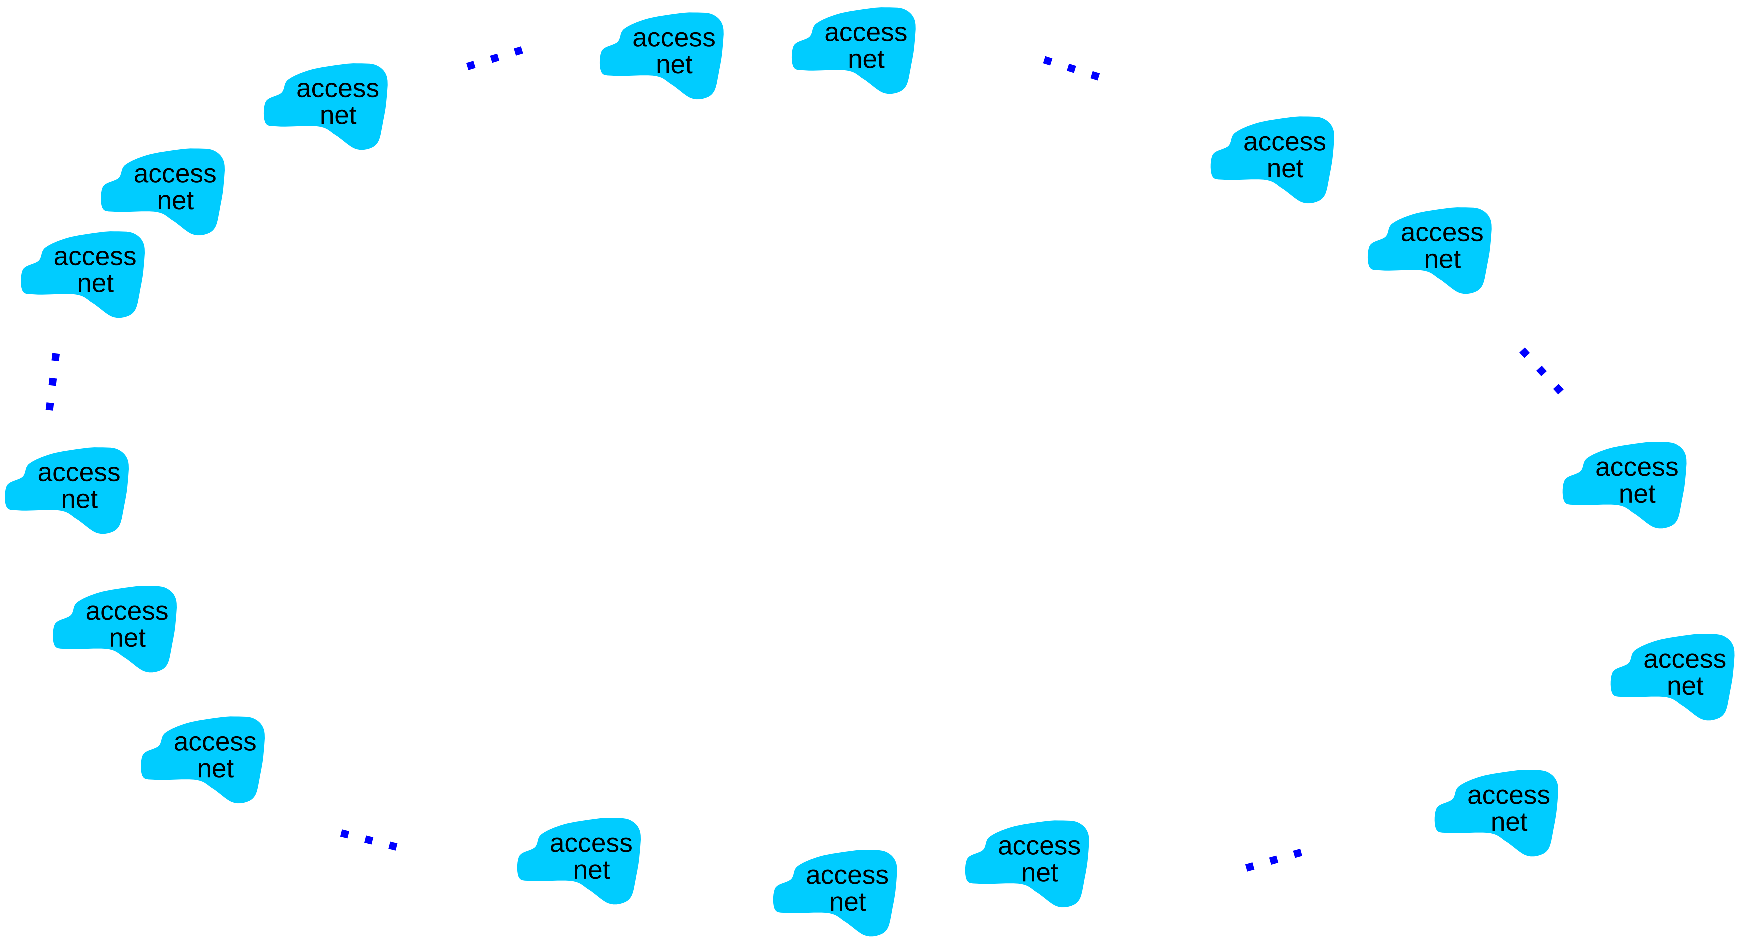

Internet Structure: Network of Networks

Option: connect each access ISP to every other access ISP?

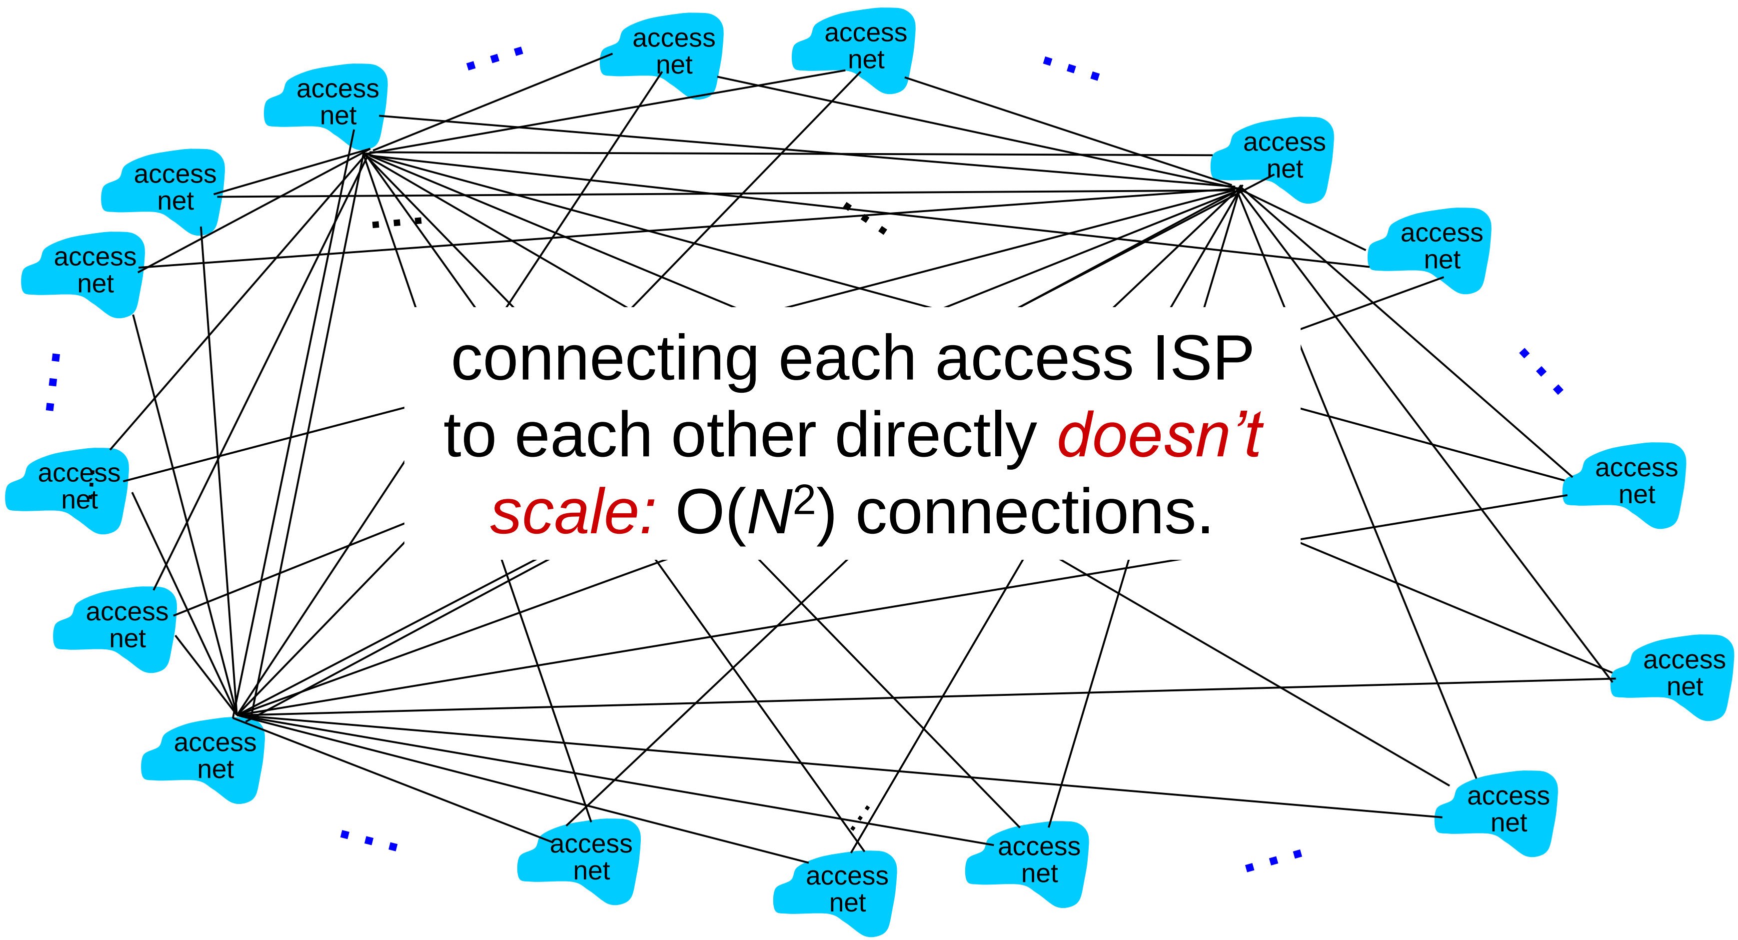

connecting each access ISP to each other directly *doesn't scale:* O(N^2) connections.

Internet Structure: Network of Networks

Option: connect each access ISP to a global transit ISP? Customer and provider ISPs have economic agreement.

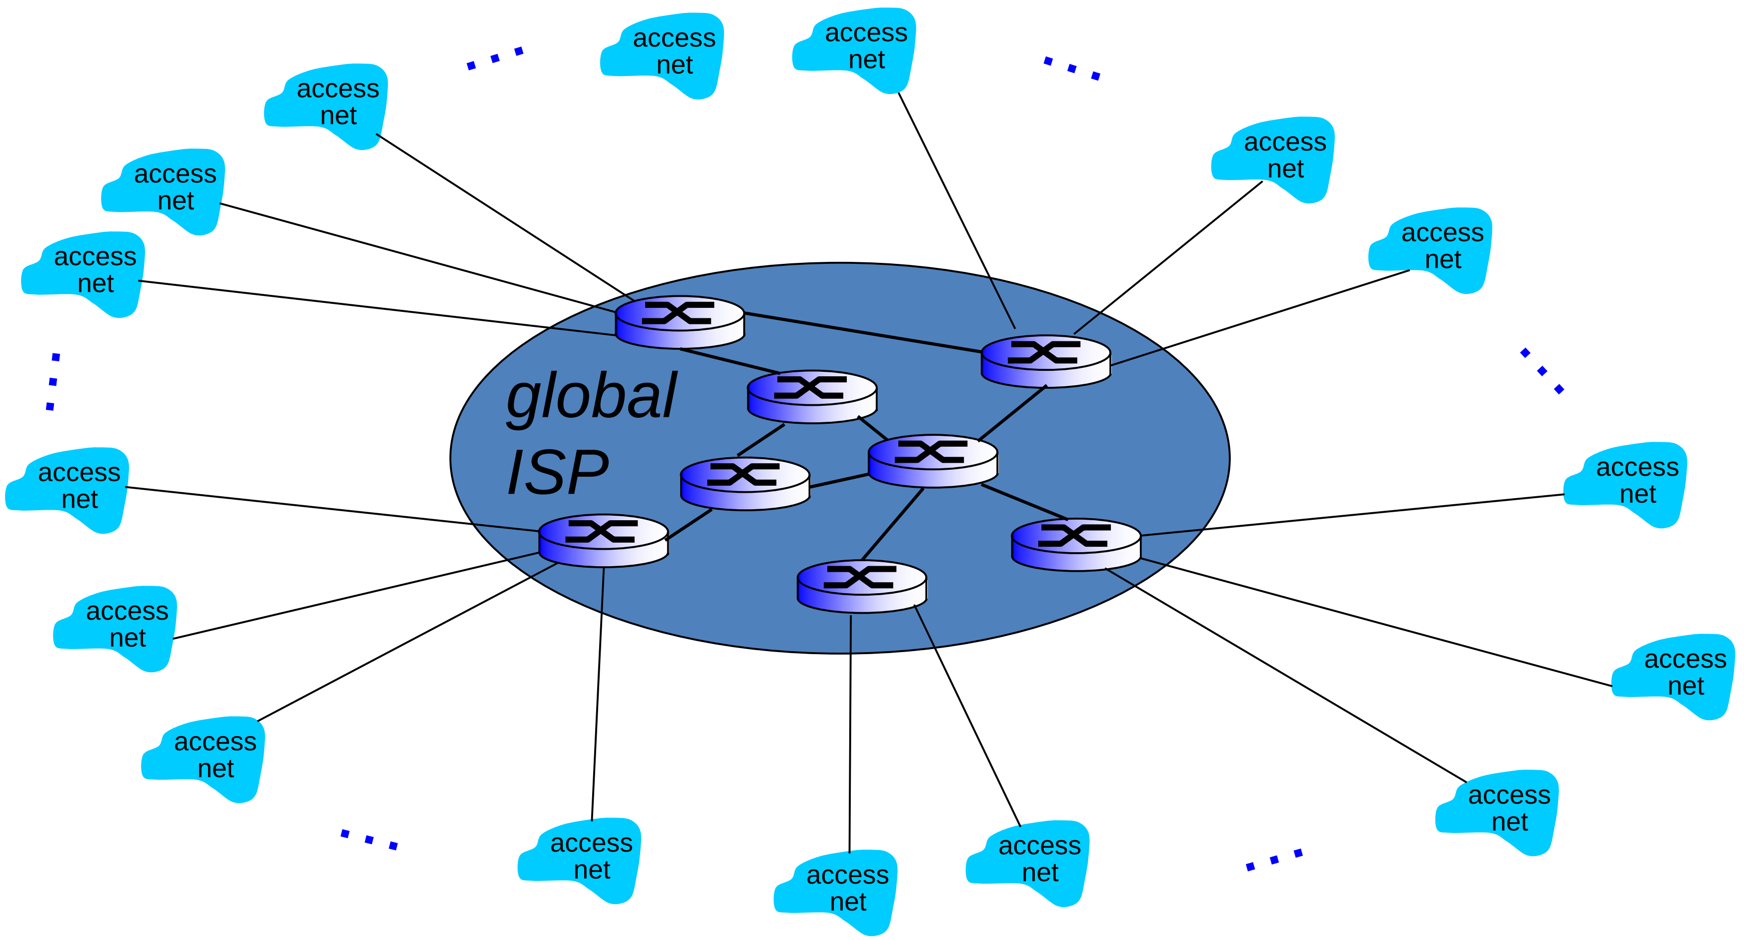

Internet Structure: Network of Networks

But if one global ISP is viable business, there will be competitors ….

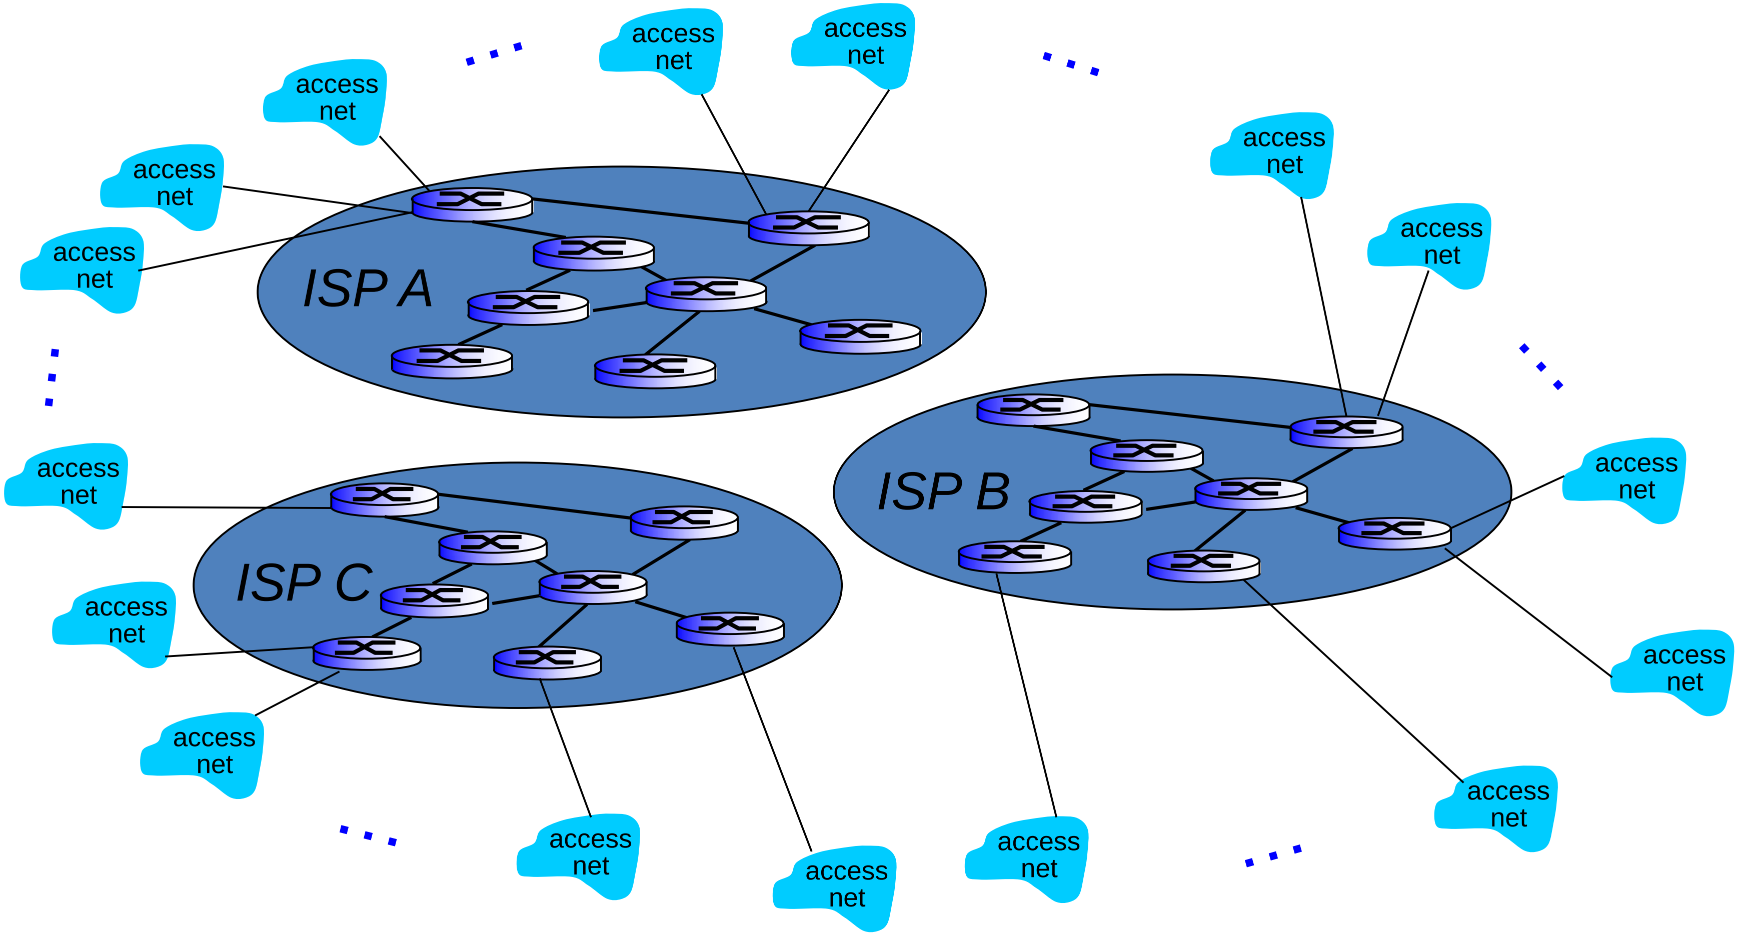

Internet Structure: Network of Networks

But if one global ISP is viable business, there will be competitors …. which must be interconnected

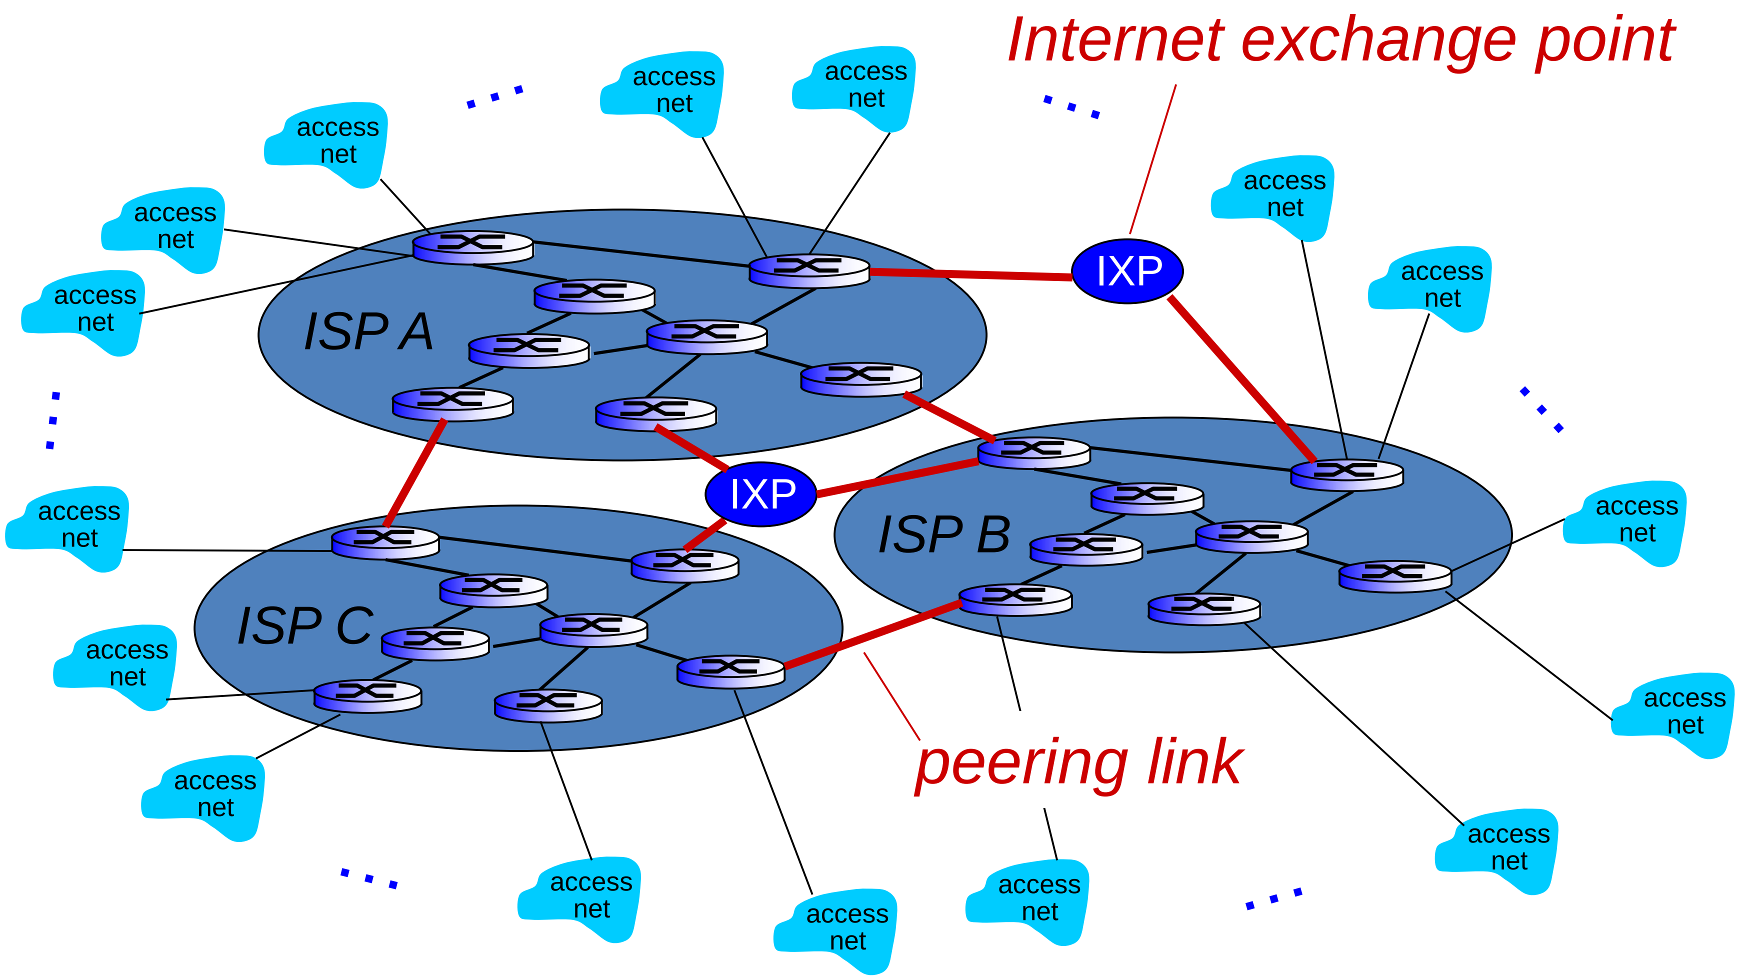

Internet Structure: Network of Networks

… and regional networks may arise to connect access nets to ISPS

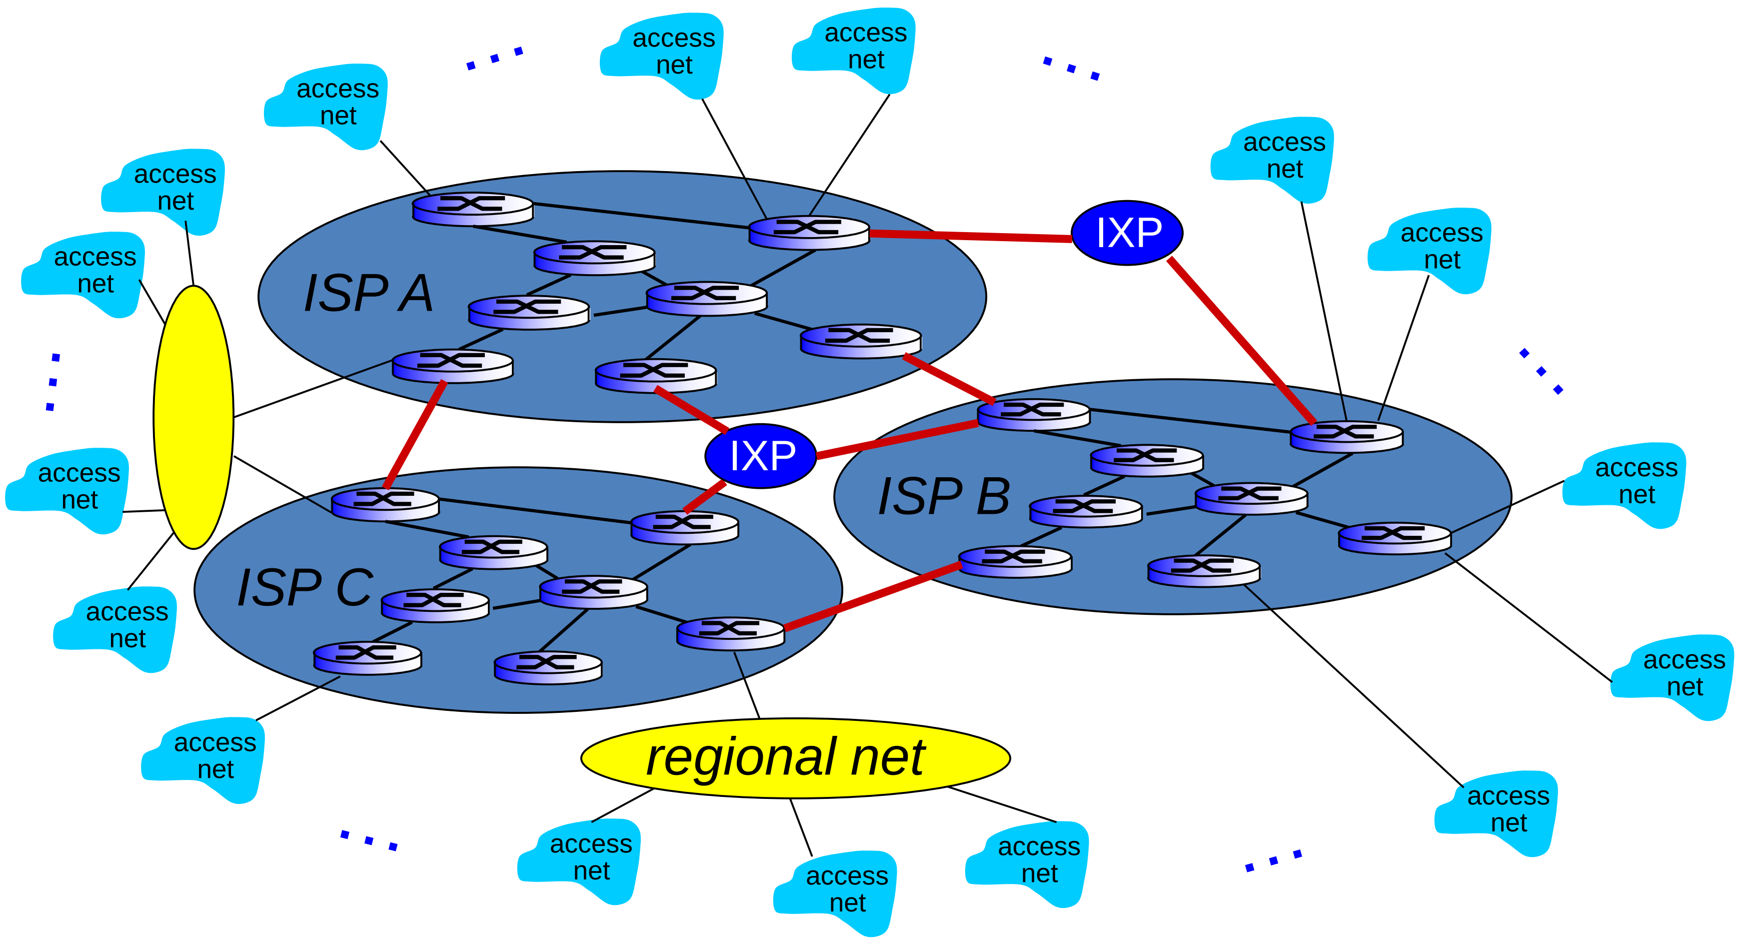

Internet Structure: Network of Networks

… and content provider networks (e.g., Google, Microsoft, Akamai) may run their own network, to bring services, content close to end users

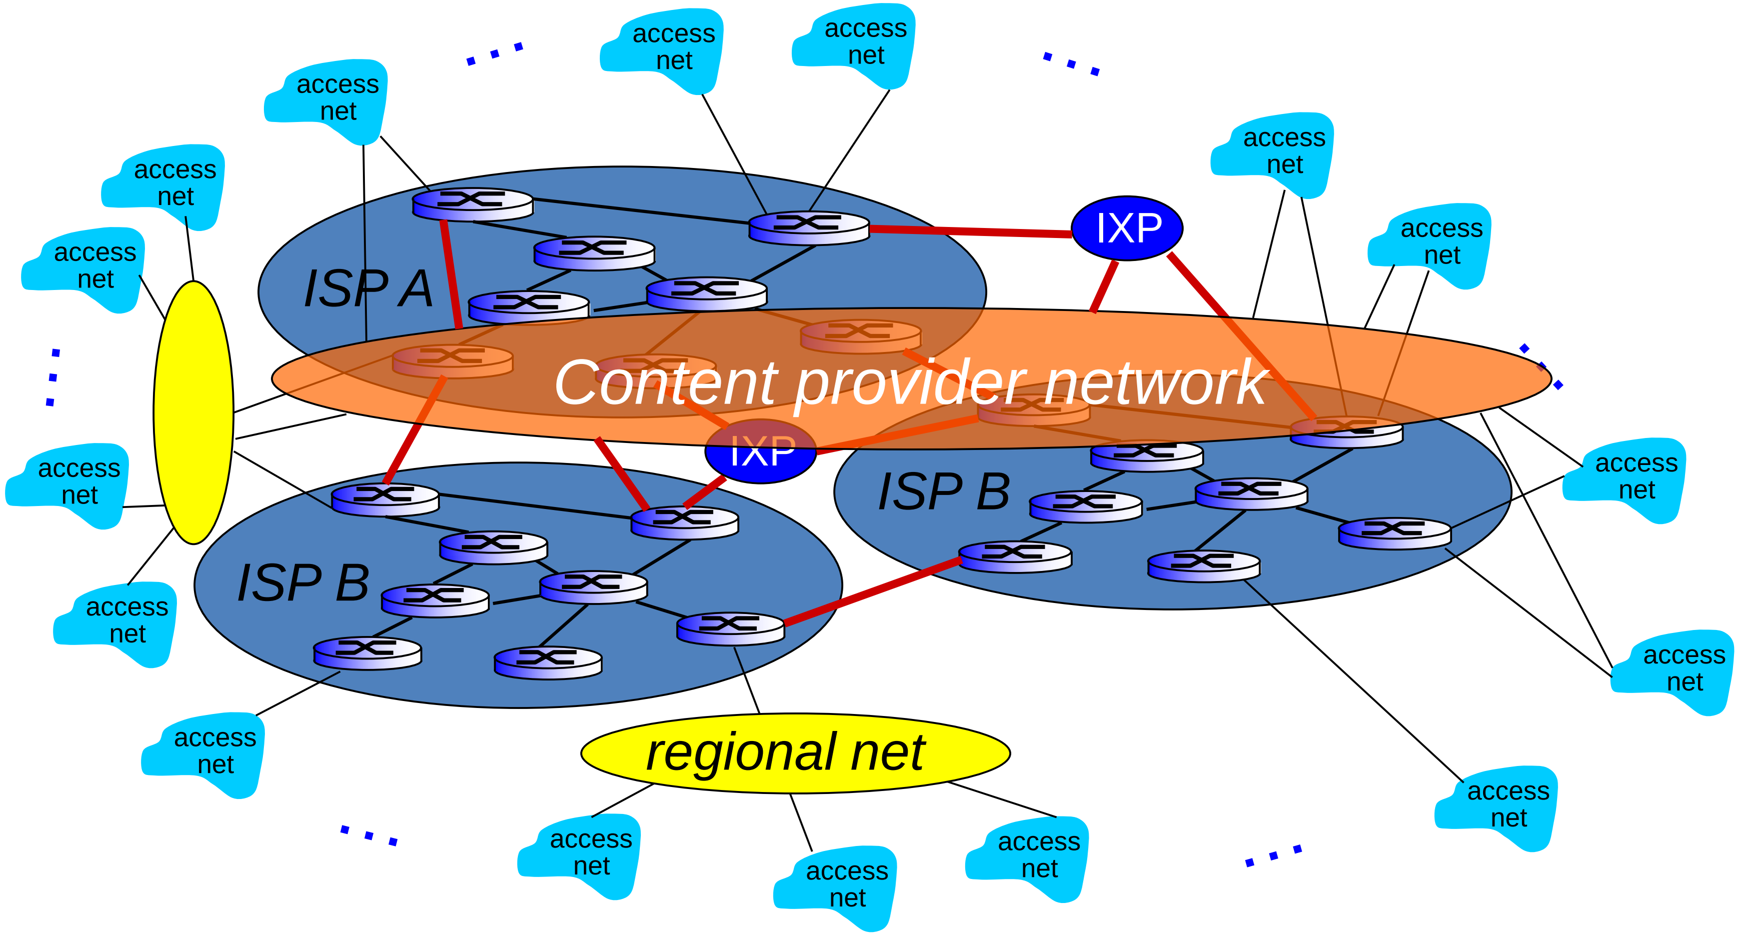

Internet Structure: Network of Networks

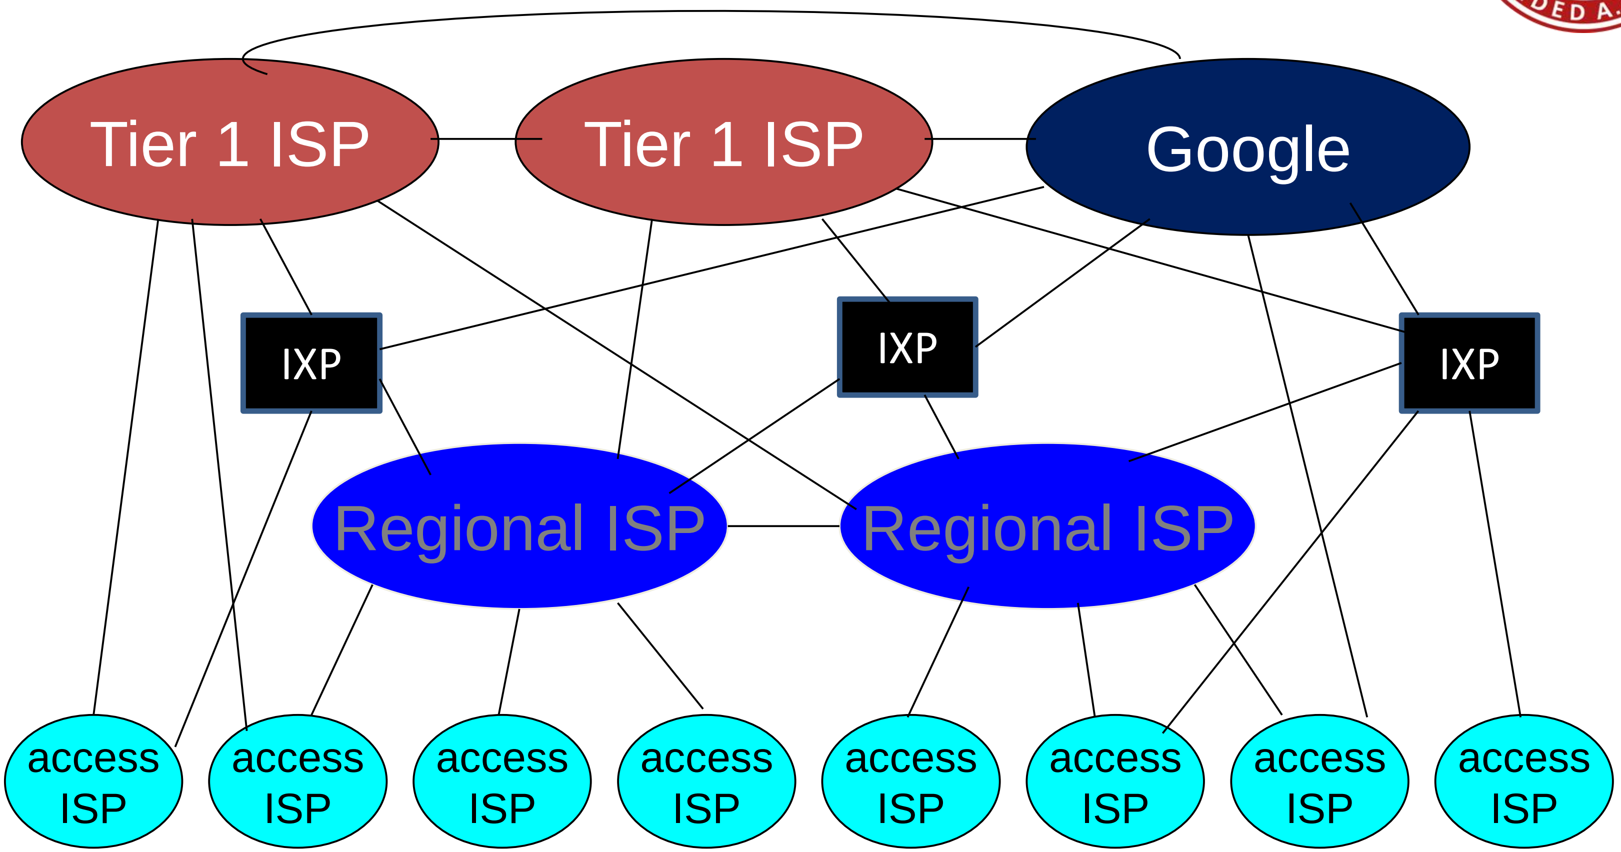

- at center: small # of well-connected large networks
  - "tier-1" commercial ISPs (e.g., Level 3, Sprint, AT&T, NTT), national & international coverage
  - content provider network (e.g, Google): private network that connects it data centers to Internet, often bypassing tier-1, regional ISPs

Internet Structure: Network of Networks

E.g. AT&T

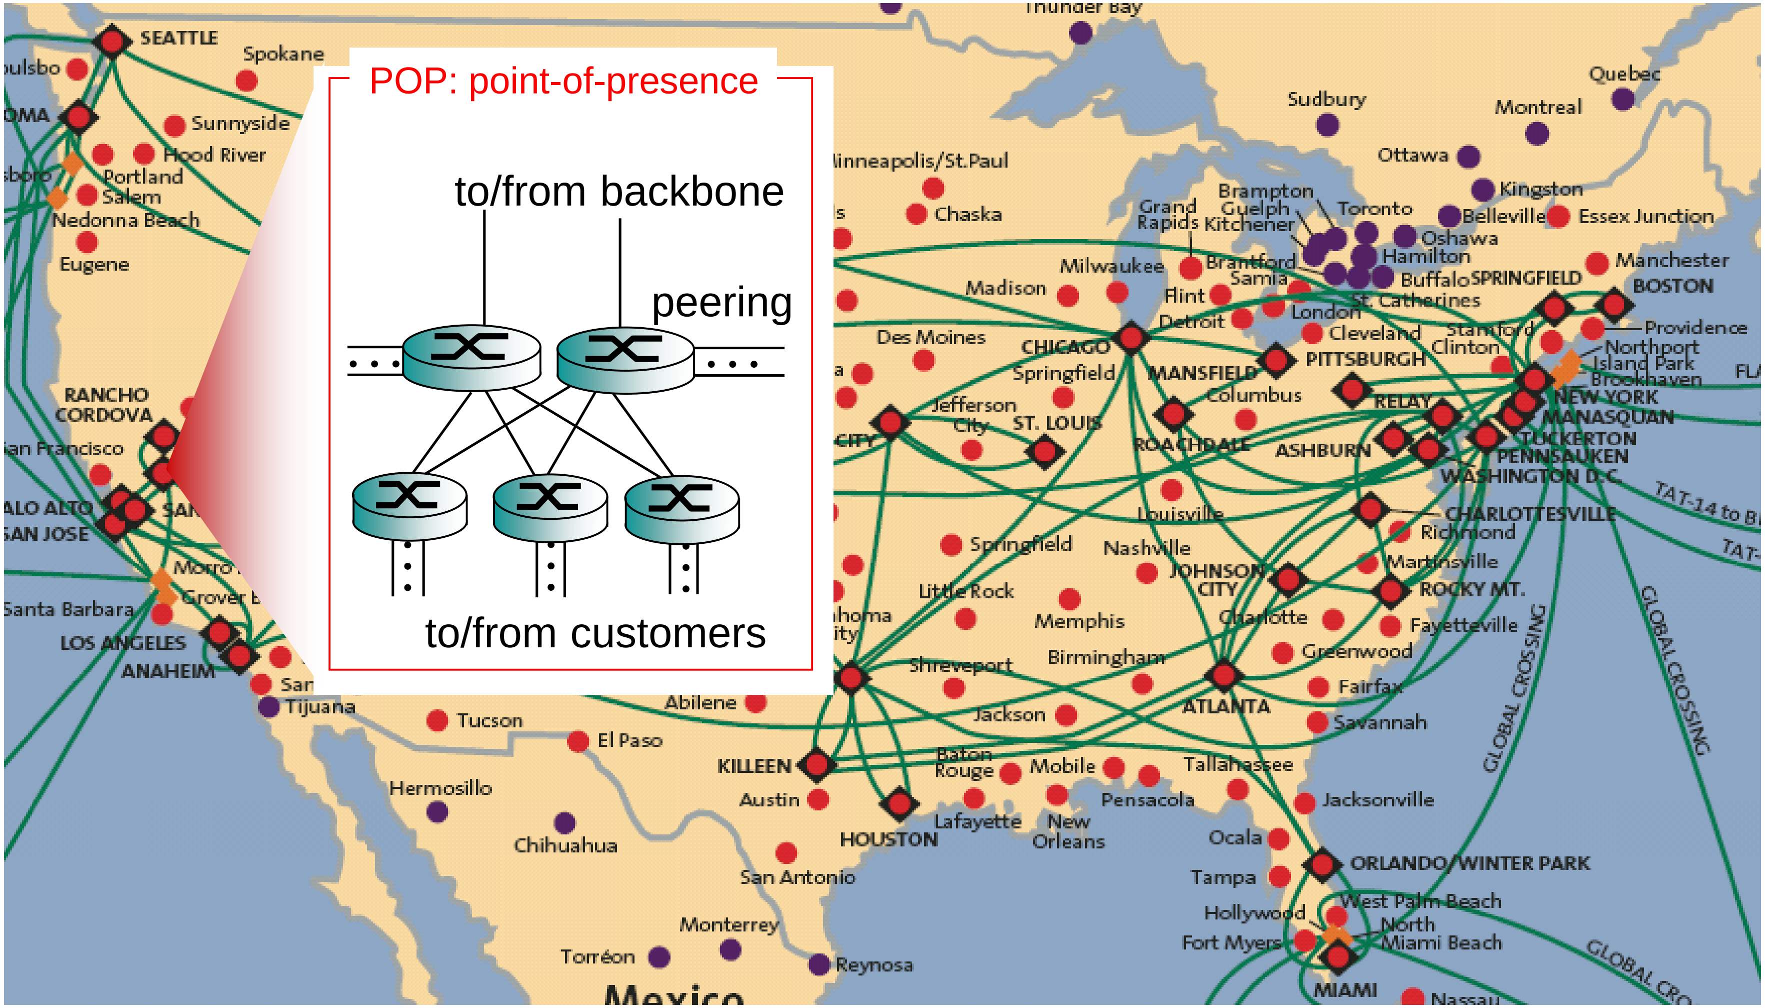

Goals for Today

- Overview
  - What is the Internet?
    - What is it and how did we get here?
  - What is about Data Centers?
- Network Structure
  - Internet
  - Data center

Networking in Data Centers

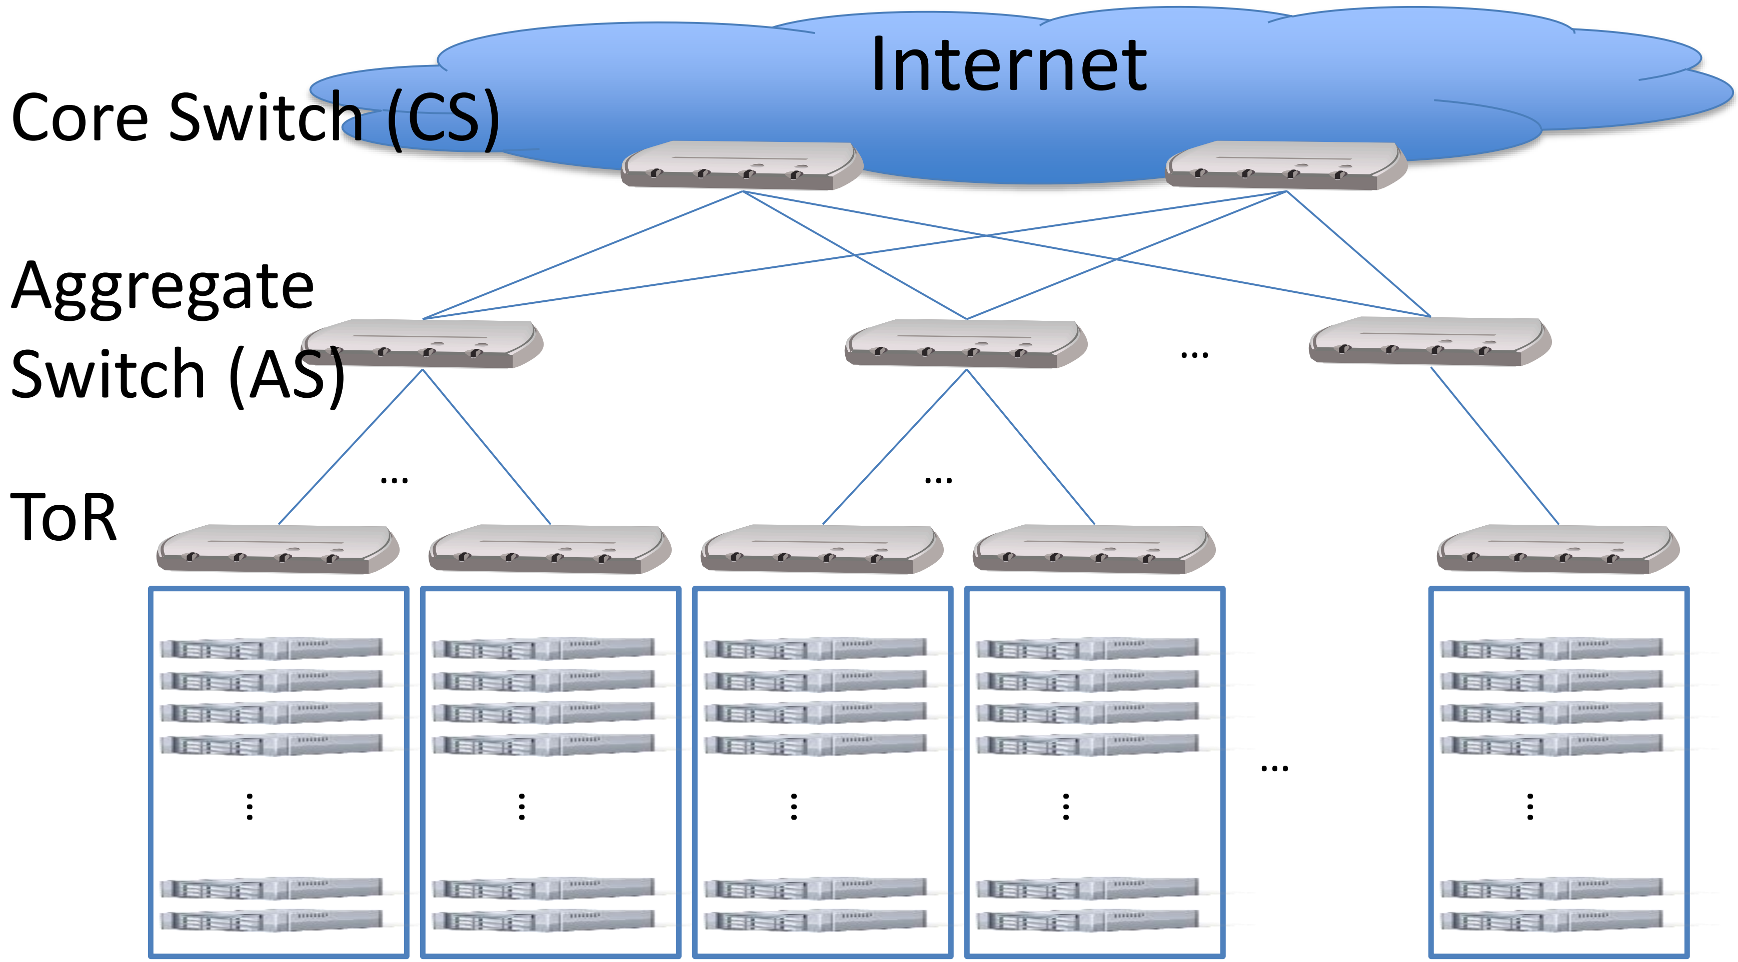

Networking in Data Centers

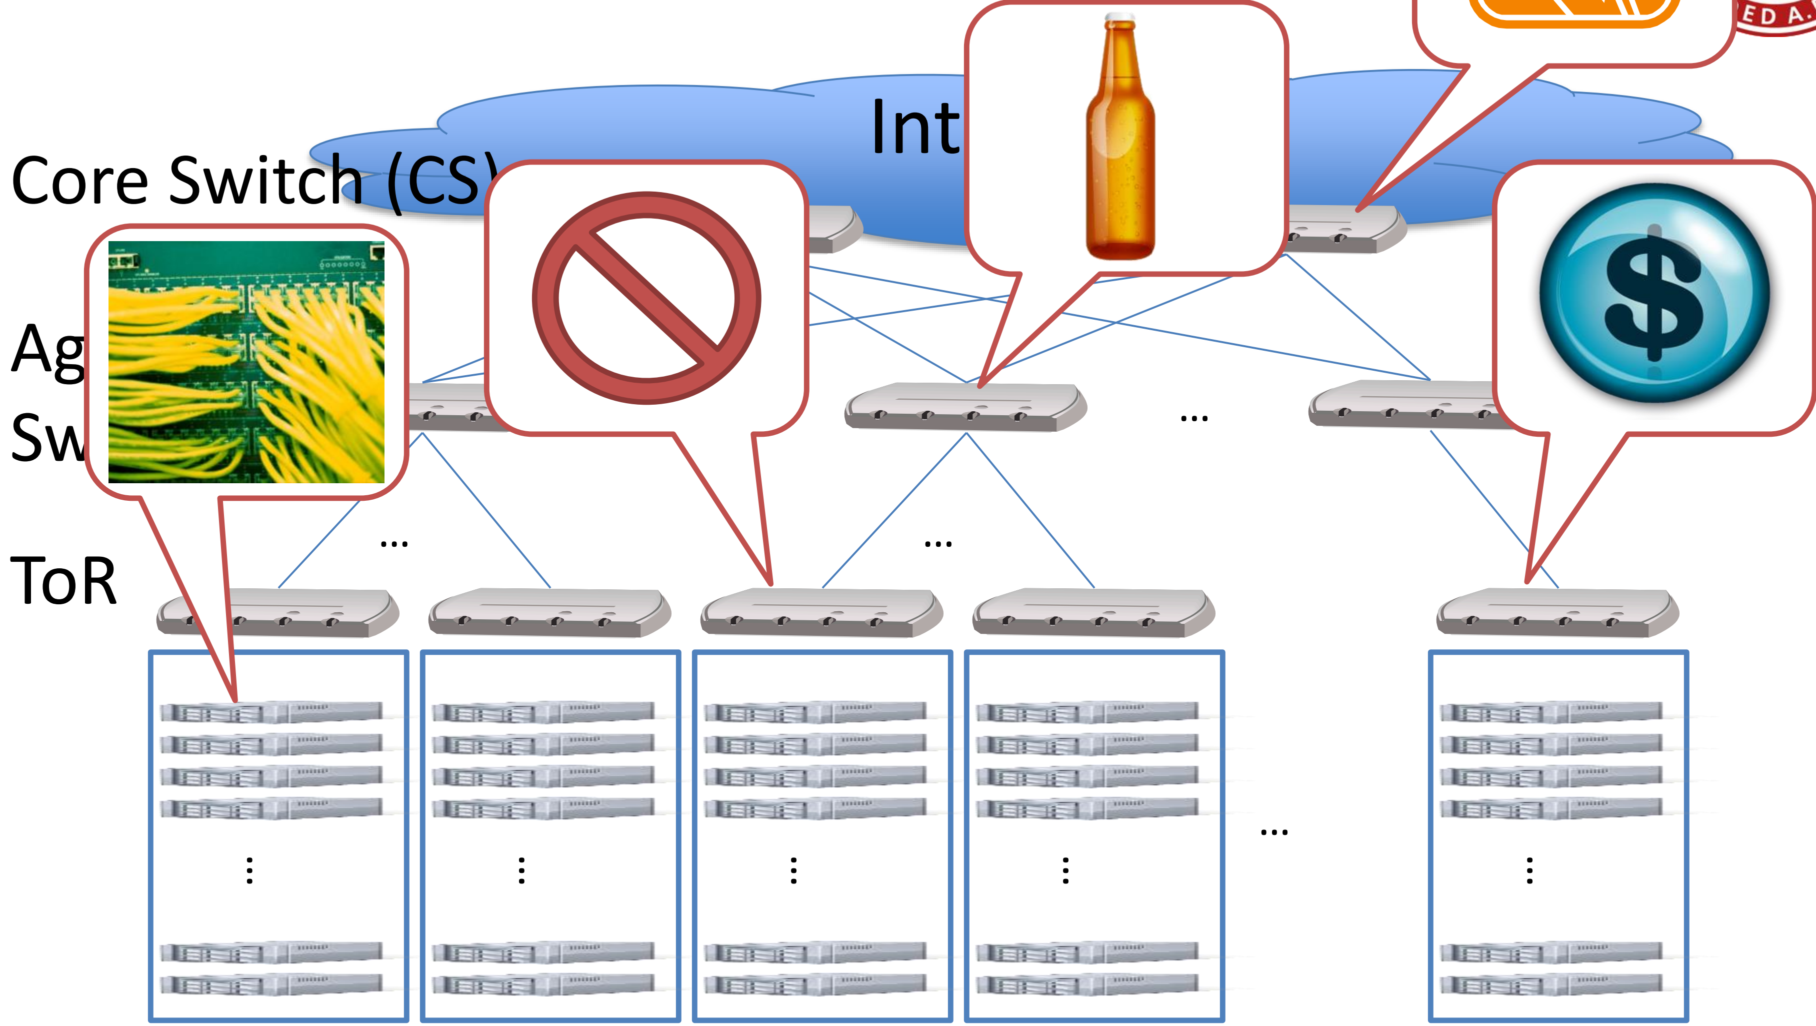

Geo-distributed Data Centers

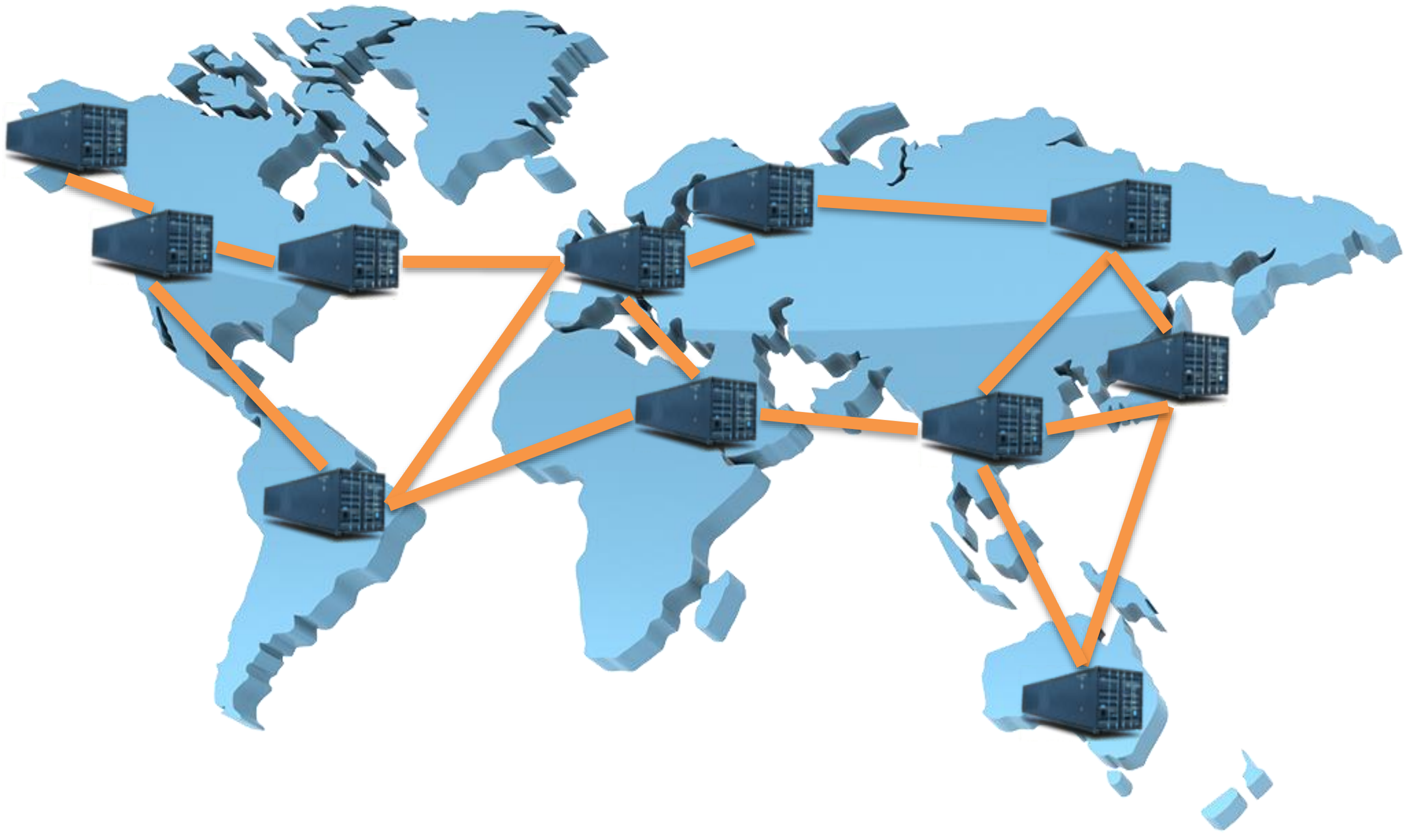

Perspective

- Large cloud service provides have deployed their own networks
  - Private networks, perhaps as large as the Internet
  - But, bypass the Internet core and connect directly with ISPs
  - Near instantaneous access betwee consumers and data centers
- Economies of scale dominate in cloud data centers

Before Next time

- No required reading and review due
- But, review chapter 2 from the book, Application Layer

- Create a project group
  - Start asking questions about possible projects

- Check website for updated schedule