

Lecture 3: Internet Edge, Core, and Structure

COMP 332, Fall 2018

Victoria Manfredi

WESLEYAN
UNIVERSITY



Acknowledgements: materials adapted from Computer Networking: A Top Down Approach 7th edition: ©1996-2016, J.F Kurose and K.W. Ross, All Rights Reserved as well as from slides by Abraham Matta at Boston University and some material from Computer Networks by Tannenbaum and Wetherall.

Today

1. Announcements

- homework 1 due Wed. by 11:59p

2. Internet organization

- edge vs. core
- Internetwork: network of networks
- IP addresses

How do you connect to Internet?

Hosts connect to edge router


- access network/ISP

Access networks

- **residential**
 - DSL (telephone), cable,
- **institutional**
 - school, company
- **mobile**

Issues

- bandwidth (bps) of access network?
- shared or dedicated?



ISP: Internet Service Provider

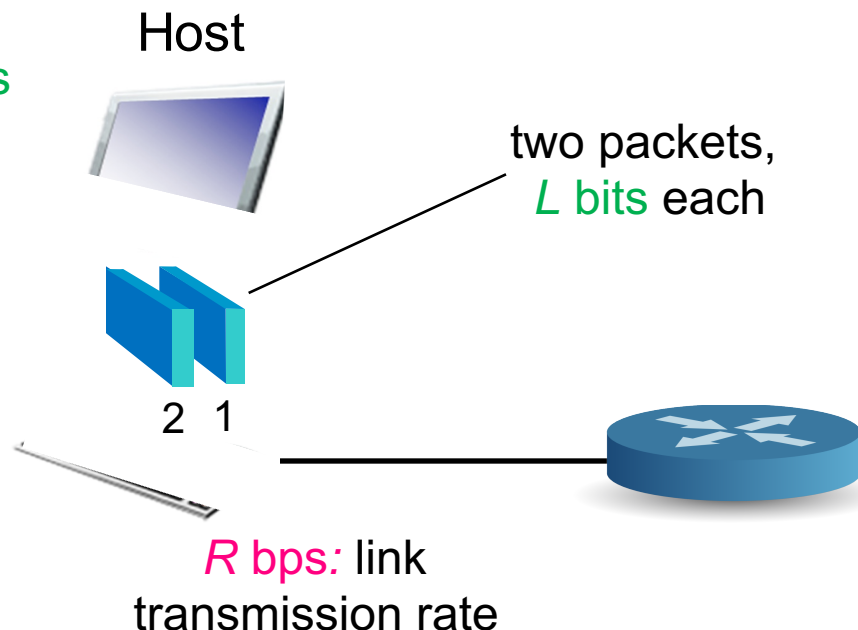
Host sends packets into access network

1. Given application message

- breaks into packets
 - smaller chunks of length L bits

2. Transmit packets into access network

- at transmission rate R
 - aka link capacity
 - aka link bandwidth



$$\text{Transmission delay} = \text{Time to transmit } L\text{-bit packet into link} = \frac{L \text{ (bits)}}{R \text{ (bits/sec)}}$$

Internet Organization

CORE

Internet Organization

EDGE

How to move data through Internet core?

Internet core

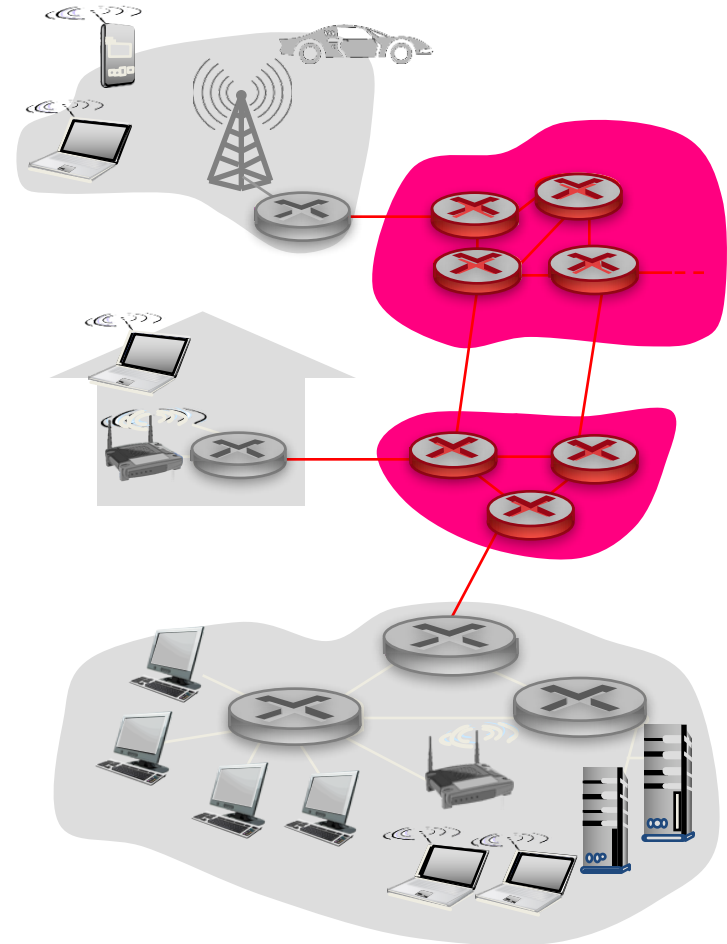
- mesh of interconnected routers

Option 1: Packet-switching

- on-demand resource allocation
- best effort service
- good bandwidth use

Option 2: Circuit-switching


- reserved resources
- guaranteed service
- may waste bandwidth



Q: human analogies of reserved versus on-demand allocation

Packet switching

1. Hosts break app-layer messages into **packets**



2. Time to transmit (push out)
 L -bit packet into R bps link:
 L / R seconds


Example

- $L = 7.5$ Mbits
- $R = 1.5$ Mbps
- 1-hop transmission delay = 5s

Packet switching

1. Hosts break app-layer messages into **packets**

3. Store-and-forward: entire packet must arrive at router before it can be transmitted on next link



2. Time to transmit (push out)
 L -bit packet into R bps link:
 L / R seconds

4. Time to transmit (push out)
 L -bit packet into R bps link:
 L / R seconds

End-end transmission delay = $2 L / R$ seconds
(assuming zero propagation, queuing, processing delay...)


Circuit switching

End-end resources allocated

- reserved for “call” between source & dest
- commonly used in traditional telephone networks

Dedicated resources

- no sharing
- circuit-like (guaranteed) performance
- circuit segment idle if not used by call (no sharing)


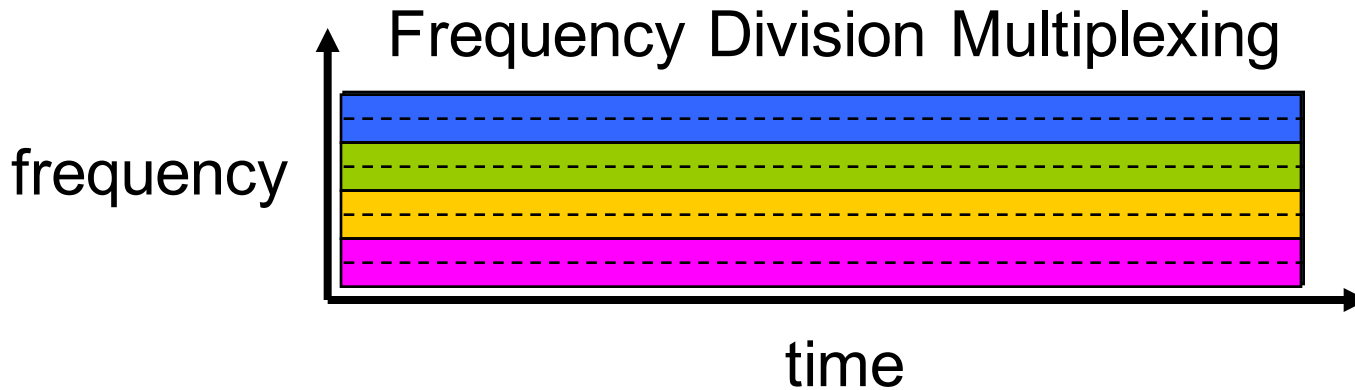


Q: what is end-end transmission delay? L / R seconds

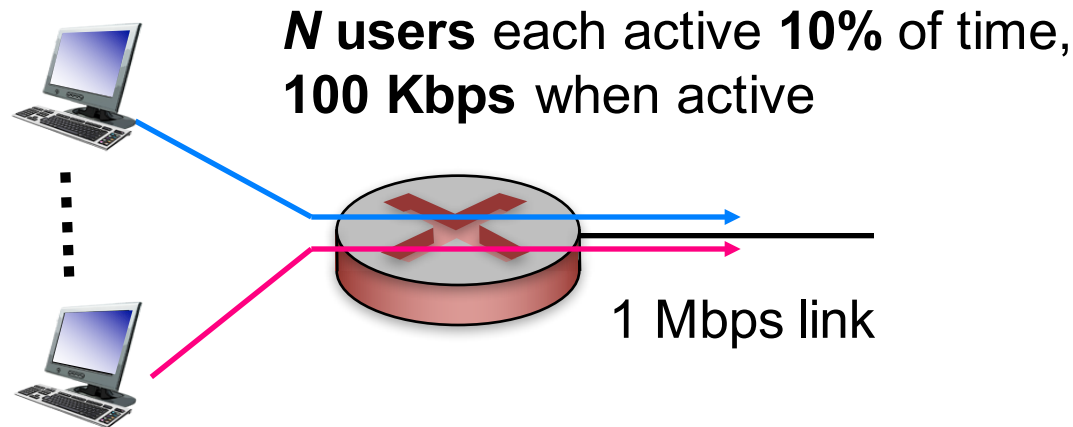
Q: what happens if there is a lull in conversation?

Circuit switching

Multiplexing 4 users



Packet switching versus circuit switching



Q: How many users can be supported?

Circuit switching

- 1 Mbps / 100 Kbps = 10
- N = 10 users

Packet switching

- N = 35 users
- prob > 10 users active simultaneously < .0004

Q: how did we get value 0.0004?

Q: what happens if > 35 users ?

Packet switching allows more users to use network!

Is packet switching always better?

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)

Binomial random variable (homework)

Suppose we do **n independent** experiments

- each experiment succeeds with **probability p**
- each experiment fails with **probability 1-p**

Independent experiment

- knowledge about one experiment occurring does not affect probability of other experiment occurring: e.g., coin toss

$$P(A \text{ and } B) = P(A) \times P(B)$$

$$P(A \text{ or } B) = P(A) + P(B)$$

$$P(X=4 \text{ and } X=5) = P(X=4) \times P(X=5)$$

$$P(X=4 \text{ or } X=5) = P(X=4) + P(X=5)$$

Binomial random variable (homework)

Suppose we do **n independent** experiments

- each experiment succeeds with **probability p**
- each experiment fails with **probability 1-p**

X = Random Variable indicating # of successes that occur in **n experiments**

Probability of i successes

$$P(X = i) = \binom{n}{i} p^i (1-p)^{n-i}$$

Probability of $n-i$ failures

n choose i
different ways to get i
successful experiments

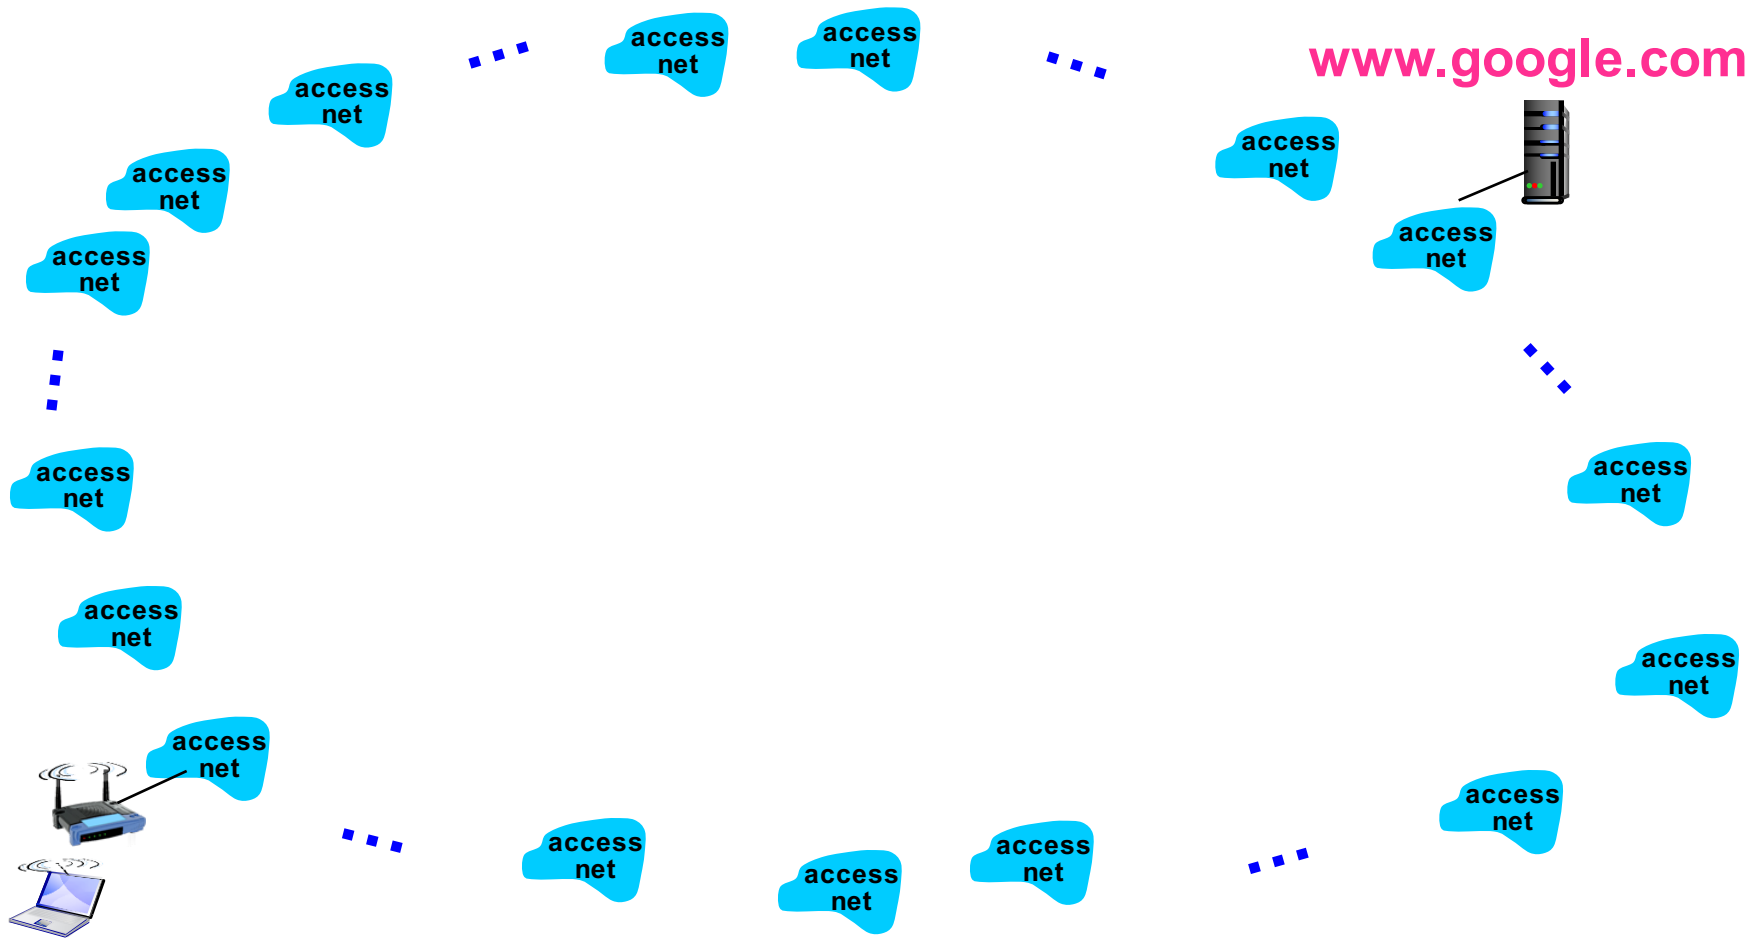
Internet Organization

A NETWORK OF NETWORKS

A network of networks

Hosts connect to Internet via access ISPs (e.g., Comcast)

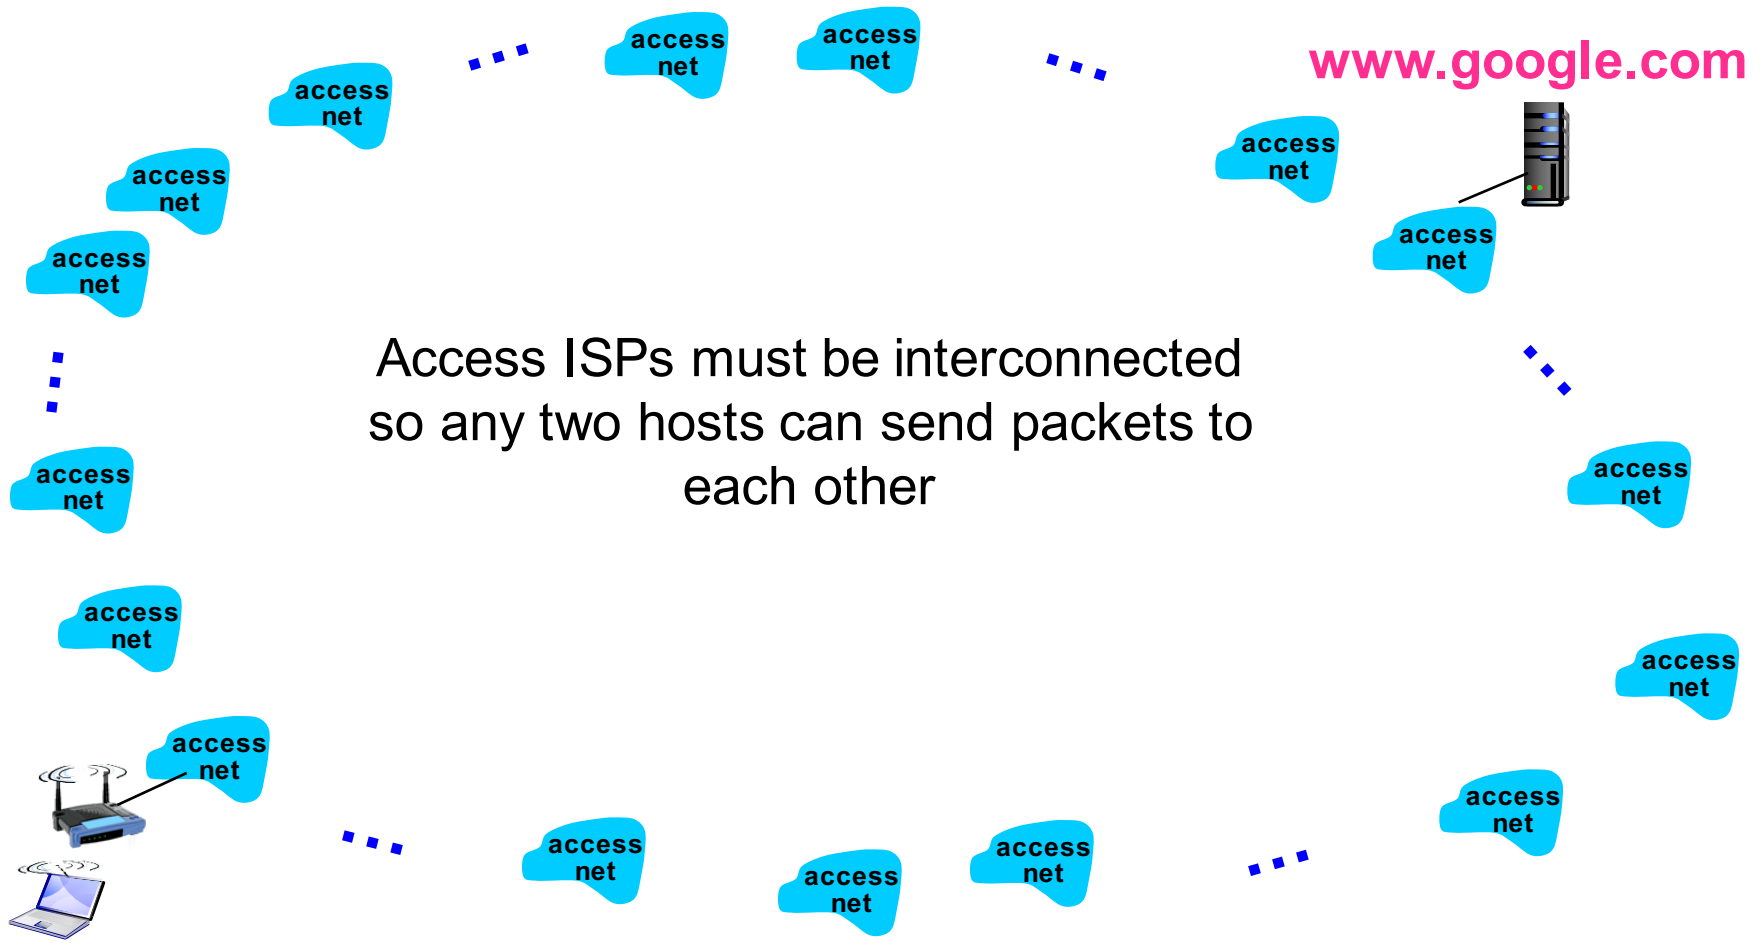
- residential, company and university ISPs



Your host

A network of networks

Q: given millions of access ISPs, how to connect together?

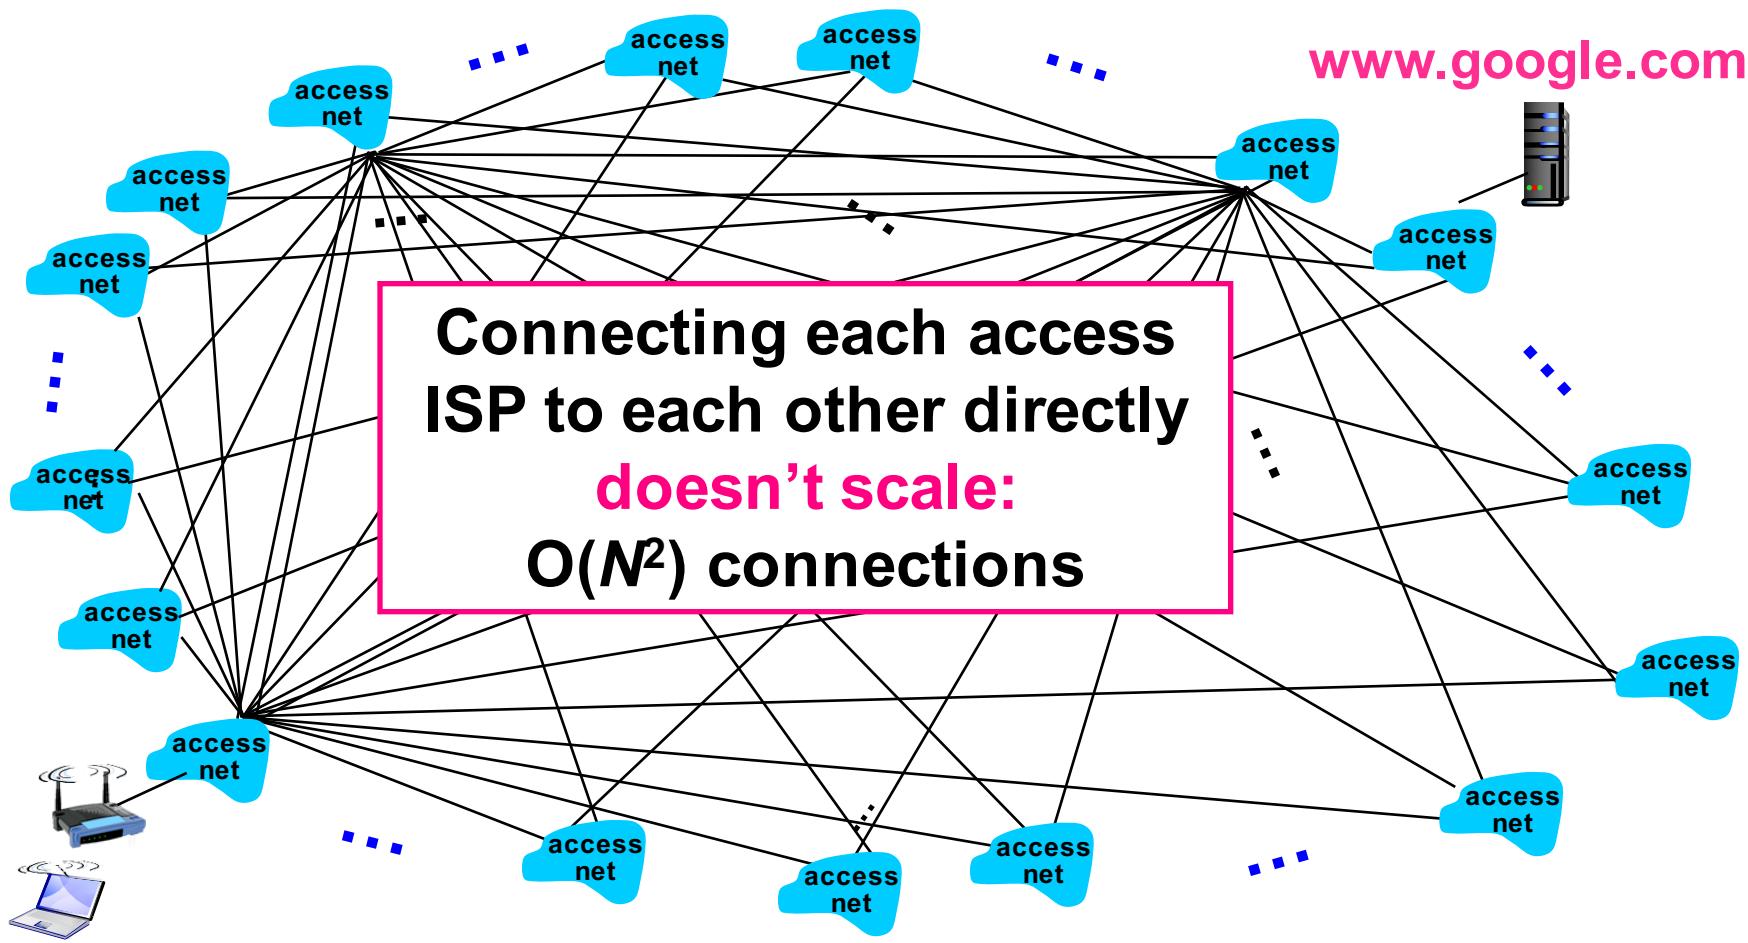


Your host

vumanfredi@wesleyan

A network of networks

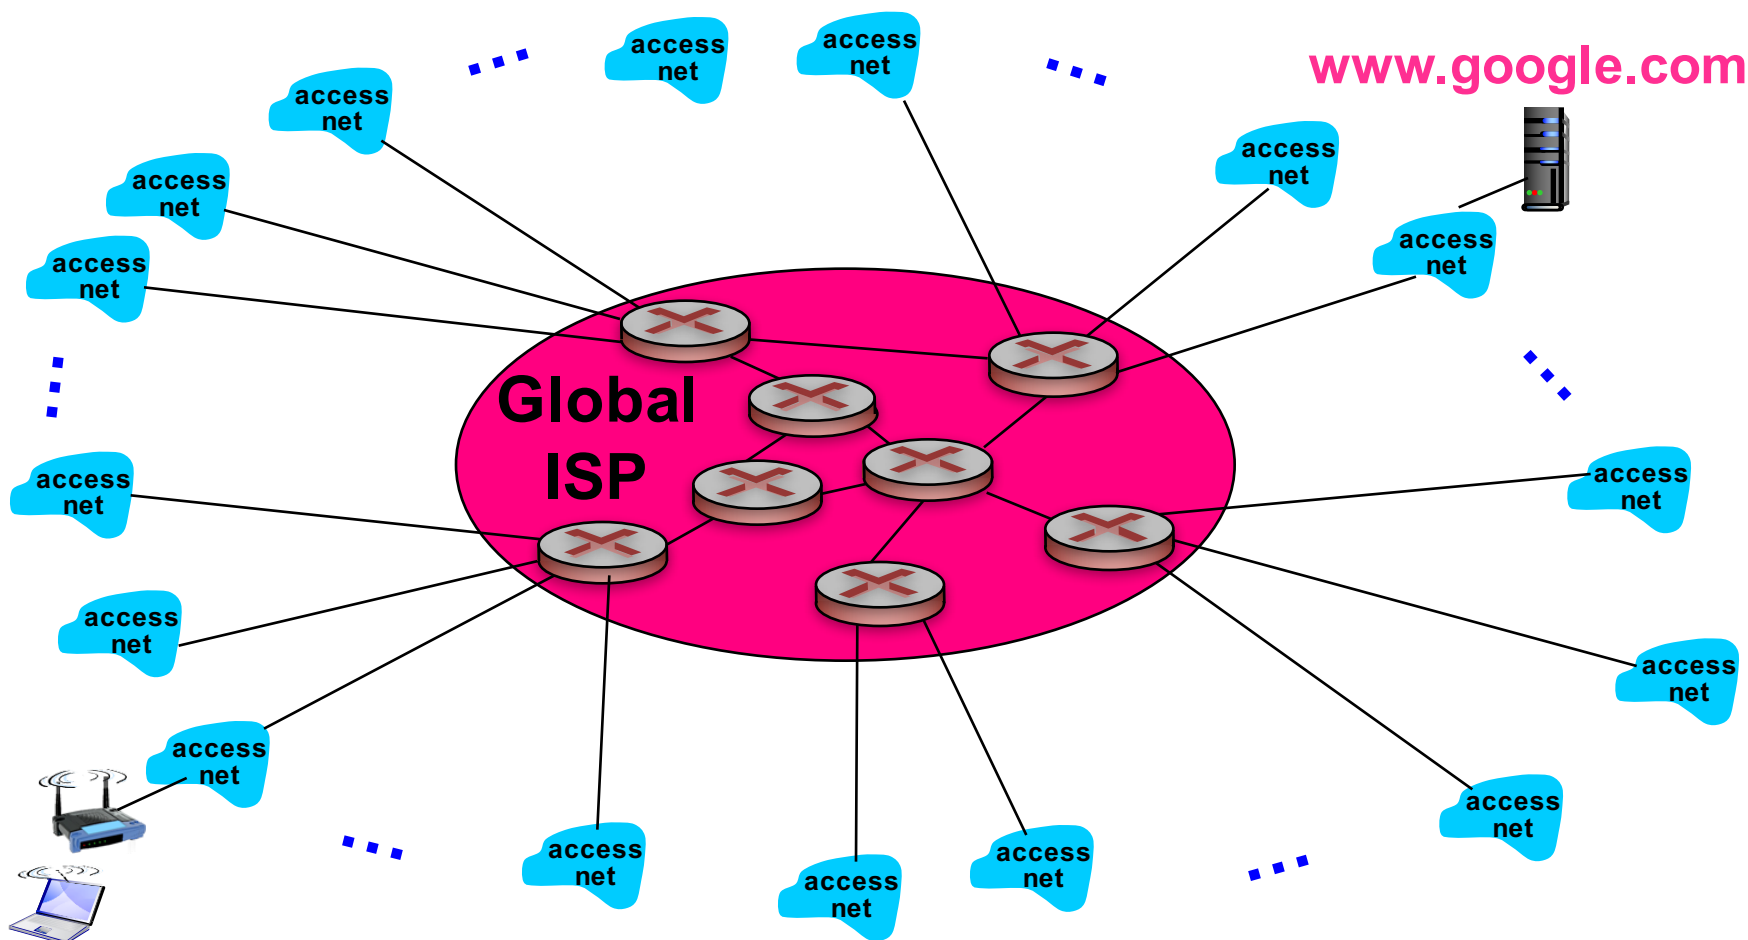
Option 1: connect each access ISP to every other access ISP?



A network of networks

Option 2: connect each access ISP to one global transit ISP?


Customer and provider ISPs have economic agreement.



Your host

A network of networks

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




IXP: Internet Exchange Point

Room full of routers

- meeting point where **multiple ISPs** can peer together

Peer

- let Internet traffic cross/transit your computer network without fee




By Fabienne Serriere [CC BY-SA 3.0
(<http://creativecommons.org/licenses/by-sa/3.0/>)],
via Wikimedia Commons

Amsterdam IXP

- where **100s** of ISPs connect
- optical fiber patch panel connecting different ISPs

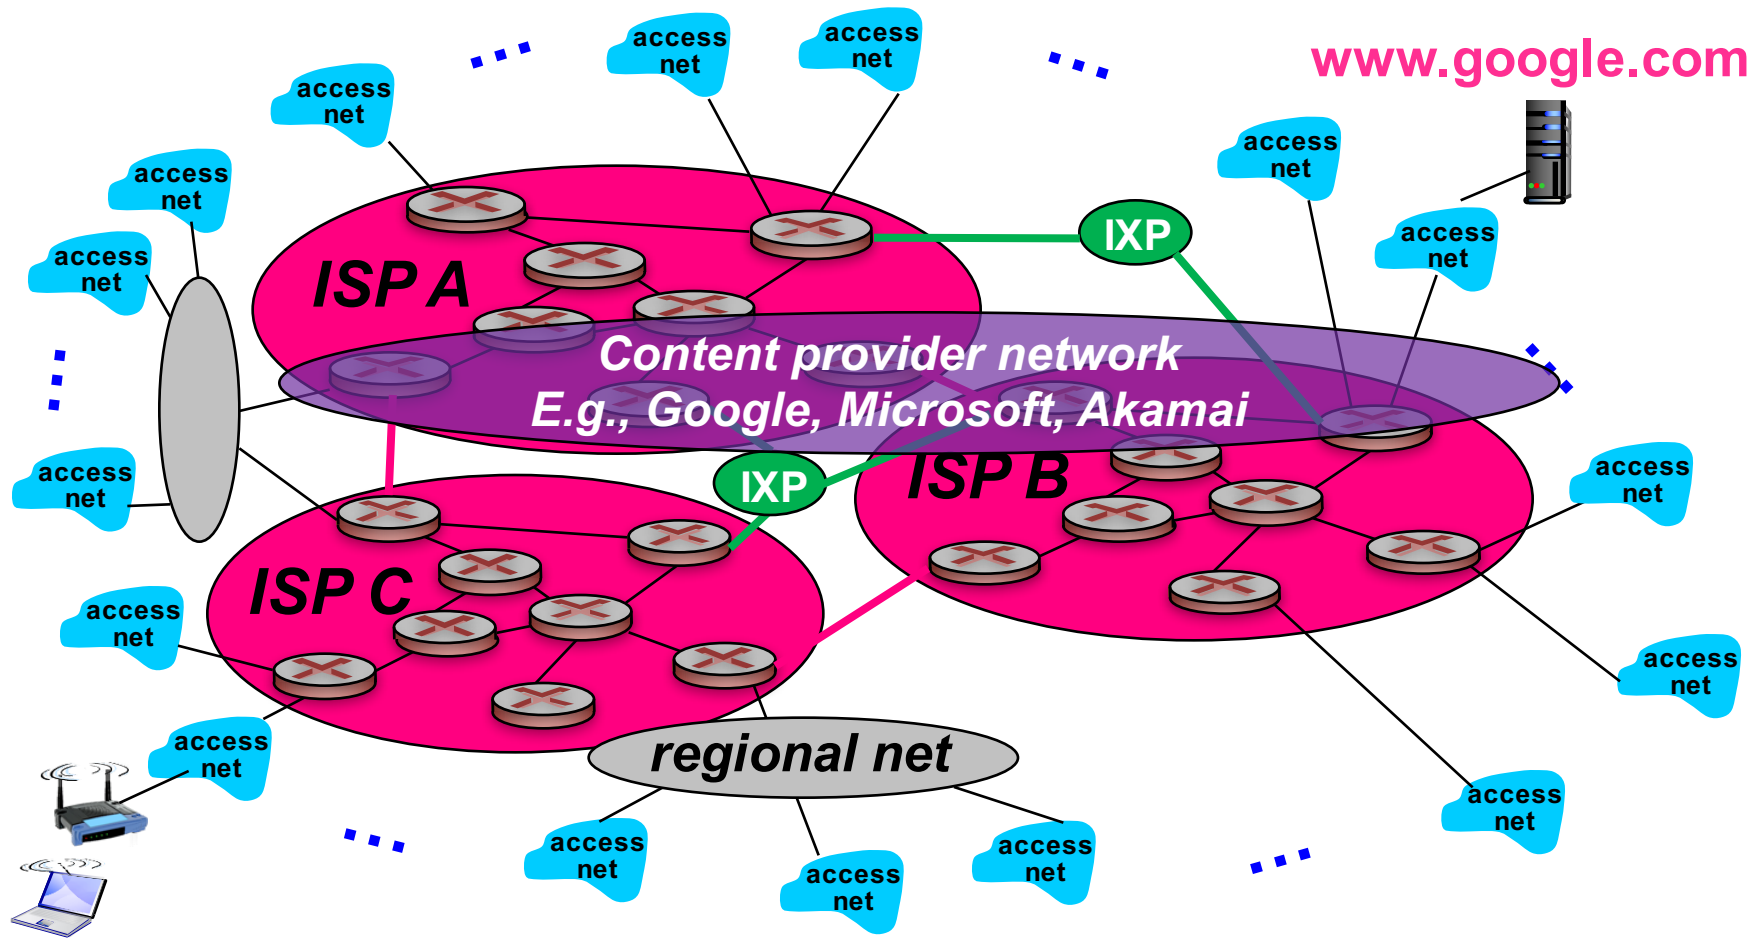
A network of networks

... and regional networks may arise to connect access nets to ISPs



A network of networks

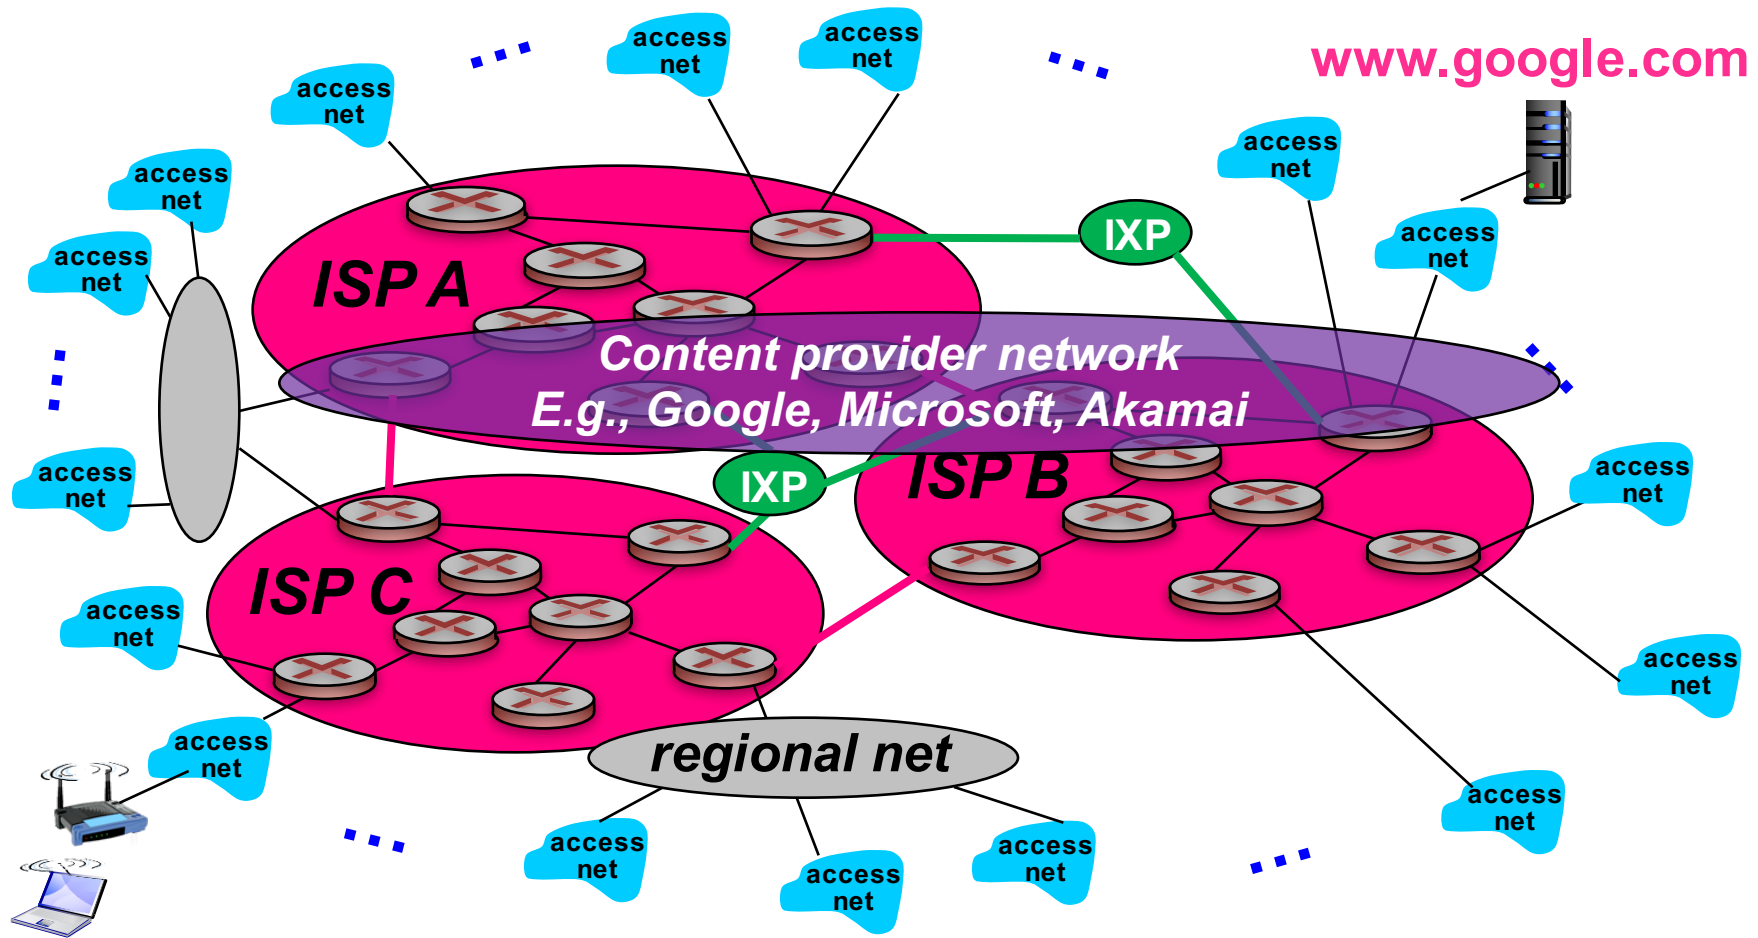
... and **content provider networks** may run their own network, to bring services, content close to end users



Your host

A network of networks

Resulting network of networks is very complex:
evolution was driven by economics and national policies



Your host

Different kinds of ISPs

Tier-1 commercial ISPs: e.g., Level 3, Sprint, AT&T, NTT

- national & international coverage, peer with other tier 1 ISPs
- can reach all of Internet via peering only, peering typically payment free

Content provider network: e.g., Google (YouTube benefits)

- private network that connects its data centers to Internet
- often bypasses tier-1, regional ISPs, may buy transit
- by not paying provider ISP, save money and better control QoS for traffic

Regional ISP: e.g., Comcast

- customer ISP of Tier 1ISP, provider ISP to access ISP
- peers with some networks but needs to purchase some IP transit to reach some parts of Internet

Access ISPs: connect end systems to Internet

- any of these could be access ISP
- company or Wesleyan may connect directly into Tier 1 or Regional ISP

Internet terminology

Internet transit

- service of letting Internet traffic cross or transit a computer network
- usually used to connect a smaller ISP to larger Internet

Multi-home

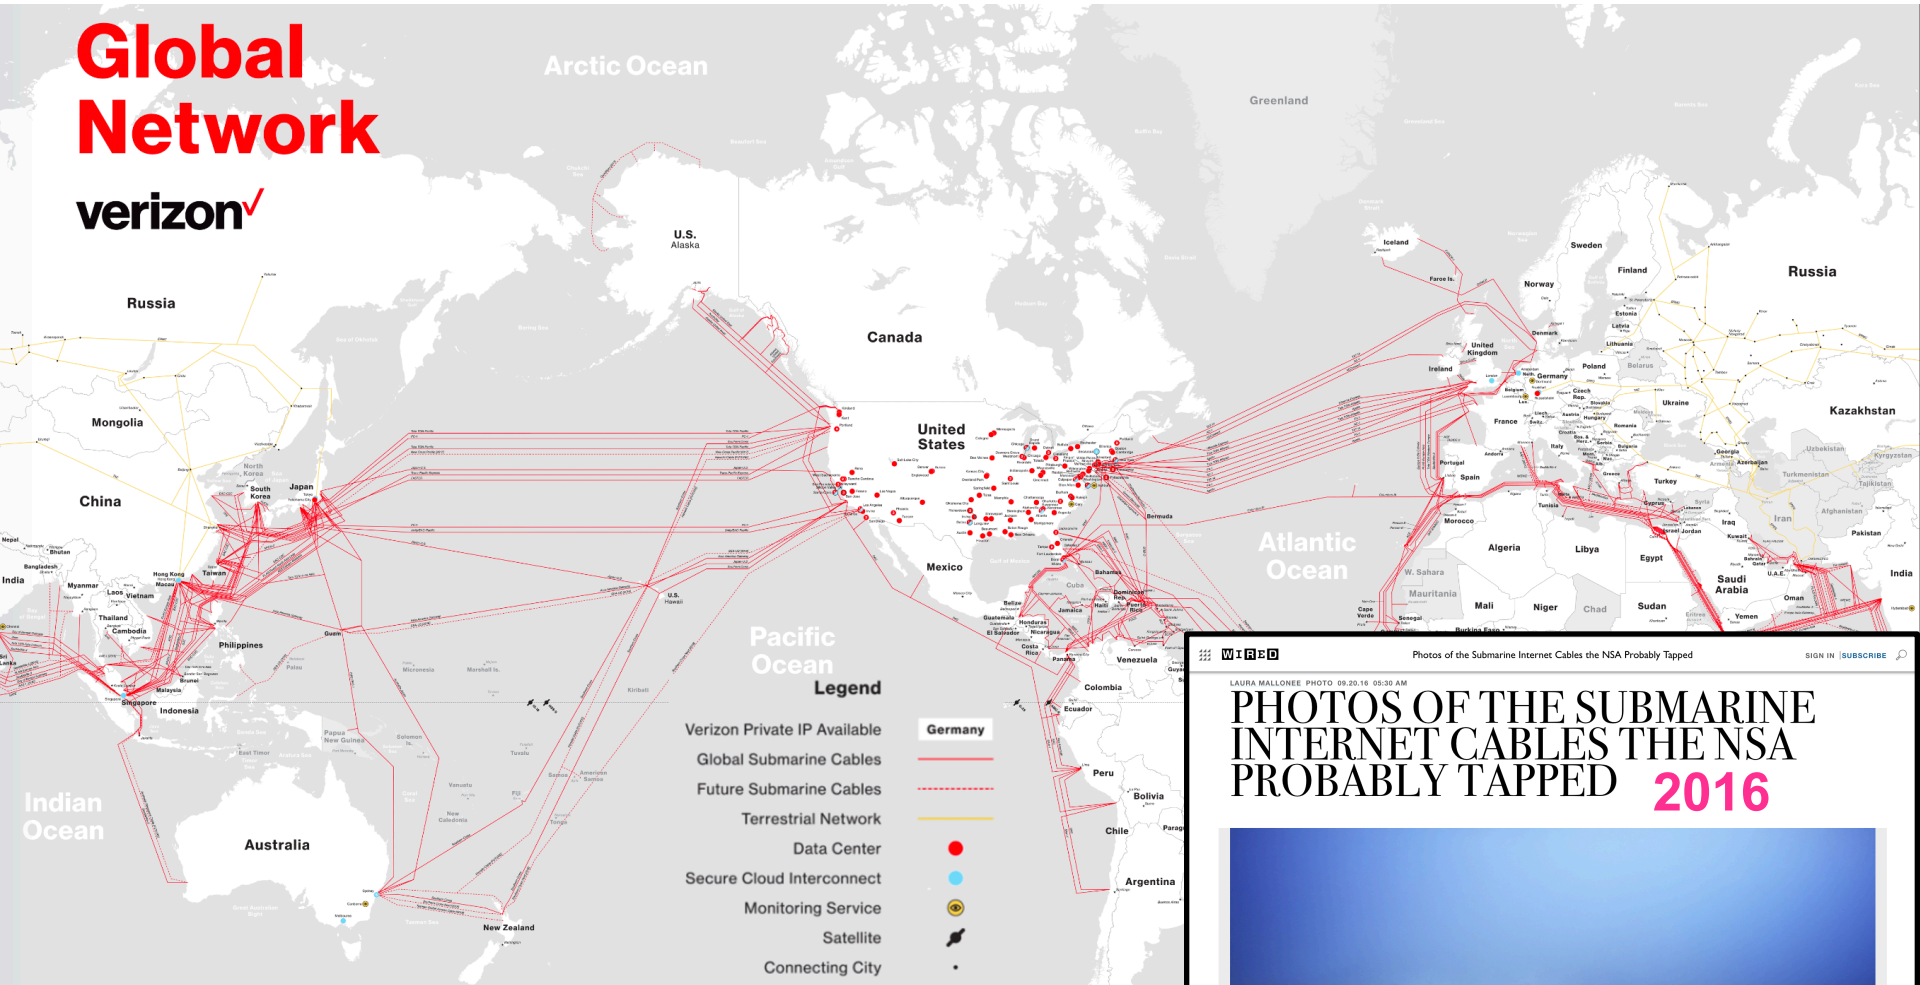
- customer ISP connects to 2 or more provider ISPs
- Why? For robustness. Can also multi-home your home network

PoP: Point-of-Presence

- where customer packets enter network
- e.g., you connecting to access ISP or ISP connecting to provider ISP

Tier-1 ISP has global reach

Global Network



<http://www.verizonenterprise.com/about/network/>

