Communication Networks

Spring 2017

Laurent Vanbever

www.vanbever.eu

ETH Zürich (D-ITET)

February, 20 2017

The Internet

An *exciting* place

8 billion

8 billion

estimated* # of Internet connected devices in 2016

* Cisco Visual Networking Index 2017

11.6 billion

estimated* # of Internet connected devices

in 2021

* Cisco Visual Networking Index 2017

~3 exabytes

estimated* daily global IP traffic

in 2016

* Cisco Visual Networking Index 2017

If 

= 1 Gigabyte

volume(Great Wall of China) = 1 exabyte

~3 exabytes

estimated* daily global IP traffic

in 2016

* Cisco Visual Networking Index 2017

~6 exabytes

estimated* daily global IP traffic

in 2020

* Cisco Visual Networking Index 2017

A *few* Internet services
you *might* have heard of…









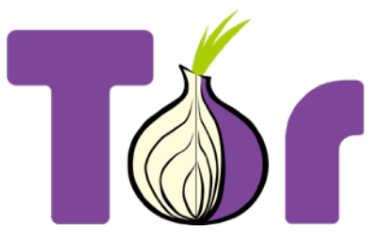

The Internet

A *tense* place

Countries get disconnected

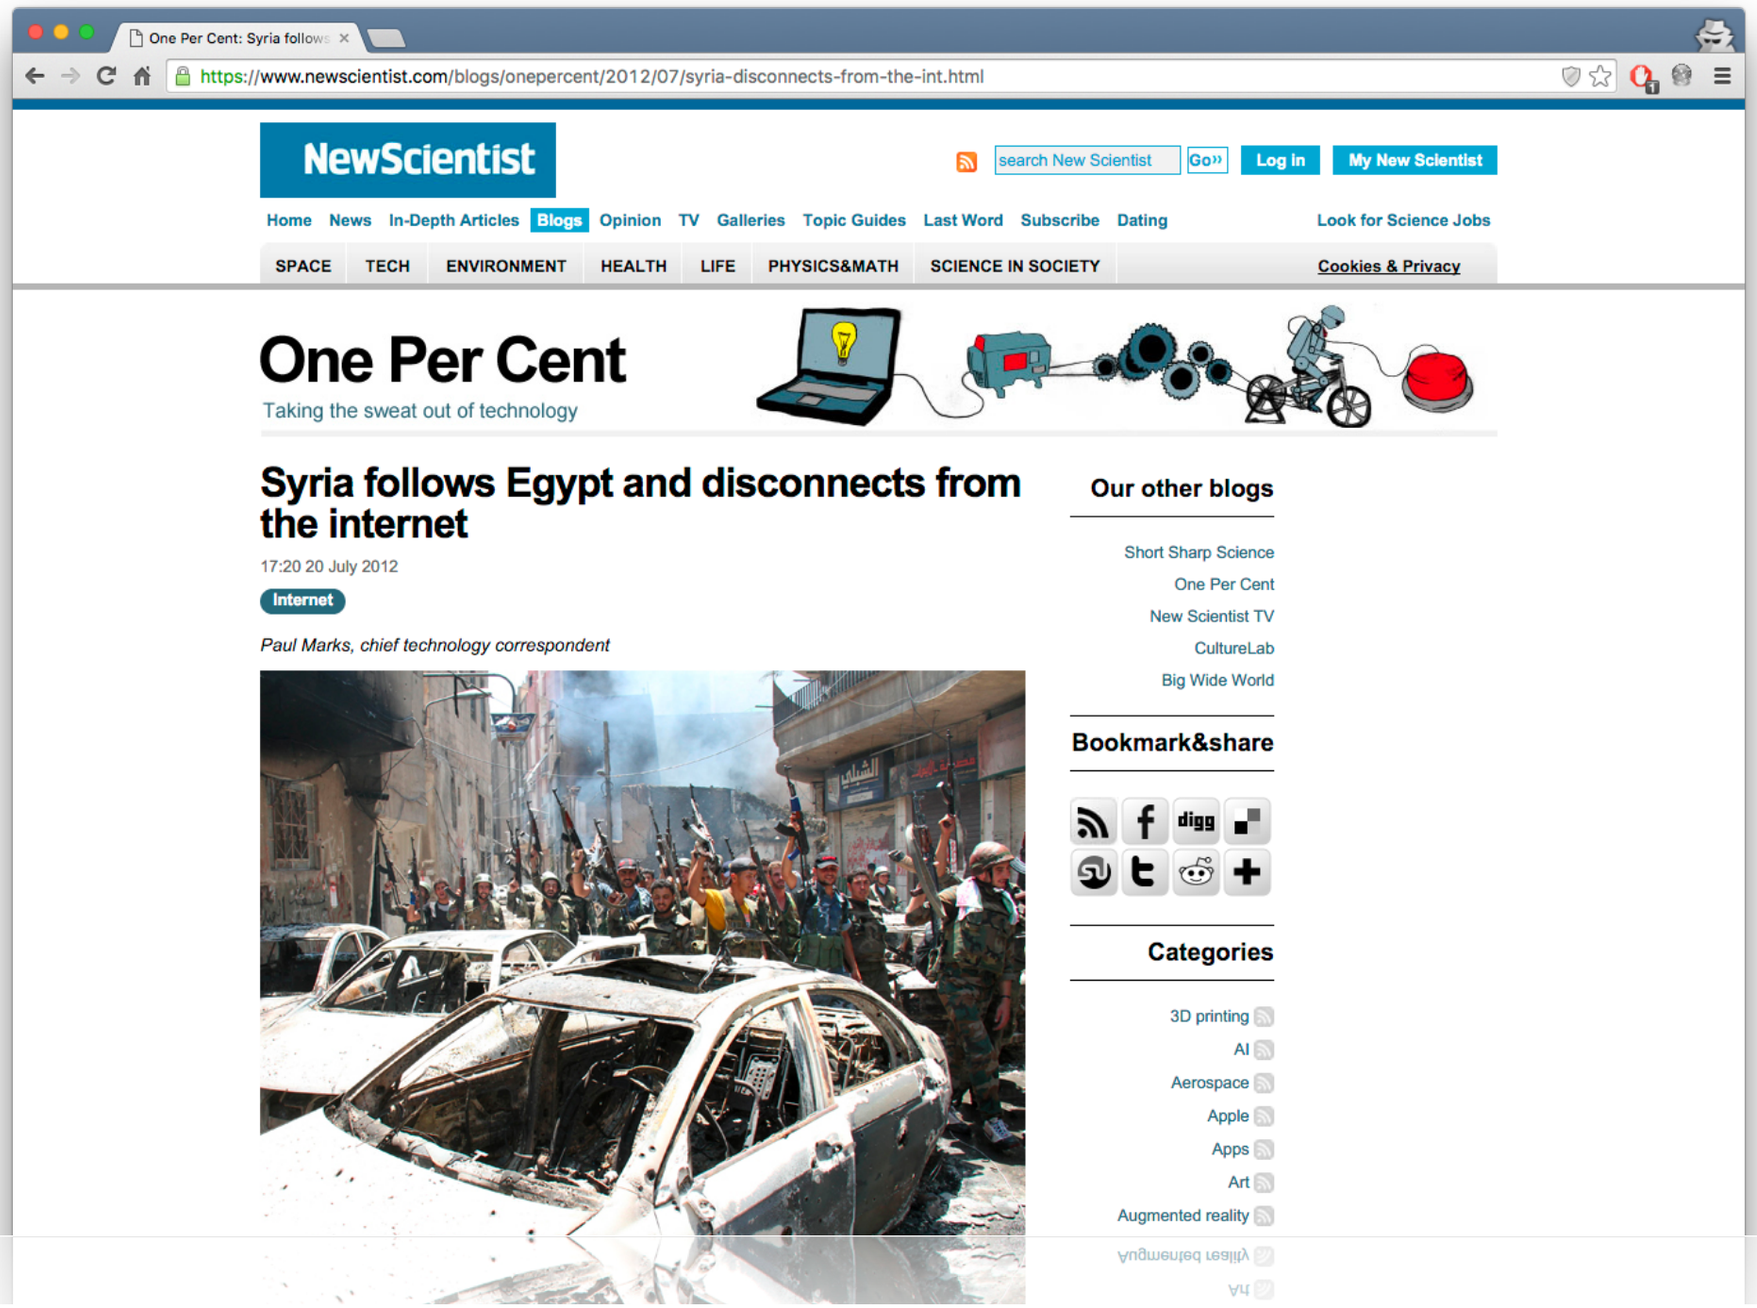

Also check: http://research.dyn.com/2012/11/could-it-happen-in-your-countr/

Internet traffic to/from Egypt in January 2011

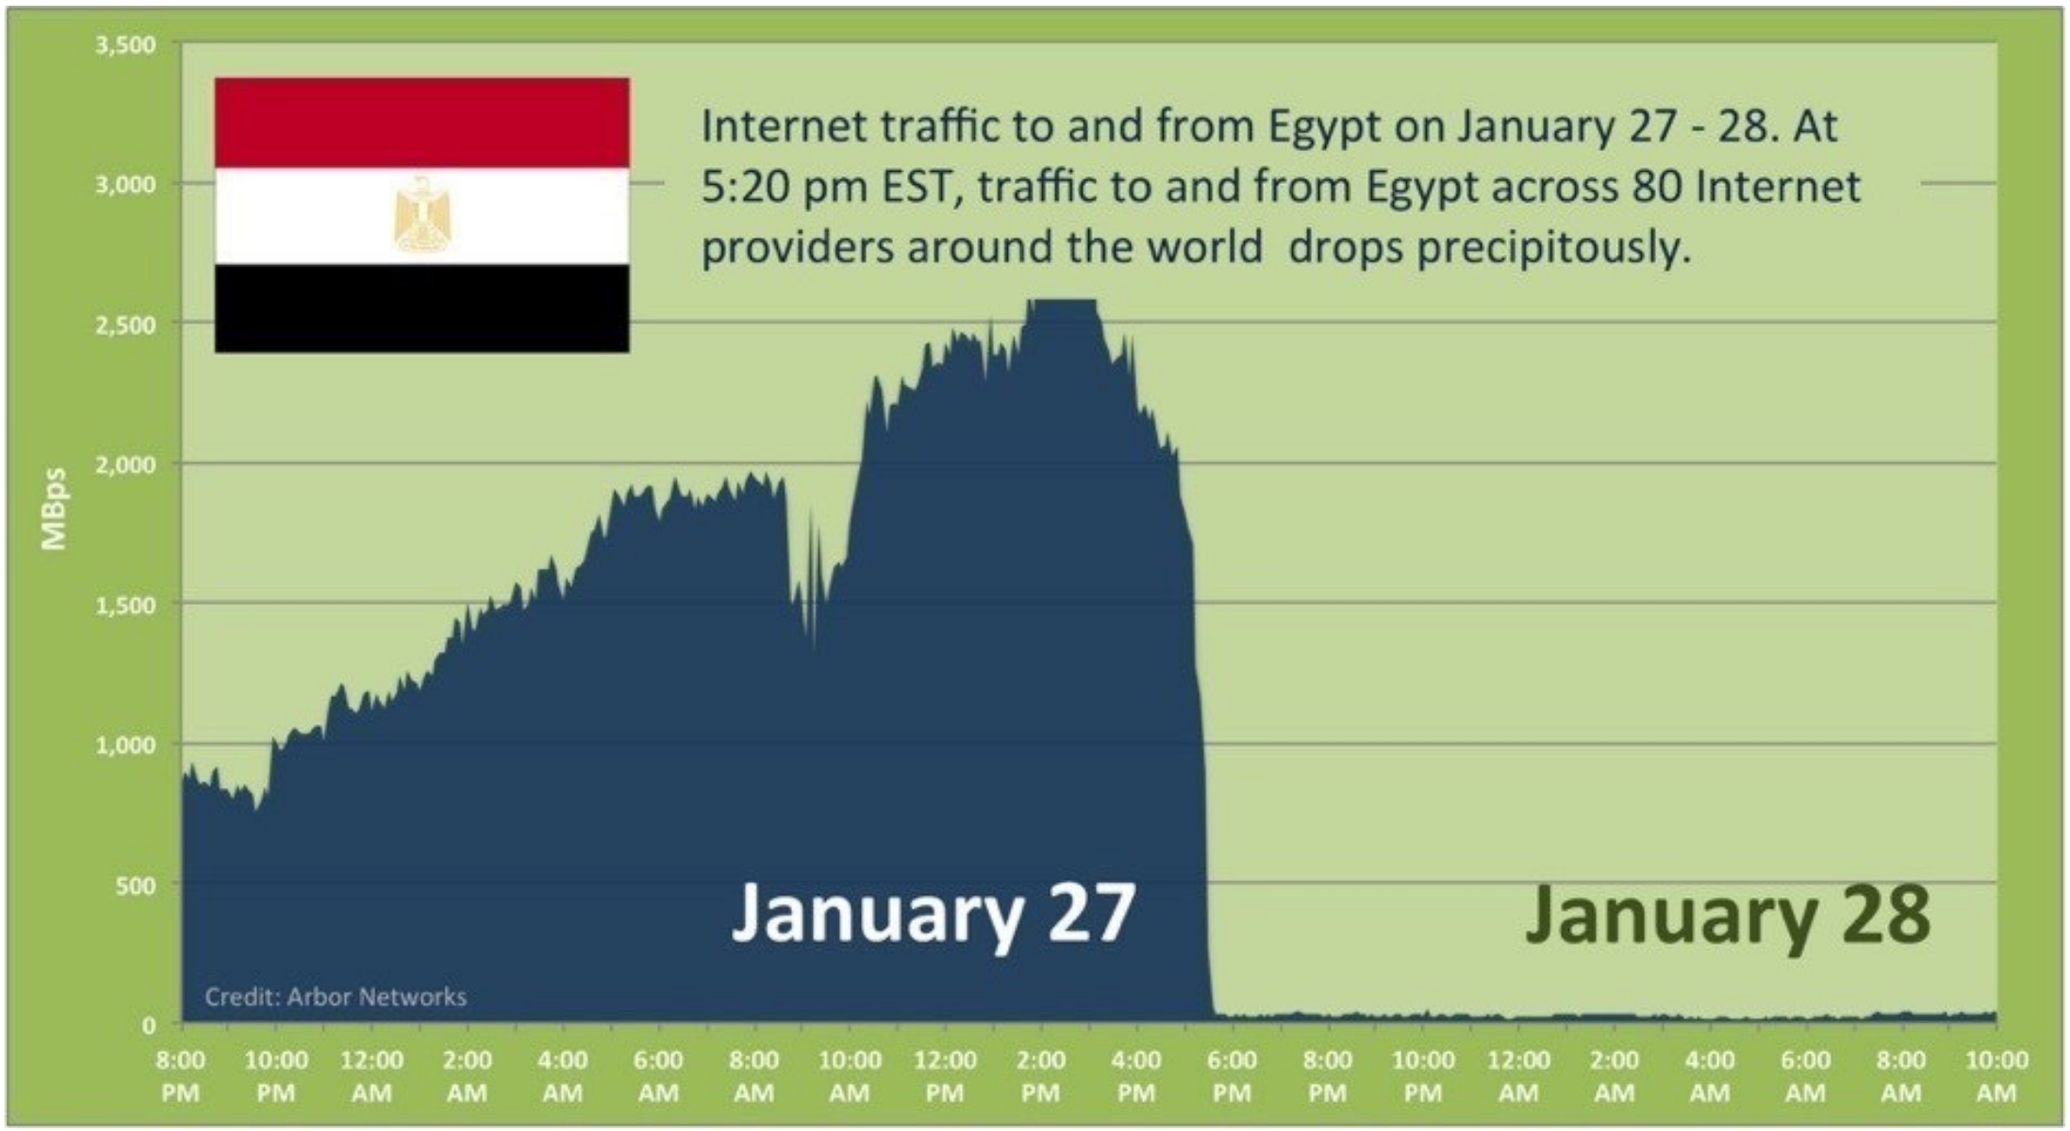

http://huff.to/1KxxoZF

Iraq Downs Internet To Combat ×

dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

Dyn PRODUCTS RESOURCES SUPPORT COMPANY Search Dyn SIGN IN

VANTAGEPOINT
IN: RESEARCH

Iraq Downs Internet To Combat Cheating…Again!

Research // Feb 6, 2017 // Doug Madory

Earlier this morning, the national fiber backbone of Iraq was taken offline in an effort to combat cheating on 6th grade placement exams. It was the fourth such outage in the past five days. 2017 marks the third year Iraq has used government-directed internet blackouts to combat cheating on student exams.

These recent outages are a continuation of a growing (and somewhat puzzling) trend by governments in many developing parts of the world to cut communications services in a desperate attempt to staunch rampant cheating on high-stakes student exams.

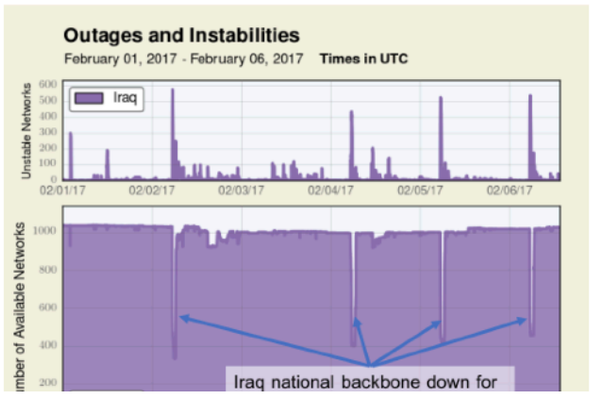

http://dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

Iraq's Exam Blackouts

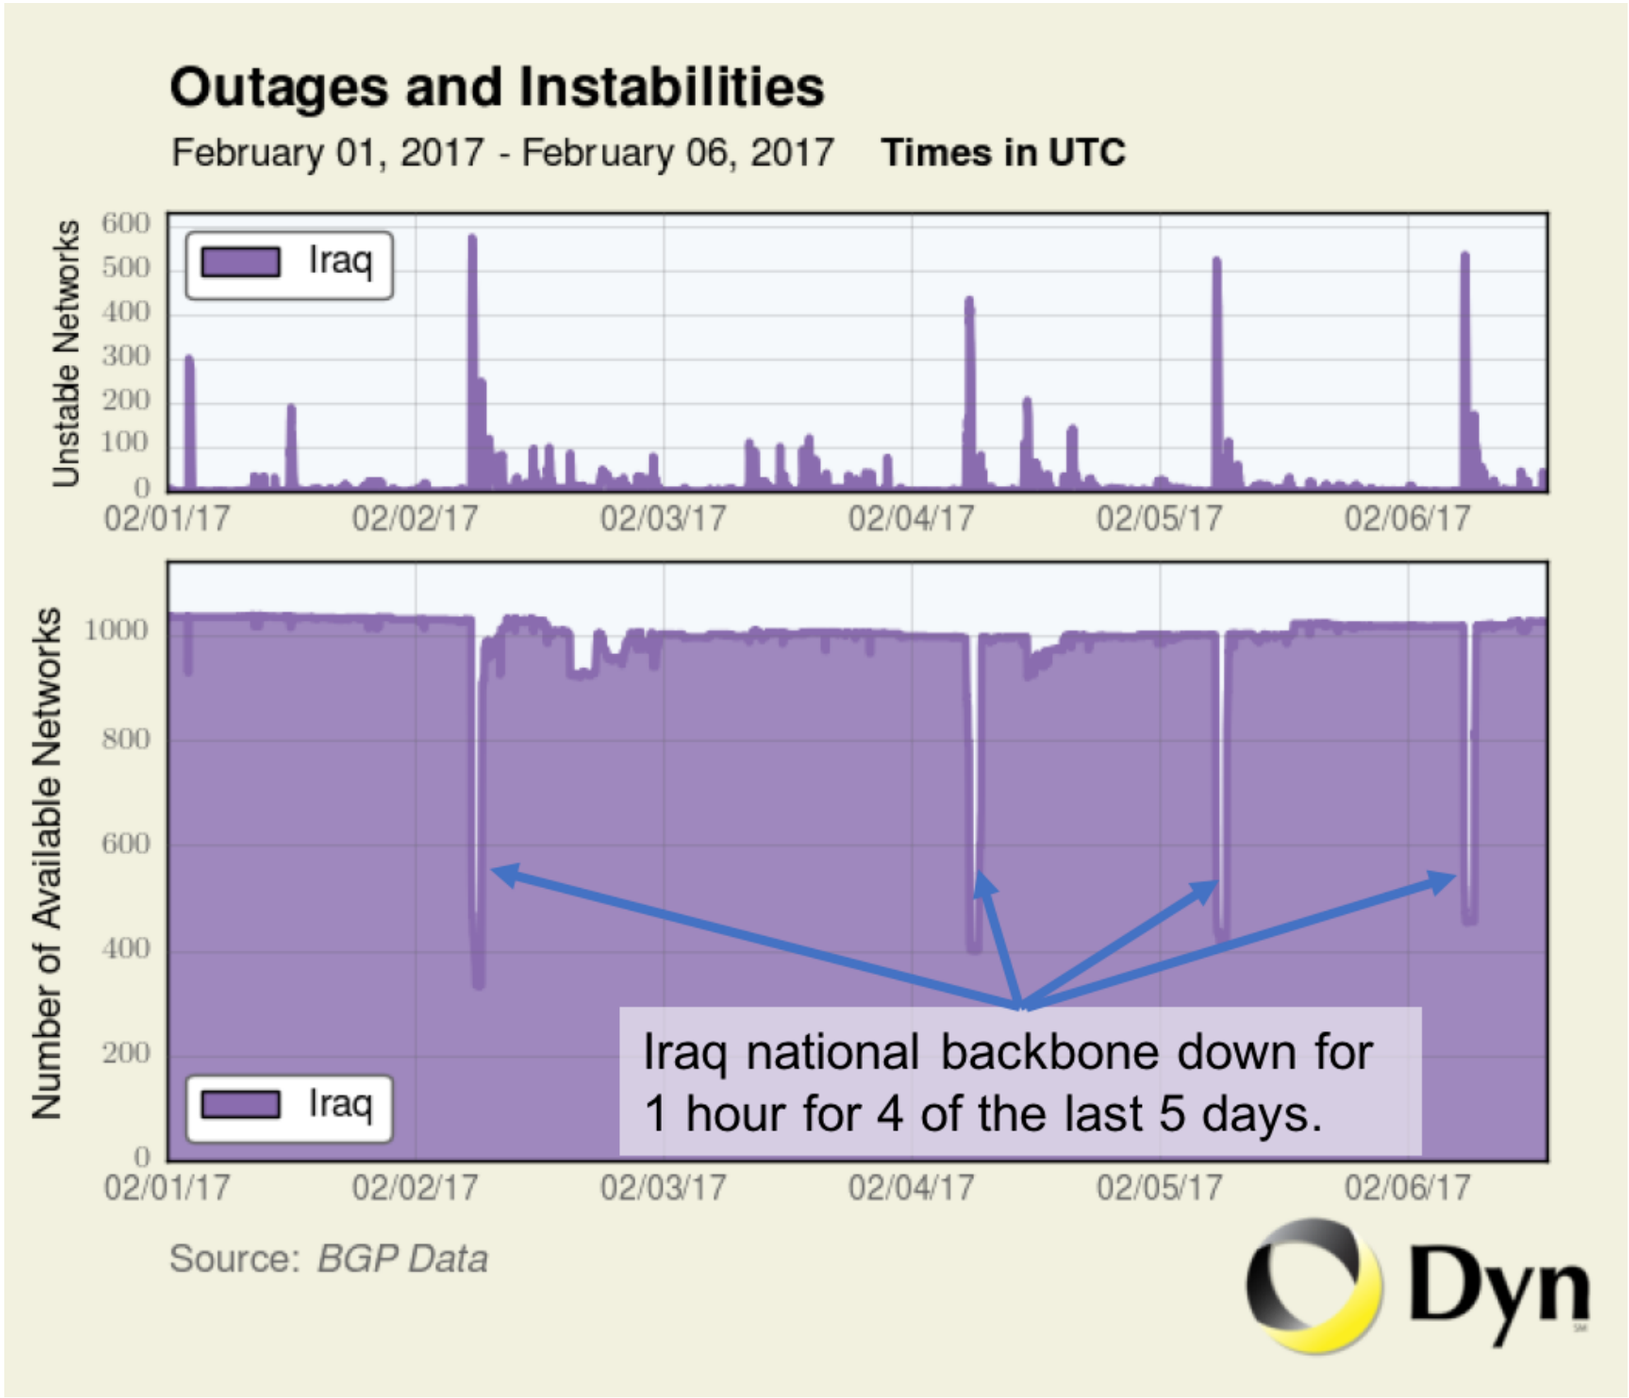

http://dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

Governments shut down the in

www.wired.co.uk/article/over-50-internet-shutdowns-2016

WIRED

Technology | Science | Culture | Video | Reviews | Magazine | More | Follow

Nintendo Switch isn't just a console, it's a 127-year saga that began with a deck of cards

—

Nintendo | 1 hour ago

Halo Wars 2 review: a solidly Spartan sequel to the real-time strategy classic

—

8/10

Halo | 56 minutes ago

Apple's WWDC 2017 will return to San Jose in June

—

WWDC | 1 hour ago

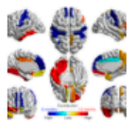

This algorithm can spot signs of autism in children a year before they're diagnosed

—

Autism | 2 hours ago

Humb collec migra

—

Humblc

Internet Freedom

Governments shut down the internet more than 50 times in 2016

Economic impact alone was £1.9bn, with greater fears over human rights and freedom of speech

http://www.wired.co.uk/article/over-50-internet-shutdowns-2016

By MATT KAMEN

—

Tuesday 3 January 2017

Communications get eavesdropped on…

http://wapo.st/1UVKamr

TOP SECRET//SI//ORCON//NOFORN

(TS//SI//NF) Introduction

U.S. as World's Telecommunications Backbone

- Much of the world's communications flow through the U.S.
- A target's phone call, e-mail or chat will take the **cheapest** path, **not the physically most direct** path – you can't always predict the path.
- Your target's communications could easily be flowing into and through the U.S.

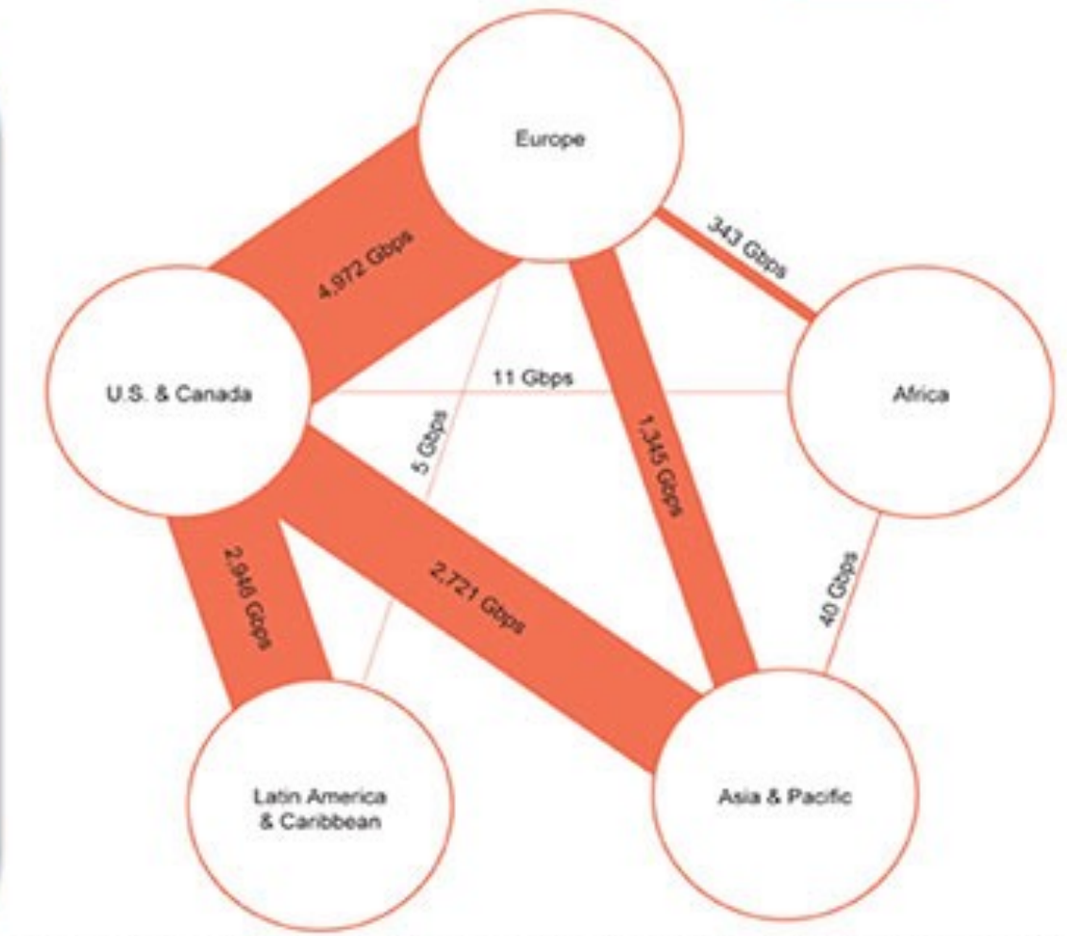

International Internet Regional Bandwidth Capacity in 2011

Source: Telegeography Research

TOP SECRET//SI//ORCON//NOFORN

http://wapo.st/1LcAw6p

Some Internet communications
are interfered against or heavily congested

Can ISPs selectively slow down traffic?

The U.S. Federal Communications Commission (FCC) ordered Comcast to stop interfering with p2p traffic

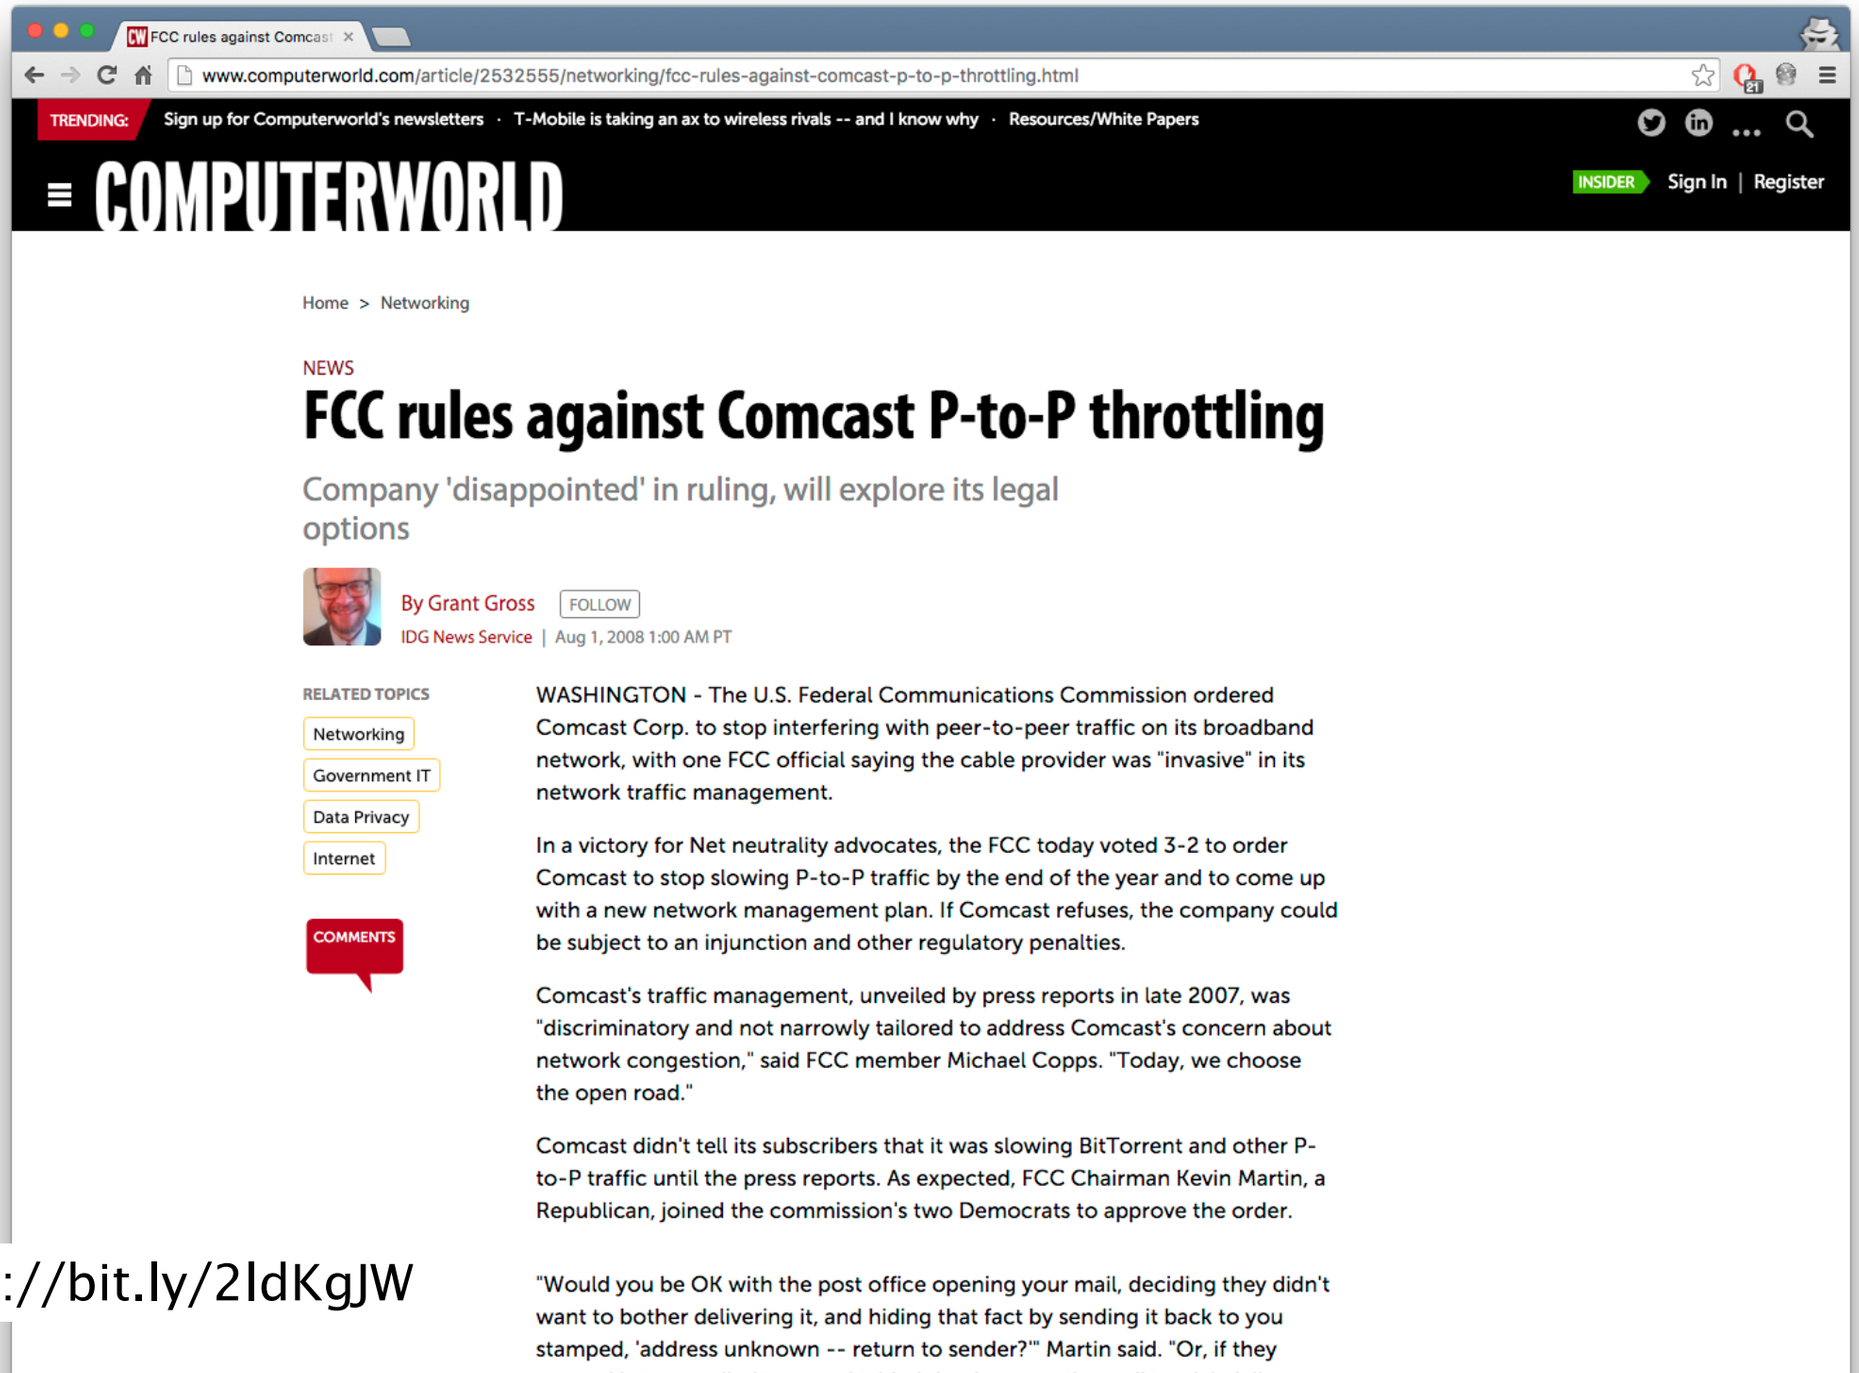

http://bit.ly/2IdKgJW


Snap Aims for Valuation of More Than $20 Billion in I.P.O.


Facebook's Zuckerberg, Bucking Tide, Takes Public Stand Against Isolationism


TECH WE'RE USING
Why I Still Love TiVo and How a Sous Vide Gadget Rescued Me


Airlines Phasing Out Screens Because You Are All on Your Devices


SOCIAL Q'S
Family Planning ... for Your Phones

TECH TIP
Adding Facebook's Birthday List to a Calendar Program


BITS
Daily Pay a Brea

TECHNOLOGY

F.C.C. Sets Net Neutrality Rules

By REBECCA R. RUIZ MARCH 12, 2015

The Federal Communications Commission is to take a more active role in regulating the Internet as a public utility, which is expected to provoke court cases from major broadband providers. By NATALIA V. OSIPOVA and CAITLIN PRENTKE on March 12, 2015. Photo by The New York Times. Watch in Times Video »

<iframe title= Embed

WASHINGTON — The Federal Communications Commission on Thursday released extensive details of how it would regulate broadband Internet providers as a public utility, producing official wording that almost certainly sets the stage for extended legal fights.

The release of the rules had been eagerly anticipated by advocates and lawmakers, as well as broadband and technology companies, since the agency approved new rules for Internet service two weeks ago. The details came in a 313-page document that included the new rules and the legal justifications for them.

The rules revealed how the strict laws would be modified for Internet providers, exempting the companies from the sort of price controls typically applied to utilities, for example. But the full text of the new order also raised uncertainties about broad and subjective regulation. One catchall provision, requiring "just and reasonable" conduct, allows the F.C.C. to decide what is acceptable on a case-by-case basis.

Opponents of the rules, including many of the leading Internet providers, spent Thursday poring over the document. It was not known who would file the first legal challenges, or exactly what legal arguments would be made. Many experts, though, said the document included plenty of opportunity for

RELATED COVERAGE


F.C.C. Approves Net Neutrality Rules, Classifying Broadband Internet Service as a Utility FEB. 26, 2015


NEWS ANALYSIS
The Push for Net Neutrality Arose From Lack of Choice FEB. 25, 2015

http://nyti.ms/2kZUnDA

TECHNOLOGY

Trump's F.C.C. Pick Quickly Targets Net Neutrality Rules

By CECILIA KANG FEB. 5, 2017

Ajit Pai, the new chairman of the Federal Communications Commission, has taken a first swipe at net neutrality. Christopher Gregory/The New York Times

WASHINGTON — In his first days as President Trump's pick to lead the Federal Communications Commission, Ajit Pai has aggressively moved to roll back consumer protection regulations created during the Obama presidency.

Mr. Pai took a first swipe at net neutrality rules designed to ensure equal access to content on the internet. He stopped nine companies from providing discounted high-speed internet service to low-income individuals. He withdrew an effort to keep prison phone rates down, and he scrapped a proposal to break open the cable box market.

In total, as the chairman of the F.C.C., Mr. Pai released about a dozen actions in the last week, many buried in the agency's website and not publicly announced, stunning consumer advocacy groups and telecom analysts. They said Mr. Pai's message was clear: The F.C.C., an independent agency, will mirror the Trump administration's rapid unwinding of government regulations that businesses fought against during the Obama administration.

"With these strong-arm tactics, Chairman Pai is showing his true stripes,"

The Trump White House
Stories about President Trump's administration.

Trump's Pick to Replace Flynn Turns Down the Job FEB 16
Did Trump Drop 2-State Policy? Mideast Experts Say Don't Be So Sure FEB 16
In 77 Chaotic Minutes, Trump Defends 'Fine-Tuned Machine' FEB 16
Trump Directs Justice Department to Investigate 'Criminal Leaks' FEB 16
House G.O.P. Leaders Outline Plan to Replace Obama Health Care Act FEB 16

See More »

http://nyti.ms/2kyROGJ

Some Internet communications
are interfered against or heavily congested

Who should pay the other for Internet connectivity?

A primer on the conflict between Netflix and Comcast

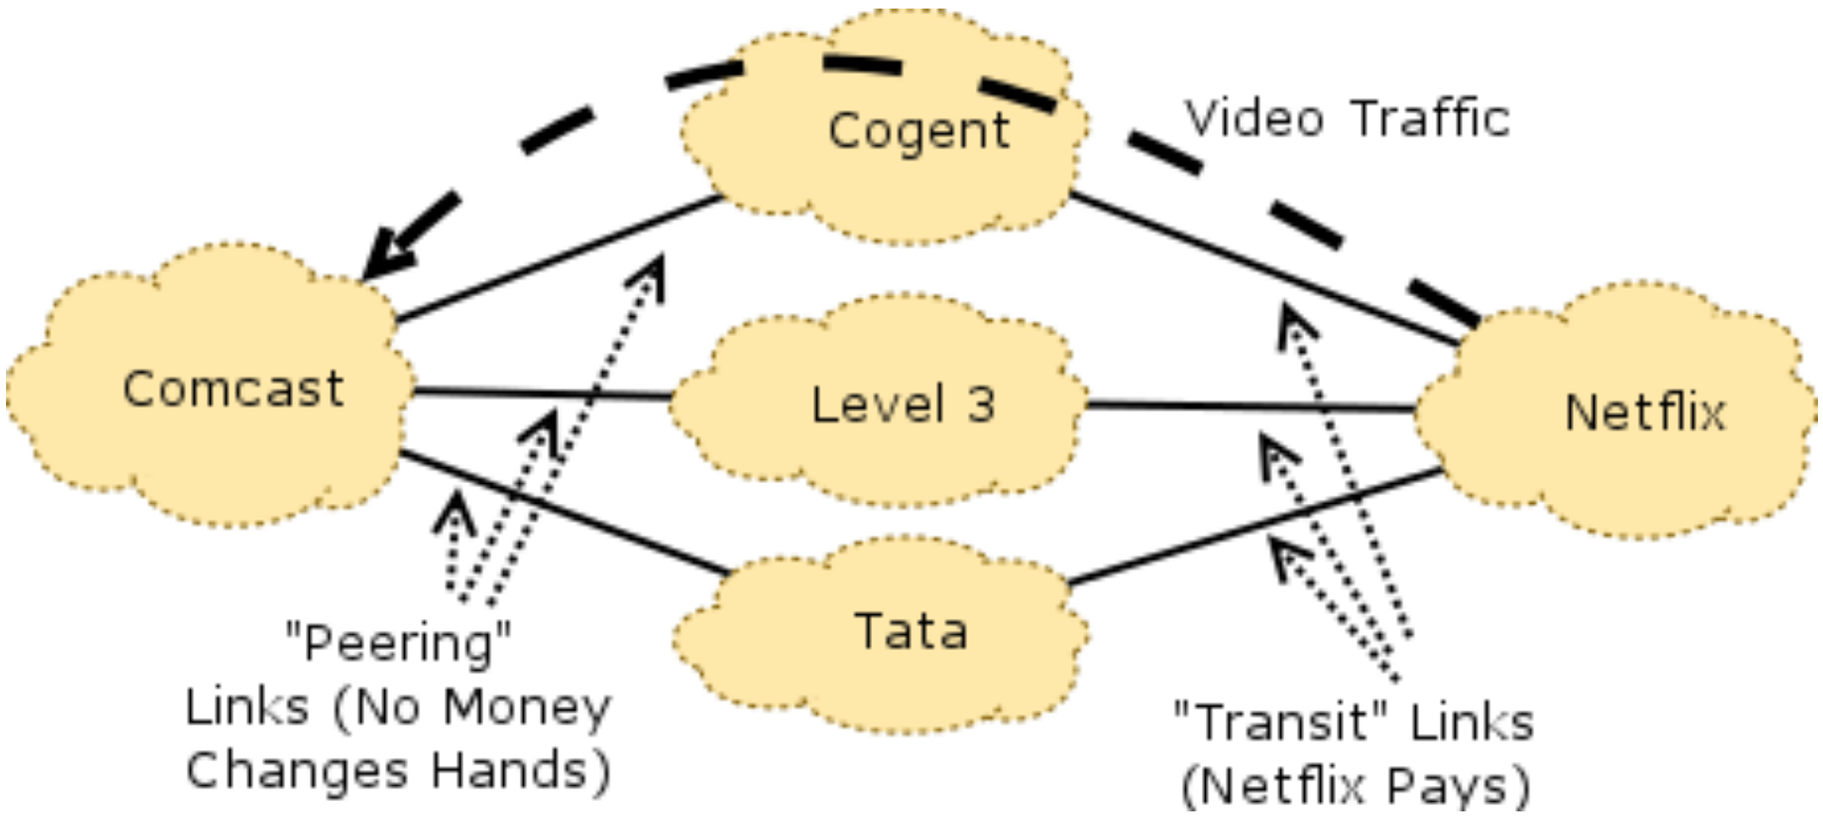

https://freedom-to-tinker.com/blog/feamster/why-your-netflix-traffic-is-slow-and-why-the-open-internet-order-wont-necessarily-make-it-faster/

Due to congestion, throughput across Cogent to Time Warner, Comcast, Verizon were miserable

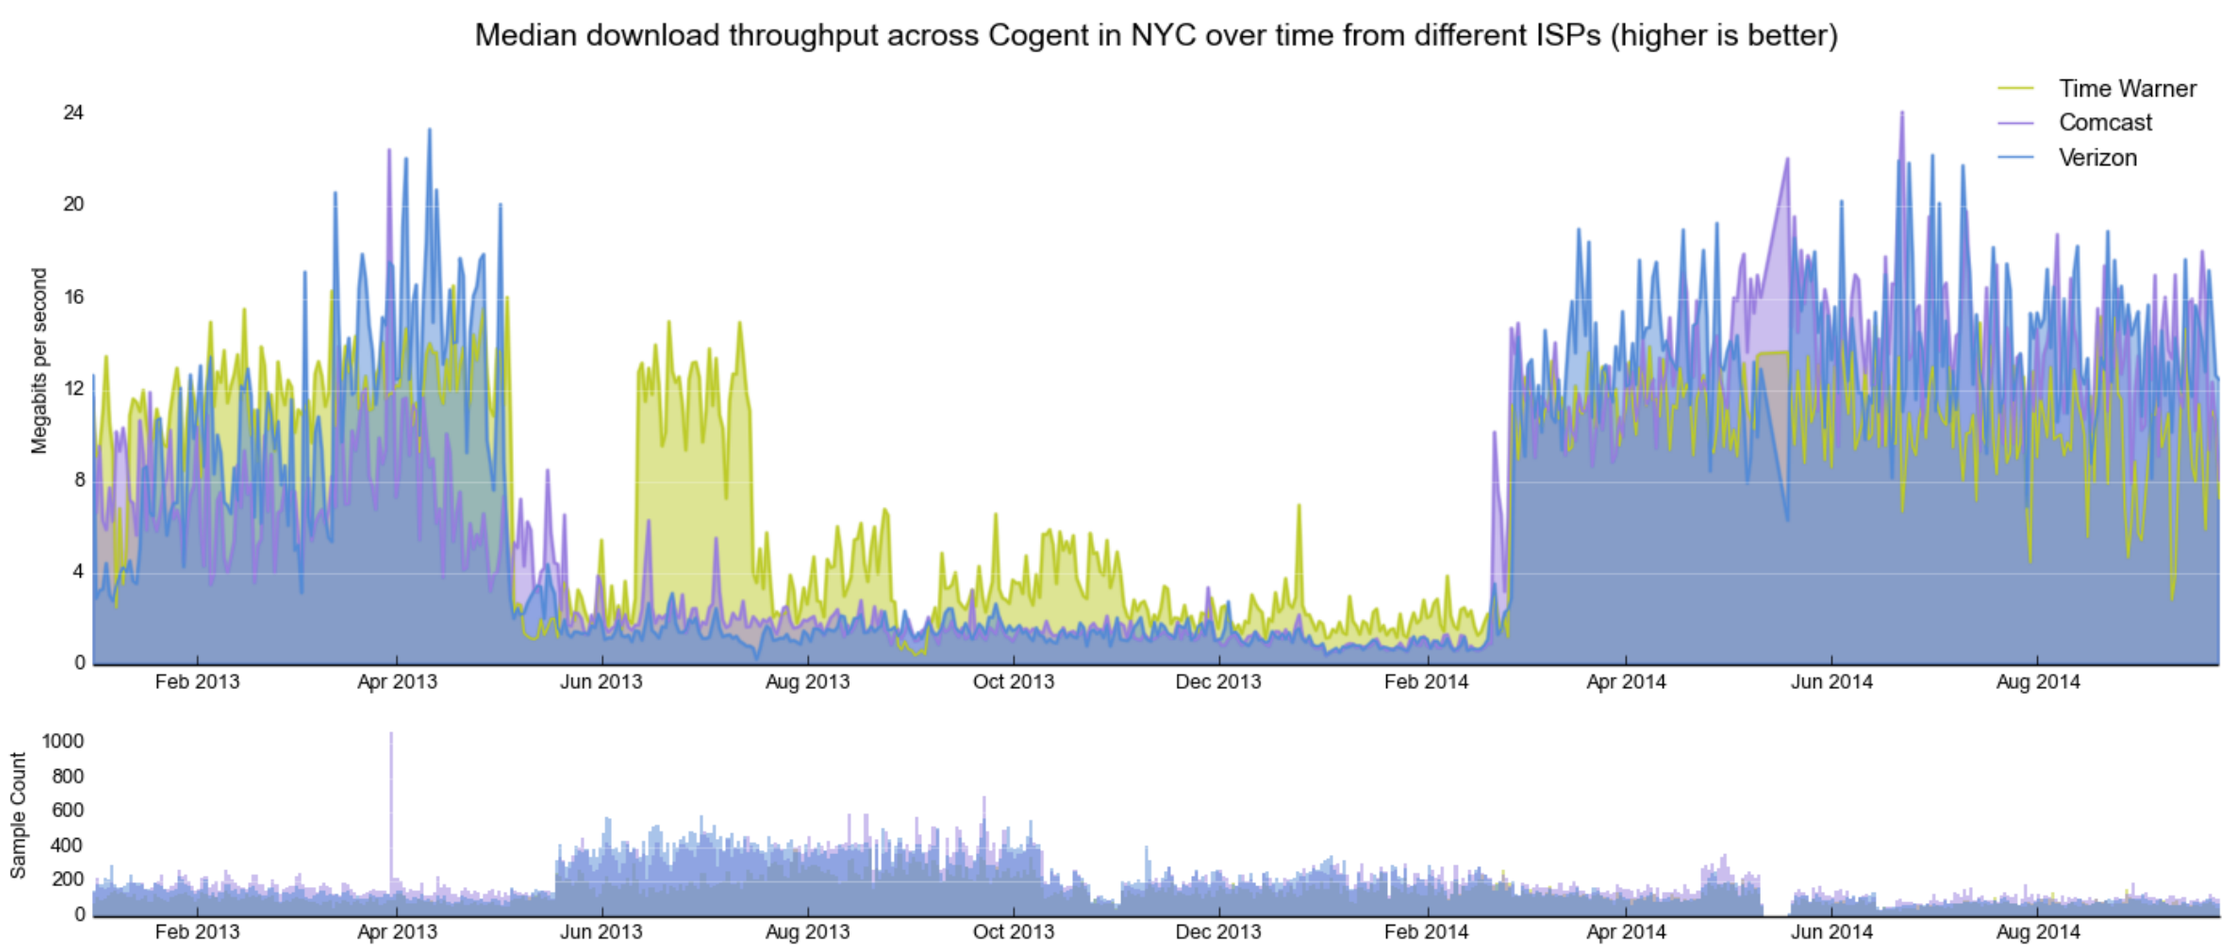

http://bit.ly/1thPzro

Situation massively improved after
Netflix agreed to paid direct connection to the providers

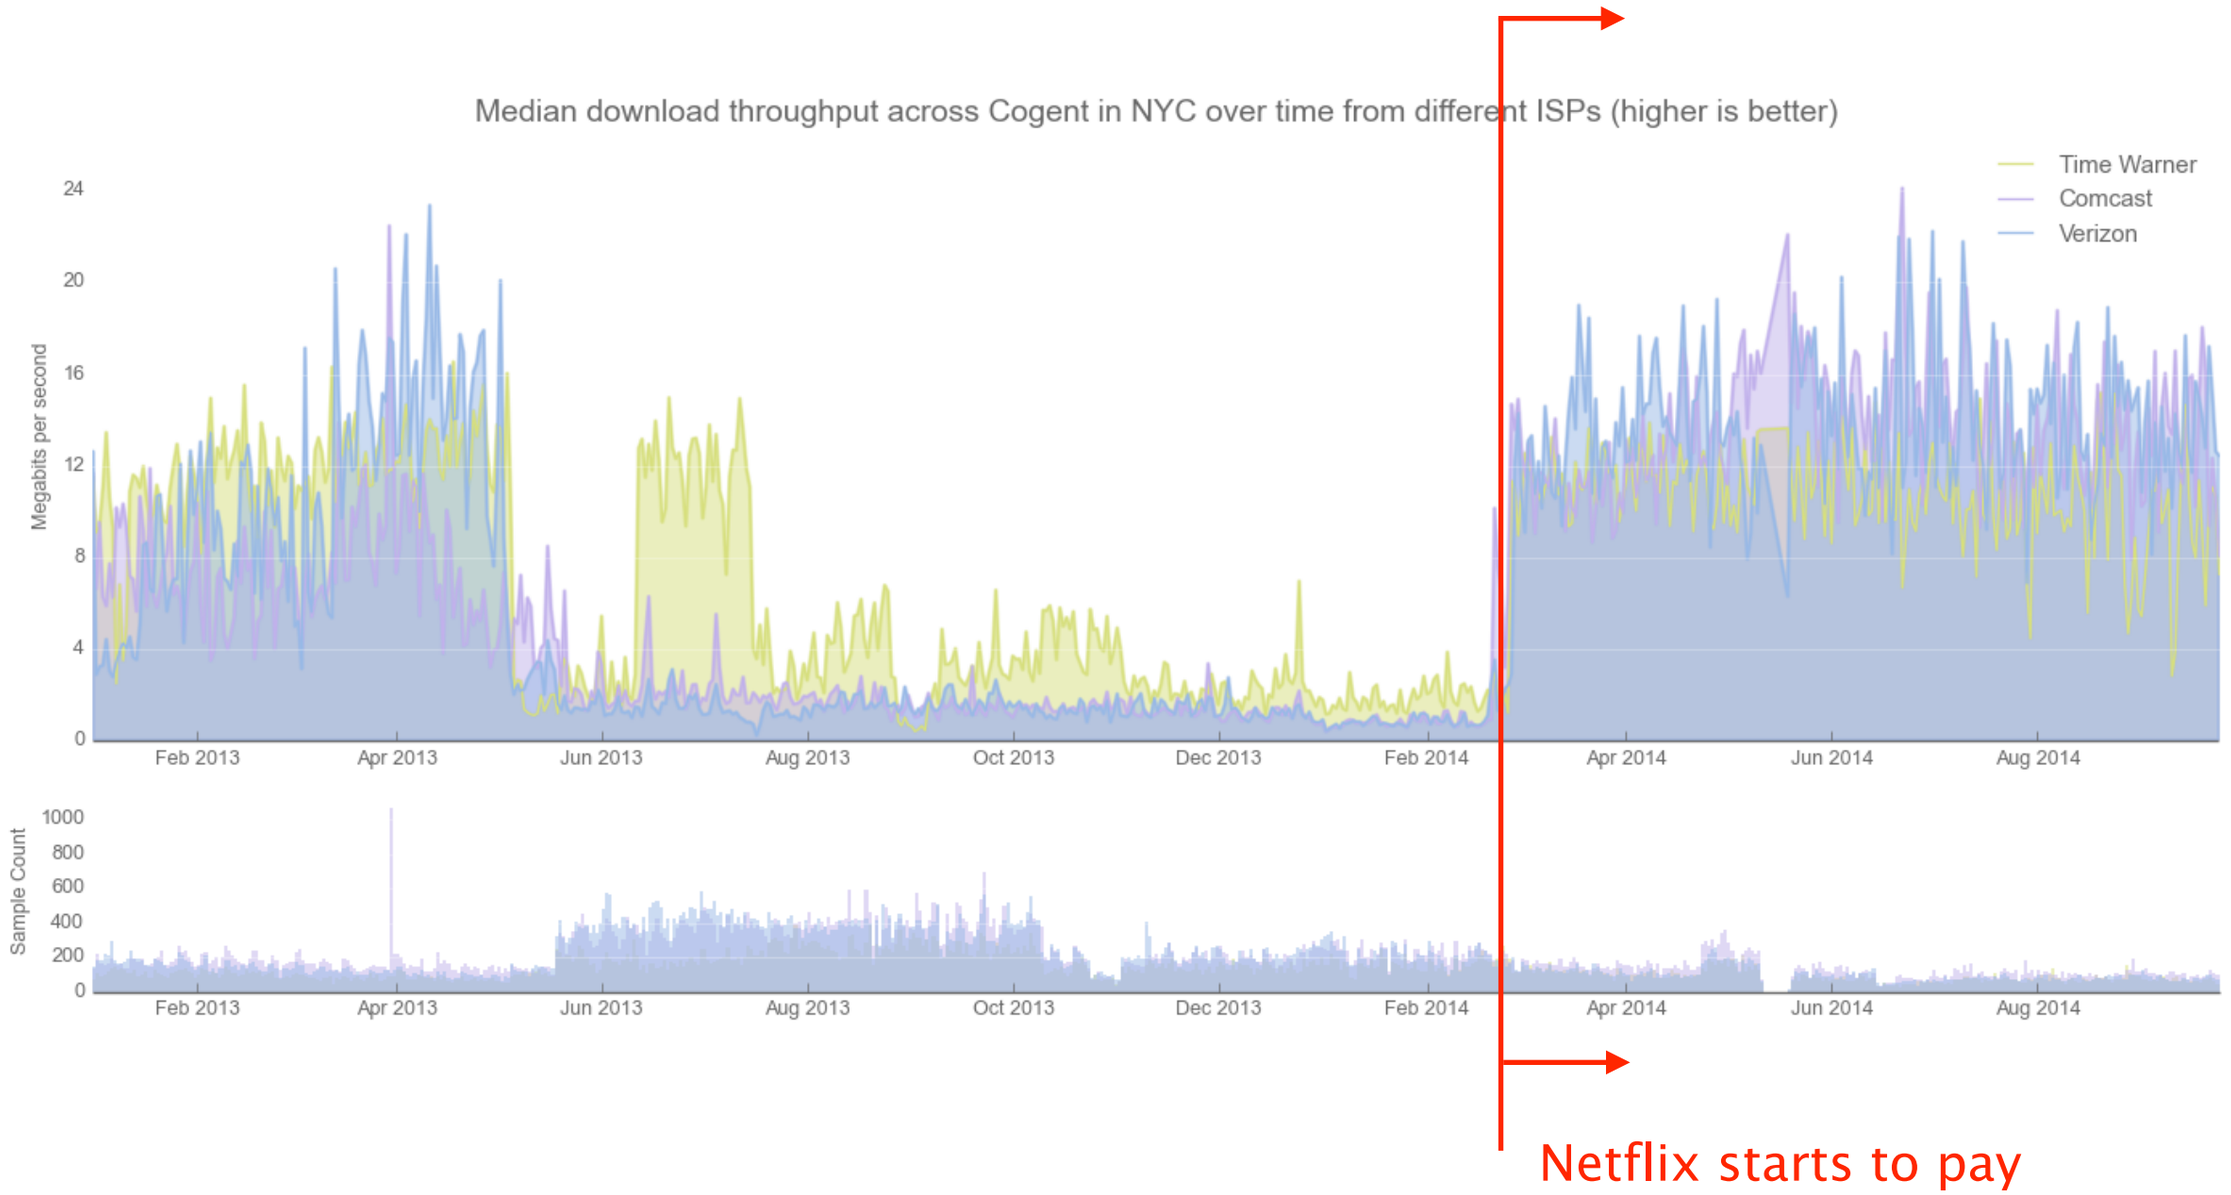

http://arstechnica.com/tech-policy/2014/04/netflix-and-verizon-reach-interconnection-deal-to-speed-up-video/

The Internet infrastructure is
a fragile environment

JUL 8, 2015 @ 03:36 PM 11,261 VIEWS

United Airlines Blames Router for Grounded Flights

Alexandra Talty, CONTRIBUTOR

I cover personal finance and travel.

FOLLOW ON FORBES (110)

Opinions expressed by Forbes Contributors are their own.

FULL BIO

After a computer problem caused nearly two hours of grounded flights for United Airlines this morning and ongoing delays throughout the day, the airline announced the culprit: a faulty router.

Spokeswoman Jennifer Dohm said that the router problem caused "degraded network connectivity," which affected various applications.

A computer glitch in the airline's reservations system caused the Federal Aviation Administration to impose a groundstop at 8:26 a.m. E.T. Planes that were in the air continued to operate, but all planes on the ground were held. There were reports of agents writing tickets by hand. The ground stop was lifted around 9:47 a.m. ET.

The outage was due to

one faulty Internet device

Facebook, Tinder, Instagram suffer widespread issues

IMAGE: GETTY IMAGES

BY JENNI RYALL
AUSTRALIA

JAN 27, 2015

***UPDATED**: Tuesday, Jan. 27 / 4:32 a.m. EST* — A Facebook spokeswoman told *Mashable* that the outage was due to a change to the site's configuration systems, and not a hacker attack. "Earlier this evening many people had trouble accessing Facebook and Instagram. This was not the result of a third party attack but instead occurred after we introduced a change that affected our configuration systems. We moved quickly to fix the problem, and both services are back to 100% for everyone.", she said.

***UPDATED**: Tuesday, Jan. 27 / 2:14 a.m. EST* — Facebook, Tinder and Twitter appear to be back to normal after a 40 minute outage and mass freak out.

The outage was due to a change to the site's configuration systems

“Human factors are responsible for 50% to 80% of network outages”

Juniper Networks, *What’s Behind Network Downtime?*, 2008

The Internet Under Crisis Conditions

Learning from September 11

Committee on the Internet Under Crisis Conditions:
Learning from September 11

Computer Science and Telecommunications Board
Division on Engineering and Physical Sciences

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

National Research Council. The Internet Under Crisis Conditions: Learning from September 11

The Internet Under Crisis Conditions

Learning from September 11

Committee on the Internet Under Crisis Conditions:
Learning from September 11

Computer Science and Telecommunications Board
Division on Engineering and Physical Sciences

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

Internet advertisements rates

suggest that

The Internet was more stable

than normal on Sept 11

Internet advertisements rates suggest that The Internet was more stable than normal on Sept 11

Information suggests that operators were watching the news instead of making changes to their infrastucture

Job Snijders
@JobSnijders

Follow

Fun fact: most BGP route leaks happen on Wednesdays, but in the weekend us humans collectively take a break! :-)

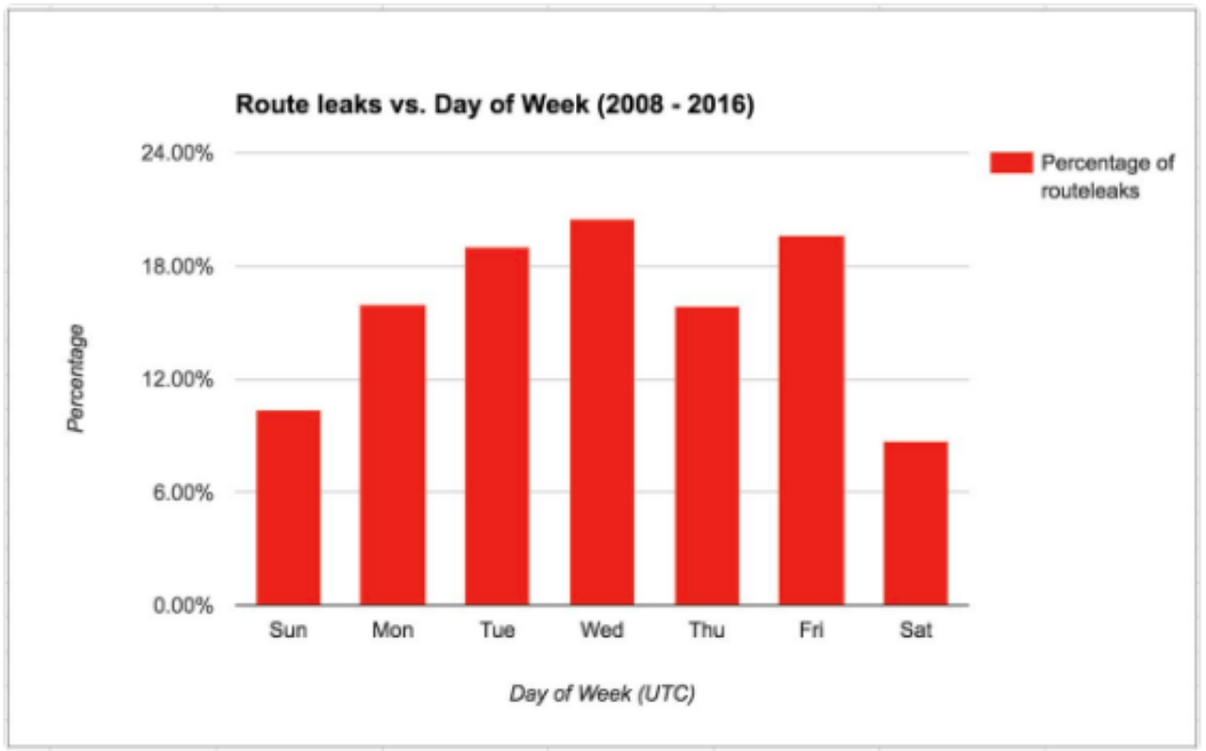

Internet scalability is at risk with
no more IPv4 addresses and a slow IPv6 deployment

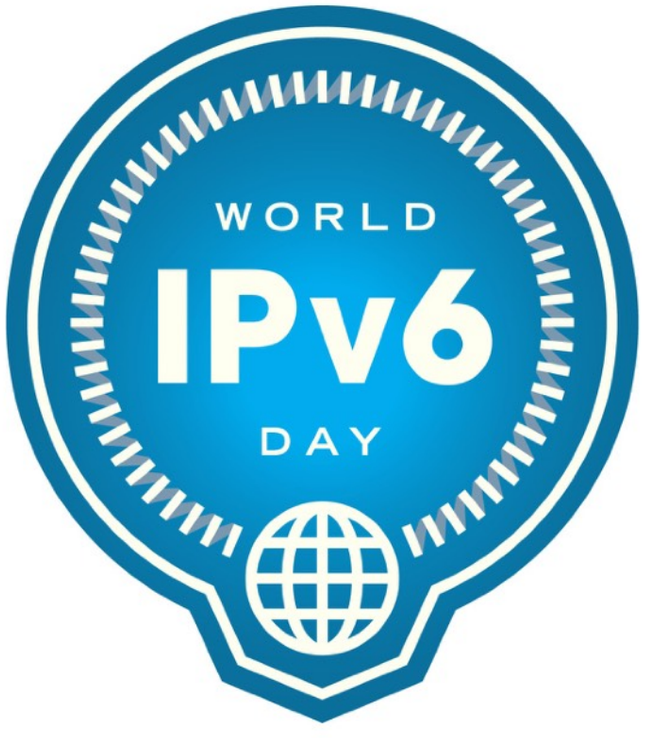

Communication Networks

Course goals

Knowledge

Understand how the Internet works and why

from your
network plug...

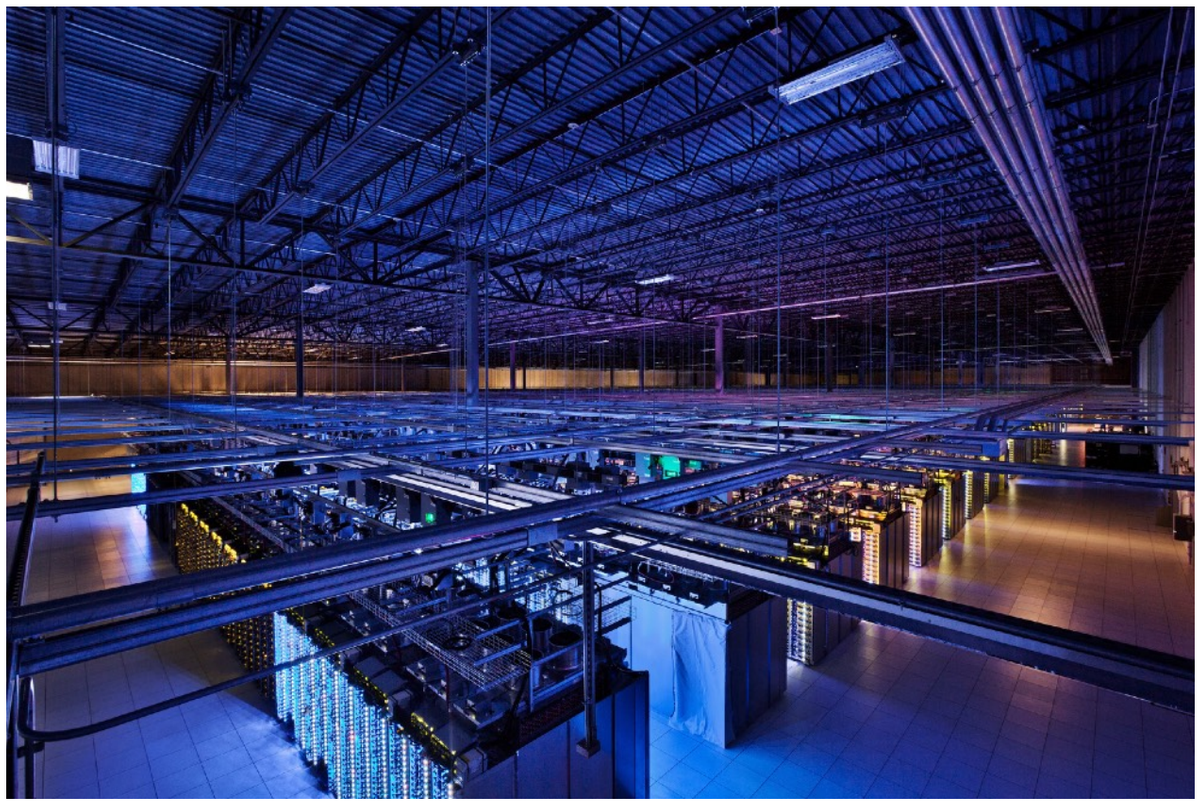

...to Google's data-center

Insight

Key concepts and problems in Networking

Naming Layering Routing Reliability Sharing

Naming Layering Routing Reliability Sharing

How do you address computers, services, protocols?

Naming Layering Routing Reliability Sharing

How do you manage complexity?

Naming Layering **Routing** Reliability Sharing

How do you **go from A to B?**

Naming Layering Routing **Reliability** Sharing

How do you communicate reliably using unreliable mediums?

Naming Layering Routing Reliability Sharing

How do you divide scarce resources among competing parties?

Skill

Build, operate and configure networks

Skill

Build, operate and configure networks

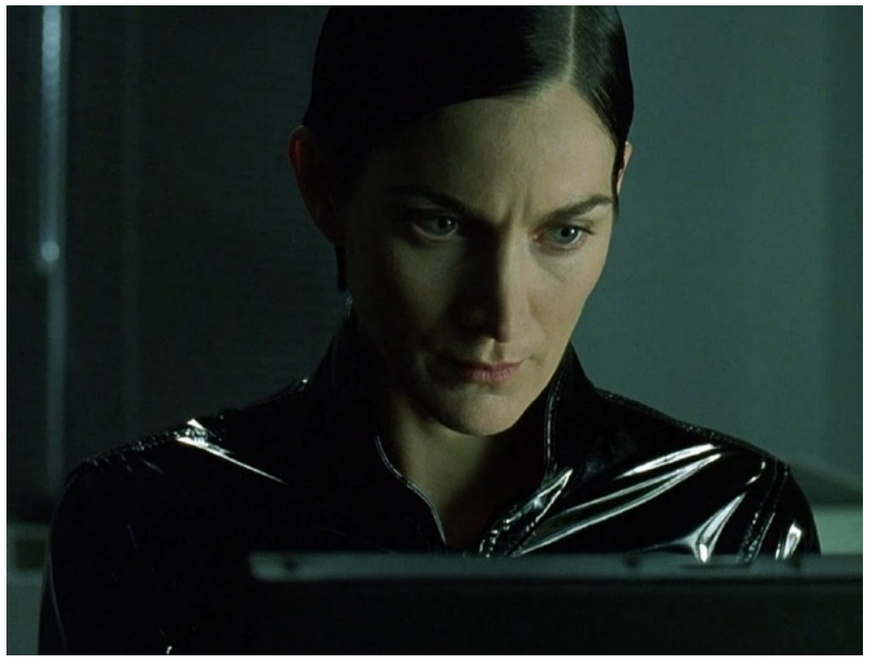

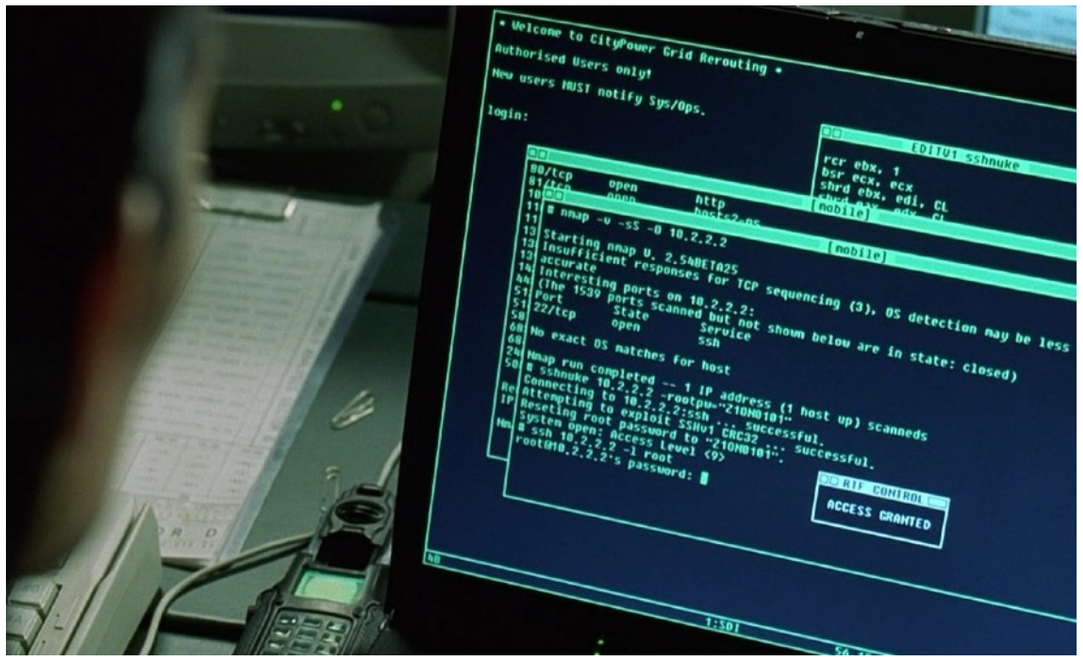

Trinity using a port scanner (nmap) in Matrix Reloaded™

Software-Defined Network

Software-Defined Network

enable network programmability

So far, network devices
have been completely locked down

closed software

closed hardware

Cisco™ device

SDN opens up the network devices, enabling network innovation

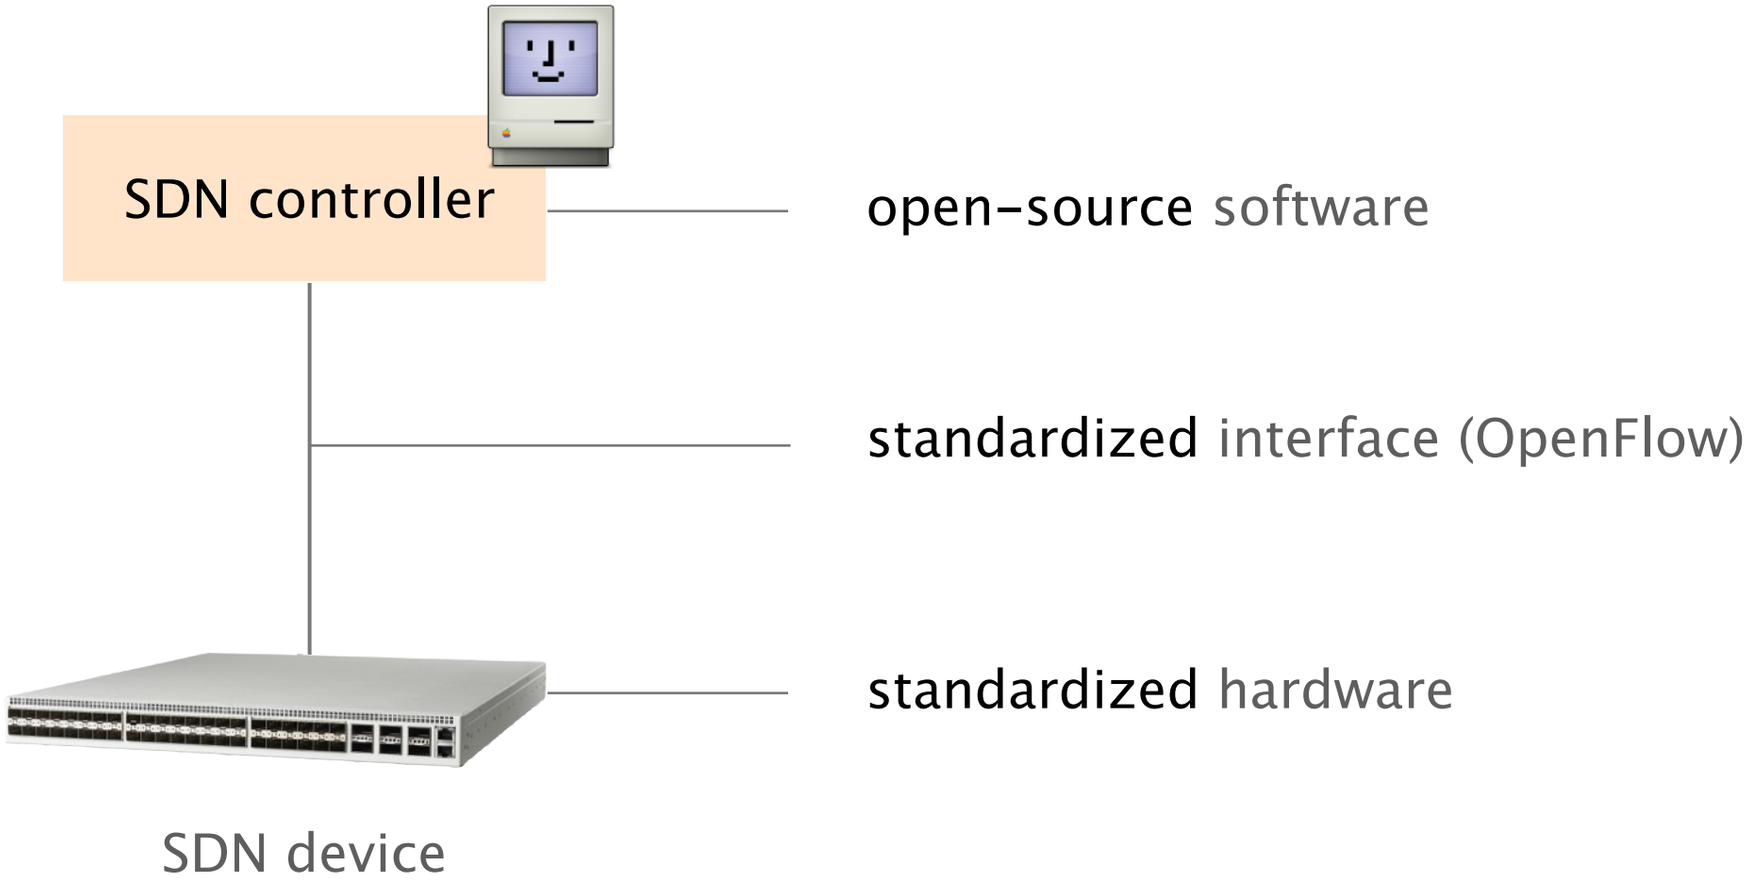

The hype around SDN is huge, both in the industry and in academia

VMware Acquires Once-Secretive Start-Up Nicira for $1.26 Billion

JULY 23, 2012 AT 1:25 PM PT

VMware, the software company best known for its virtualization technology that forms the backbones of so-called cloud computing today, said it will pay $1.26 billion for Nicira, a networking start-up that has sought to do to networks what VMware has done to computers.

The news comes on the same day that VMware was to report quarterly earnings. And while I don't usually cover VMware's earnings, I may as well mention the results: The company reported revenue for the quarter ended June rose to $1.12 billion, while earnings on a per-share basis were 68 cents. Analysts had been expecting sales of $1.12 billion and earnings of 66 cents.

Nicira had been running in stealth mode for quite awhile; I got to reveal its plans to the world last February.

The deal amounts to a nice payoff for Nicira's investors including Andreessen Horowitz, Lightspeed Venture Partners and NEA, as well as VMware founder Diane Greene and venture capitalist Andy Rachleff.

With $600M Invested in SDN Startups, the Ecosystem Builds

Scott Raynovich, June 10, 2014

More than $600 million has been invested in at least two dozen software-defined networking (SDN) startups so far, according to Rayno Report research. You can expect that to continue to climb. With the SDN ecosystem starting to take hold with a broad range of alliances and distribution partnerships, we're just getting started.

The Arista IPO will help build visibility for next-generation, software-driven networking. But Arista is selling its own hardware and is not an SDN pure-play. A new line of SDN startups, with a more radical approach to software-based networking, is building momentum. These newer SDN startups are just getting their gear into customers' hands and starting to build sales channels, so you can expect a long revenue ramp.

This excitement is boosting startup valuations, according to Rayno Report research. There are now at least ten SDN startups with valuations over $100 million. As I reported in April, a recent investment in Cumulus Networks pushed up the valuation of the private company north of $300 million, according to industry sources. Big Switch, which did a deal in 2012 valuing it near $170 million, took money from Intel in 2013, most likely boosting its valuation to over $200 million, according to several sources.

Related Articles

How to Effectively Embed SDN in the Enterprise

NFV and SDN: What's the Difference Two Years Later?

sFlow Creator Peter Phaal On Taming The Wilds Of SDN & Virtual Networking

Featured Article: Bringing Data-Driven SDN to the Network Edge

NFV Delivers Pervasive Intelligence for MNOs

Communication Networks

Course organization

Our website: http://comm-net.ethz.ch

check it out regularly

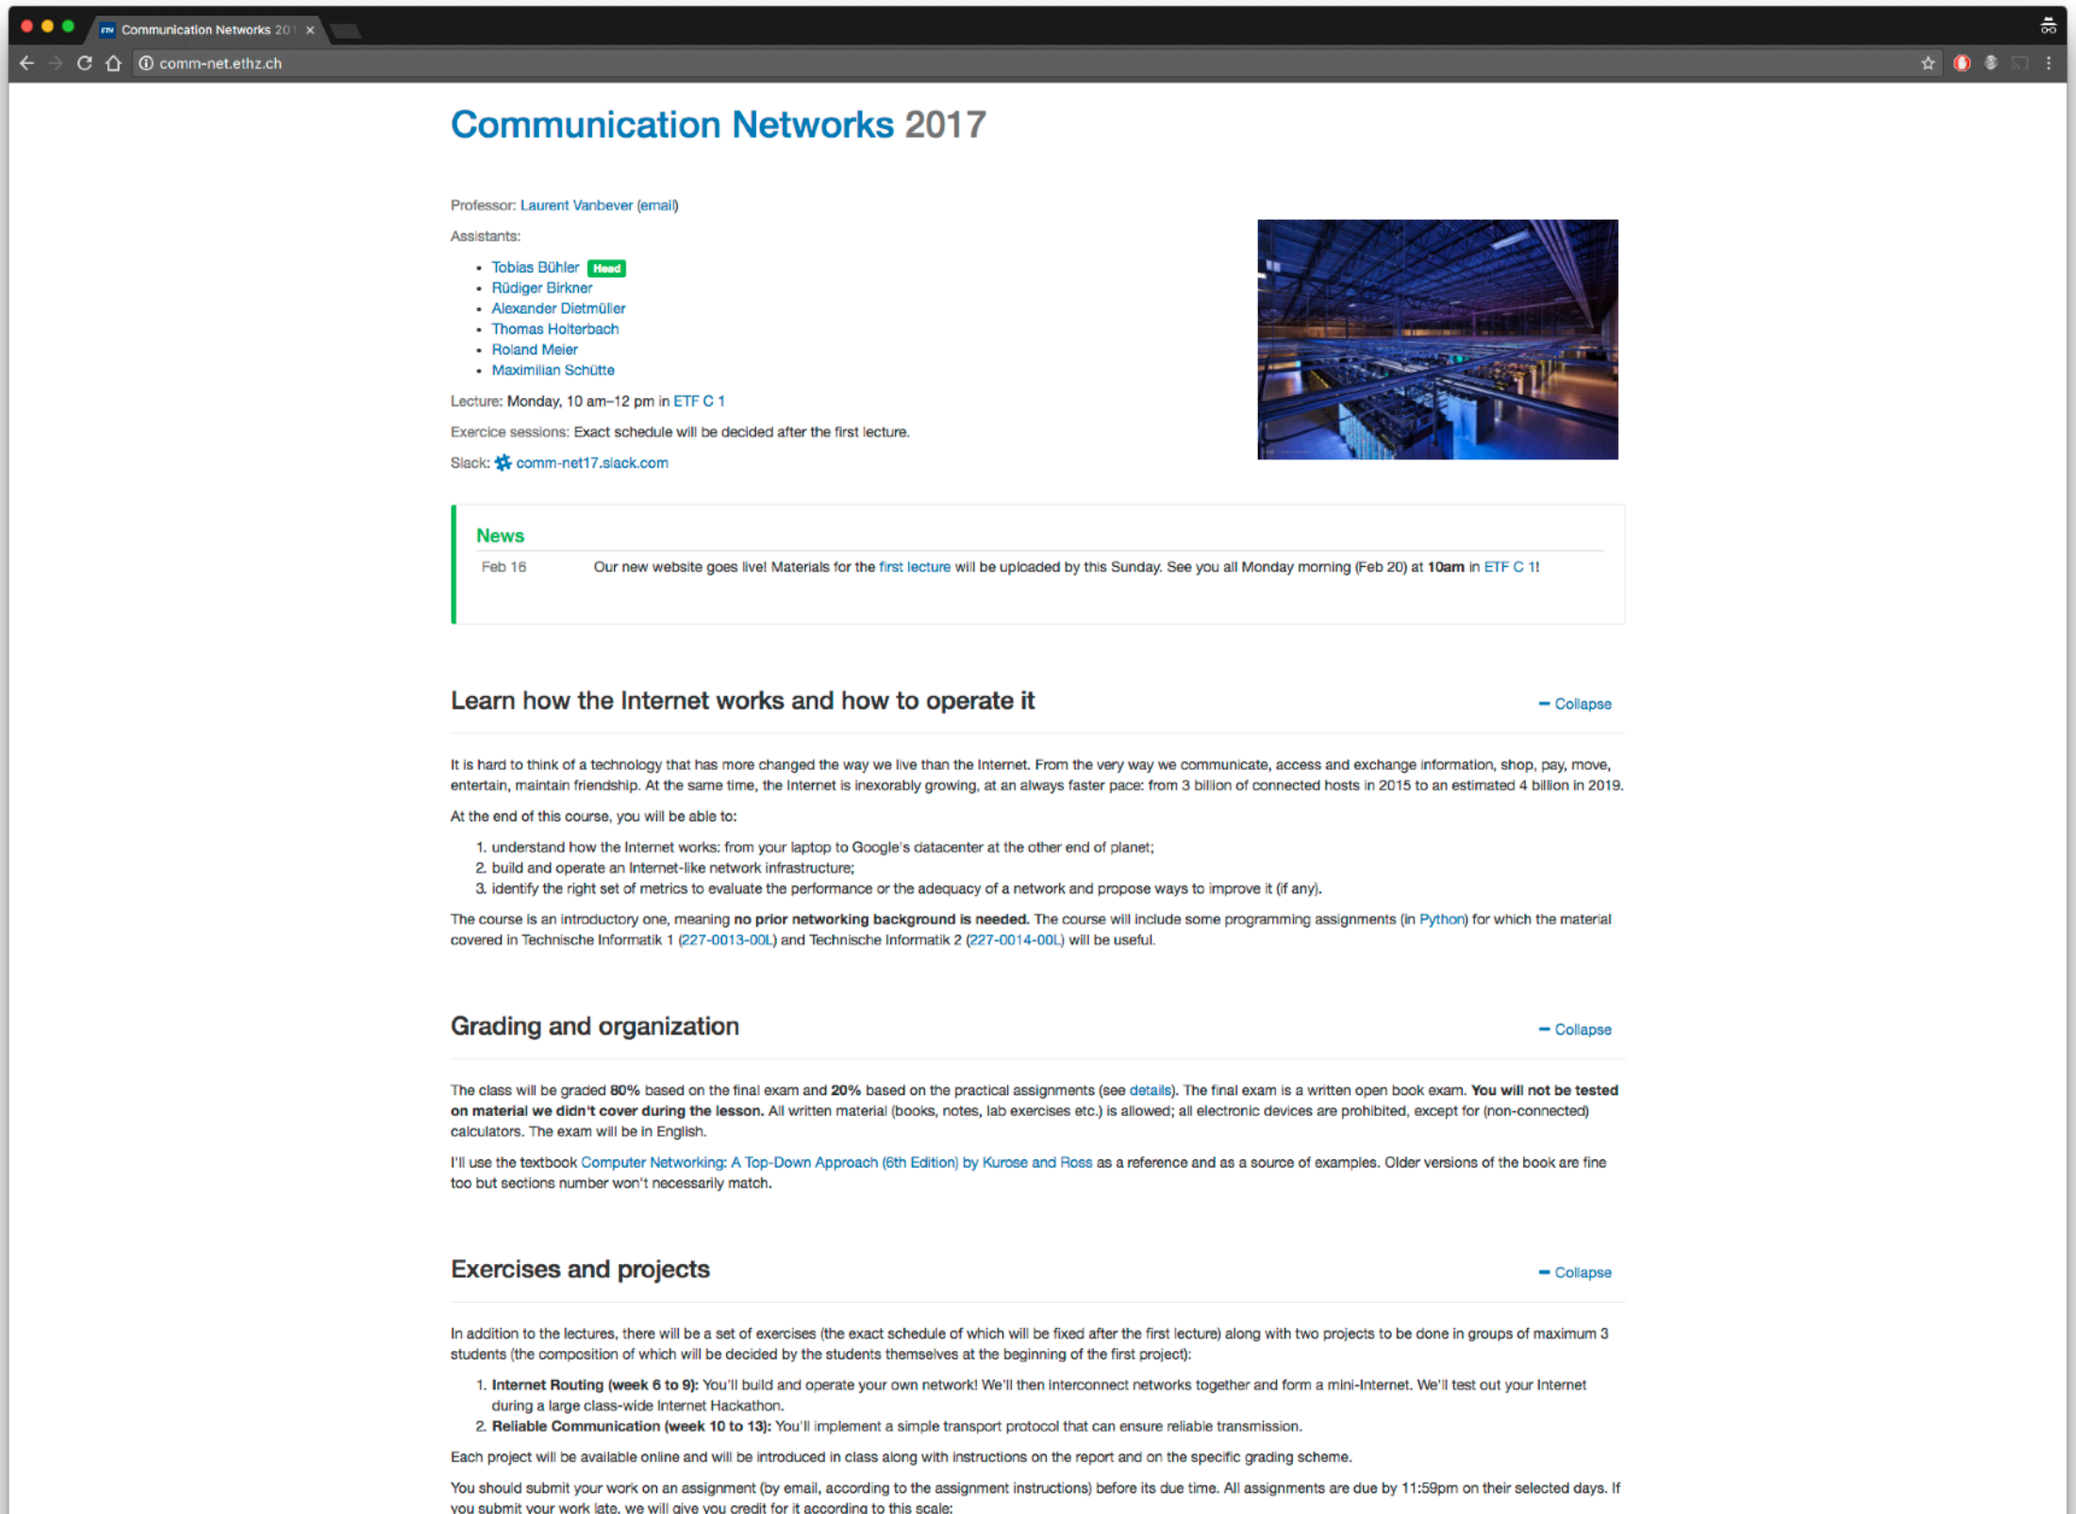

The course will be split into three parts

Part 1	Part 2	Part 3
Overview	Concepts	Today's Internet
2 lectures	2 lectures	~9 lectures

The lectures will be accompanied by exercises,
there will be two sessions per week (to ease scheduling)

available slots	Tue	8-10	vote ***today*** on comm-net.ethz.ch
		10-12	
		13-15	
		15-17	
	Wed	13-15	
		15-17	
	Thur	13-15	
		15-17	

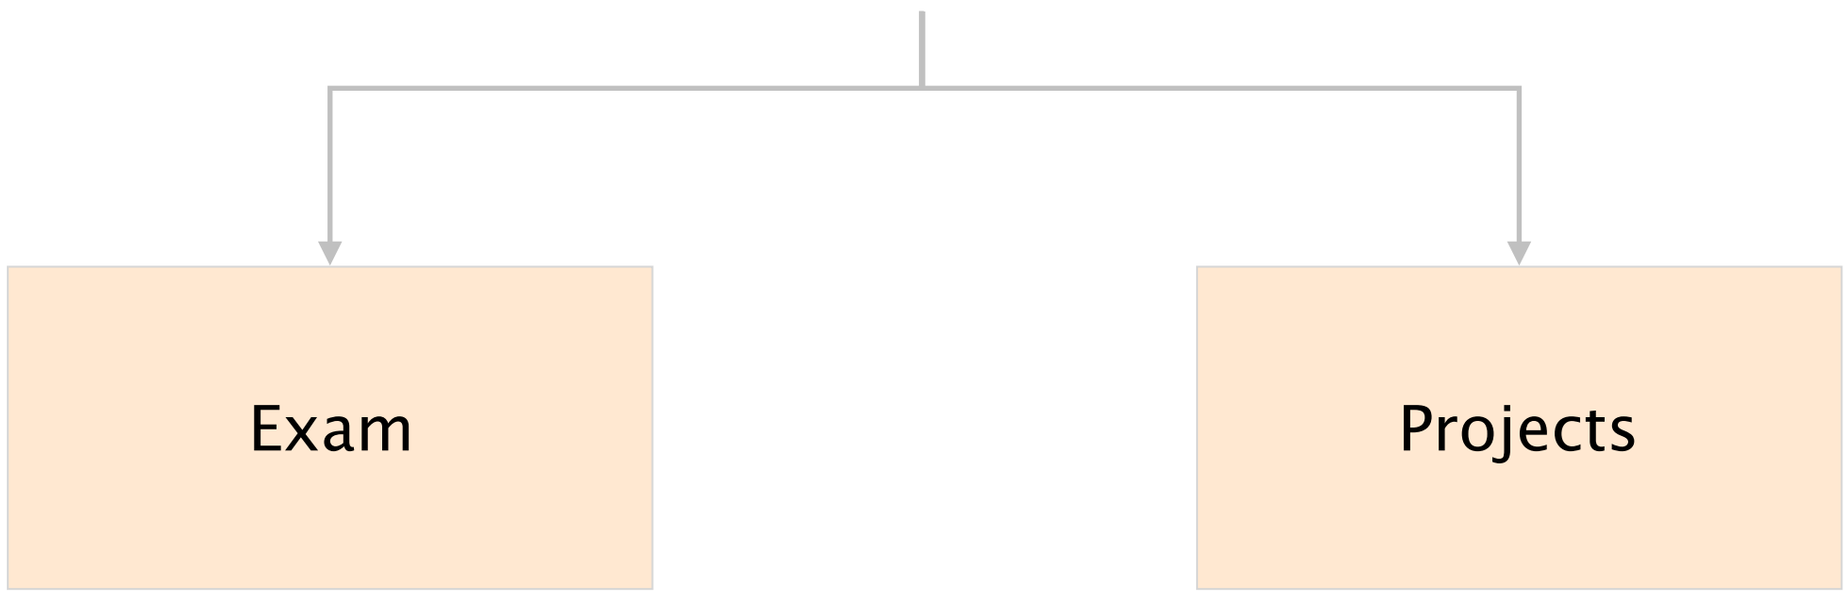

80%

written, open book

20%

mandatory

Your final grade

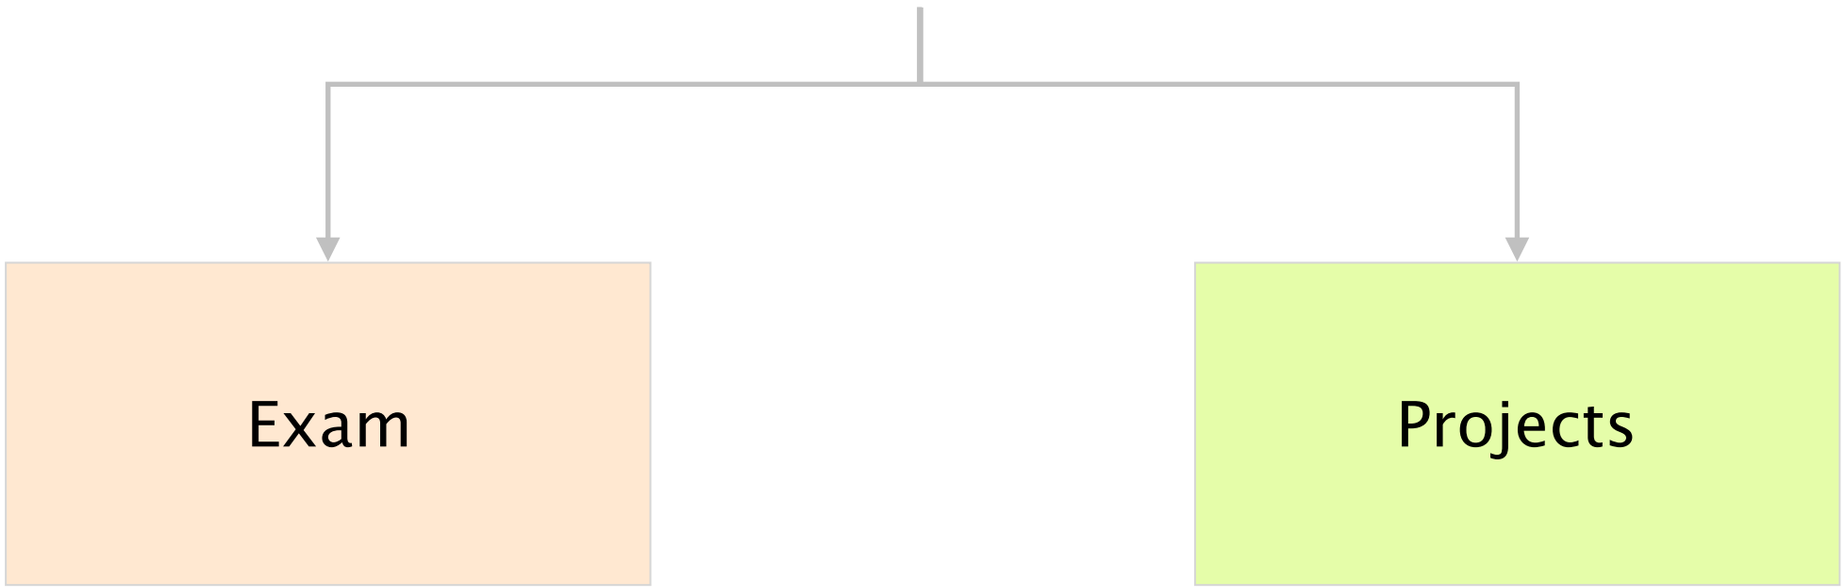

80%

written, open book

20%

mandatory

There will be two practical projects,
to be done in group of maximum three students

#1 Build and operate a small Internet

#2 Implement an interoperable reliable protocol

Detailed instructions will follow

“Internet Hackathon”

sometime in week 8

2016 edition

The course follows the textbook

Computer Networking: a Top-Down Approach

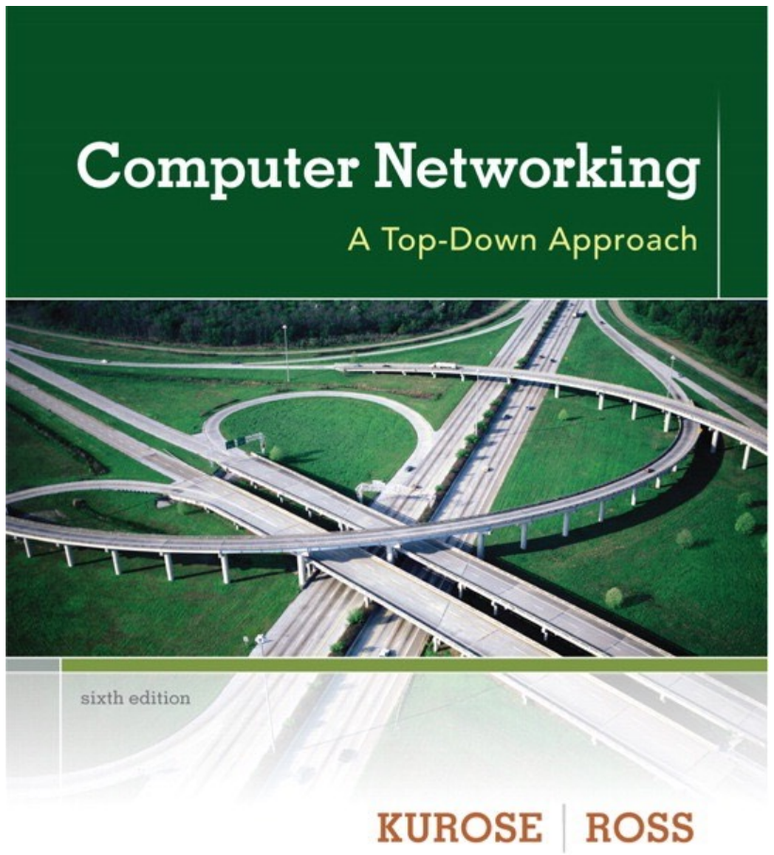

6th edition

ok to use the 5th

see sections indicated

on comm-net.ethz.ch

We'll use Slack (chat client)

to discuss about the course and assignments

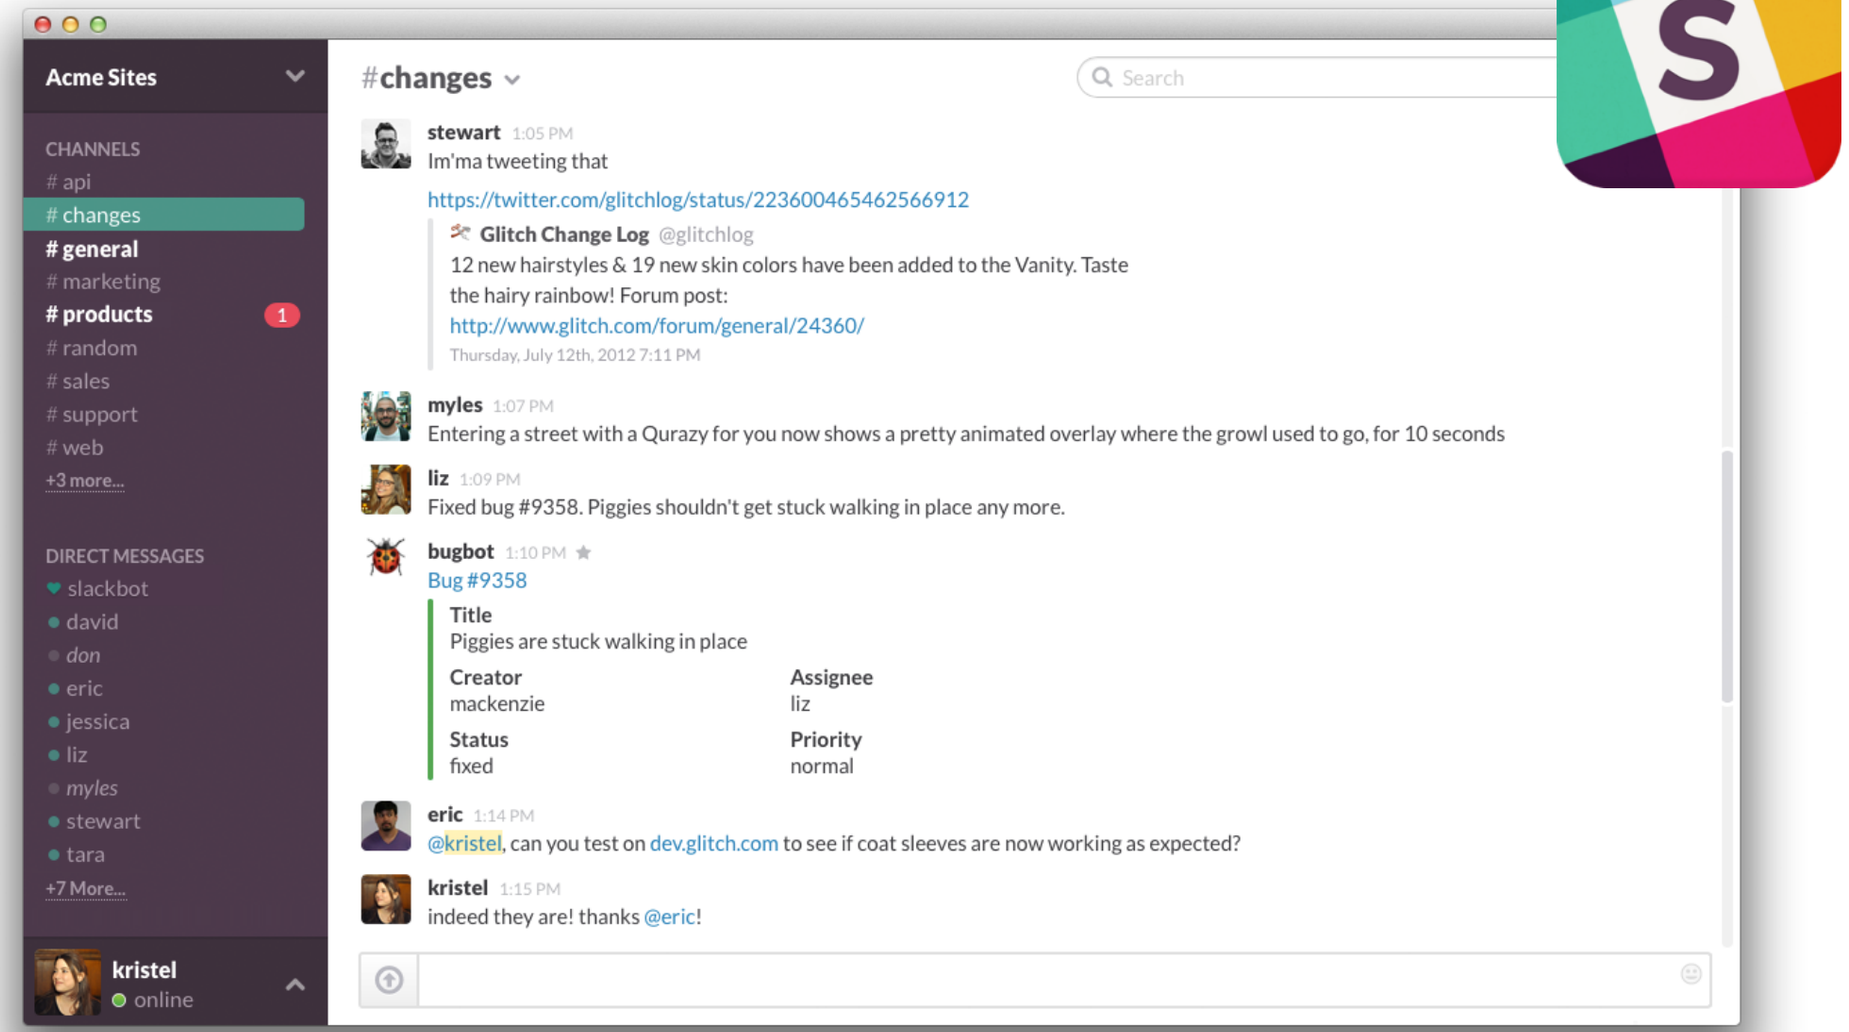

Web, smartphone and desktop clients available

Register today using your *real* name

> https://comm-net17.slack.com/signup

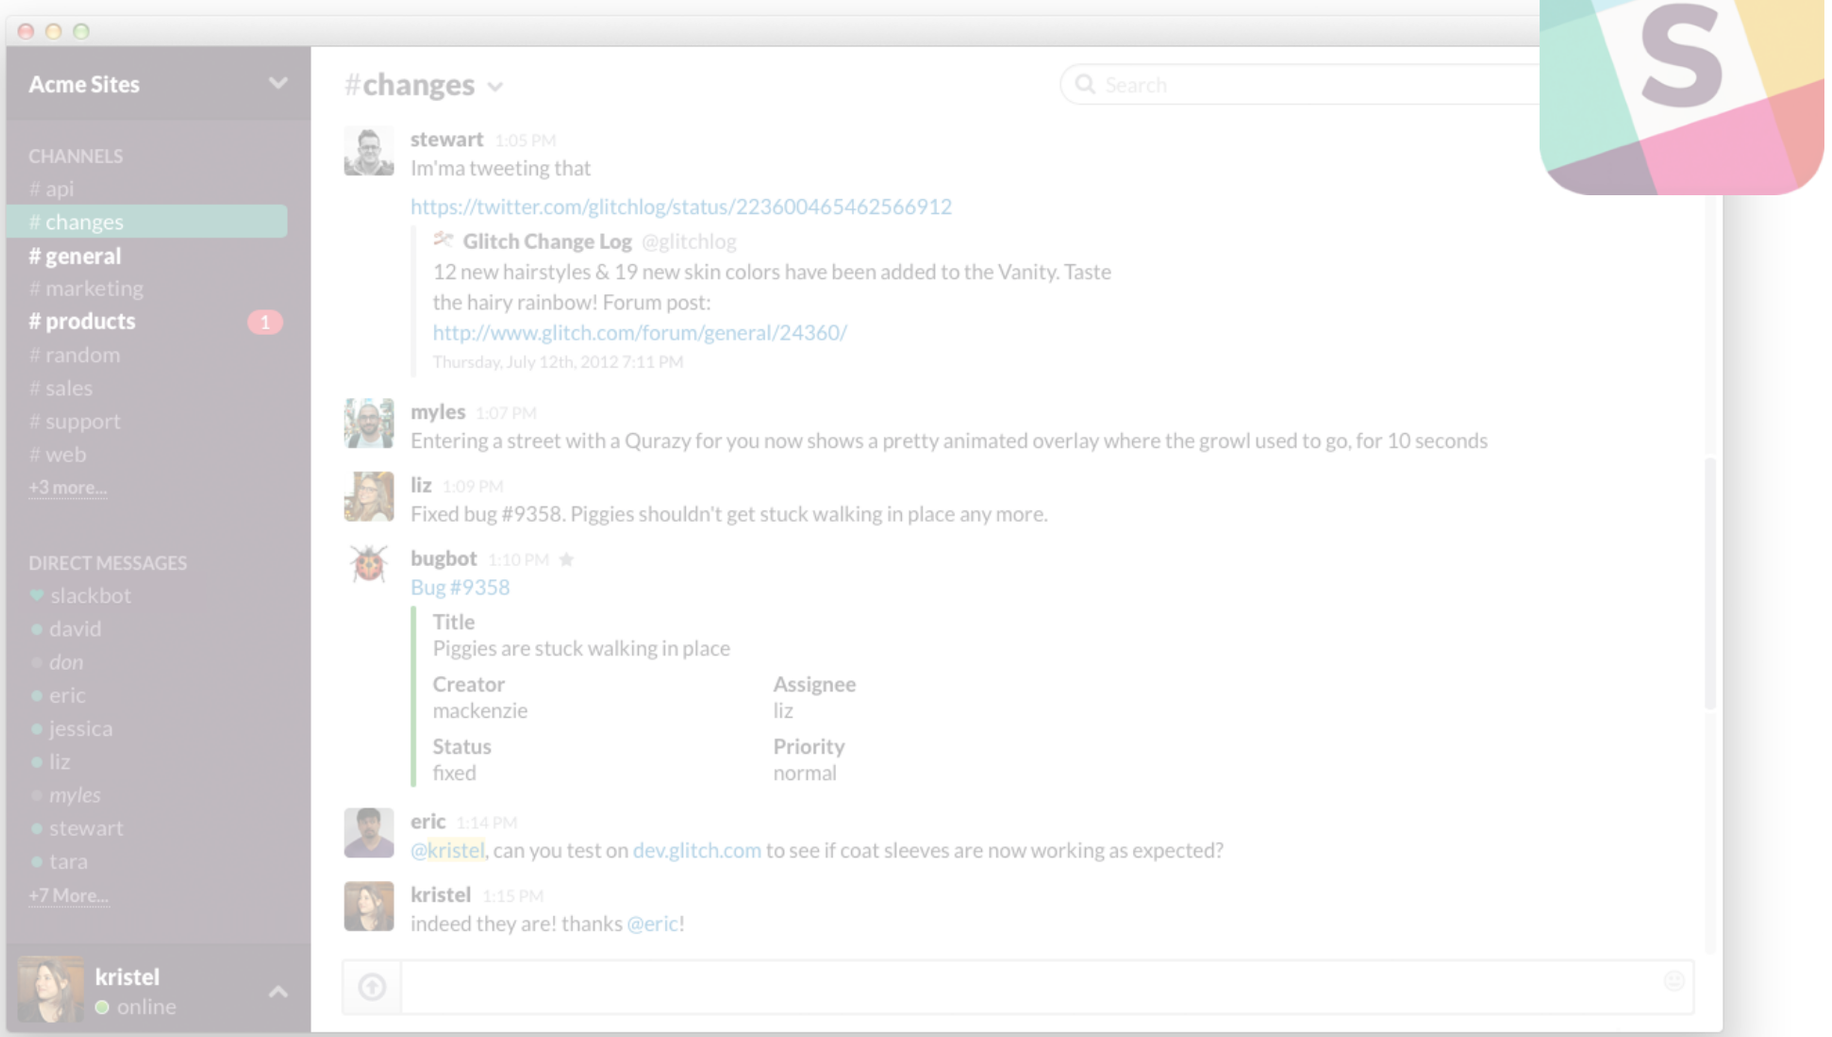

Web, smartphone and desktop clients available

Communication Networks

List any

technologies, principles, applications...

used after typing in:

> www.google.ch

and pressing enter in your browser

Communication Networks

Part 1: Overview

#1 What is a network made of?

#2 How is it shared?

#3 How is it organized?

#4 How does communication happen?

#5 How do we characterize it?

Communication Networks

Part 1: Overview

#1 What is a network made of?

How is it shared?

How is it organized?

How does communication happen?

How do we characterize it?

Networks are composed of three basic components

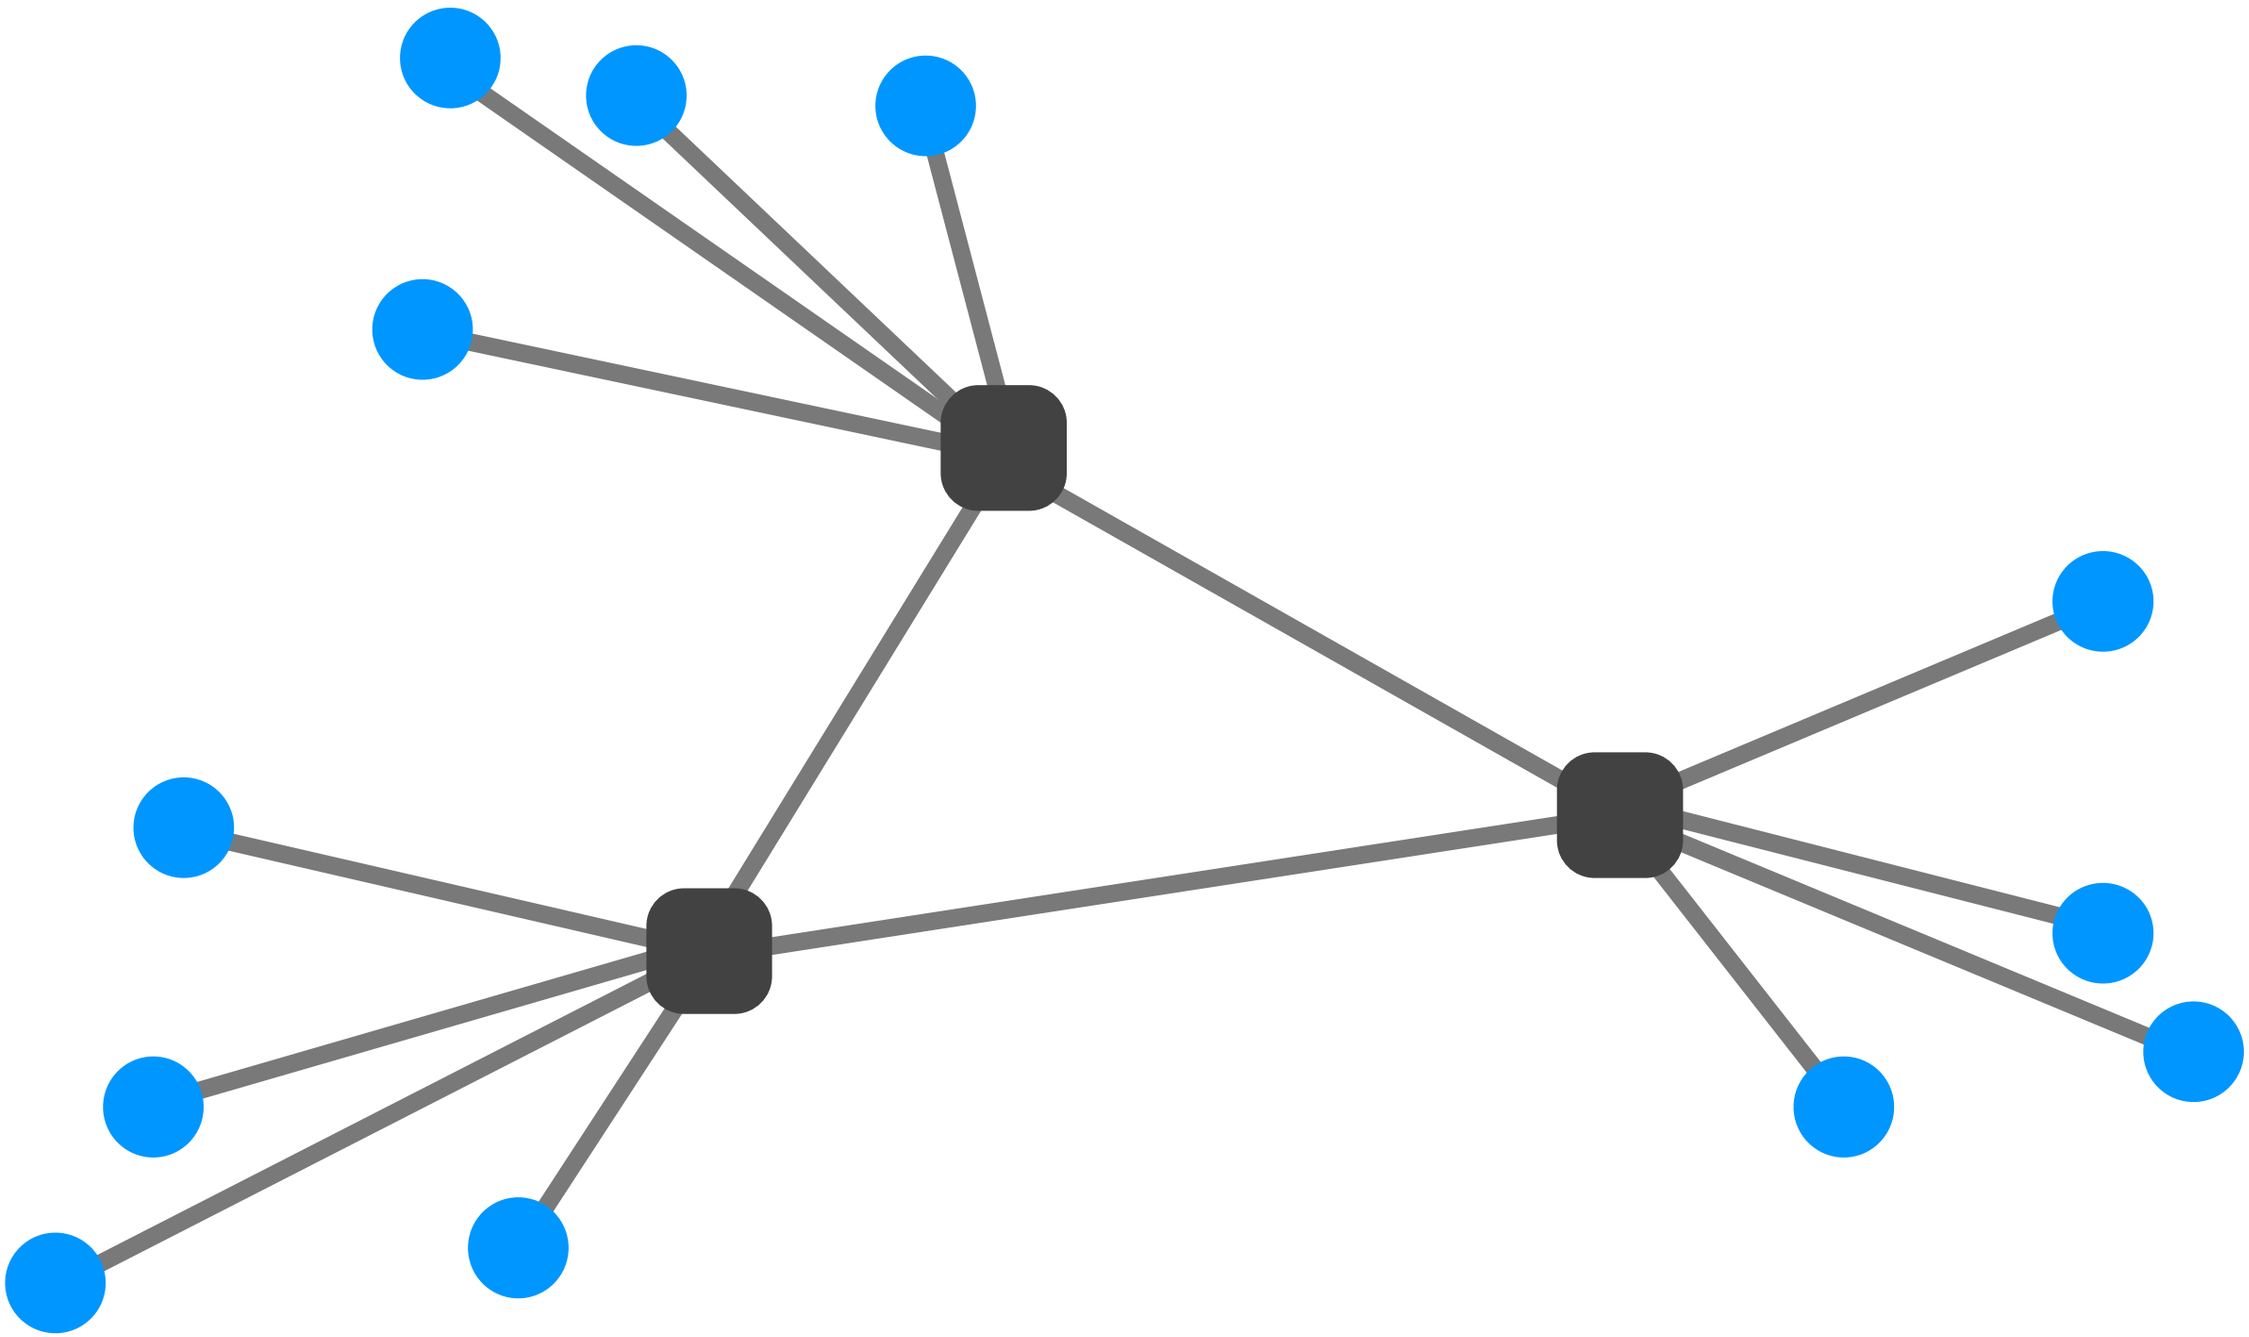

End-systems send & receive data

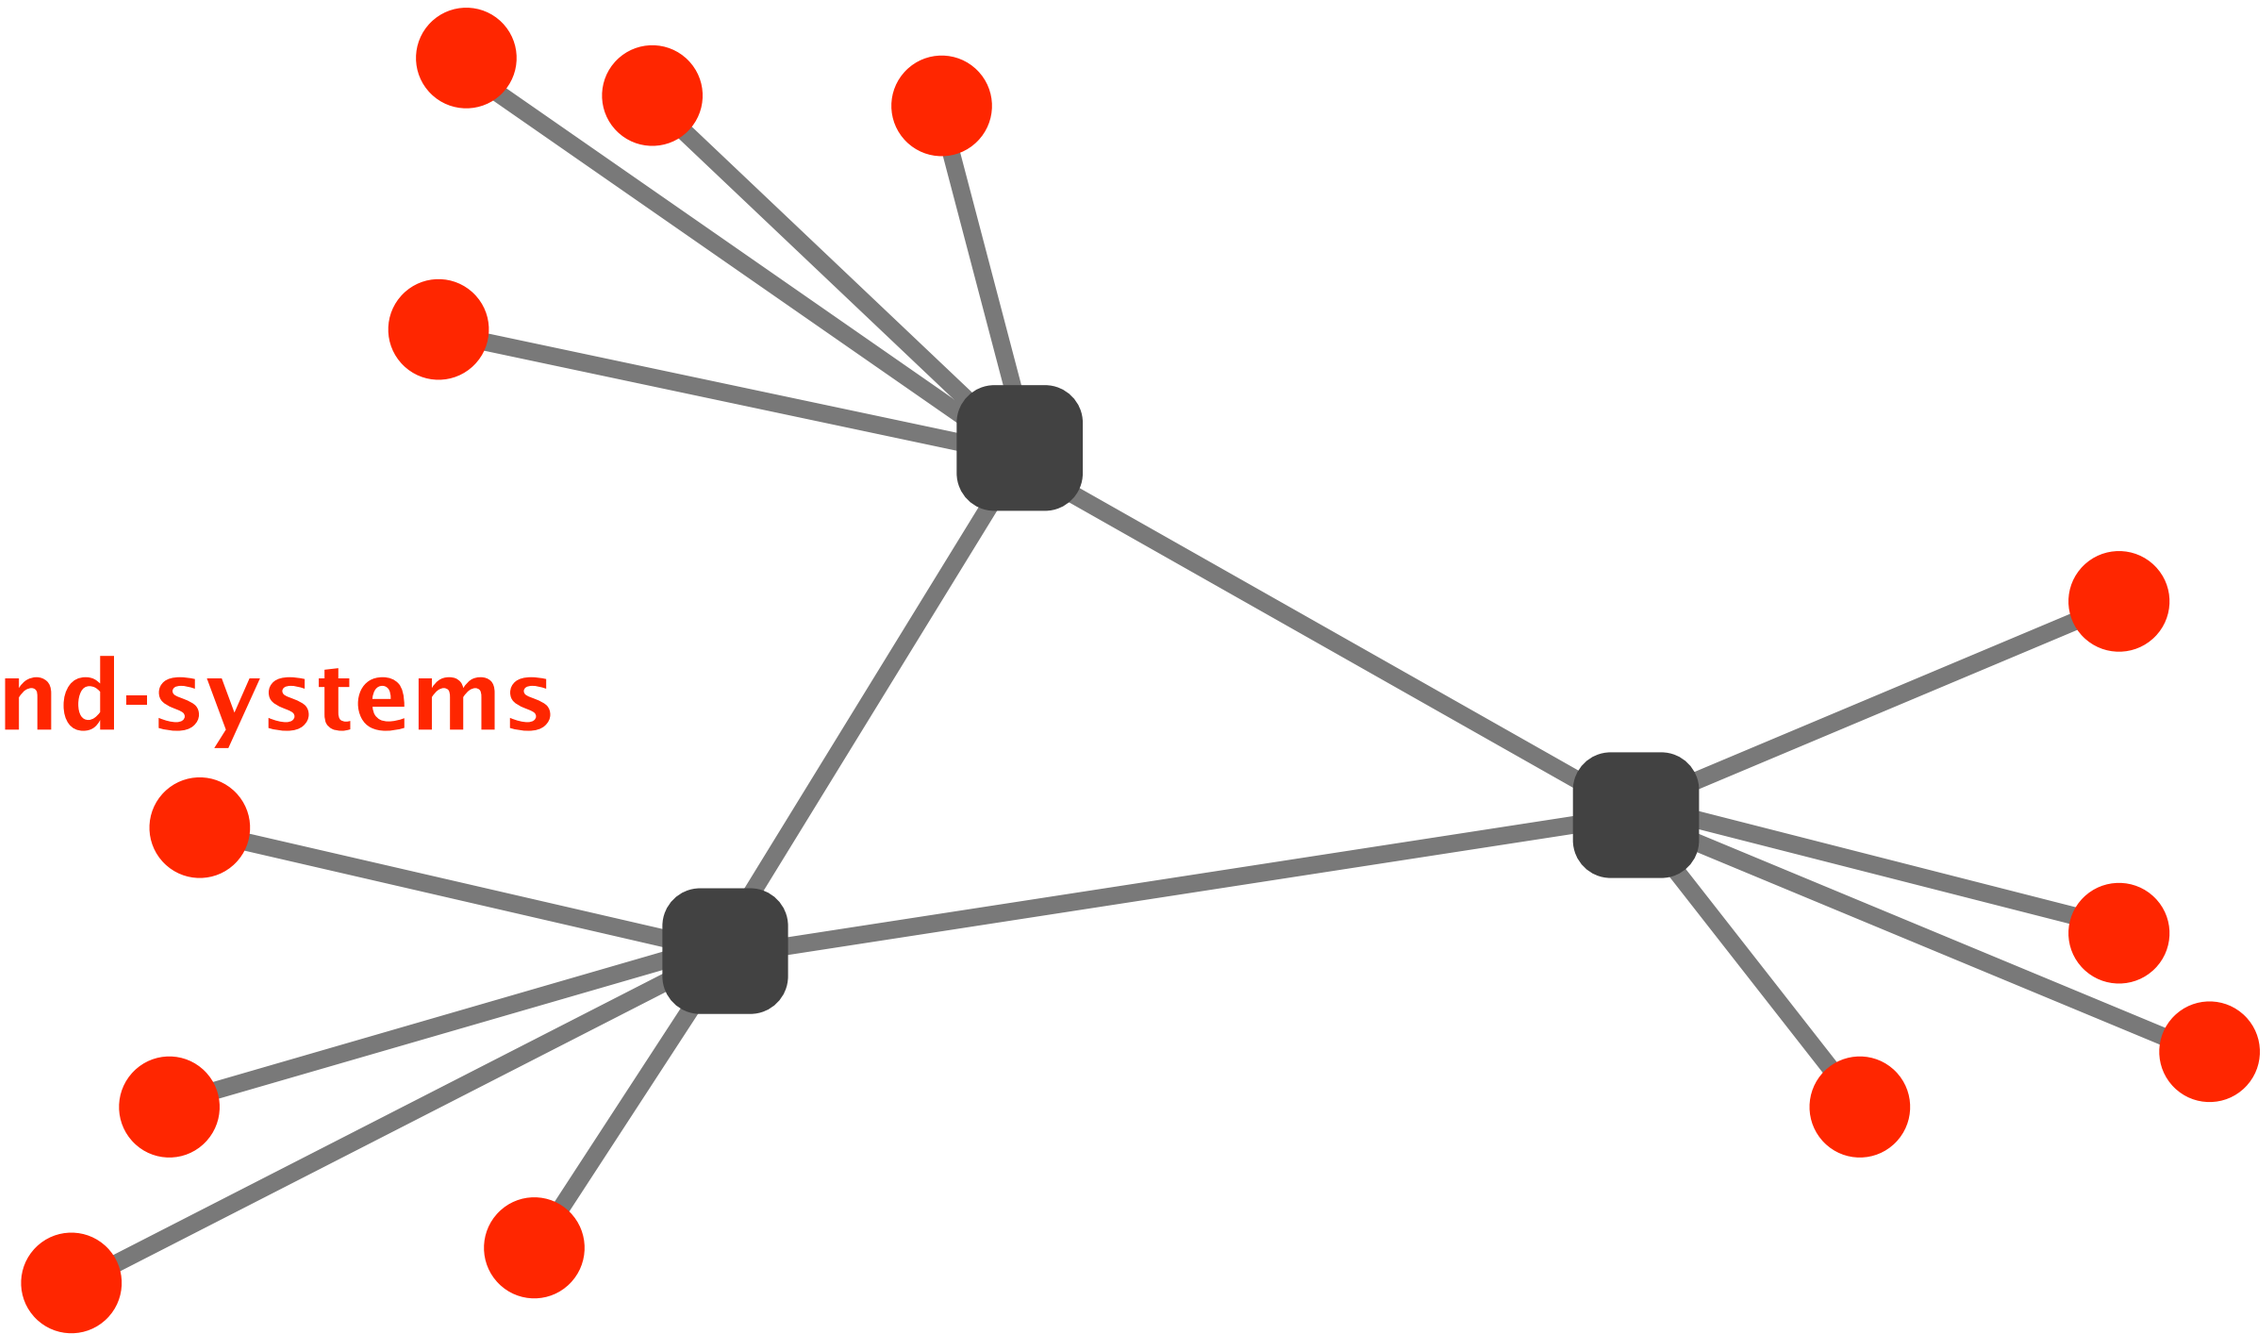

End-systems come in a wide-variety

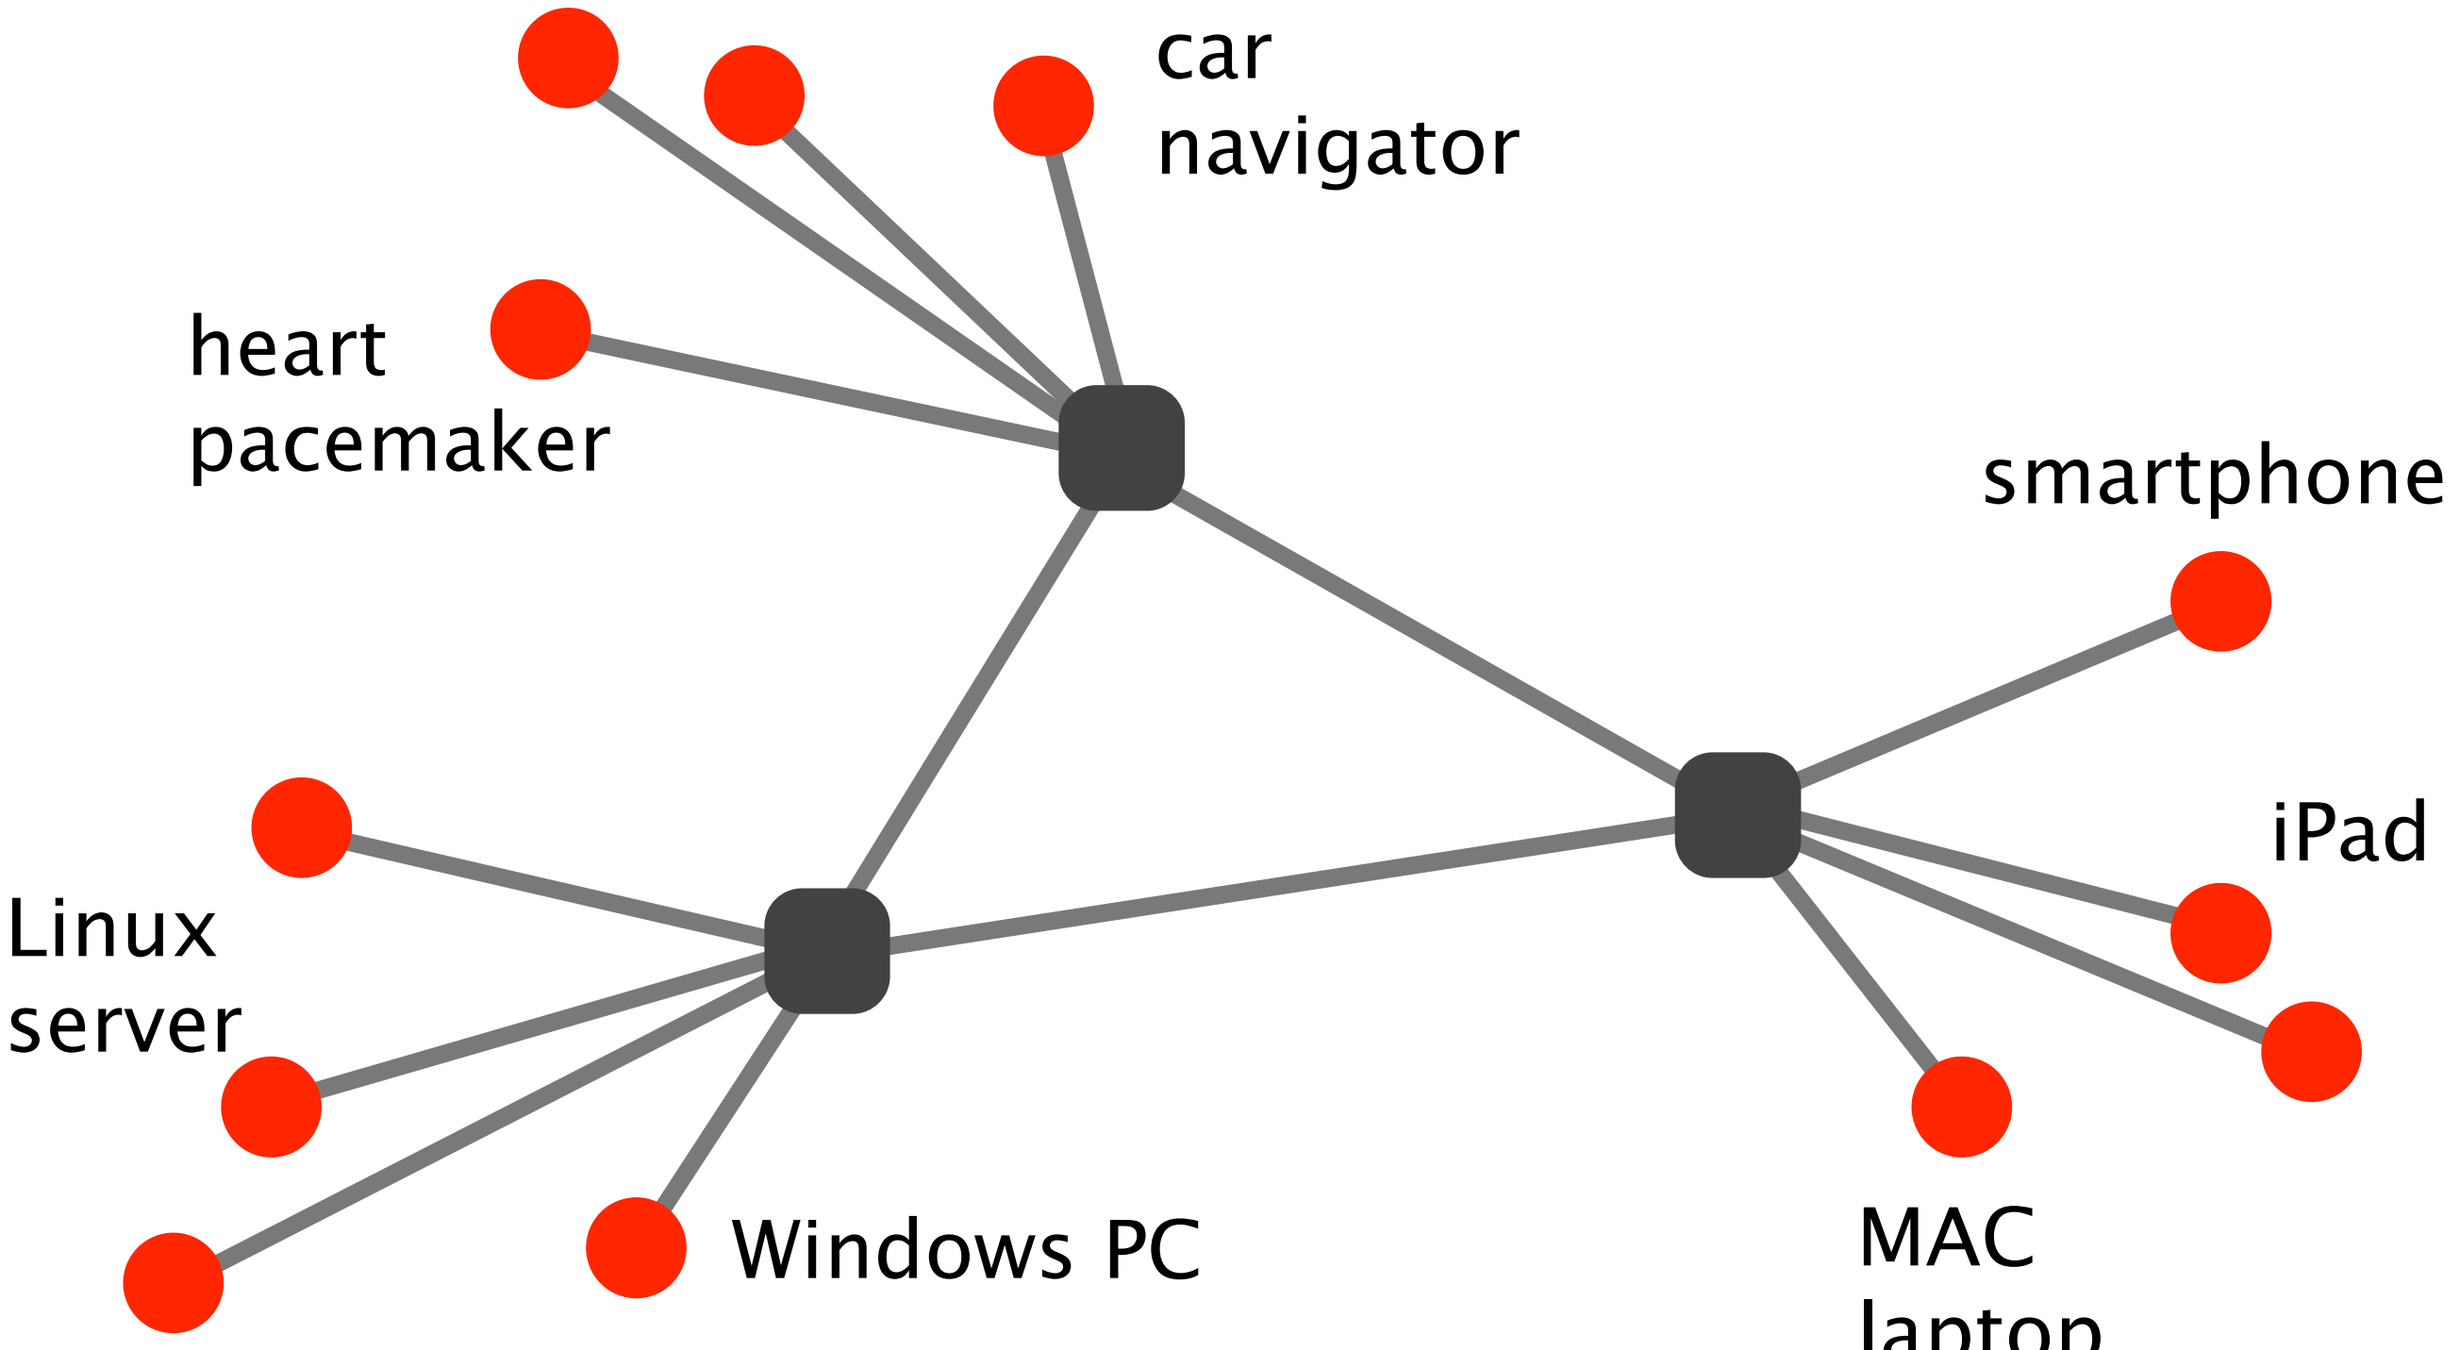

Switches & routers forward data to the destination

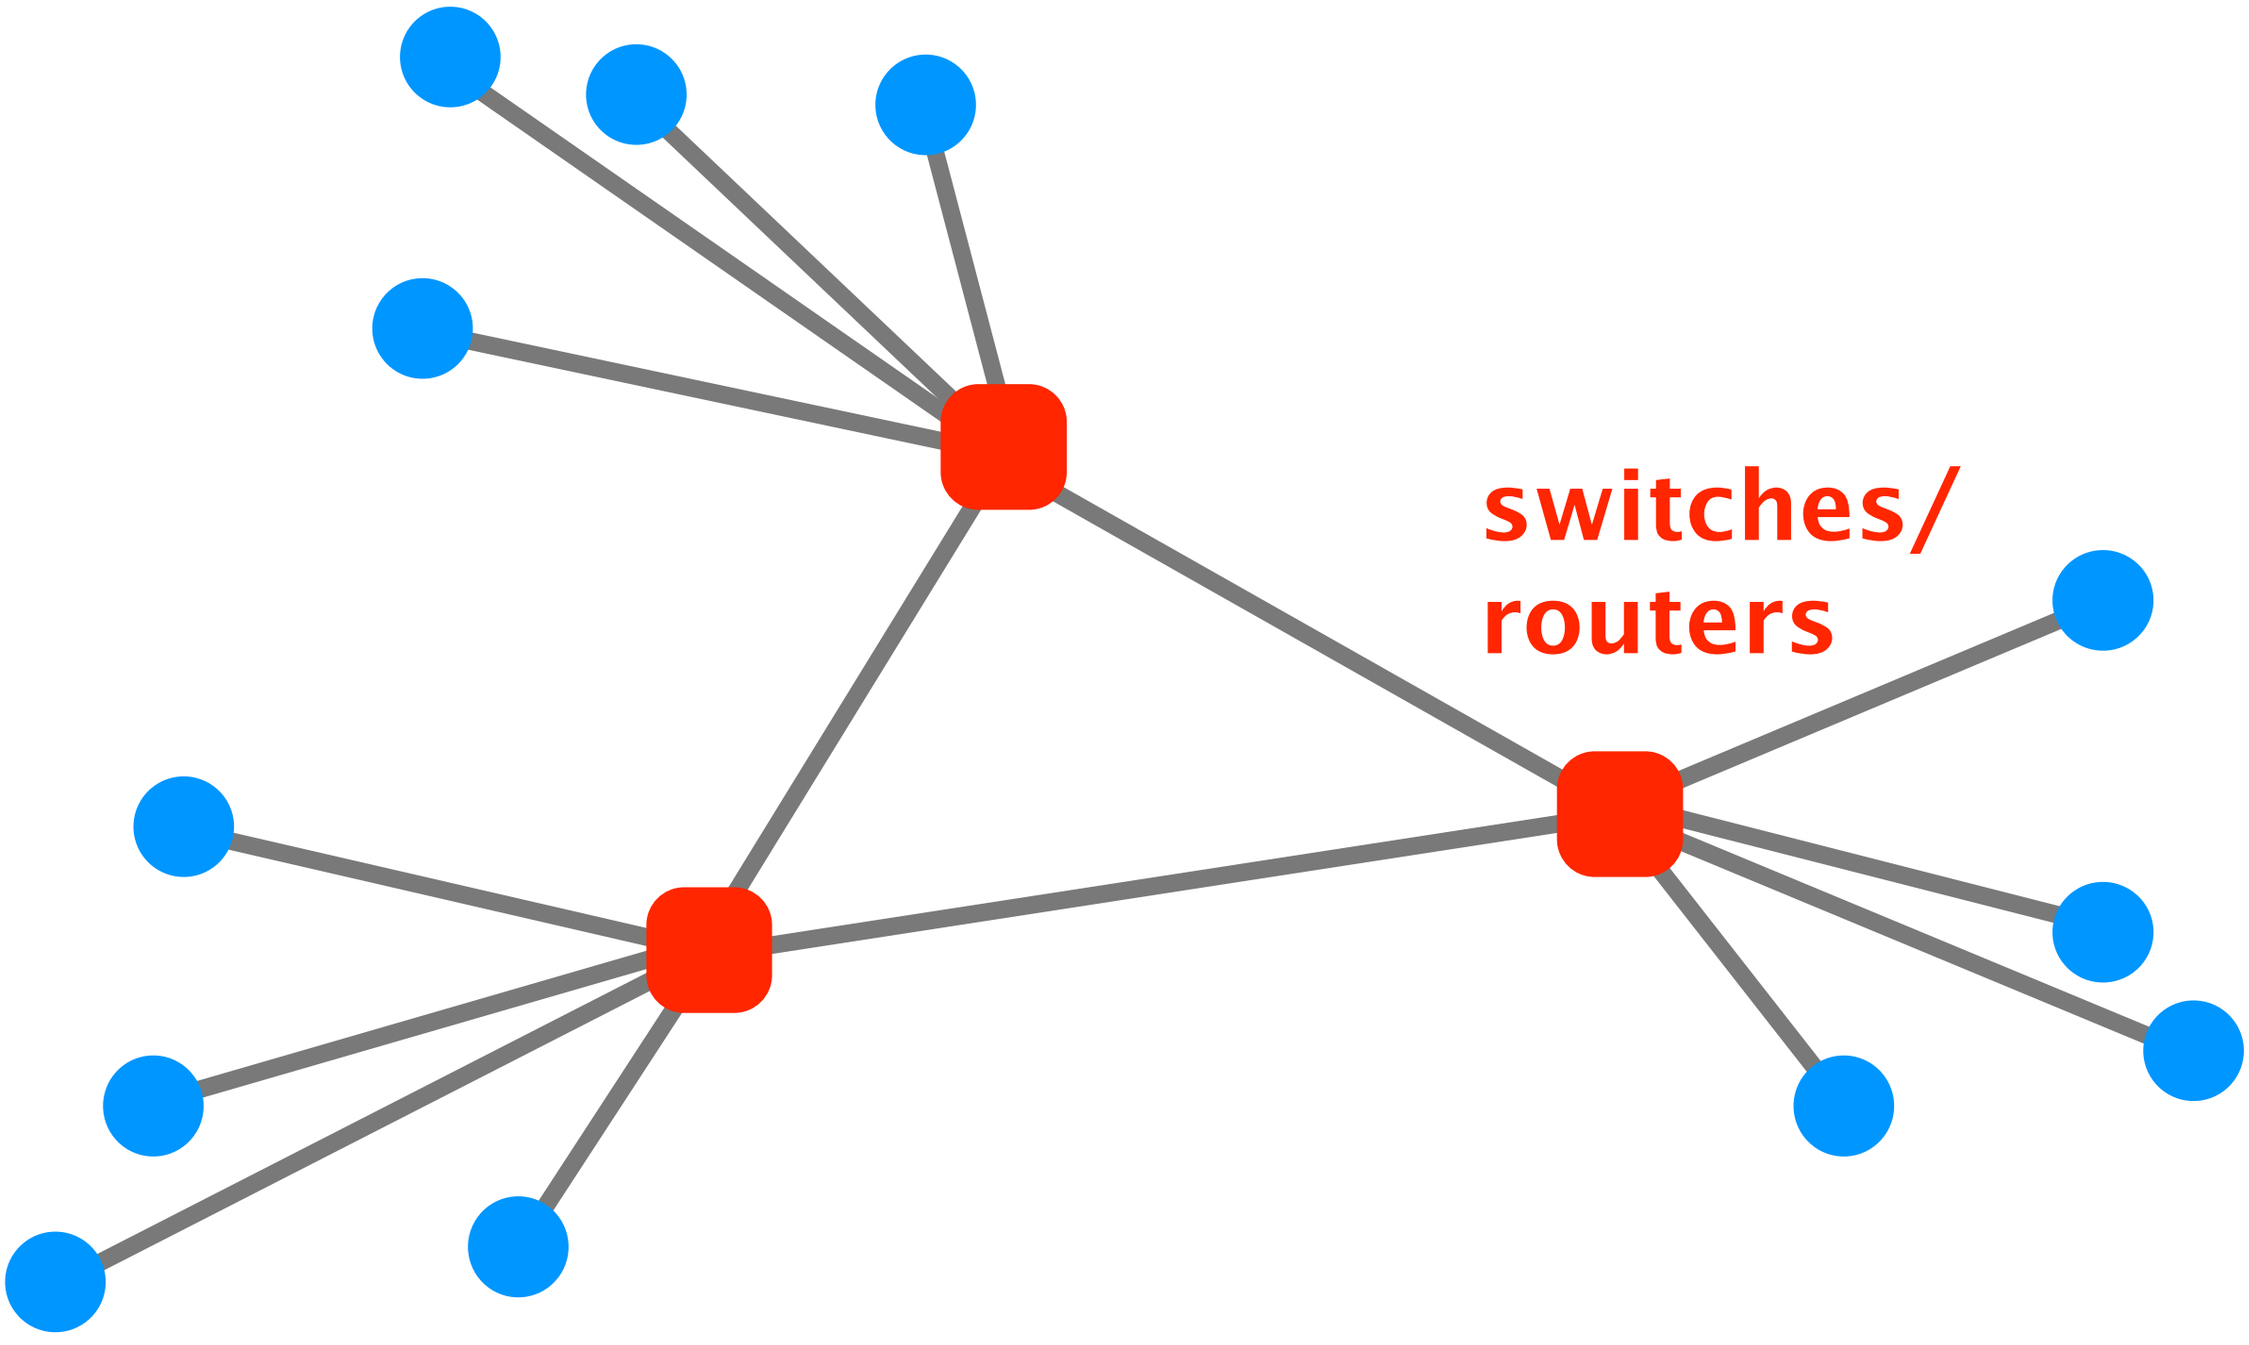

Routers vary in size and usage

Home router

~20 cm

0,5 kg

1 Gbps

Internet core router

>200cm

700kg

1.2 Tbps

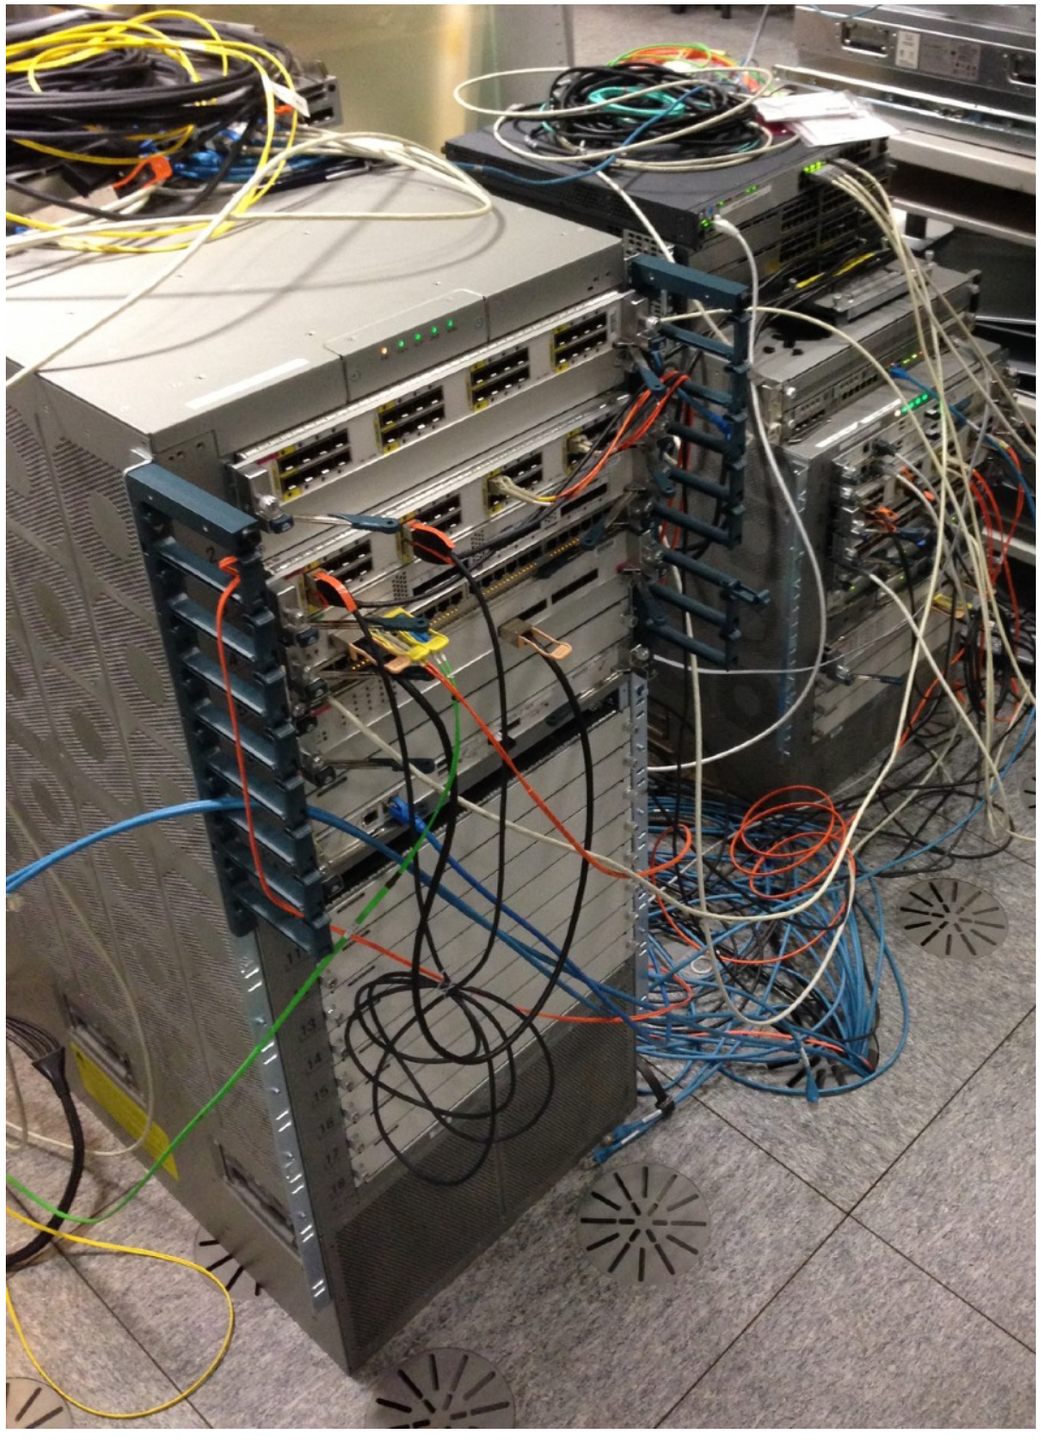

Cisco Nexus 7k

Routers @ETHZ

~25 deployed

Links connect end-systems to switches and switches to each other

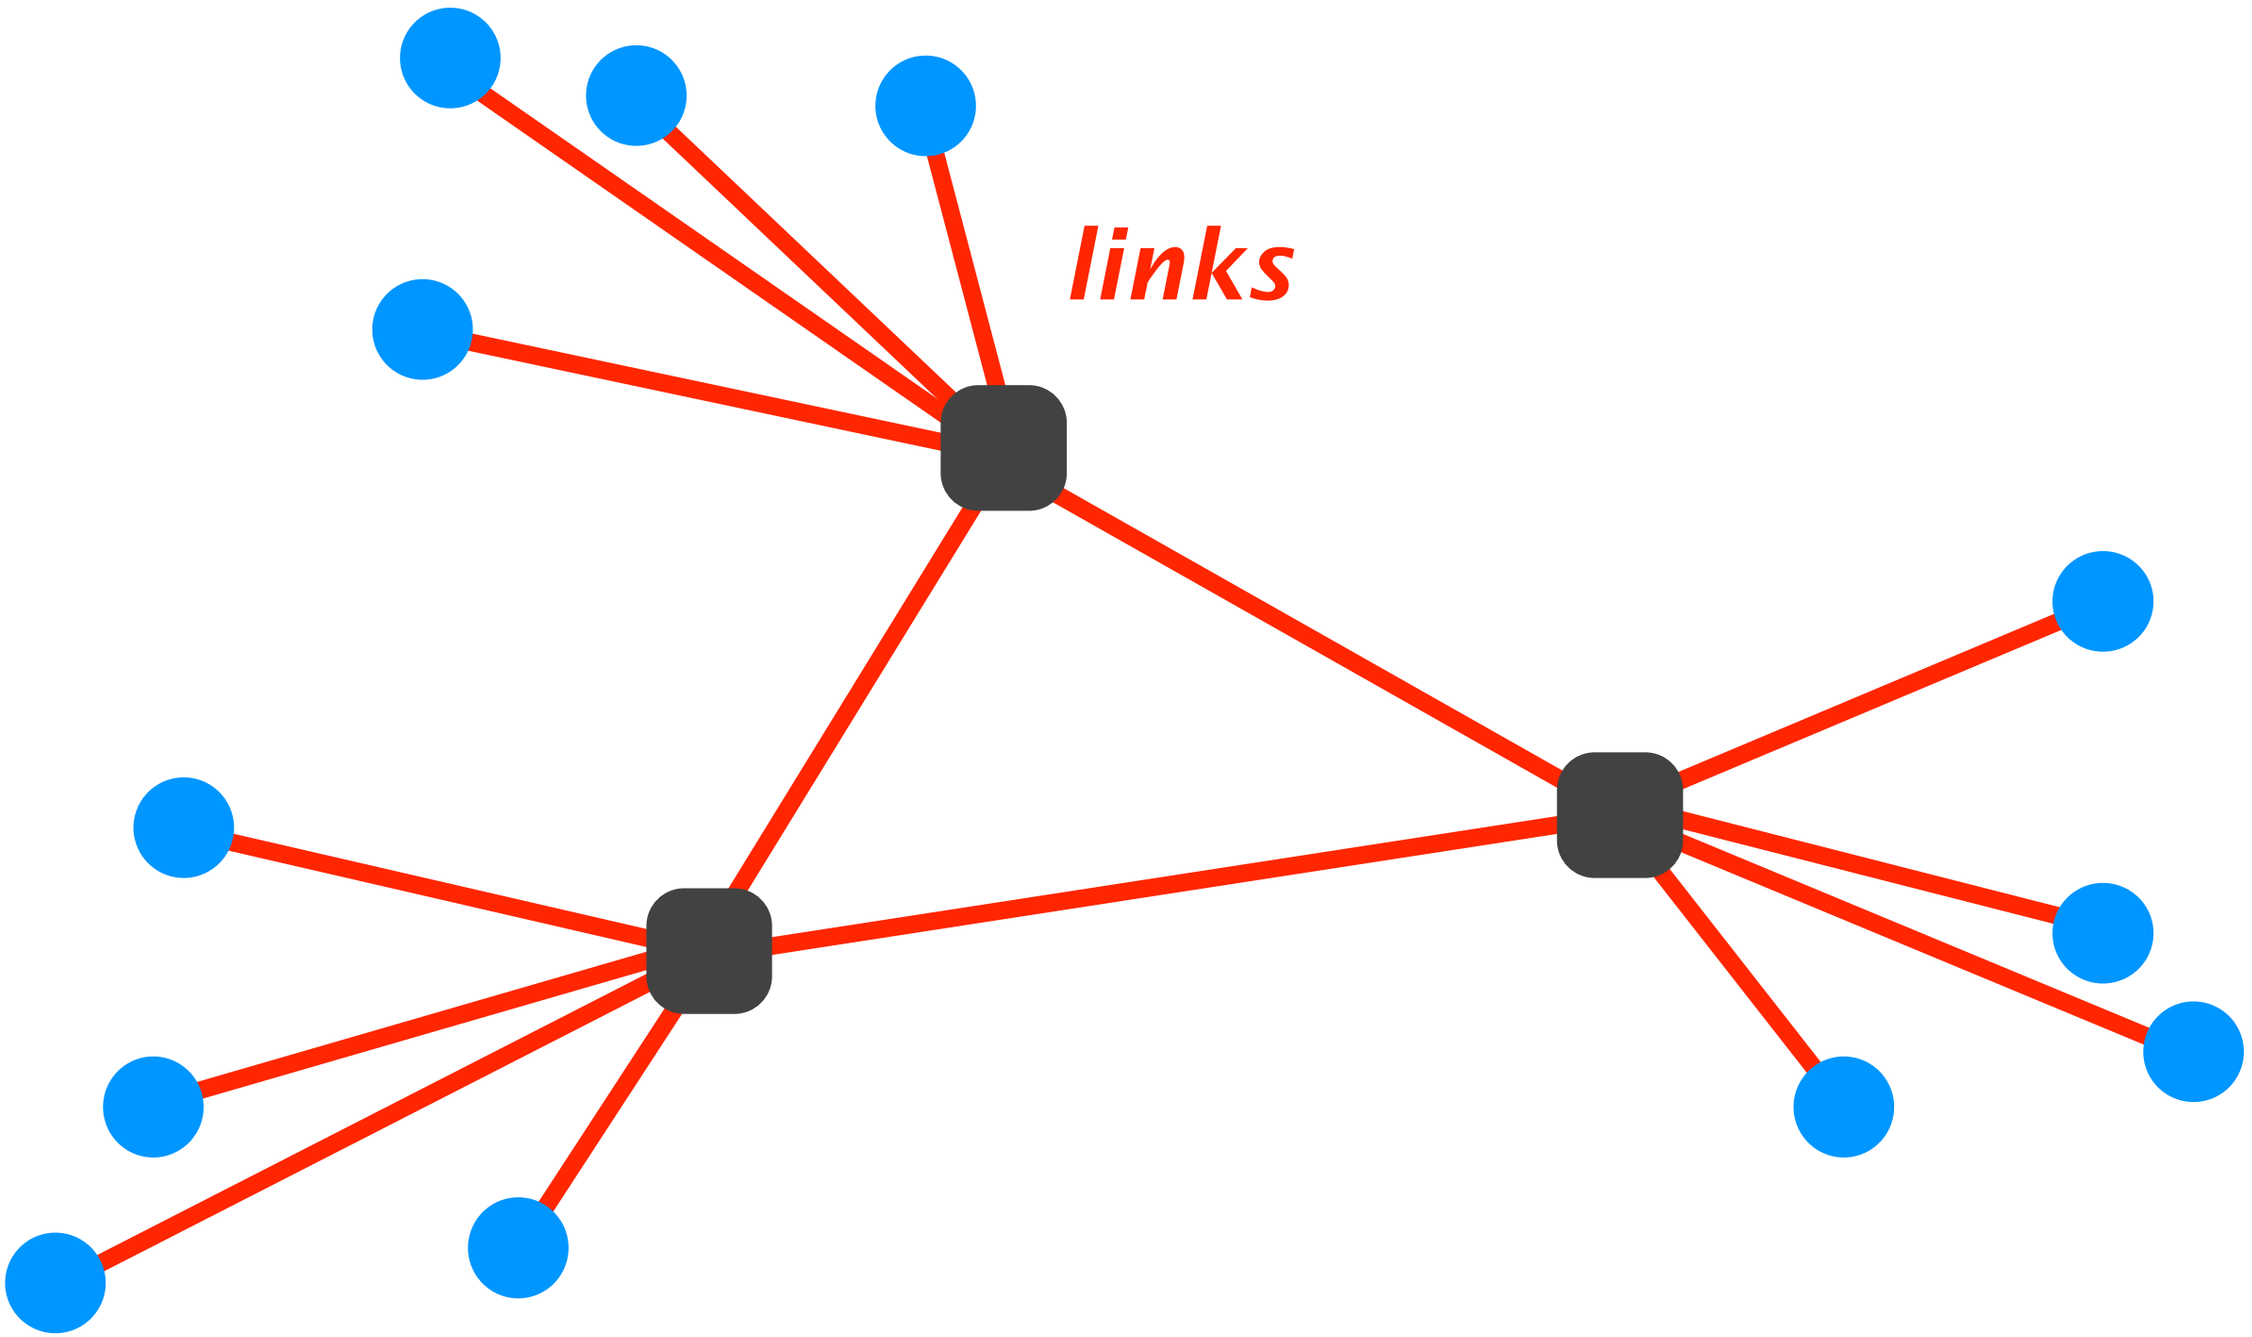

Links, too, vary in size and usage

Copper

ADSL, RJ-45,…

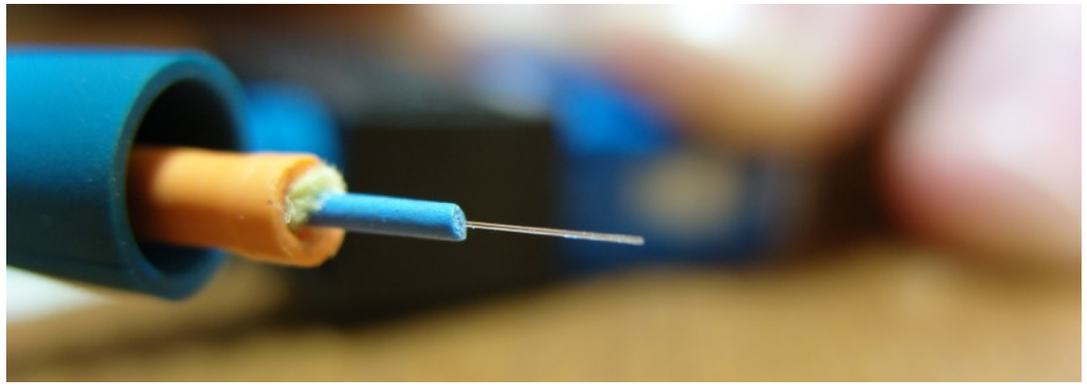

Optical fibers

Wireless link

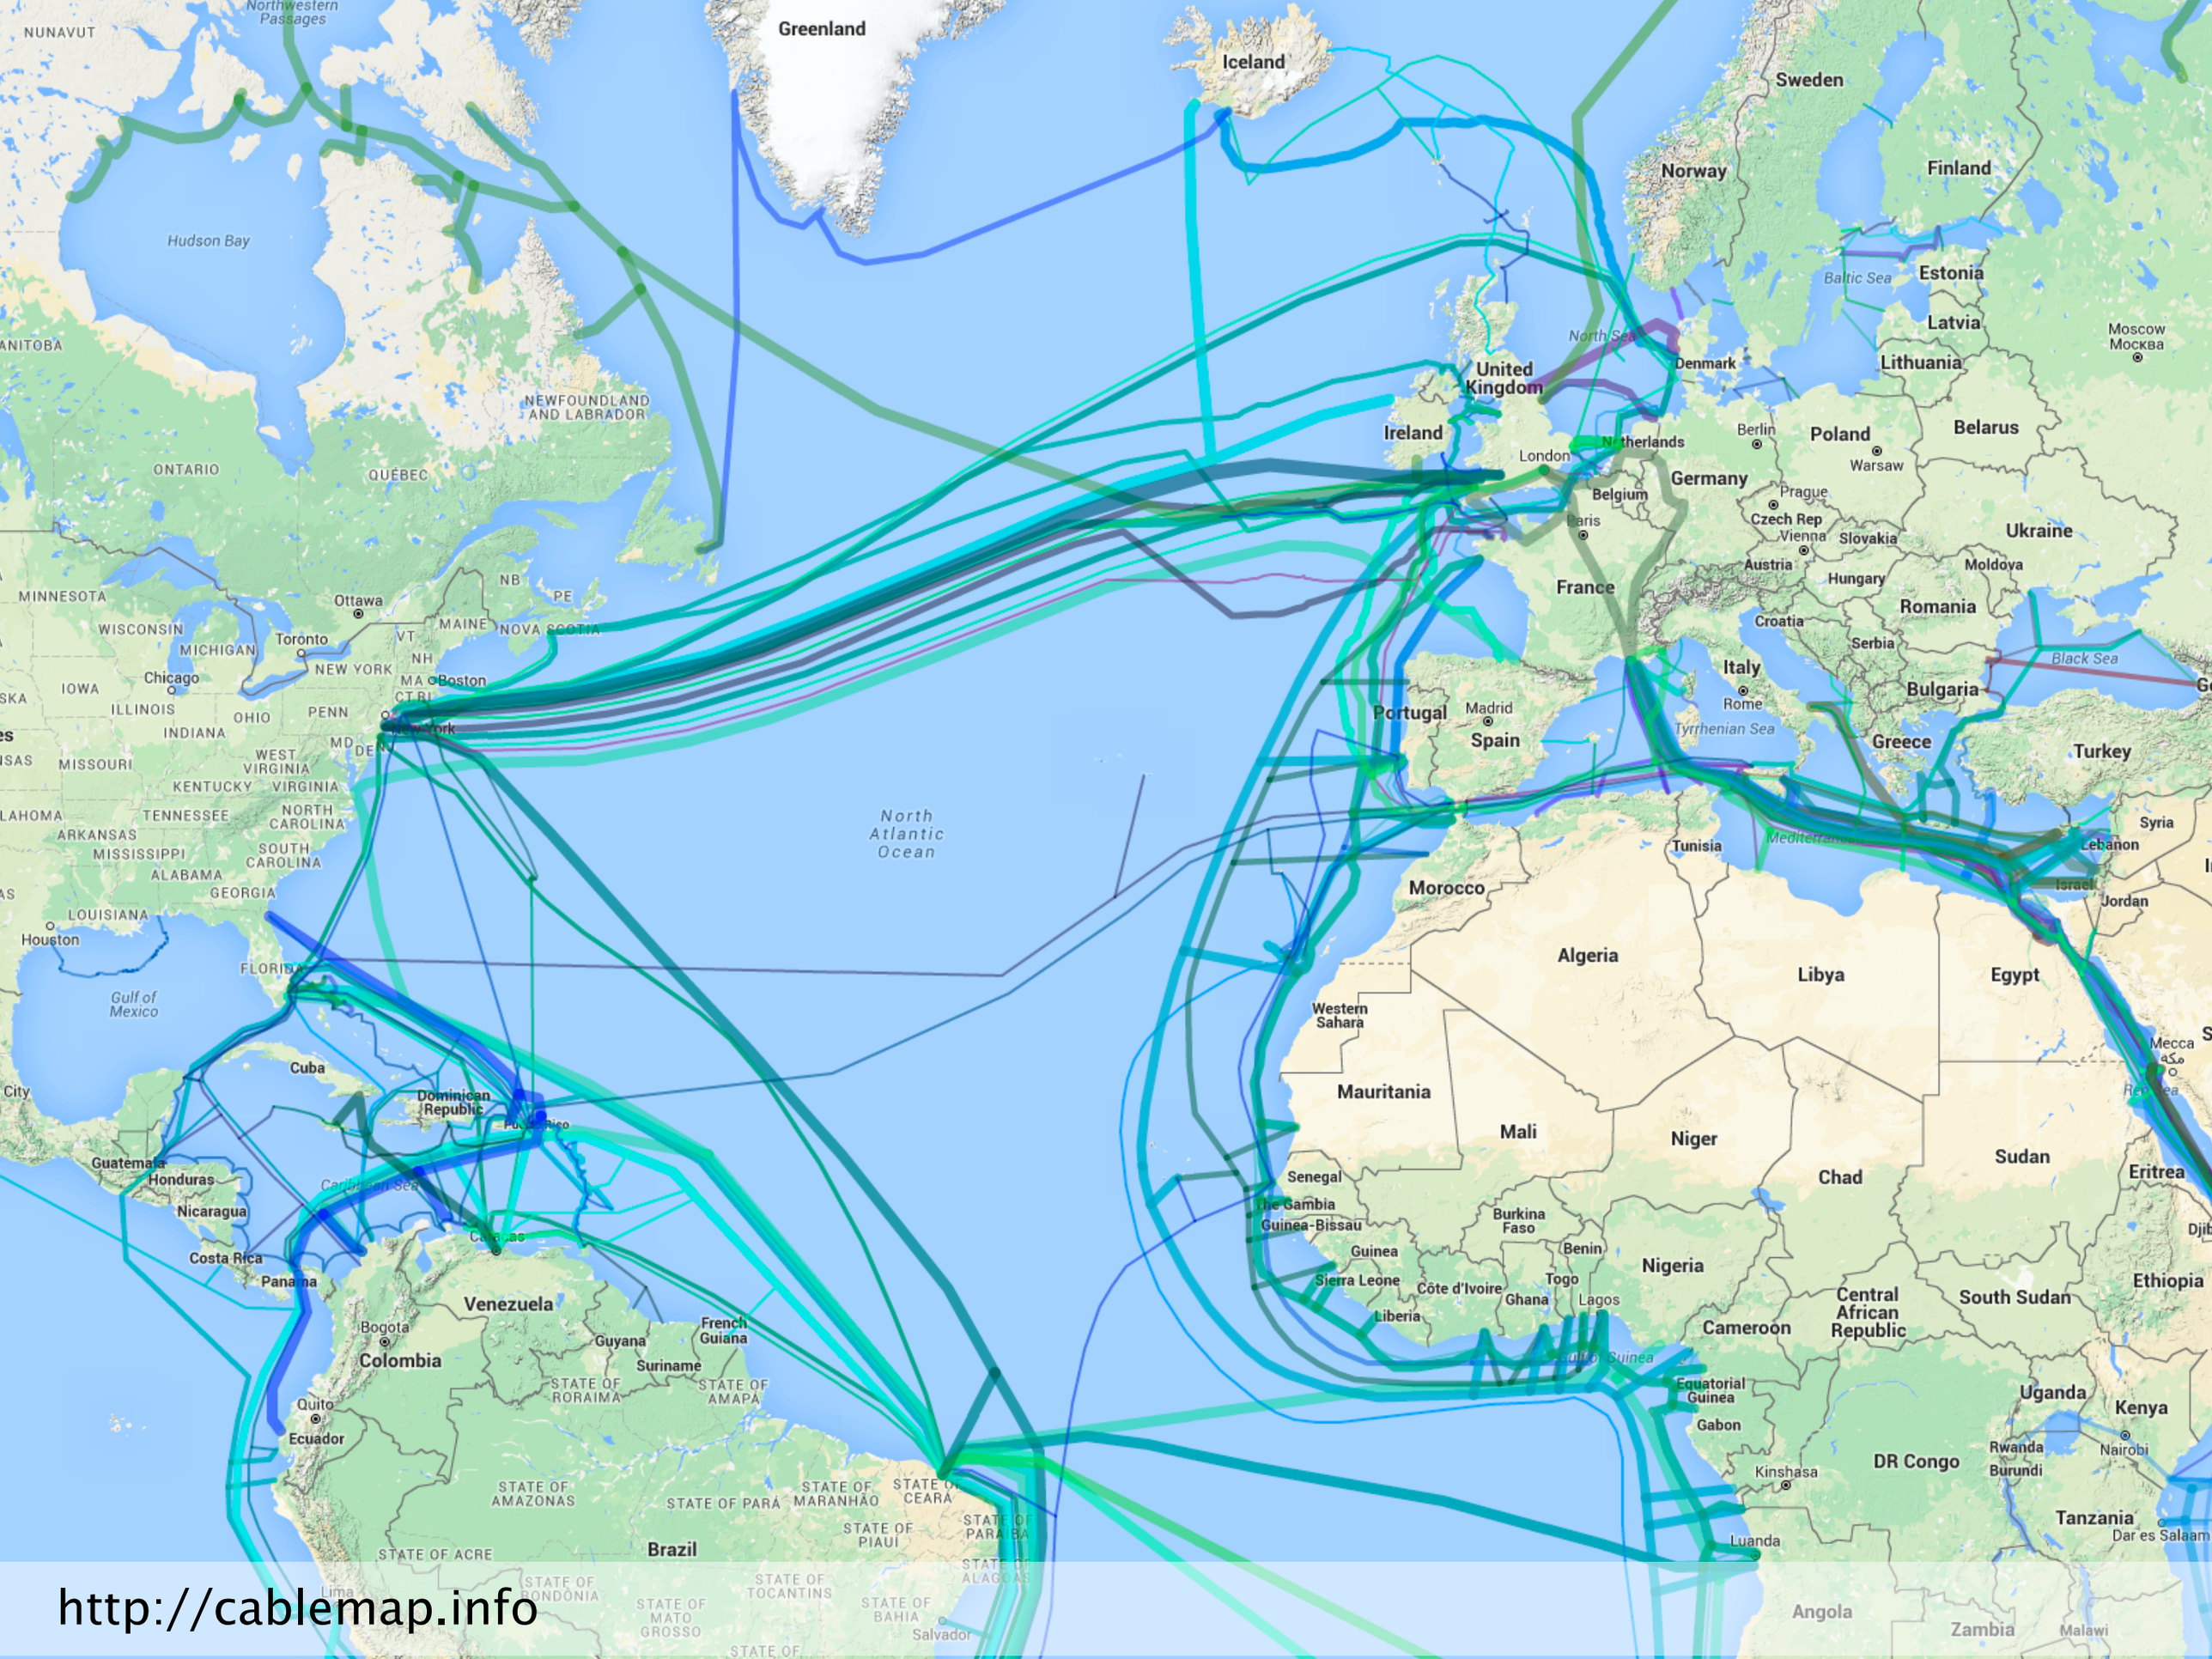

http://cablemap.info

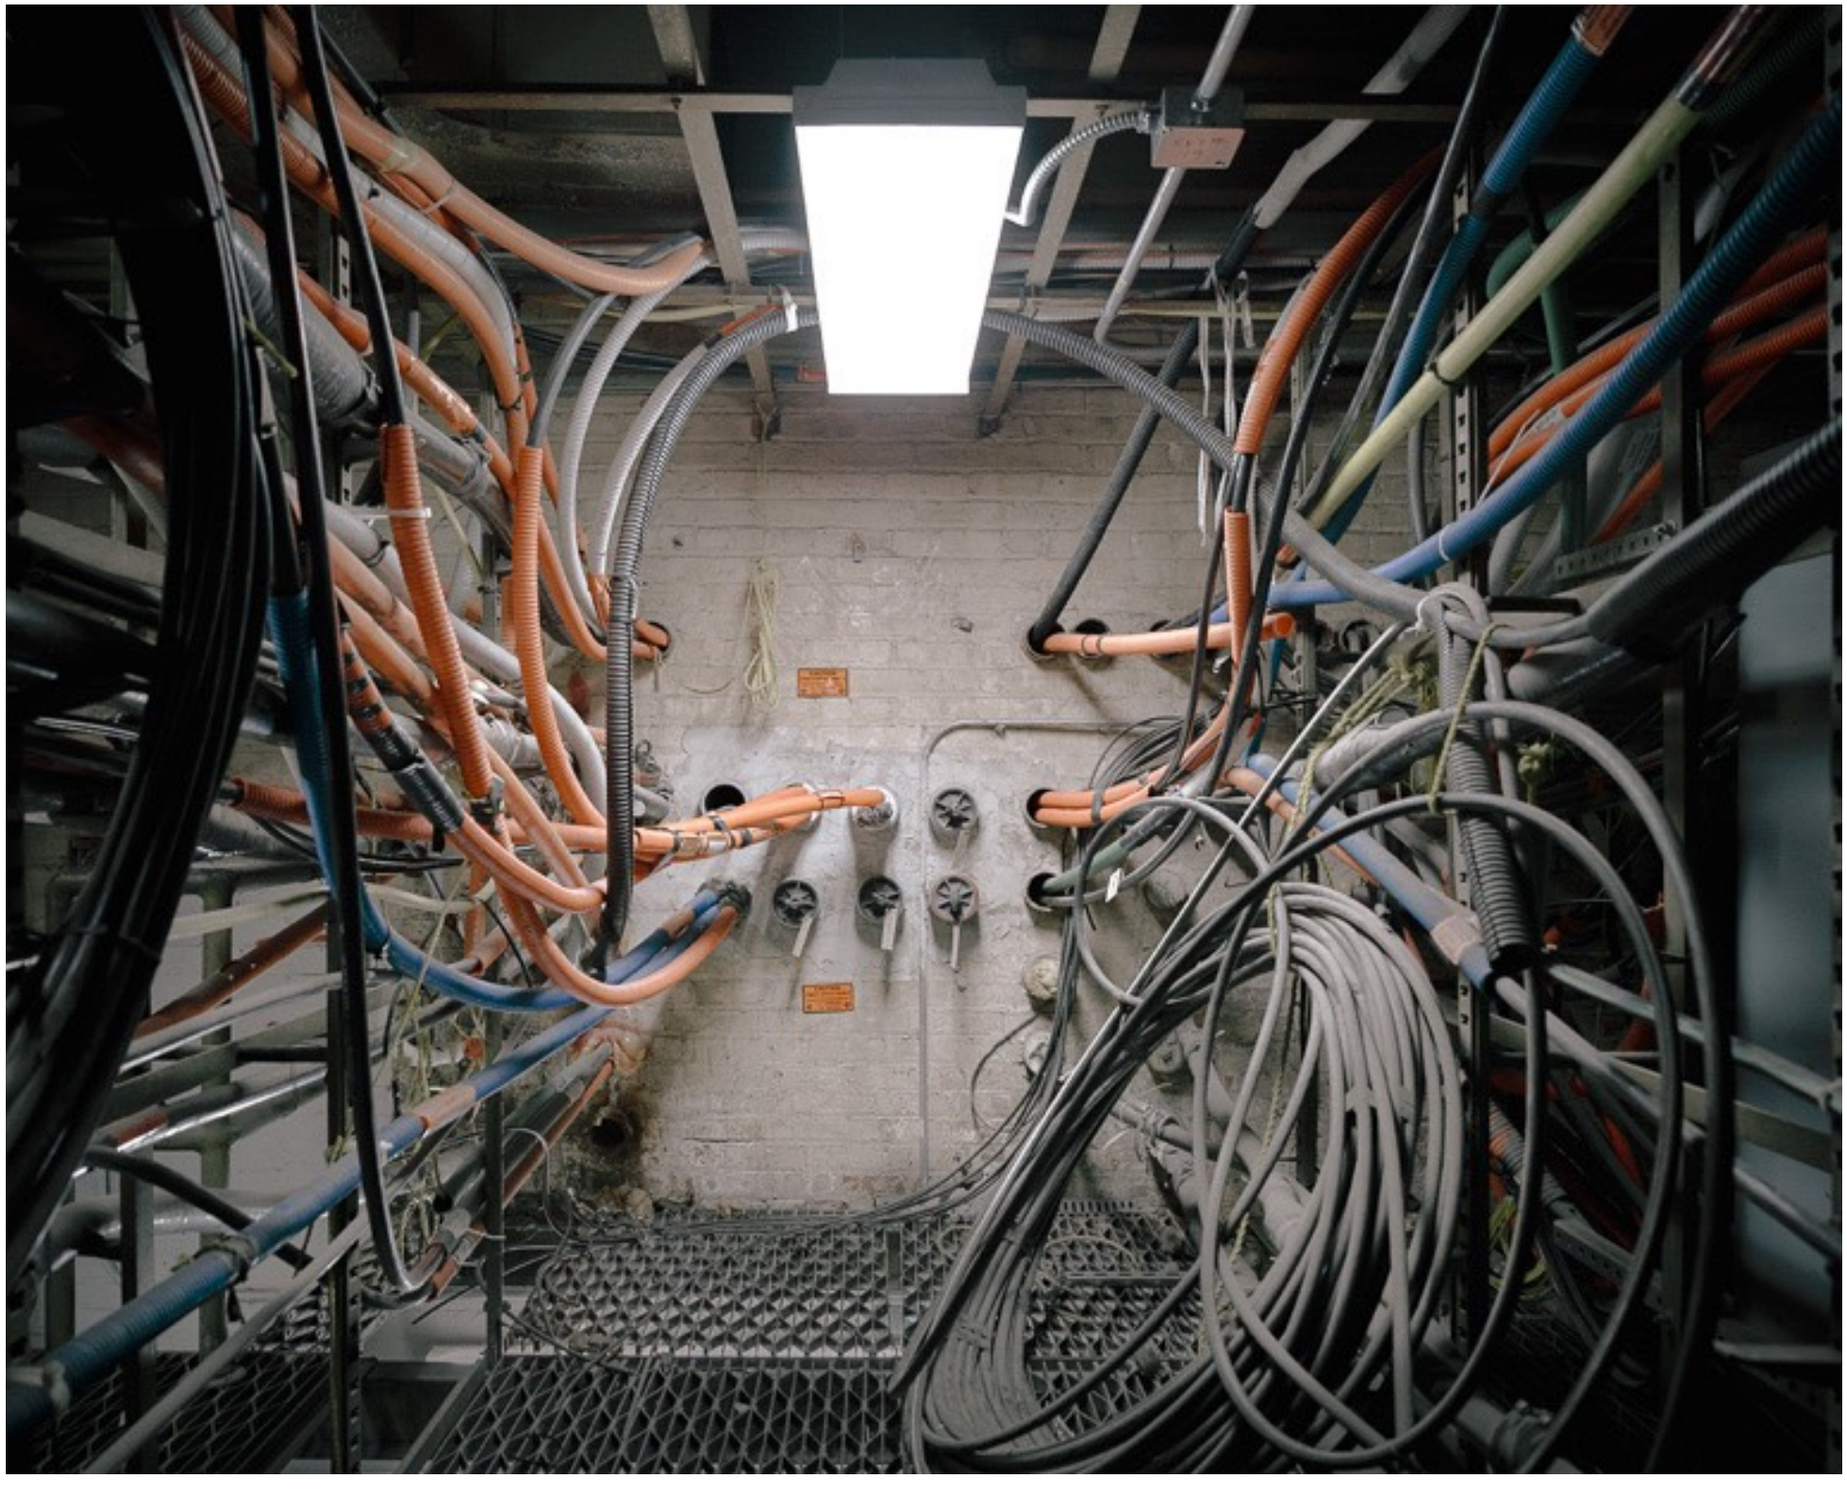

Somewhere in Manhattan… http://www.petergarritano.com/the-internet.html

The *Inter*net is a network of networks

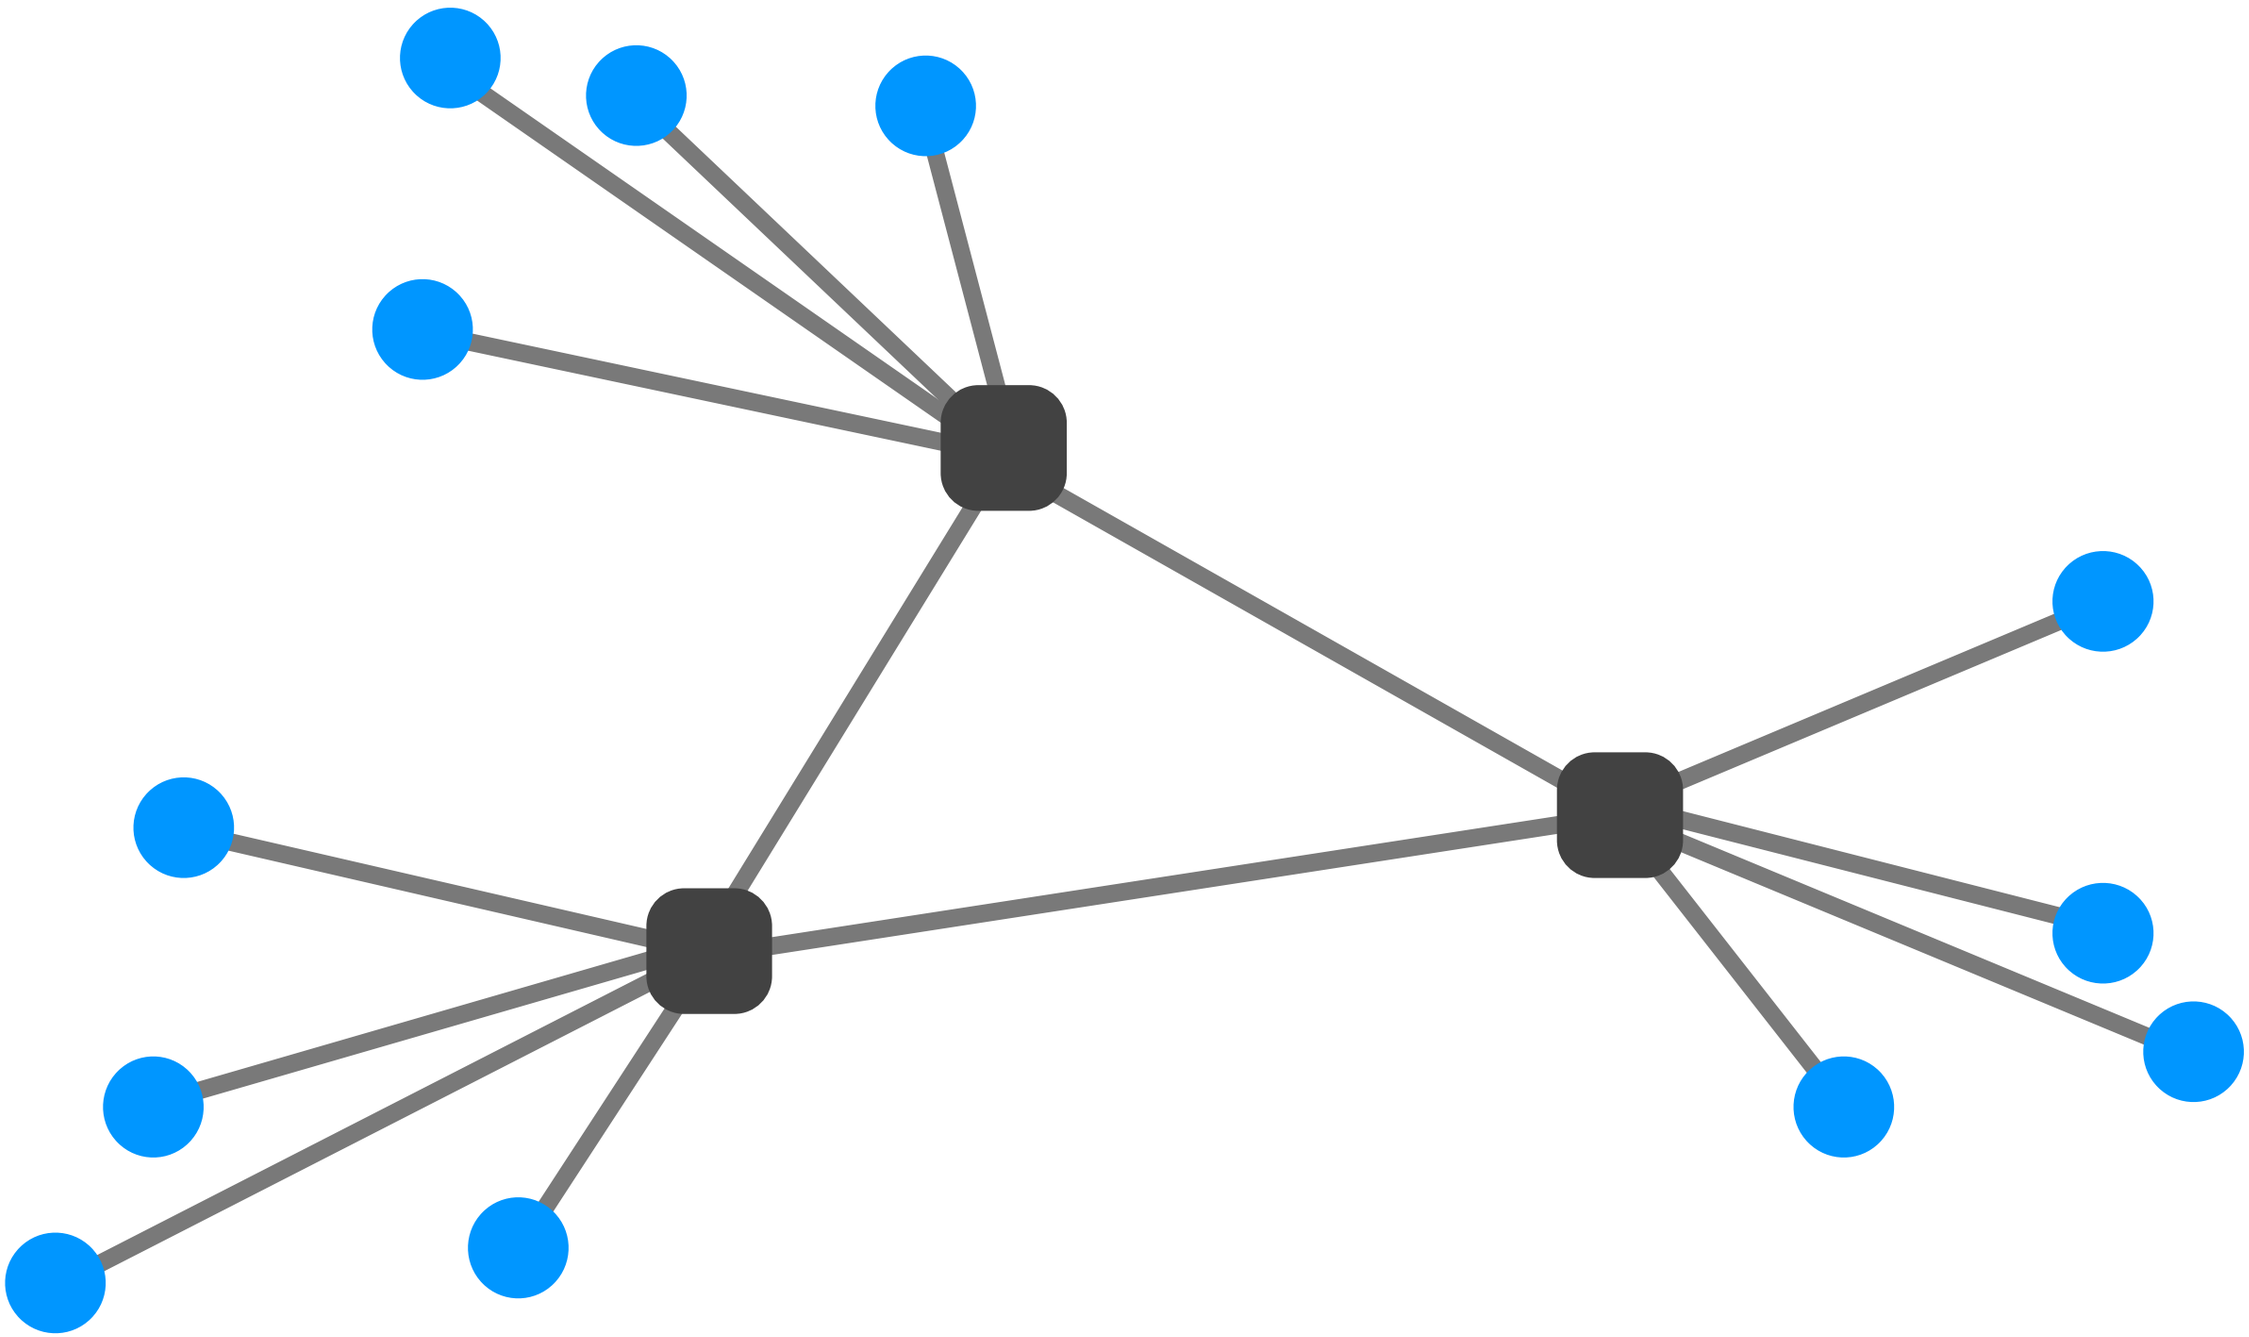

Internet Service Providers

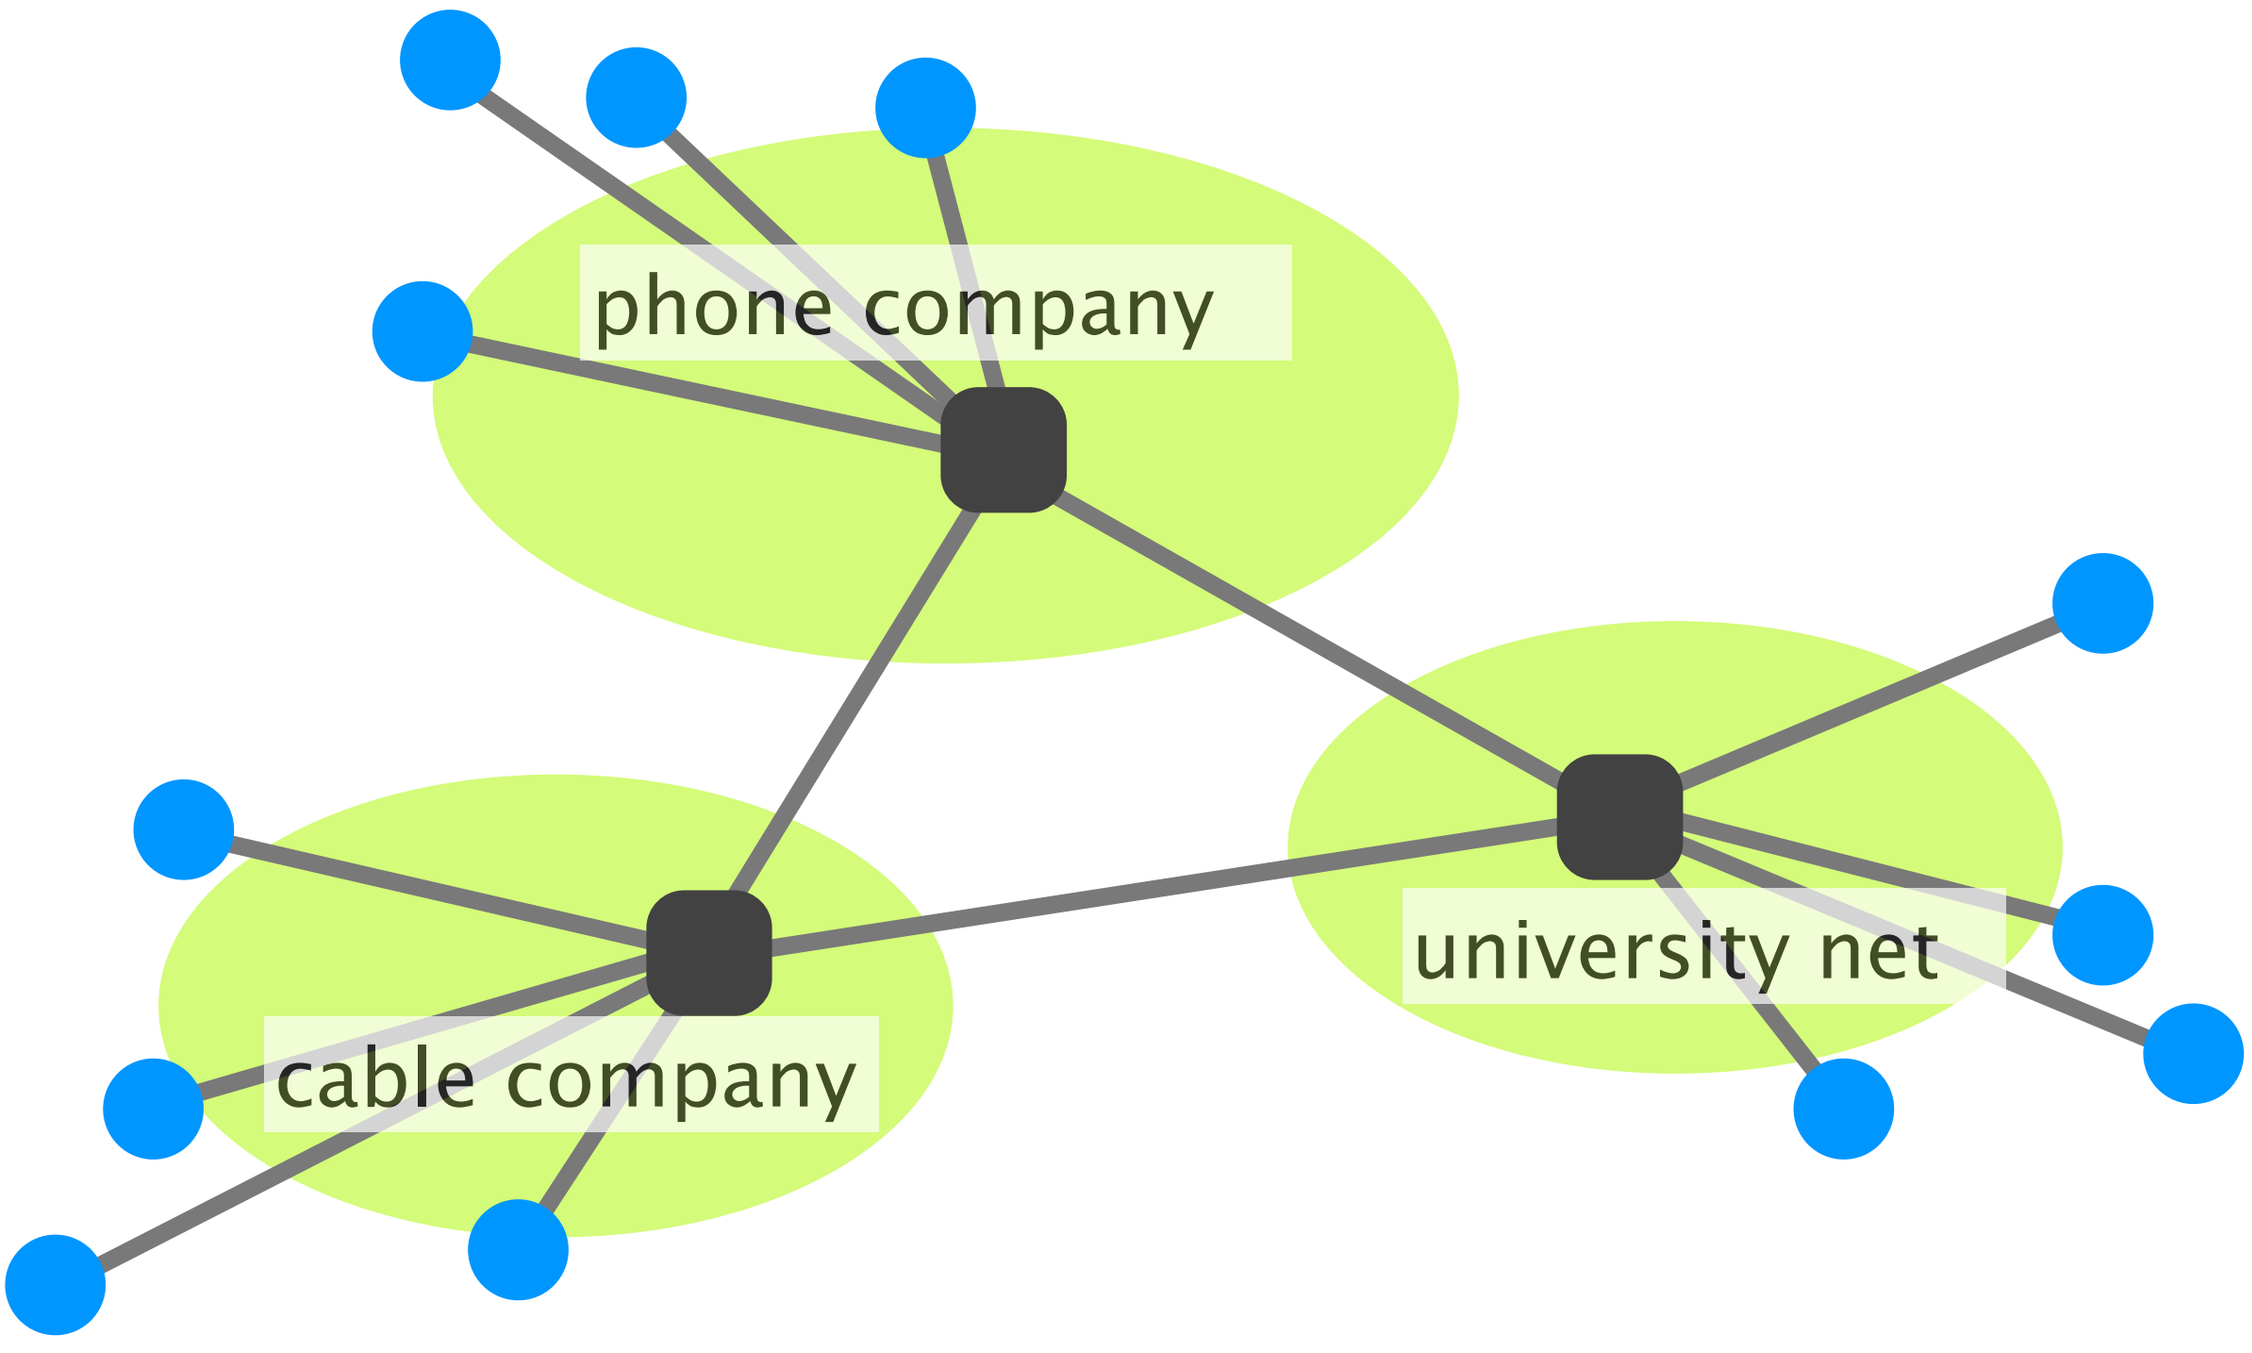

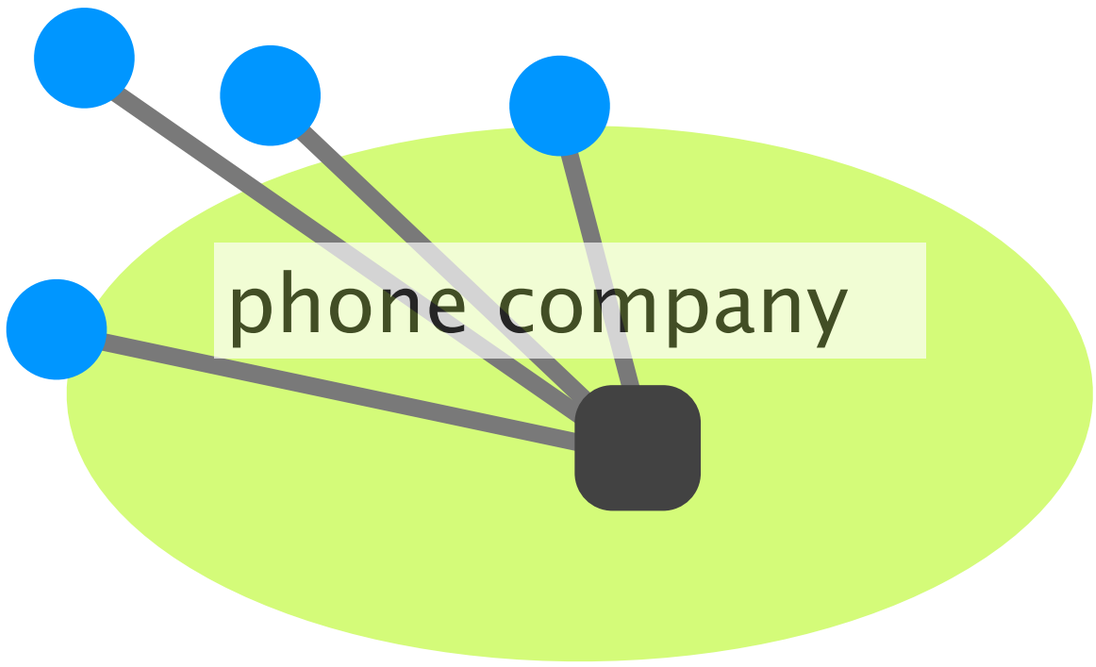
phone company

Conceptually, the last mile of the Internet looks like this

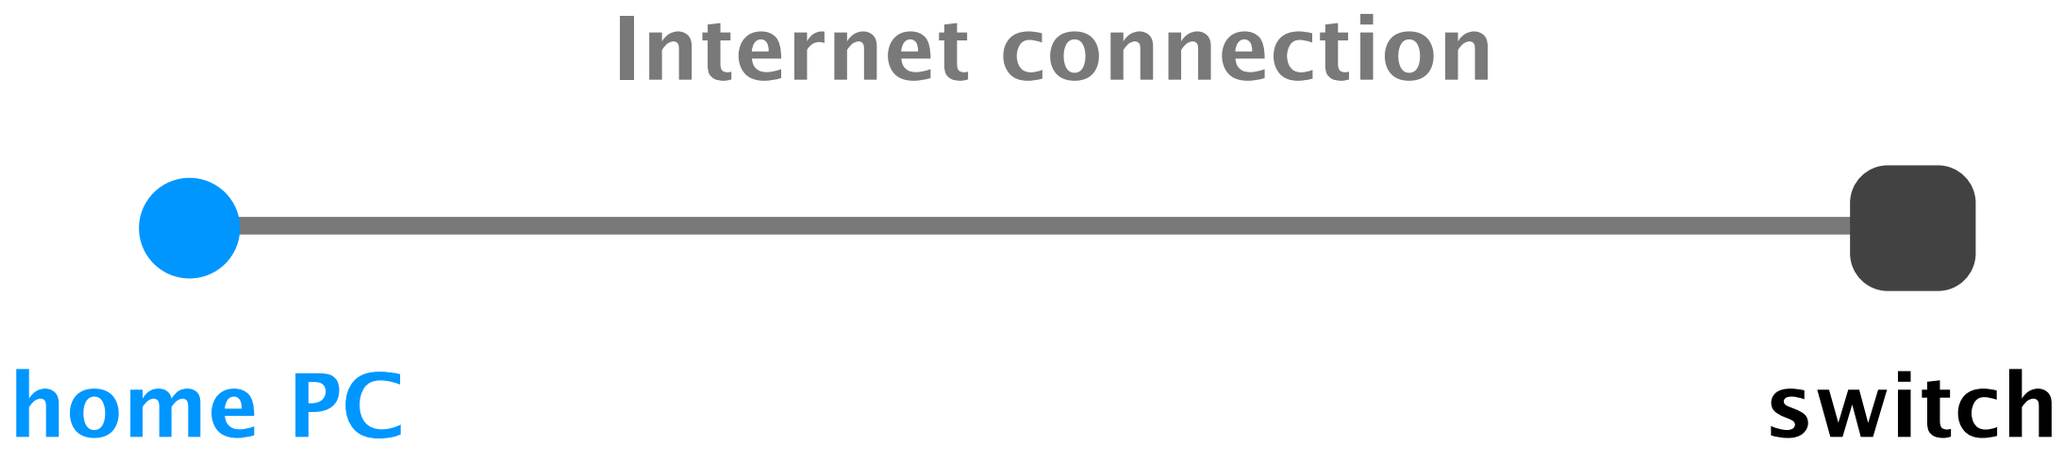

In practice, it looks more like this…

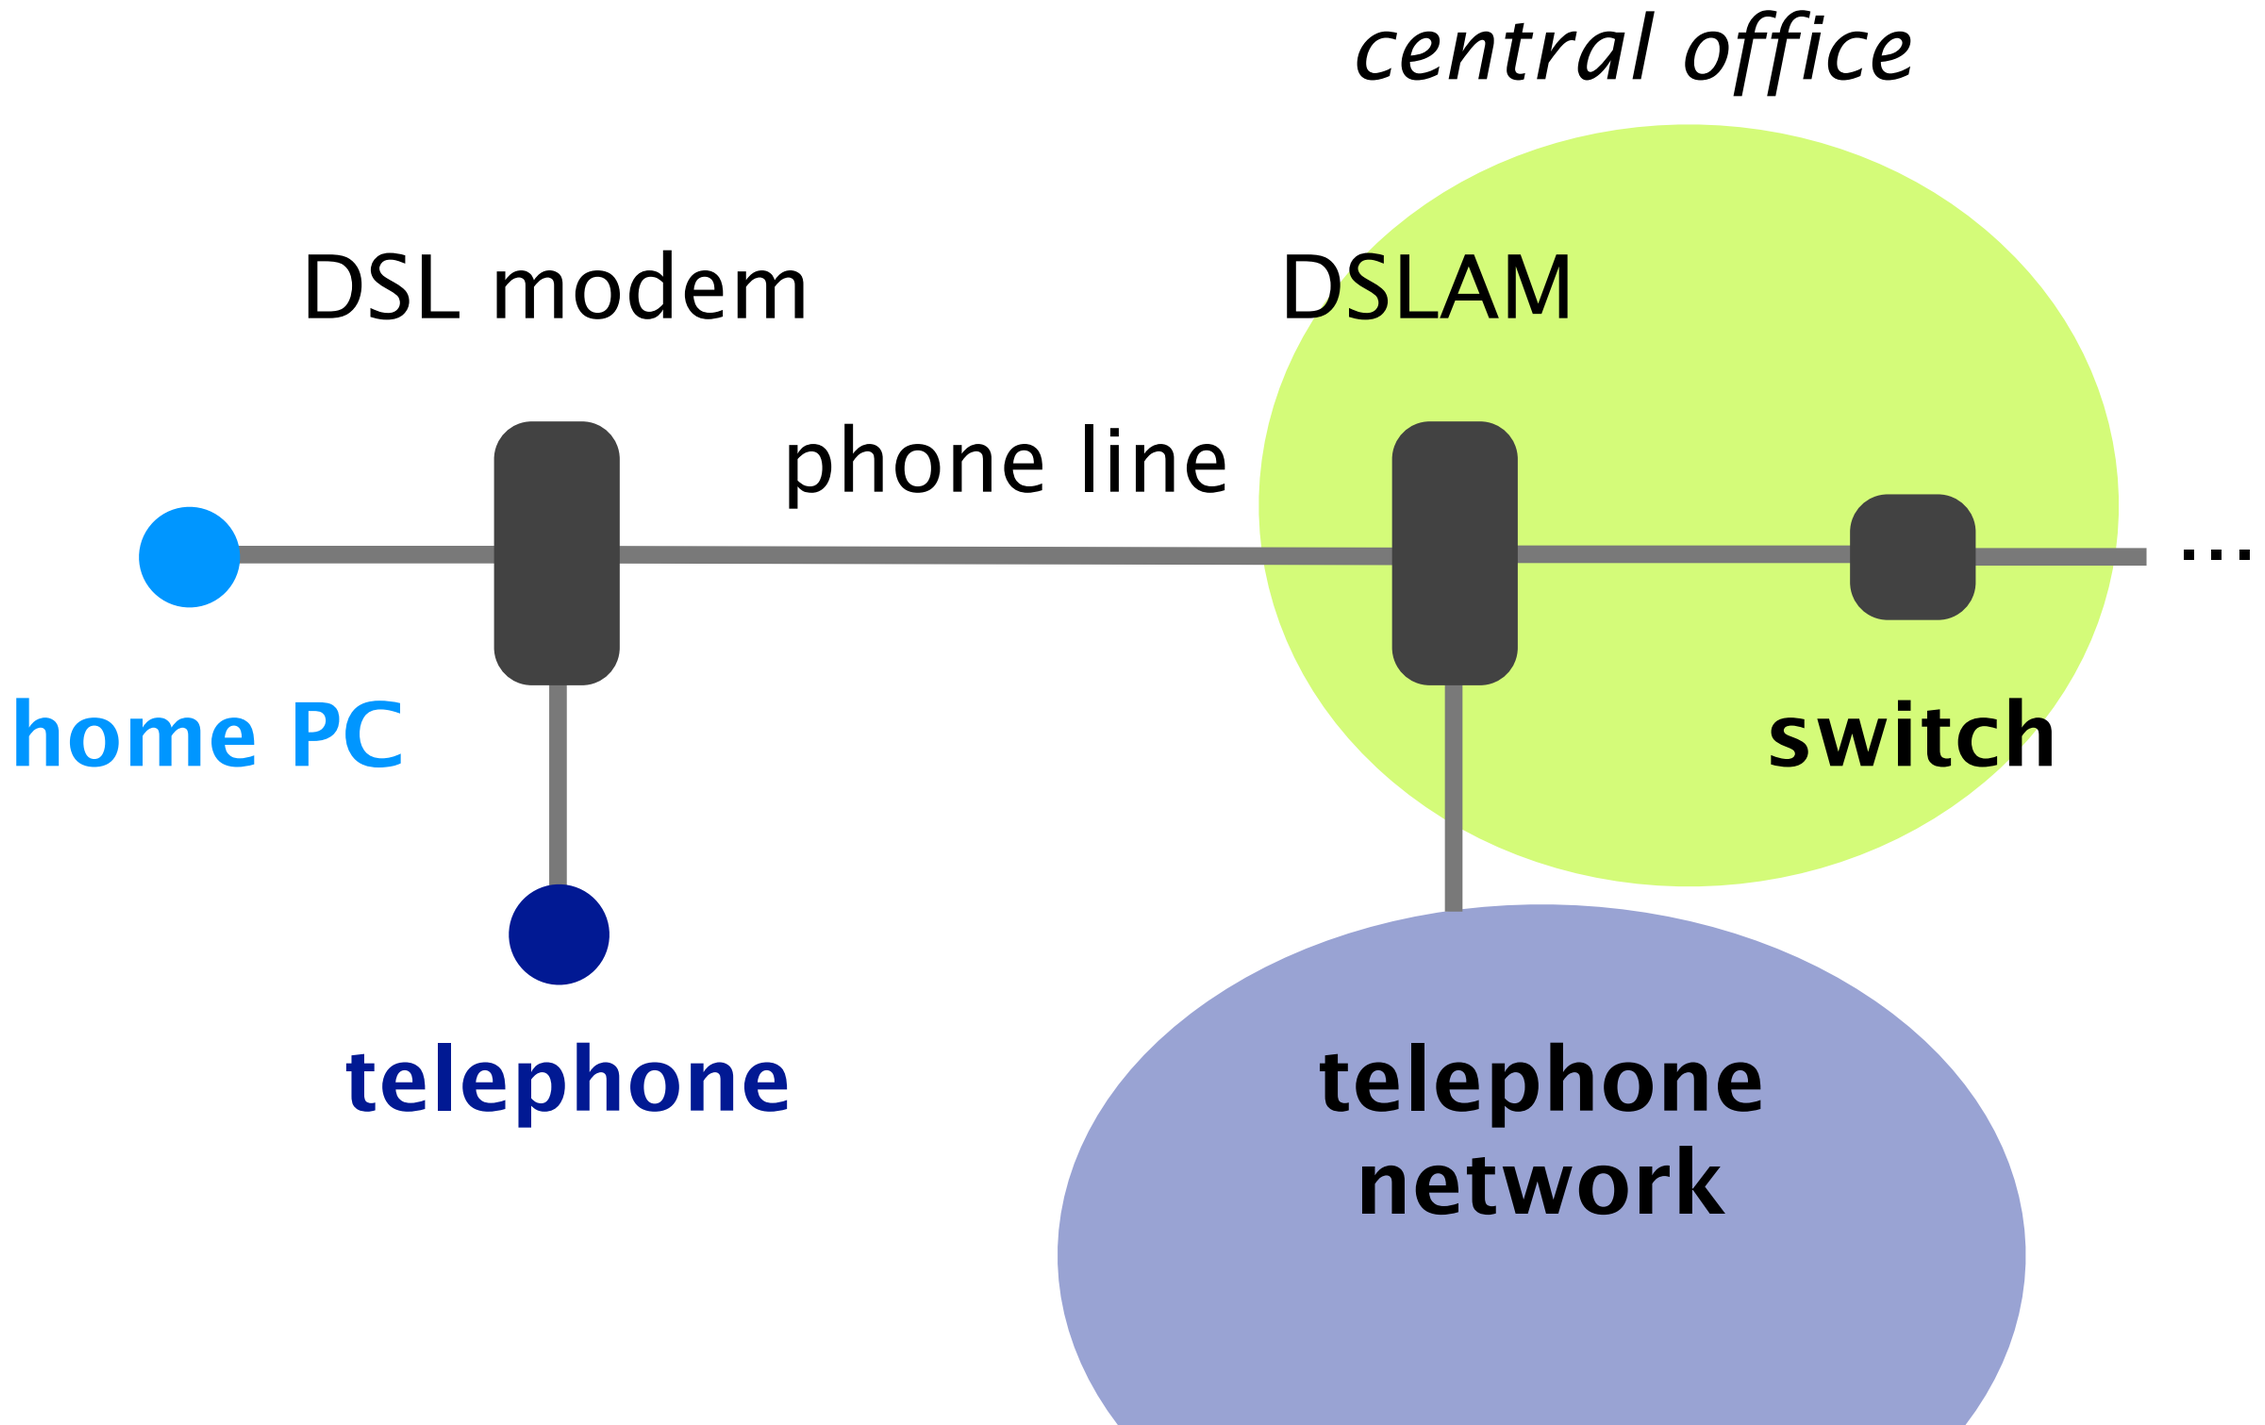

Digital Subscriber Line (DSL) brings
high BW to households over phone lines

Digital Subscriber Line (DSL) brings high BW to households over **phone lines**

↓

Why?

Digital Subscriber Line (DSL) brings high BW to households over phone lines

DSL is composed of 3 channels:

- downstream data channel — tens to few hundred Mbps
- upstream data channel — few Mbps to few tens Mpbs
- 2-ways phone channel

DSL is composed of 3 channels:

- downstream data channel
- upstream data channel
- 2-ways phone channel

tens to few hundred Mbps

few Mbps to few tens Mpbs

Why is there such an asymmetry?

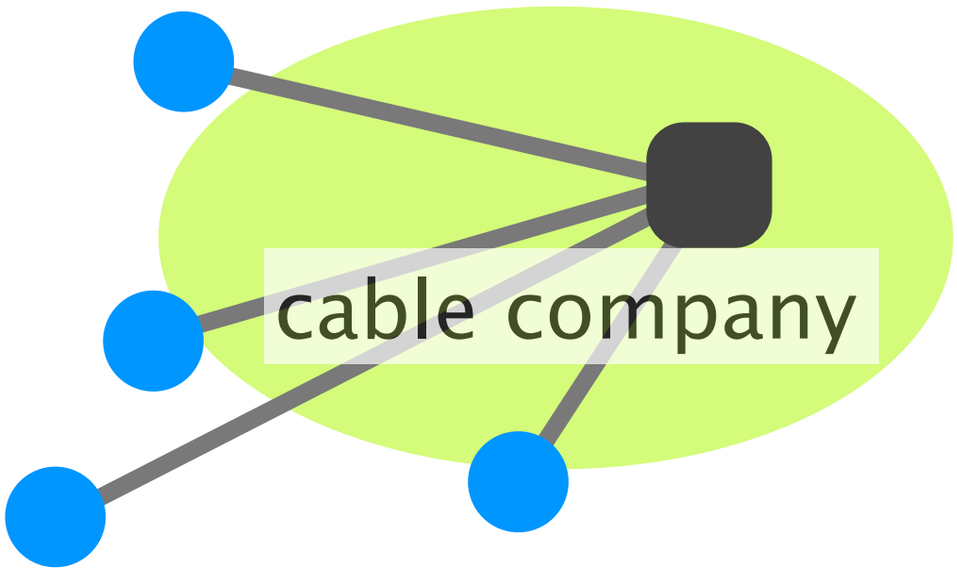
cable company

Conceptually, the last mile of the Internet looks like this

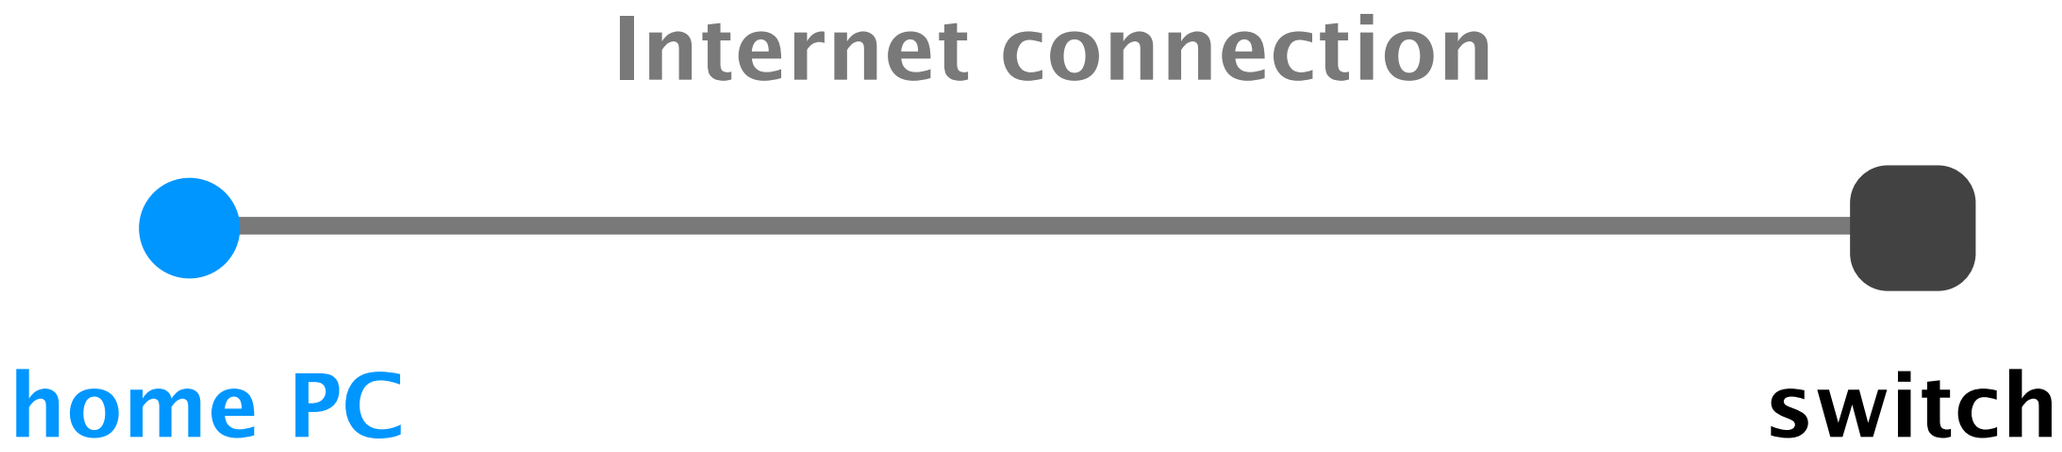

In practice, it looks more like this…

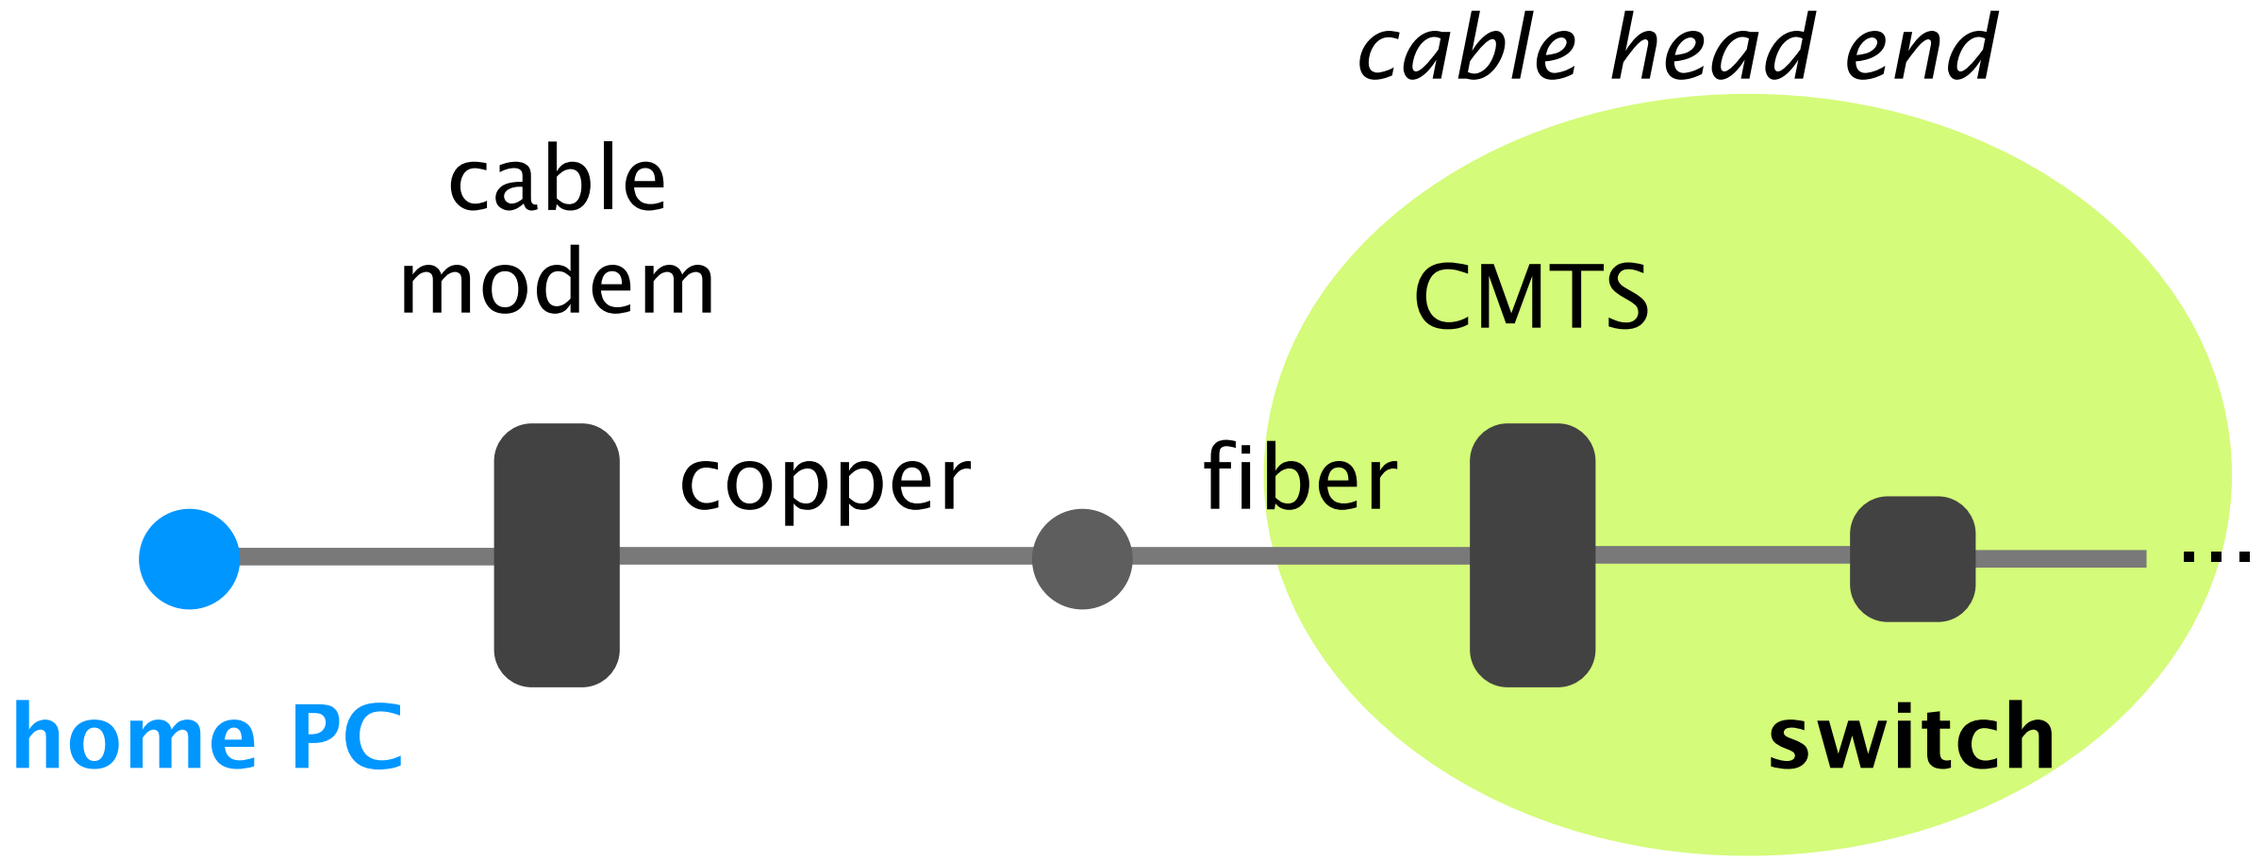

Many households share the same access

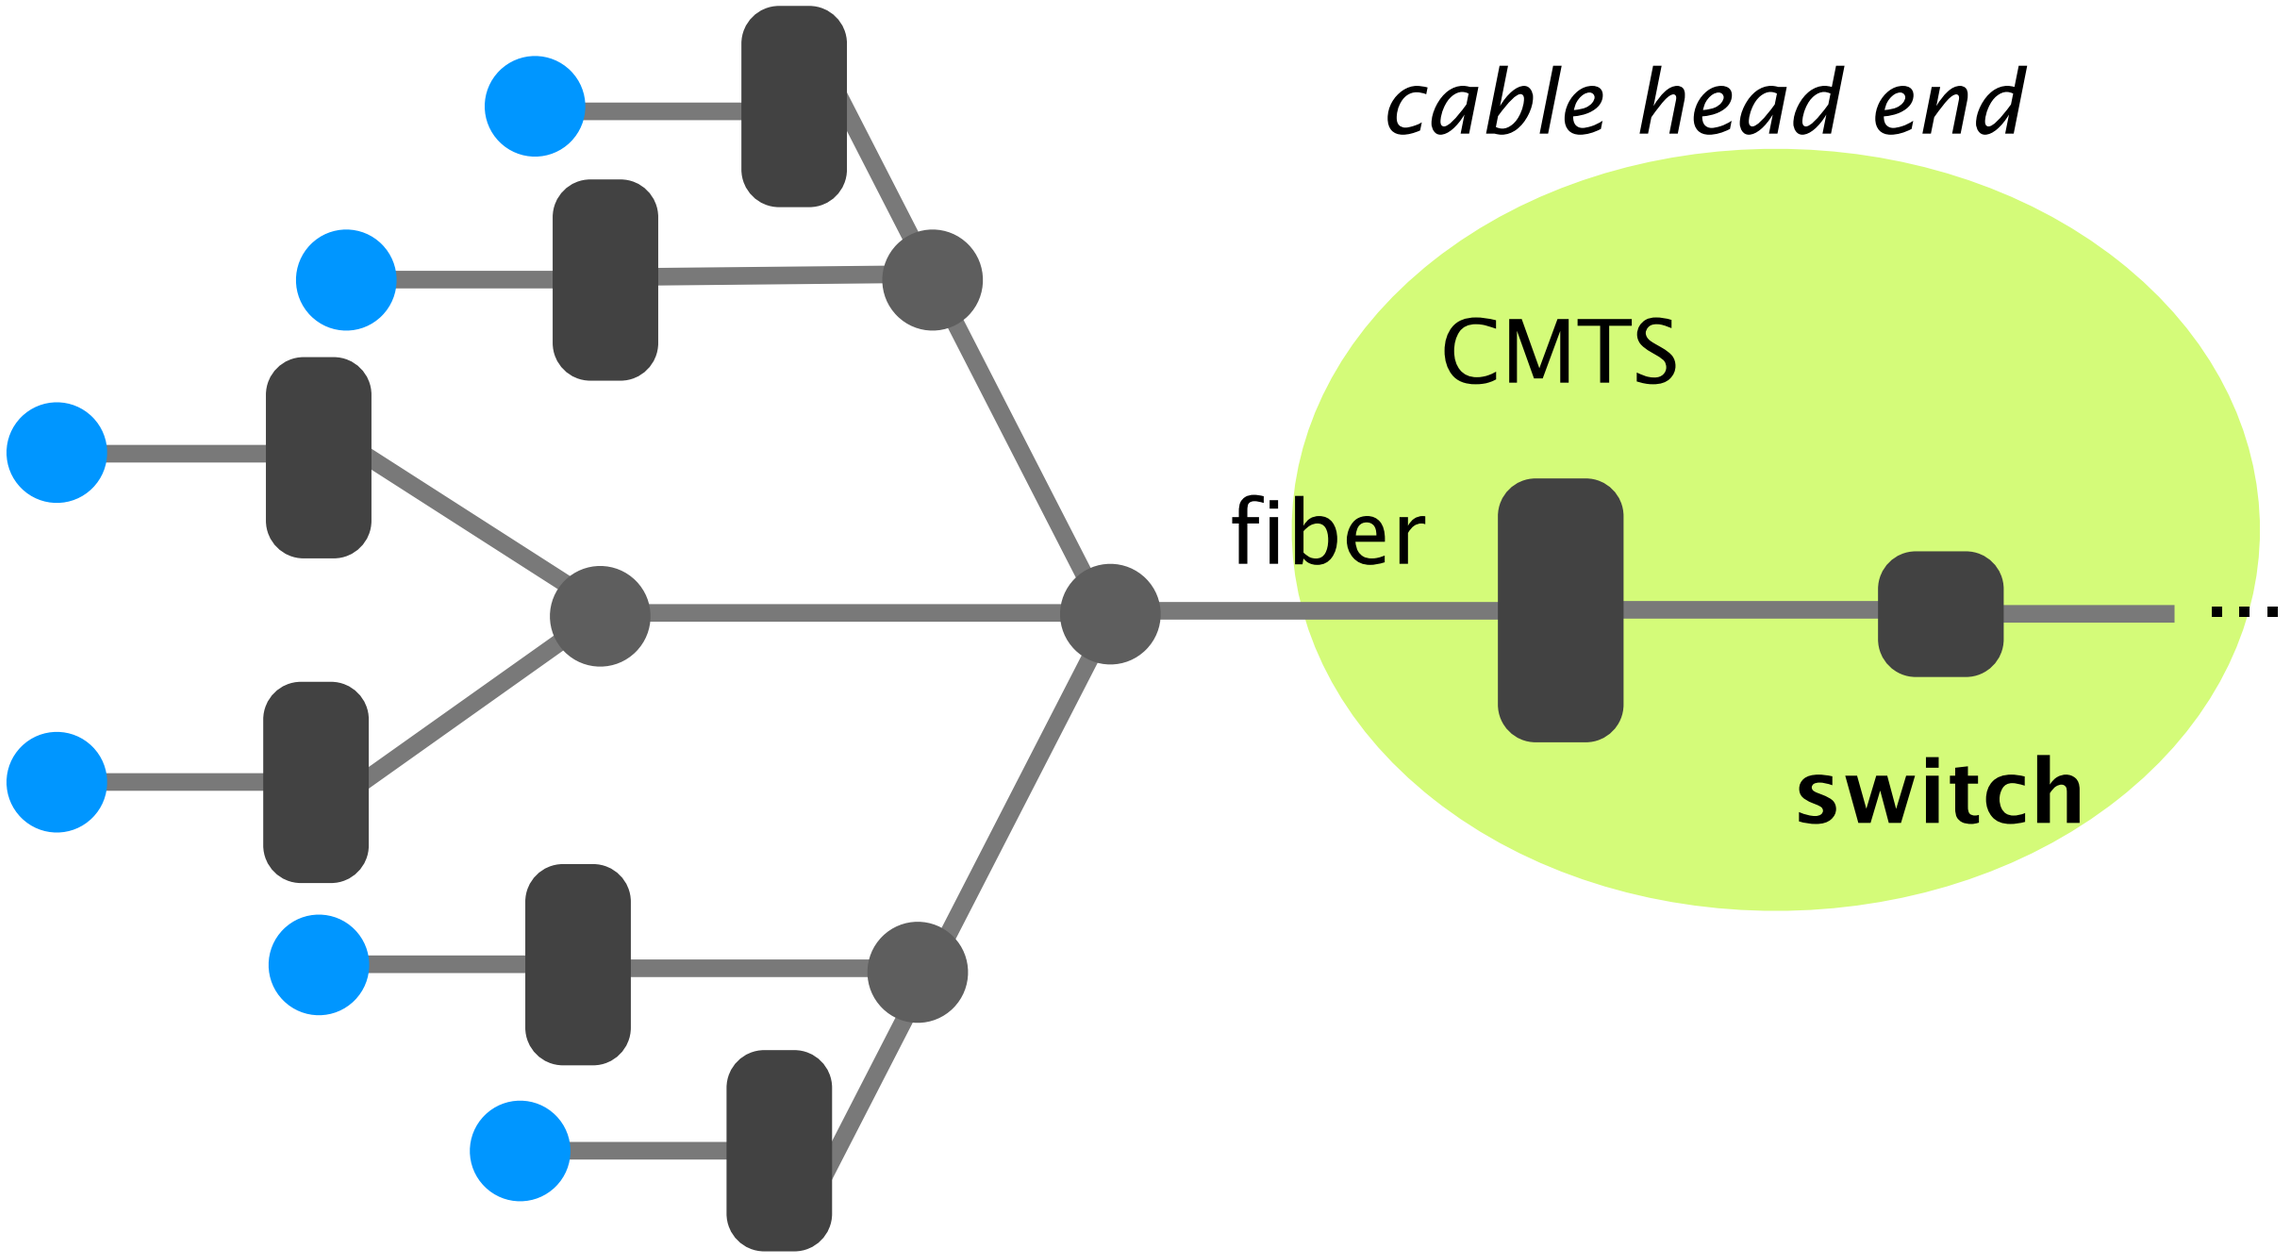

Cable Access Technologies (CATV) brings high BW to the households via cable TV

coaxial copper & fiber

- downstream data channel — tends to hundreds of Mbps
- upstream data channel — tens of Mbps

Unlike ADSL, the medium is shared between households

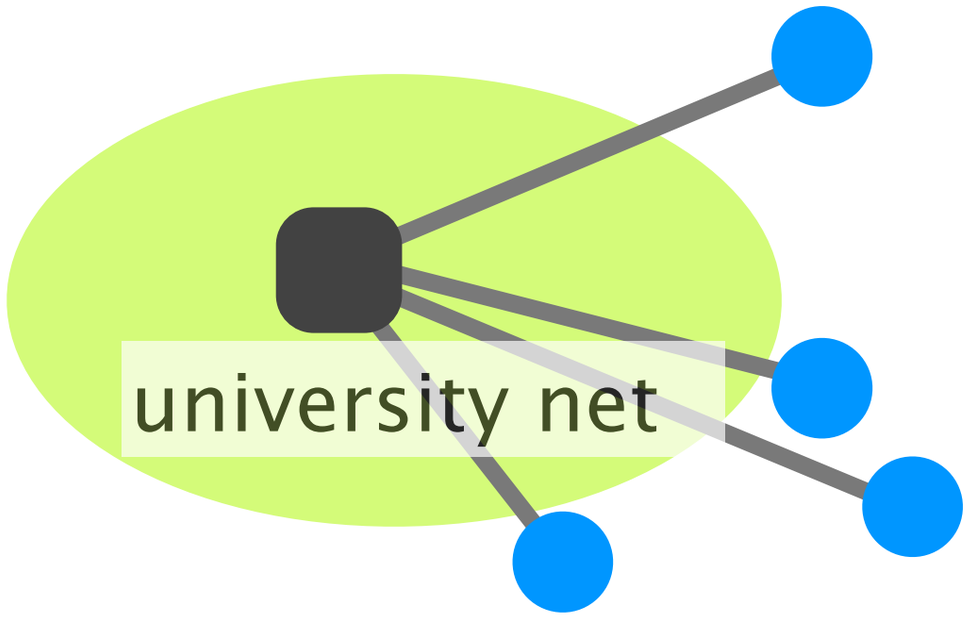
university net

Internet connection

workstation

switch

With respect to DSL and cable providers, enterprise access networks is *much* simpler

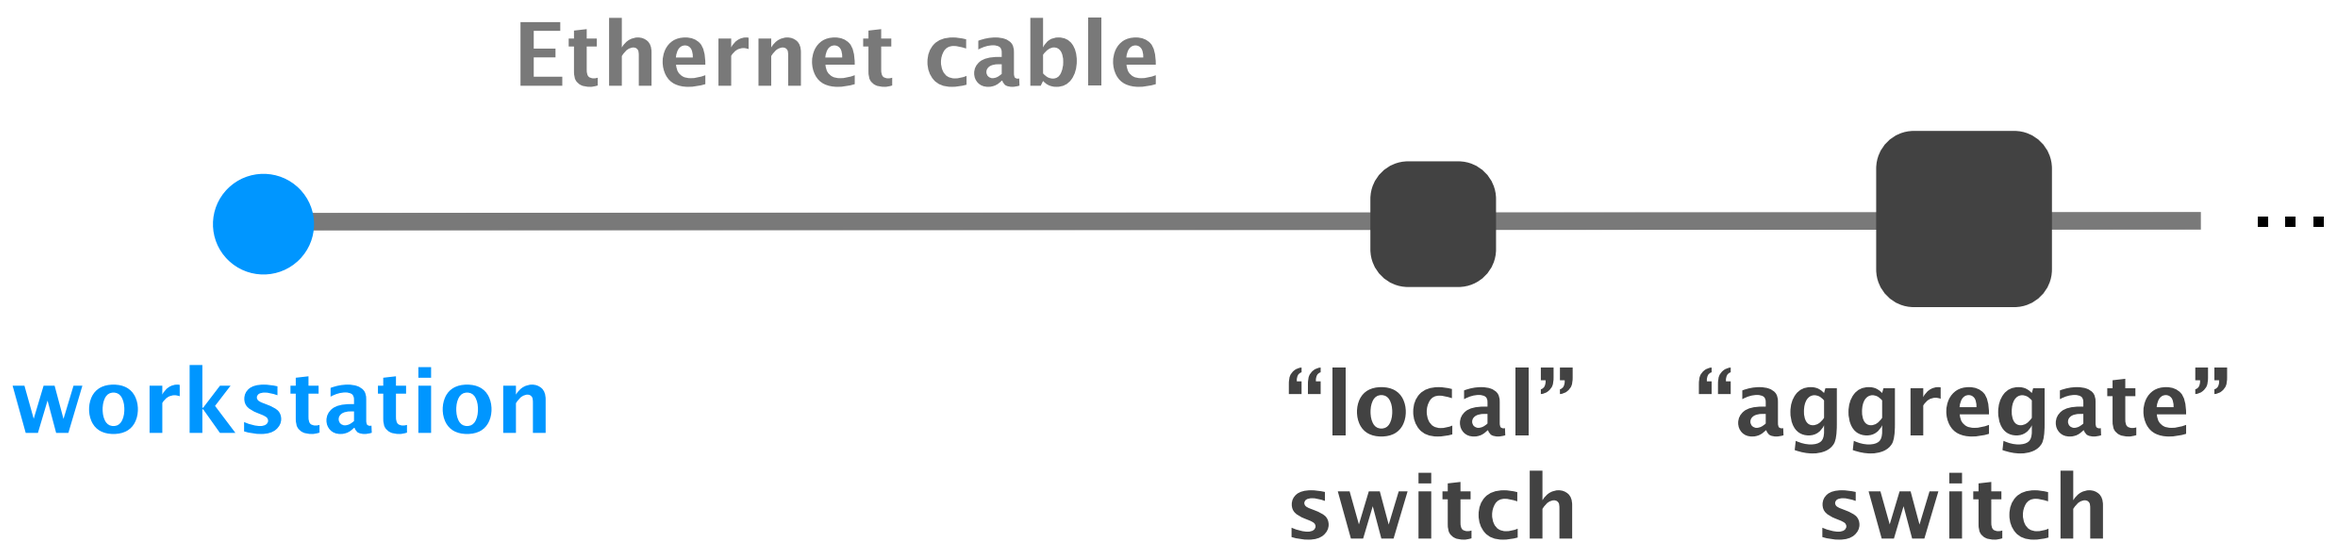

Ethernet is the most widely used Local Area Network technology

Twisted pair copper

1Gbps, 10 Gbps, 40 Gbps, 100 Gbps, …

symmetric

ADSL, CATV and Ethernet are only few examples of access technologies…

Cellular	smart phones
Satellite	remote areas
FTTH	household
Fibers	Internet backbone
Infiniband	High performance computing

…

Communication Networks

Part 1: Overview

What is a network made of?

#2 How is it shared?

How is it organized?

How does communication happen?

How do we characterize it?

Up to now, we've seen

what the last mile of the Internet looks like

What about the rest of the network?

3 requirements for a network topology

Should tolerate failures

several paths between each source and destination

Possess enough sharing to be feasible & cost-effective

number of links should not be too high

Provide adequate per-node capacity

number of links should not be too small

Compare these three designs in terms of sharing, resiliency, and per-node capacity

design	full-mesh	chain	bus
	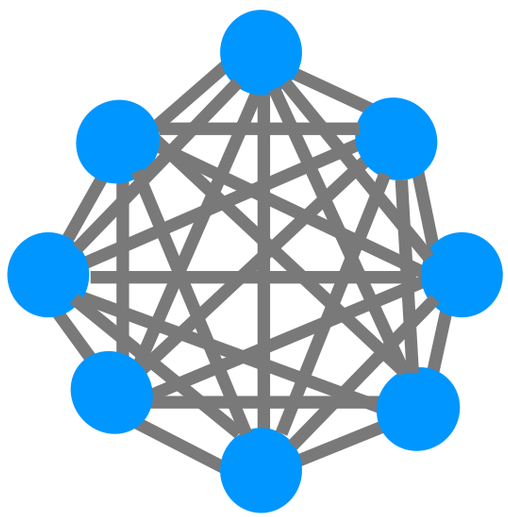	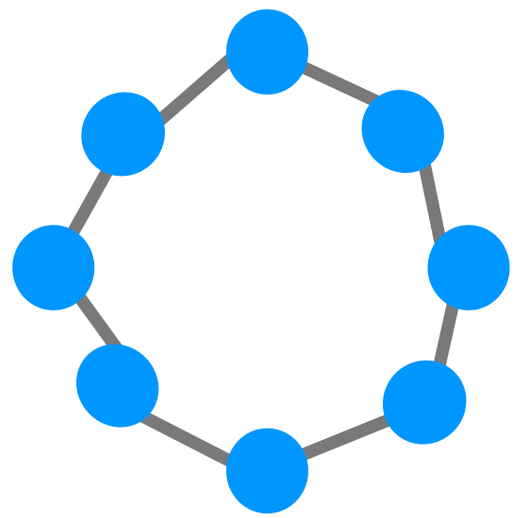	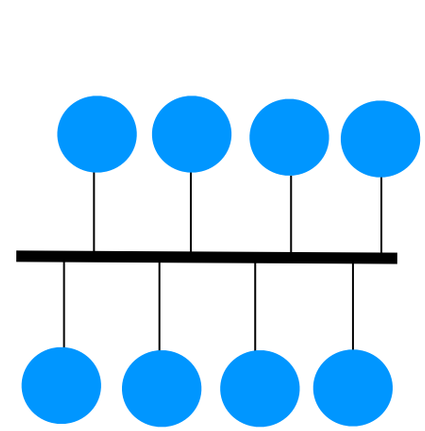
advantages			
disadvantages			

Switched networks provide

reasonable and flexible compromise

design

switched

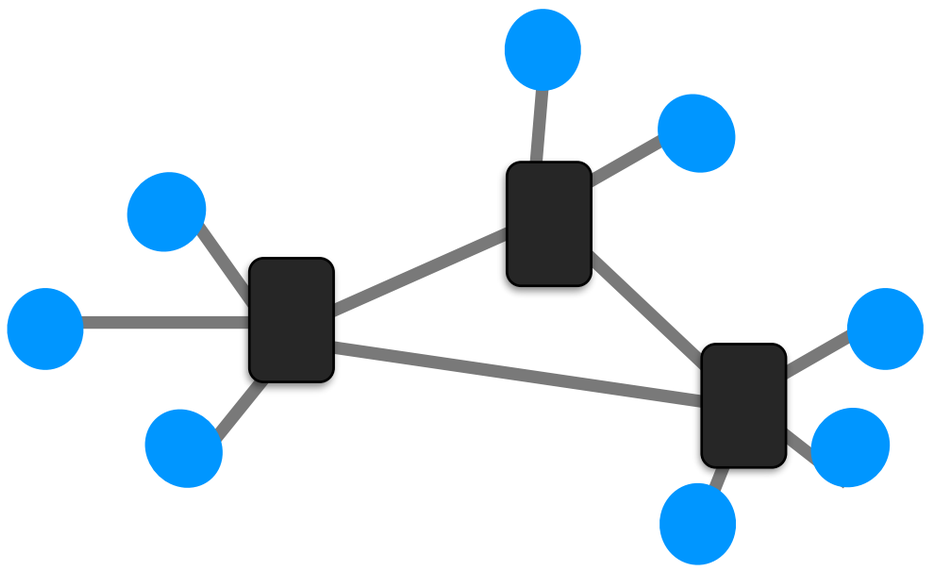

advantages

sharing and per-node capacity can be adapted to fit the network needs

disadvantages

require smart devices to perform:
forwarding, routing, resource allocation

Links and switches are shared between flows

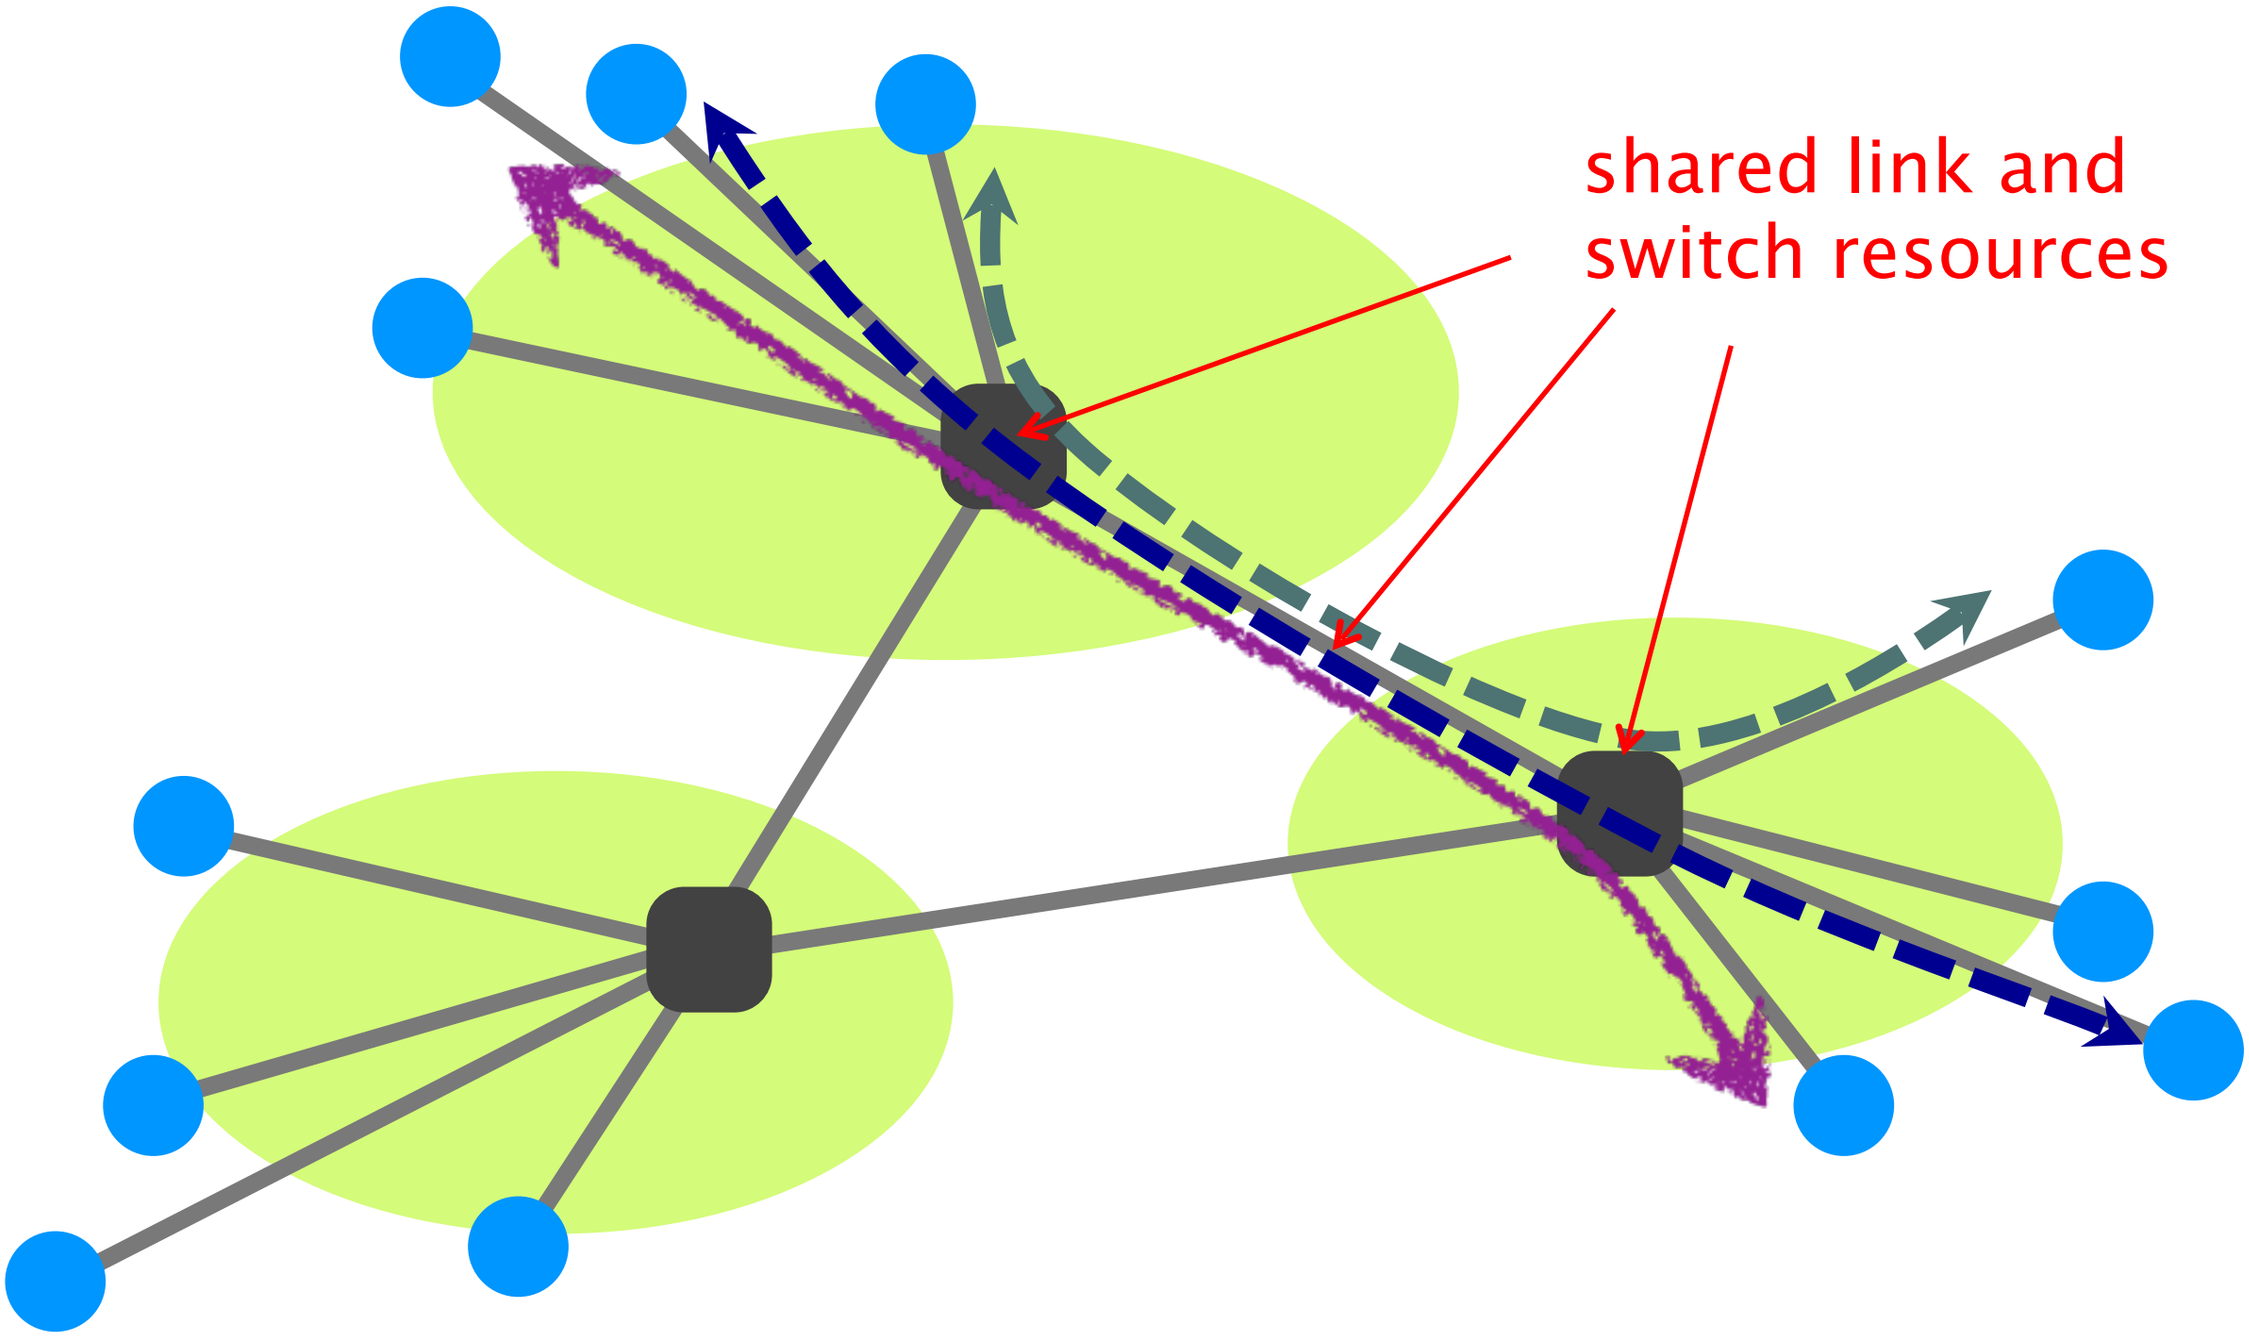

There exist two approaches to sharing:

reservation and on-demand

	Reservation	On-demand
principle	reserve the bandwidth you need in advance	send data when you need

Both are examples of statistical multiplexing

	Reservation	On-demand
multiplexing	at the flow-level	at the packet-level

Between reservation and on-demand:

Which one do you pick?

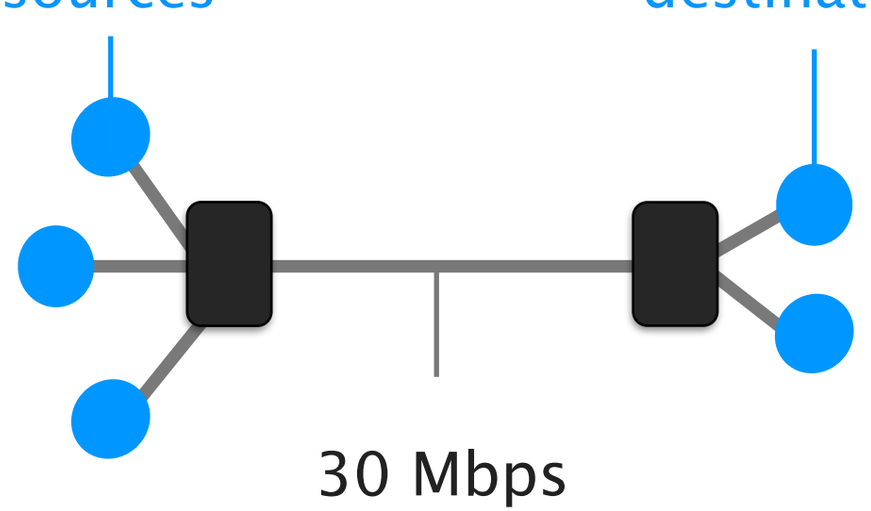

Consider that each source needs 10 Mbps

What do they get with:

- reservation
- on-demand

Assume the following peak demand and flow duration

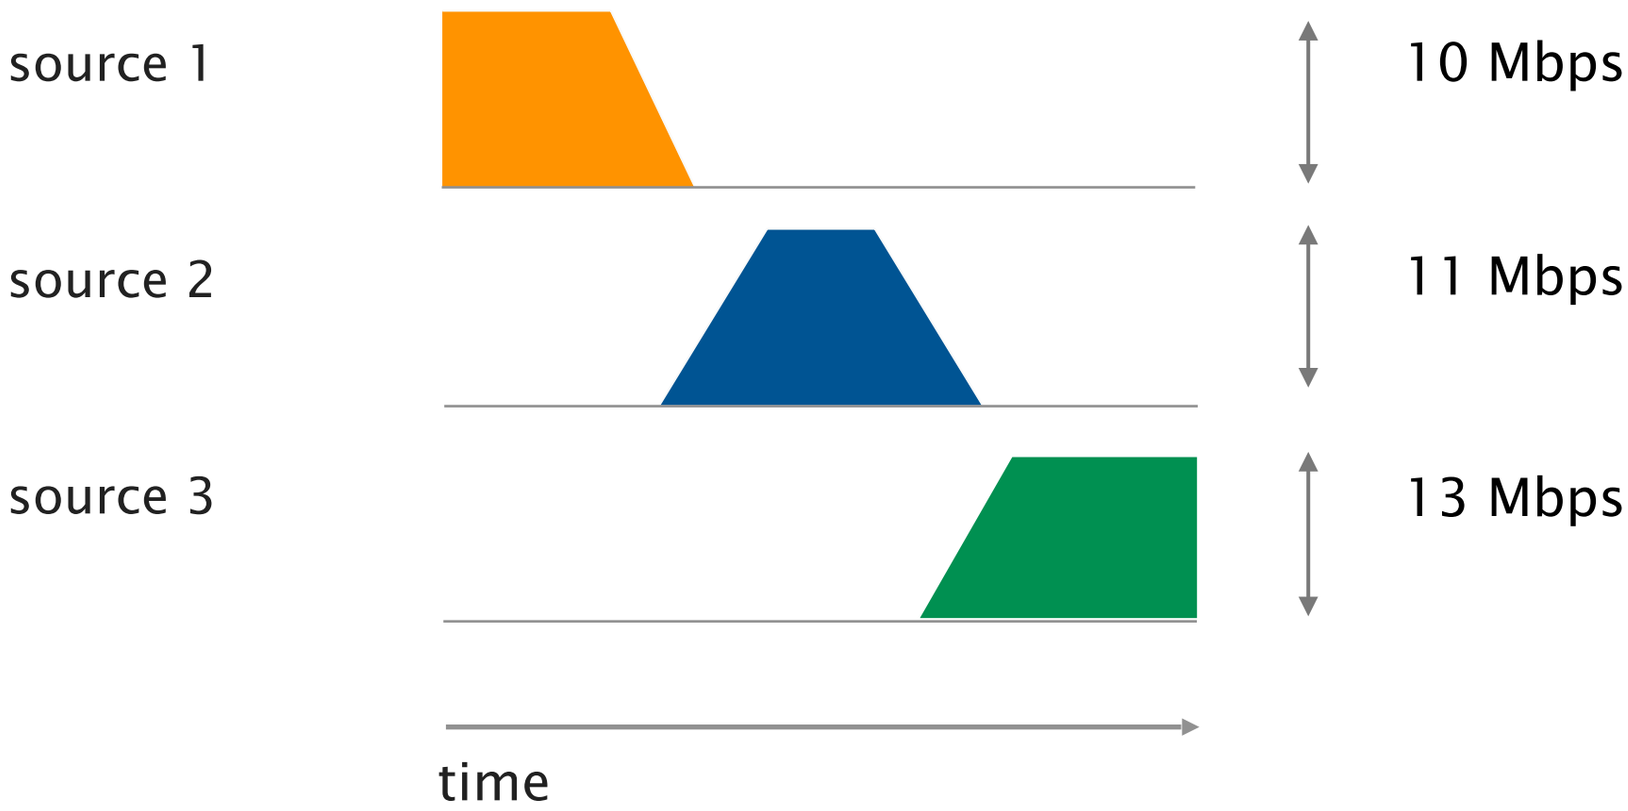

Assume the following peak demand and flow duration

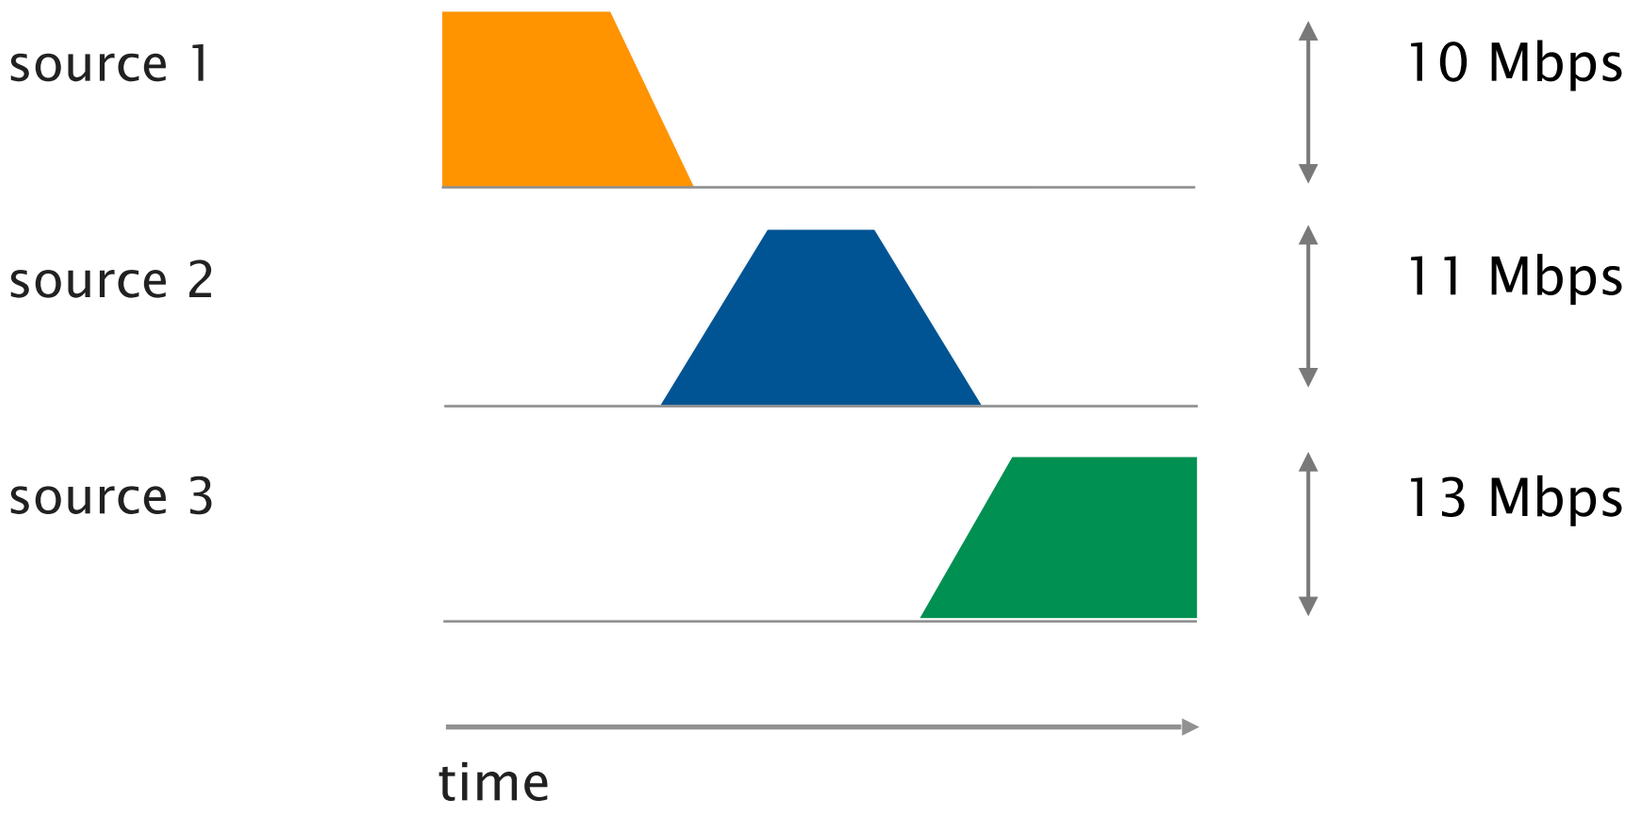

What does each source get with reservation and on-demand?

- first-come first-served
- equal (10 Mbps)

Peak vs average rates

Each flow has

- Peak rate — P
- Average rate — A

Reservation must reserve *P*, but level of utilization is *A/P*

e.g., *P*=100 Mbps, *A*=10 Mbps, level of utilization=10%

On-demand can achieve higher level of utilizations

depends on degree of sharing and burstiness of flows

Ultimately, it depends on the application

Reservation makes sense when P/A is small

voice traffic has a ratio of 3 or so

Reservation wastes capacity when P/A is big

data applications are bursty, ratios >100 are common

Ultimately, it depends on the application

Reservation makes sense when P/A is small

voice traffic has a ratio of 3 or so

Reservation wastes capacity when P/A is big

data applications are bursty, ratios >100 are common

That's why the phone network used reservations

… and why the Internet does not!

In practice, the two approaches are implemented using circuit-switching or packet-switching, respectively

	Reservation	On-demand
implem.	circuit-switching	packet-switching

	Reservation	On-demand
implem.	circuit-switching	packet-switching

Circuit switching relies on the Resource Reservation Protocol

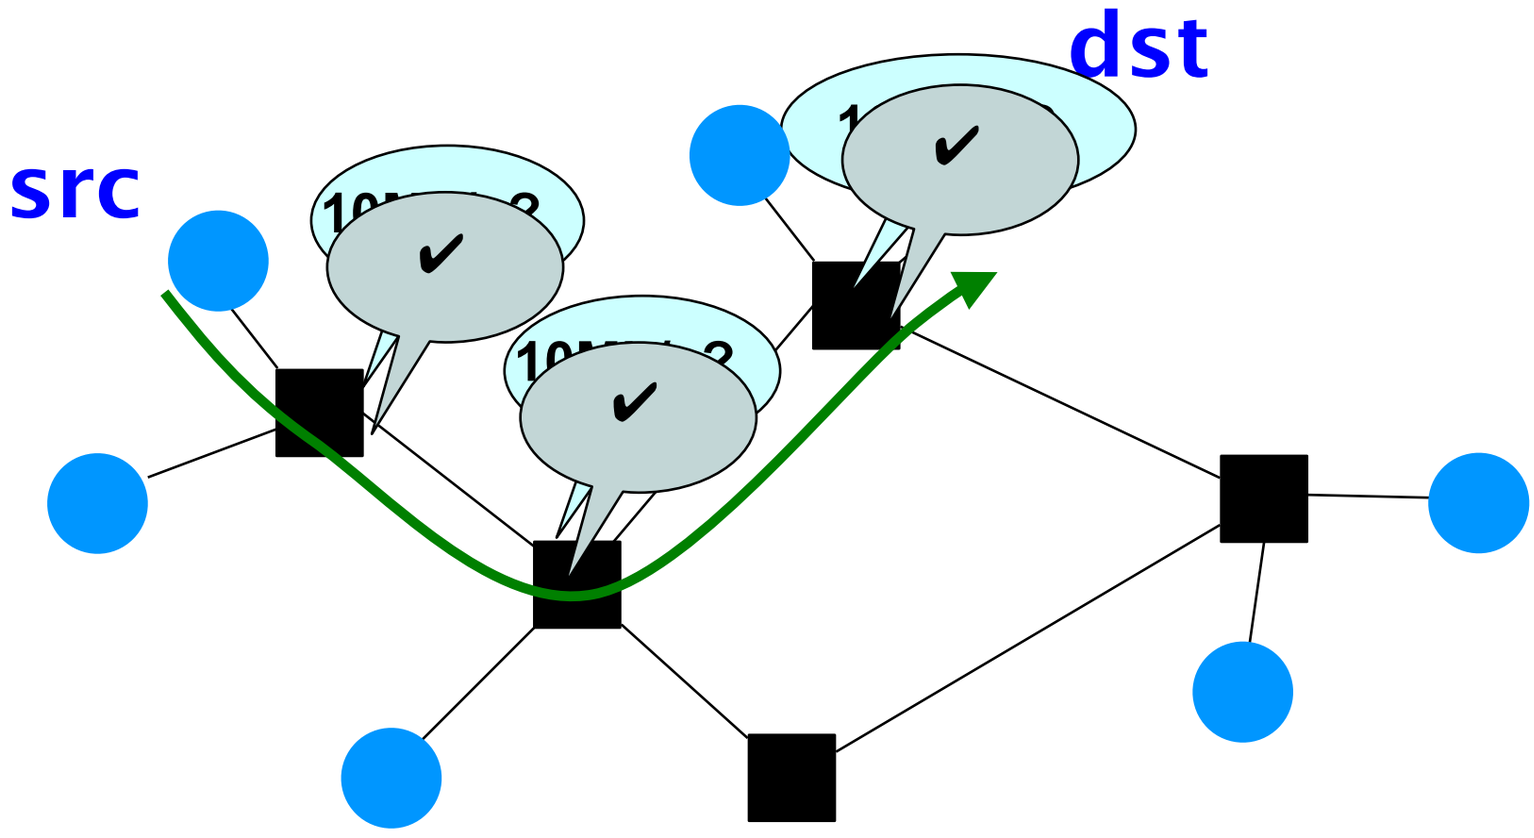

(1) src sends a reservation request for 10Mbps to dst
(2) switches “establish a circuit”
(3) src starts sending data
(4) src sends a “teardown circuit” message

The Resource Reservation Protocol establishes a circuit within each switch

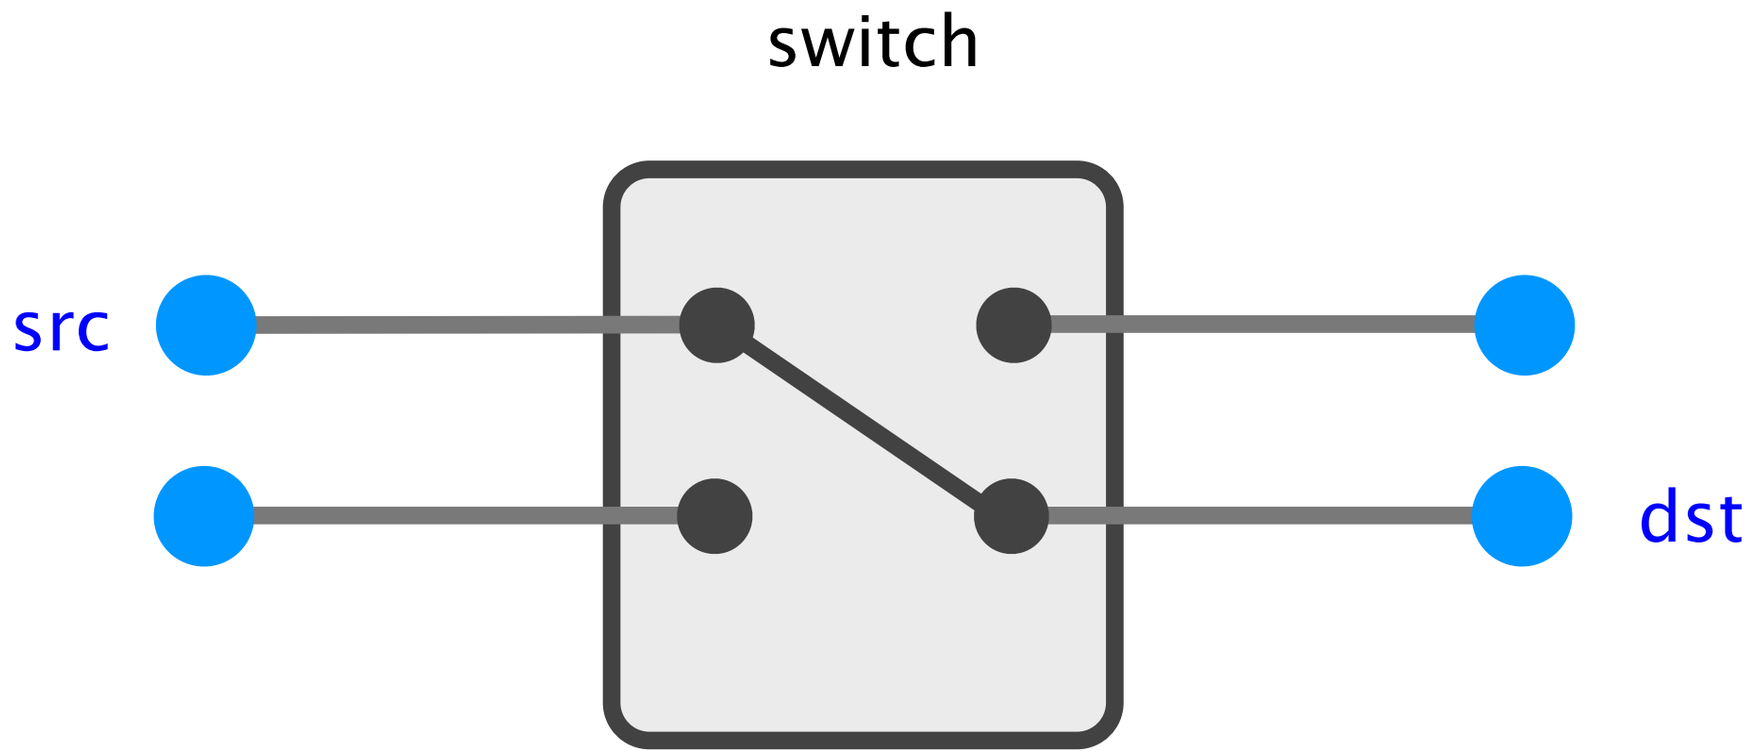

There exist many kinds of circuits

Time-based multiplexing

- divide time in slots
- allocate one slot per circuit

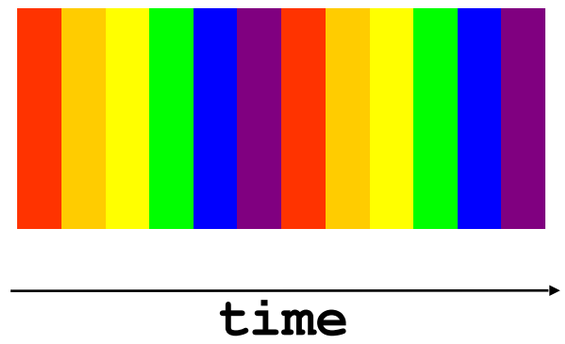

Frequency-based multiplexing

- divide spectrum in frequency bands
- allocate one band per circuit

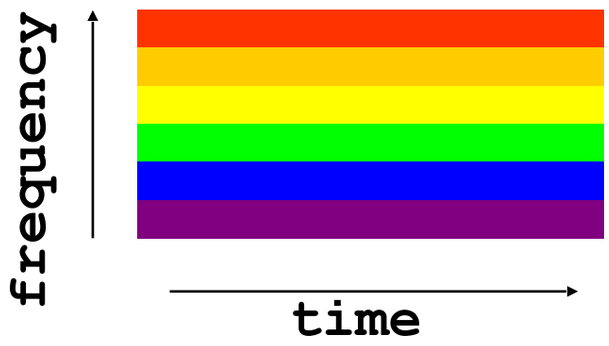

Let’s walk through example of data transfer using circuit switching

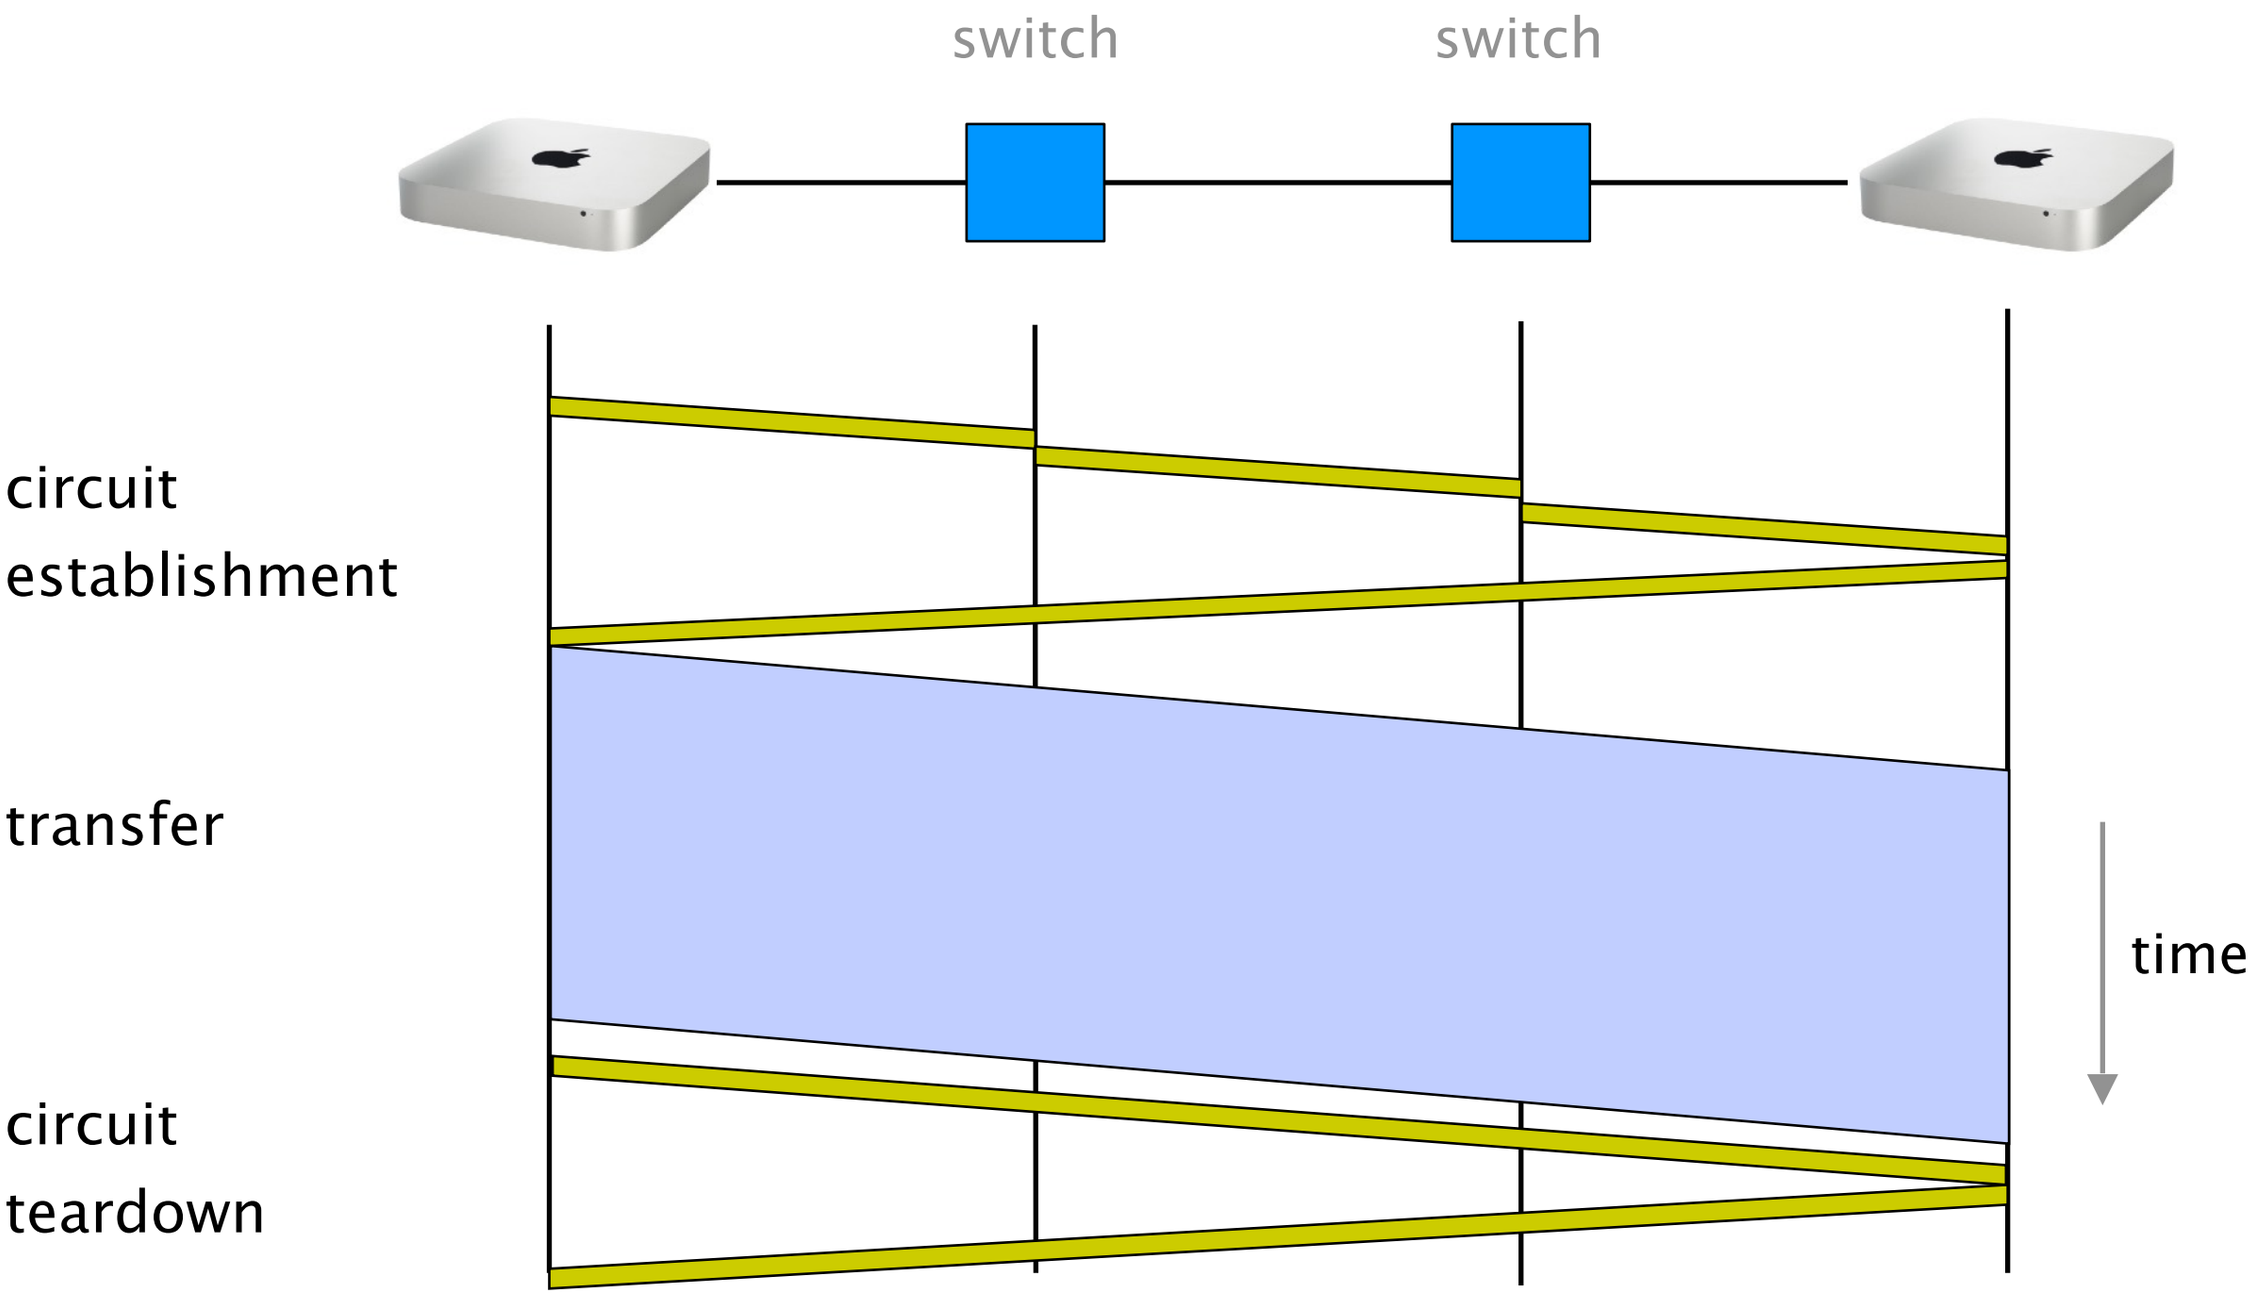

The efficiency of the transfer depends on how utilized the circuit is once established

This is an example of poor efficiency.

The circuit is mostly idle due to traffic bursts

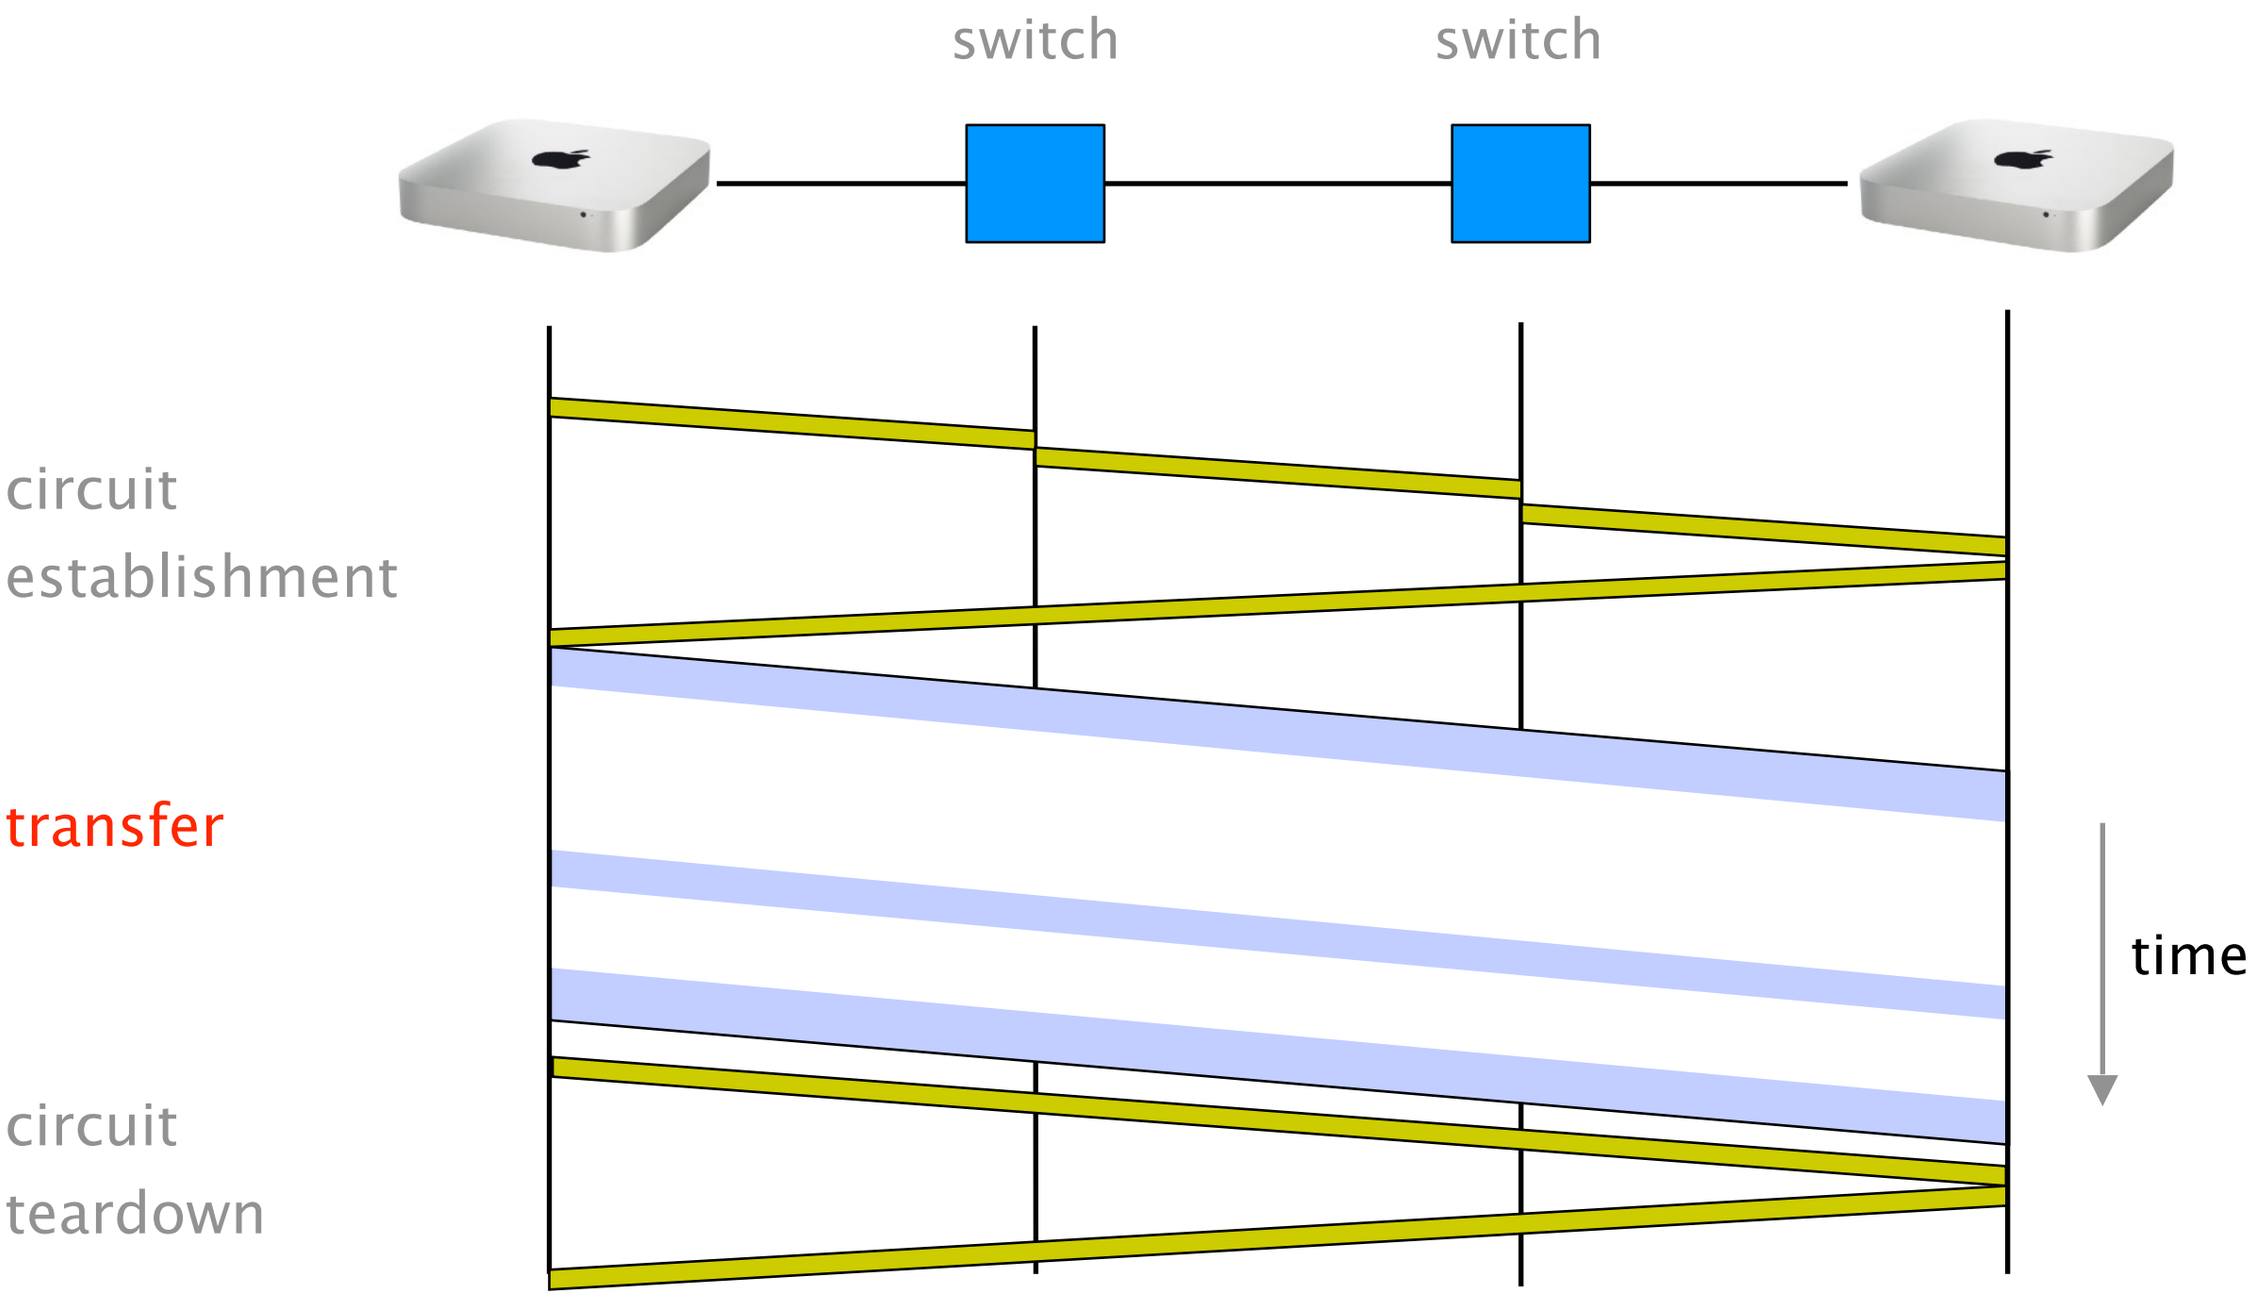

This is another example of poor efficiency.

The circuit is used for a short amount of time

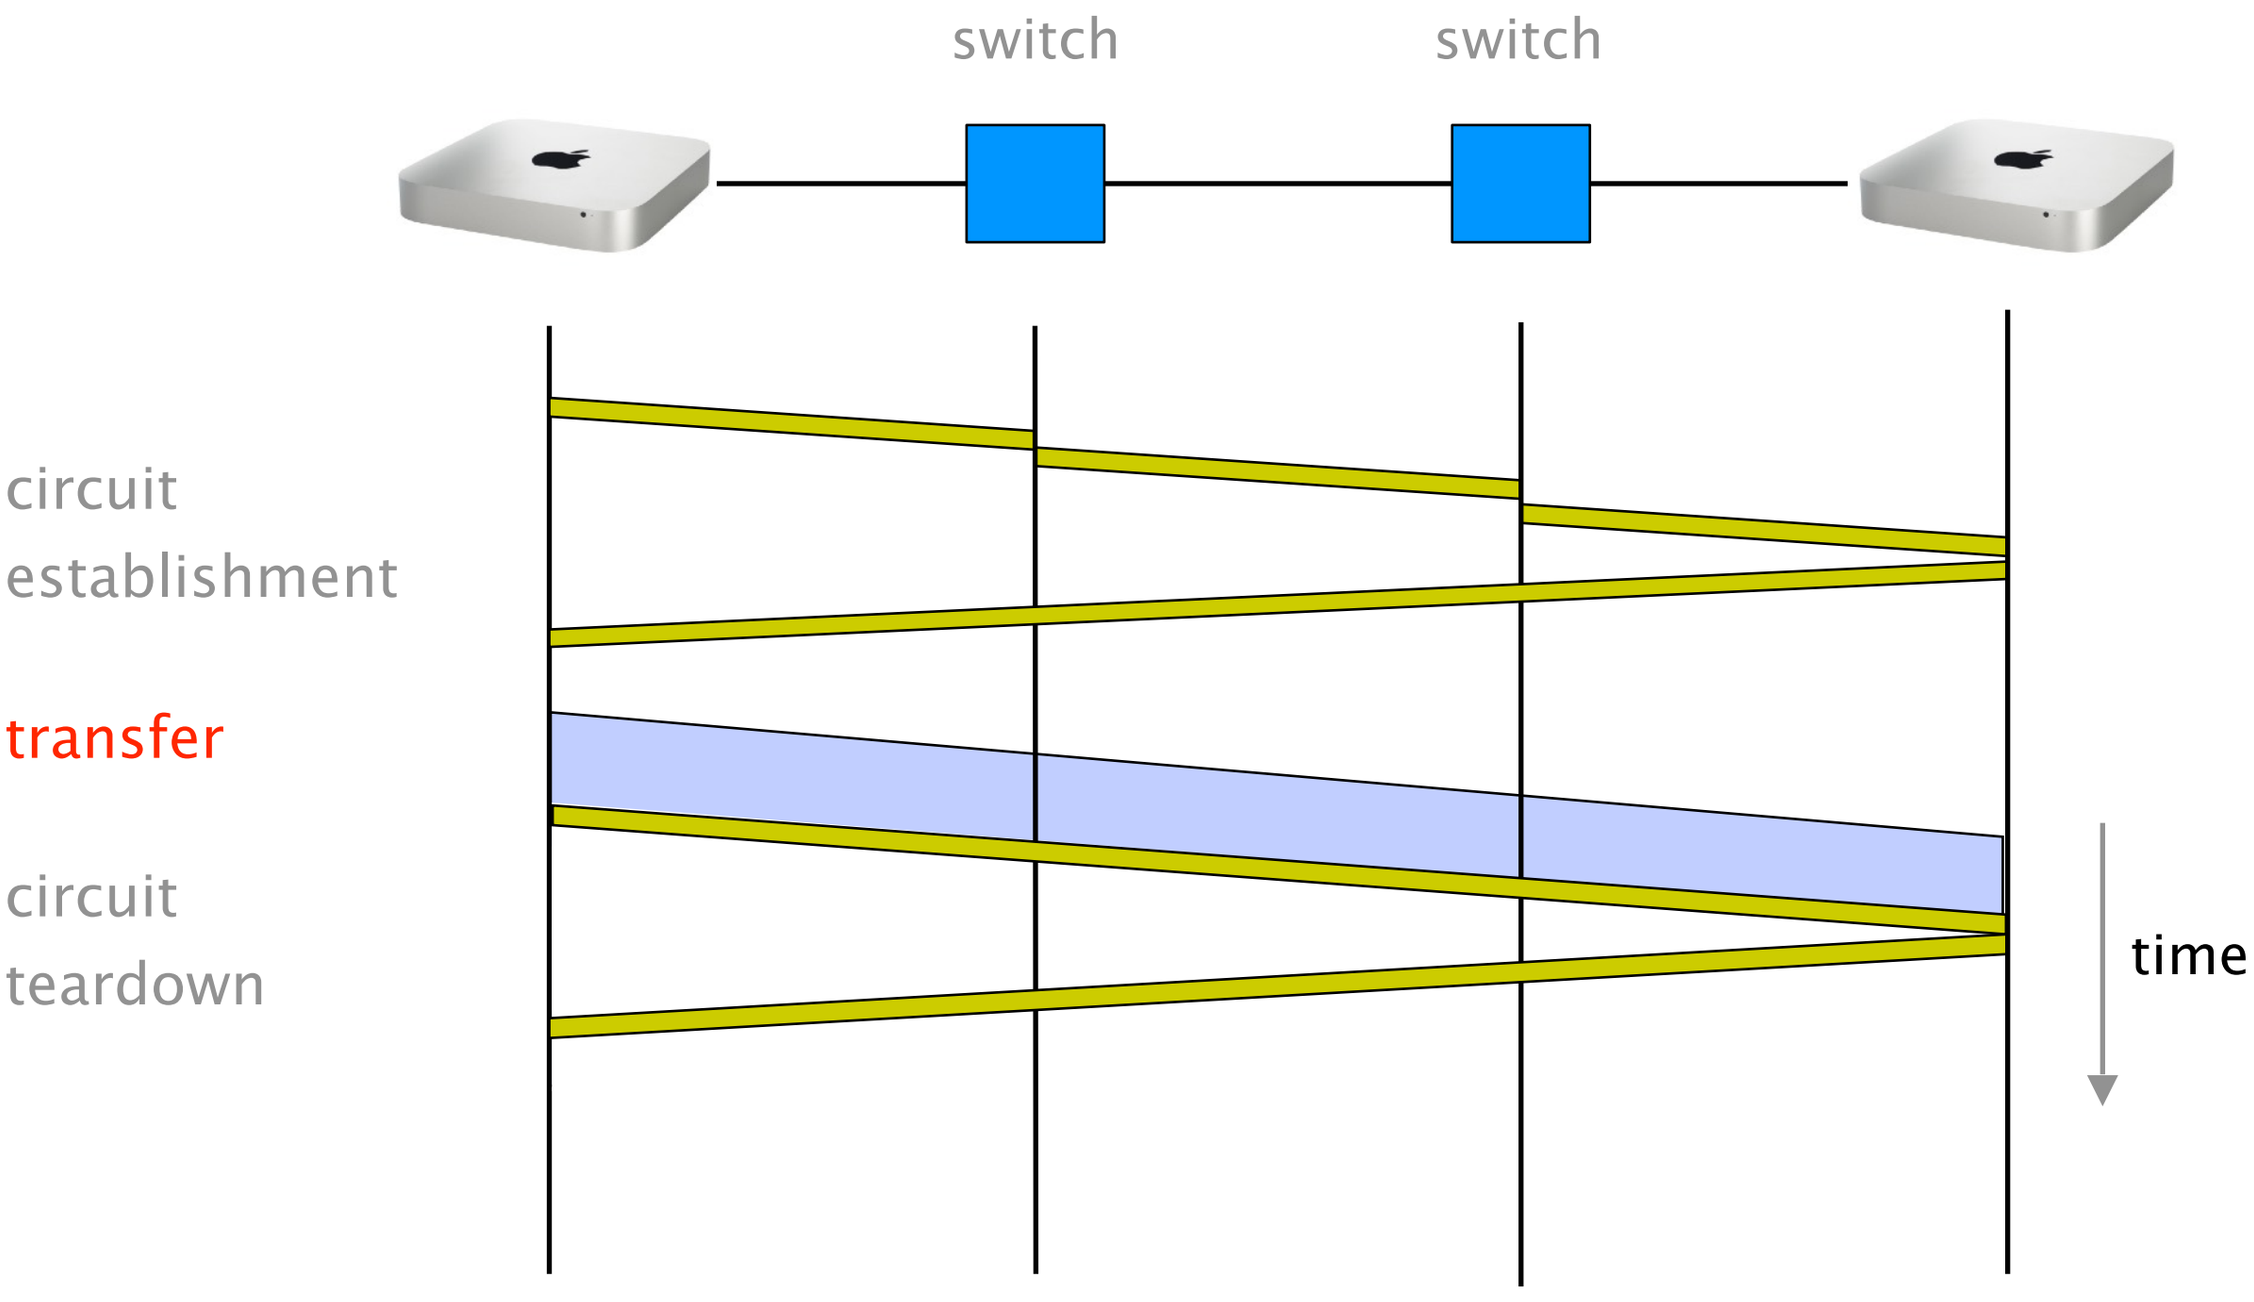

Another problem of circuit switching is that it doesn't route around trouble

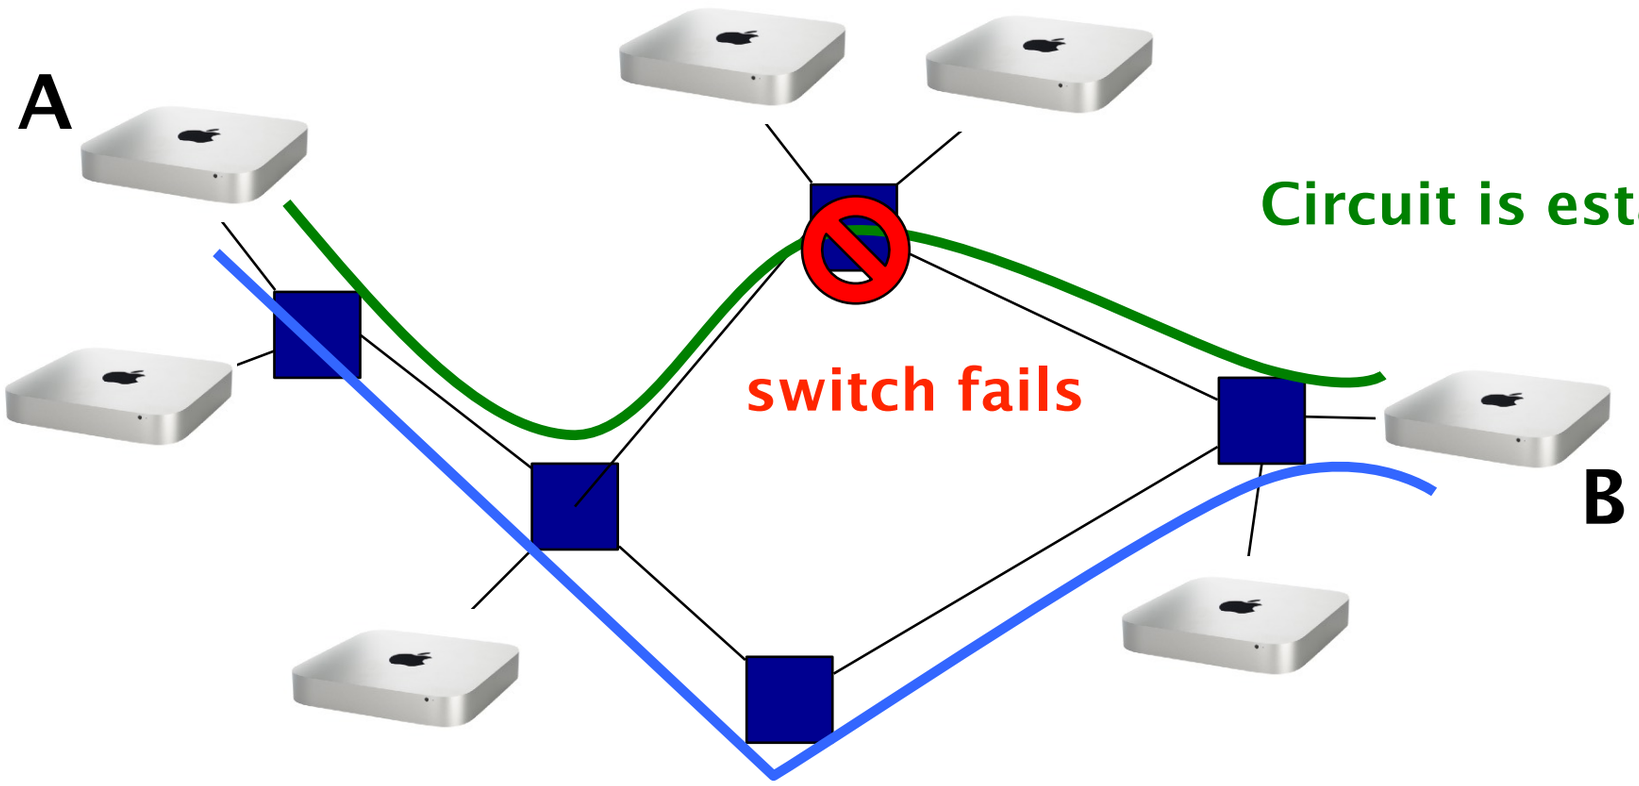

Circuit is established

A is forced to signal a new circuit to restore communication

Pros and cons of circuit switching

advantages	disadvantages
predictable performance	inefficient if traffic is bursty or short
simple & fast switching once circuit established	complex circuit setup/teardown which adds delays to transfer
	requires new circuit upon failure

What about packet switching?

Reservation

On-demand

circuit-switching

packet-switching

In packet switching,

data transfer is done using independent packets

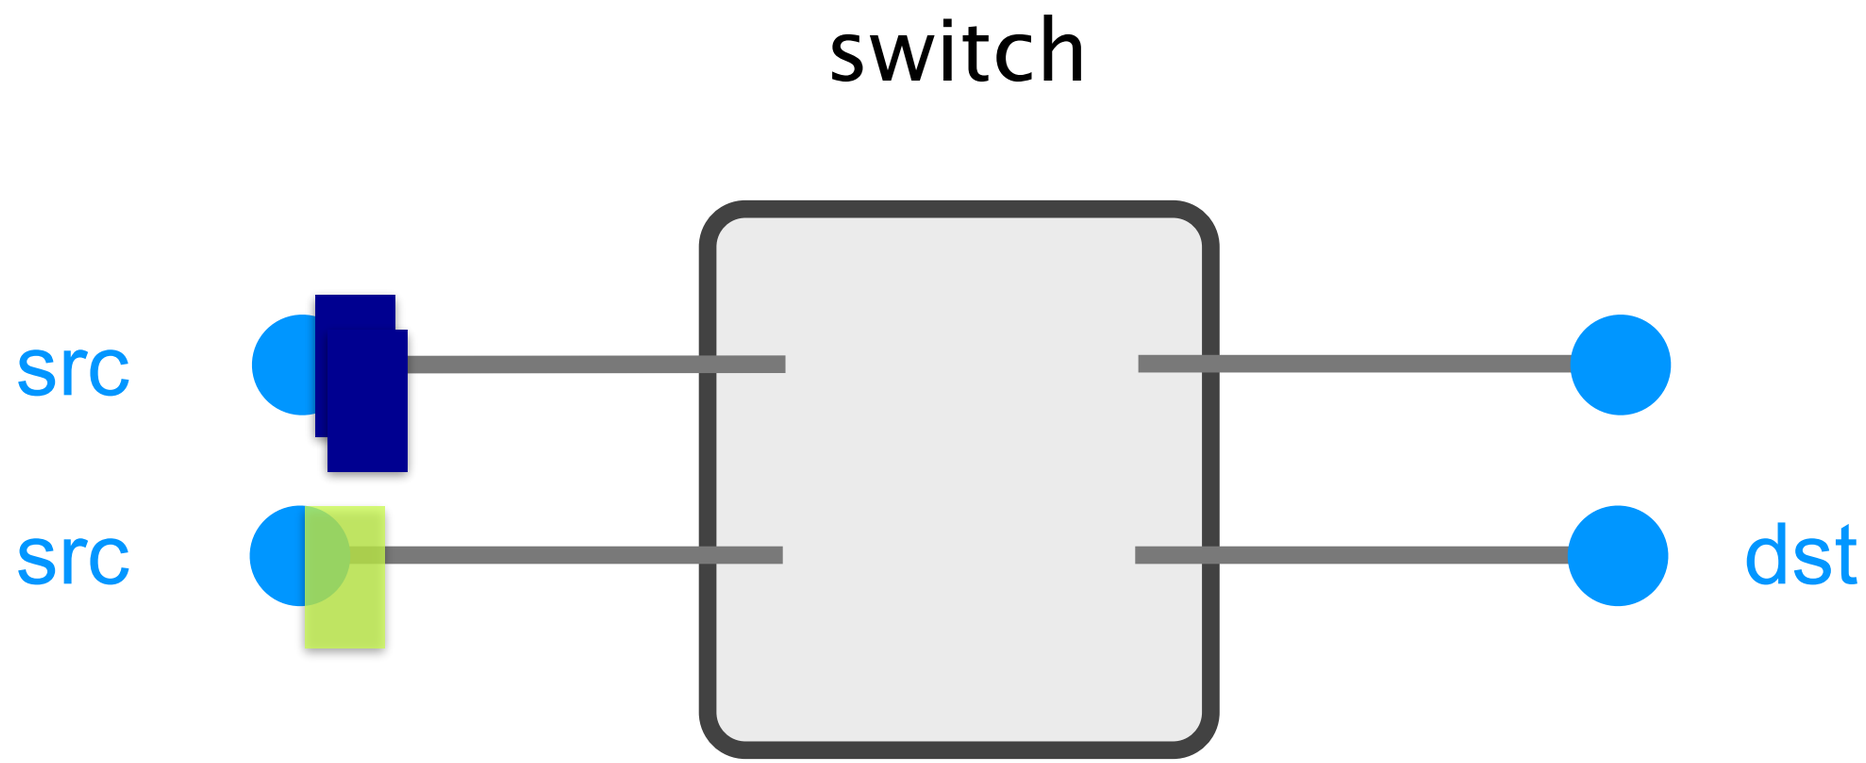

Each packet contains destination (dst)

Since packets are sent without global coordination, they can “clash” with each other

To absorb transient overload,
packet switching relies on buffers

To absorb transient overload,
packet switching relies on buffers

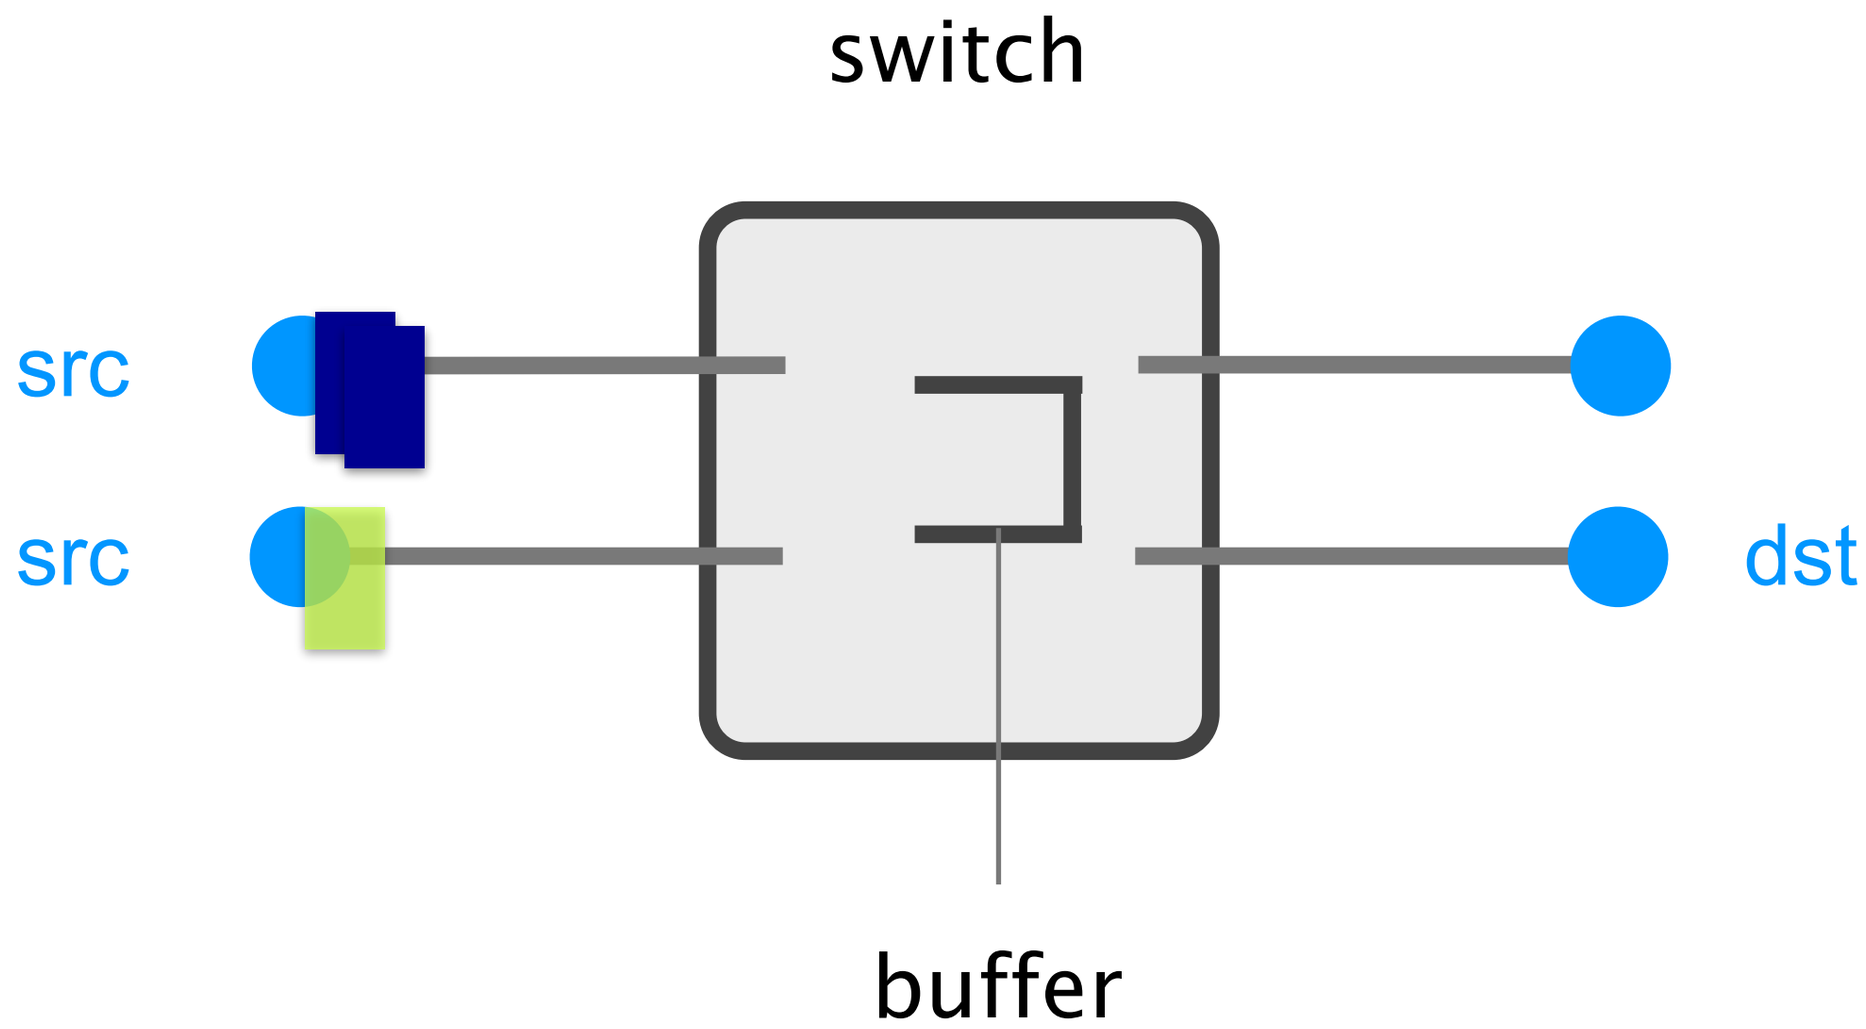

Packet switching routes around trouble

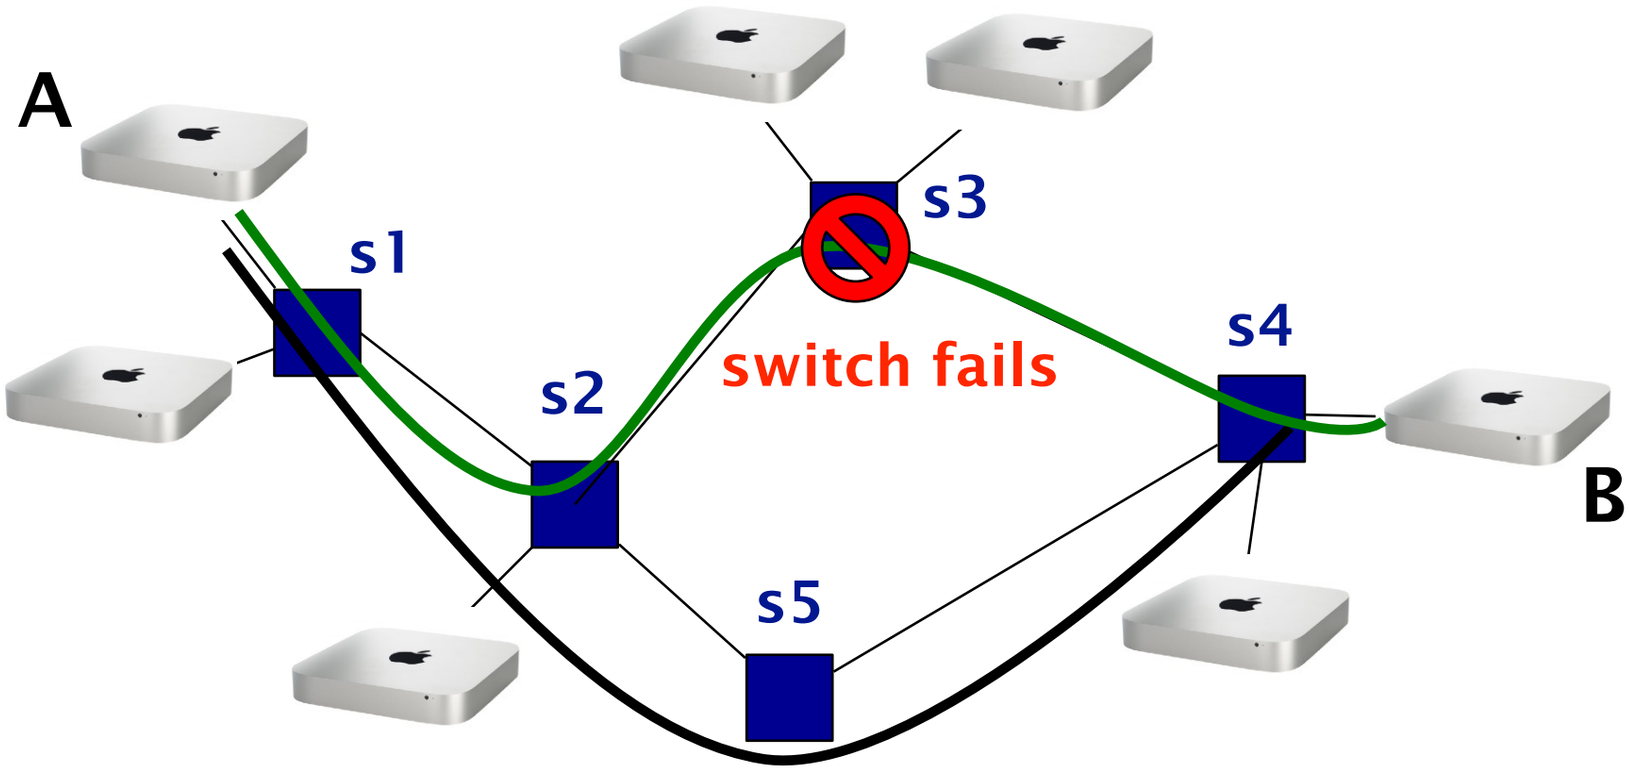

route recomputed
on the fly by s2

Pros and cons of packet switching

advantages	disadvantages
efficient use of resources	unpredictable performance
simpler to implement than circuit switching	requires buffer management and congestion control
route around trouble	

Packet switching beats Circuit switching
with respect to *resiliency* and *efficiency*

Internet  packets

Packet switching will be our focus for the rest of the course

Communication Networks

Part 1: Overview

What is a network made of?

How is it shared?

#3 How is it organized?

How does communication happen?

How do we characterize it?

So far, this is our vision of the Internet…

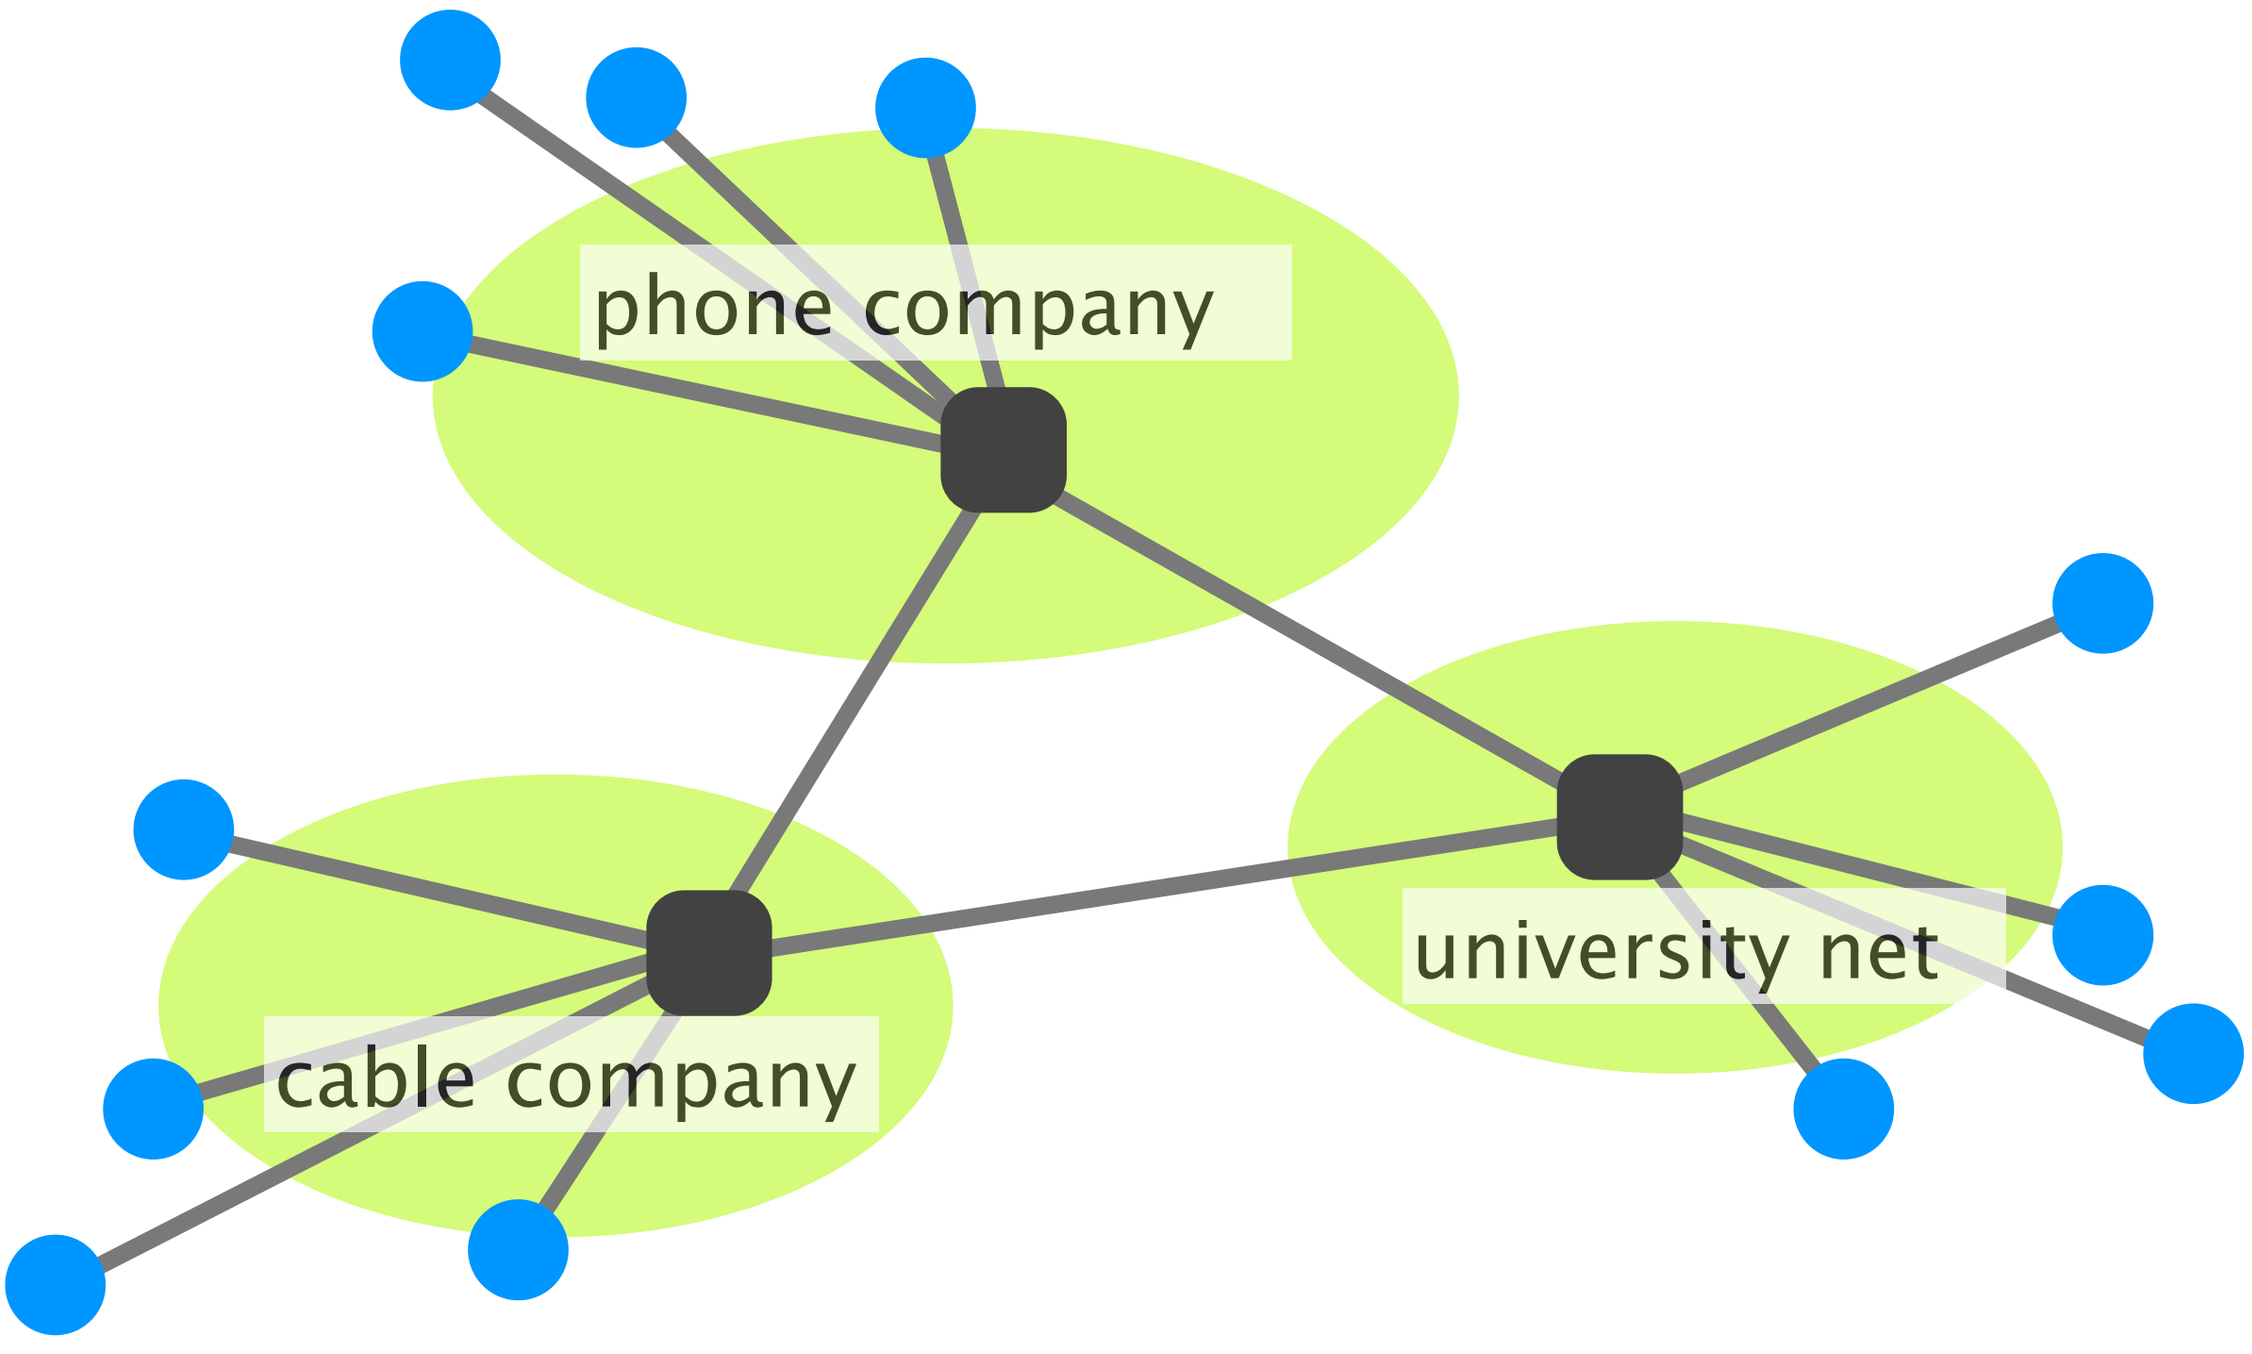

In practice, the Internet is a “bit” more complex

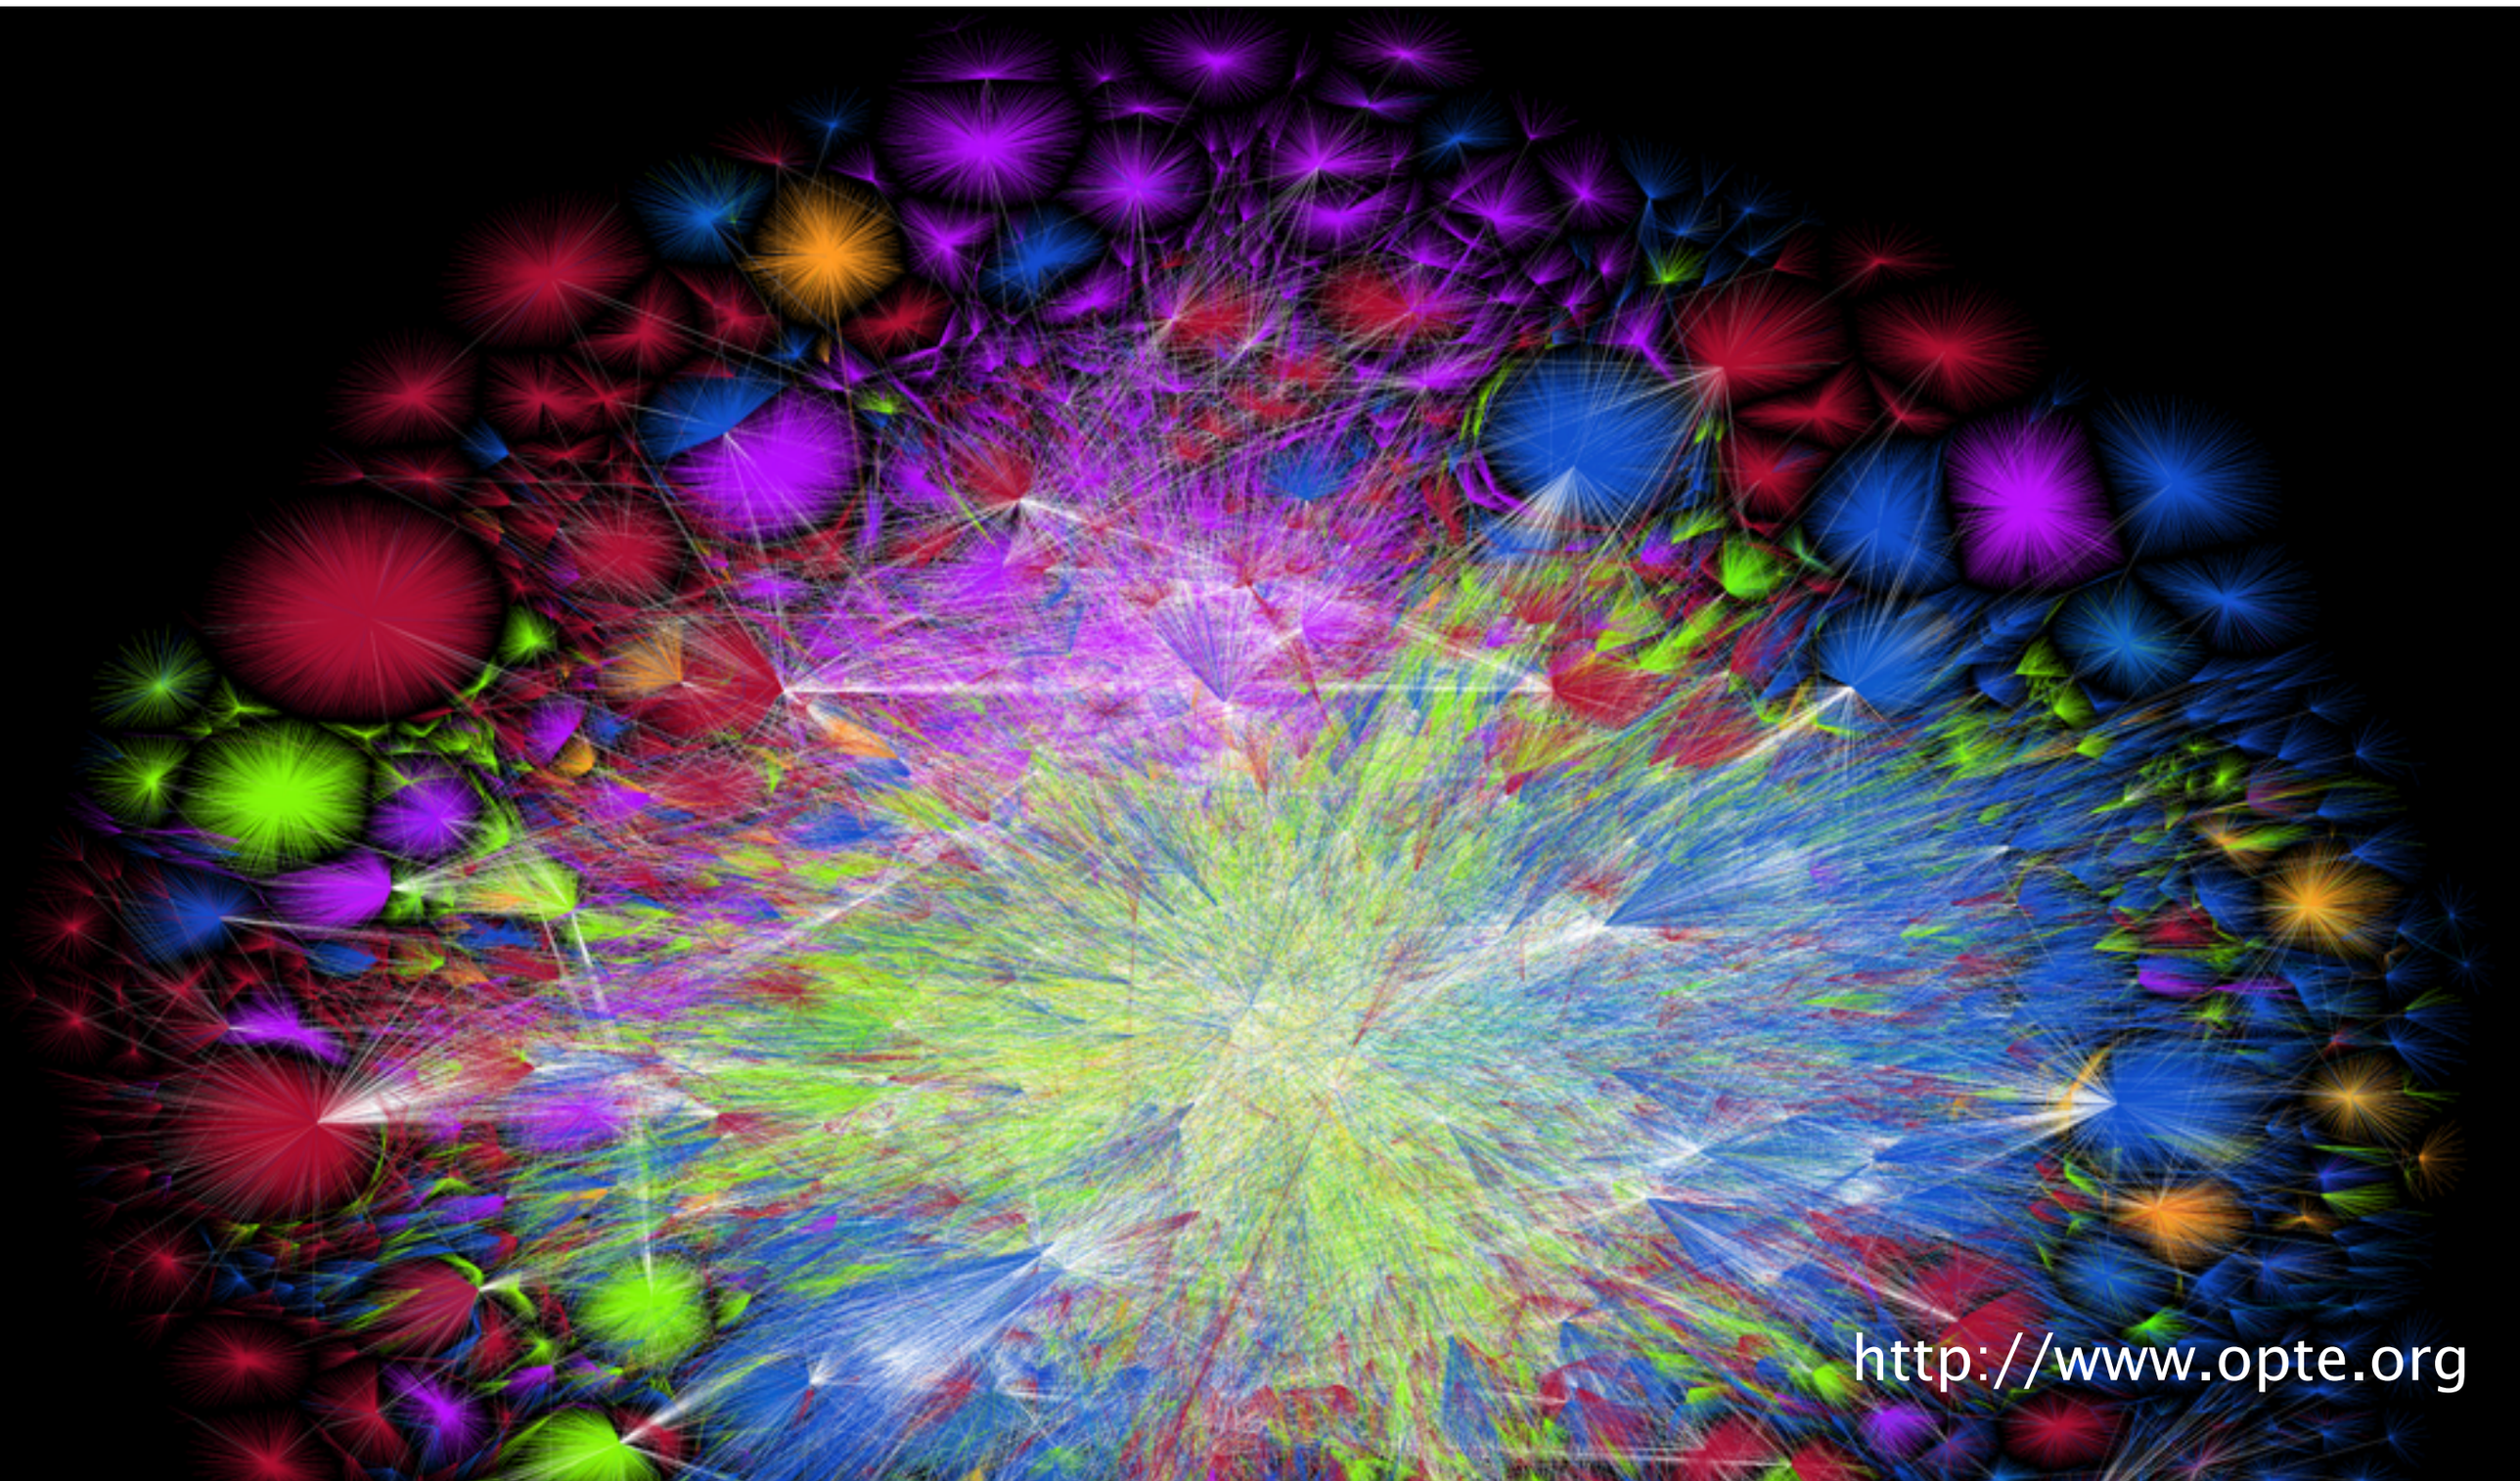

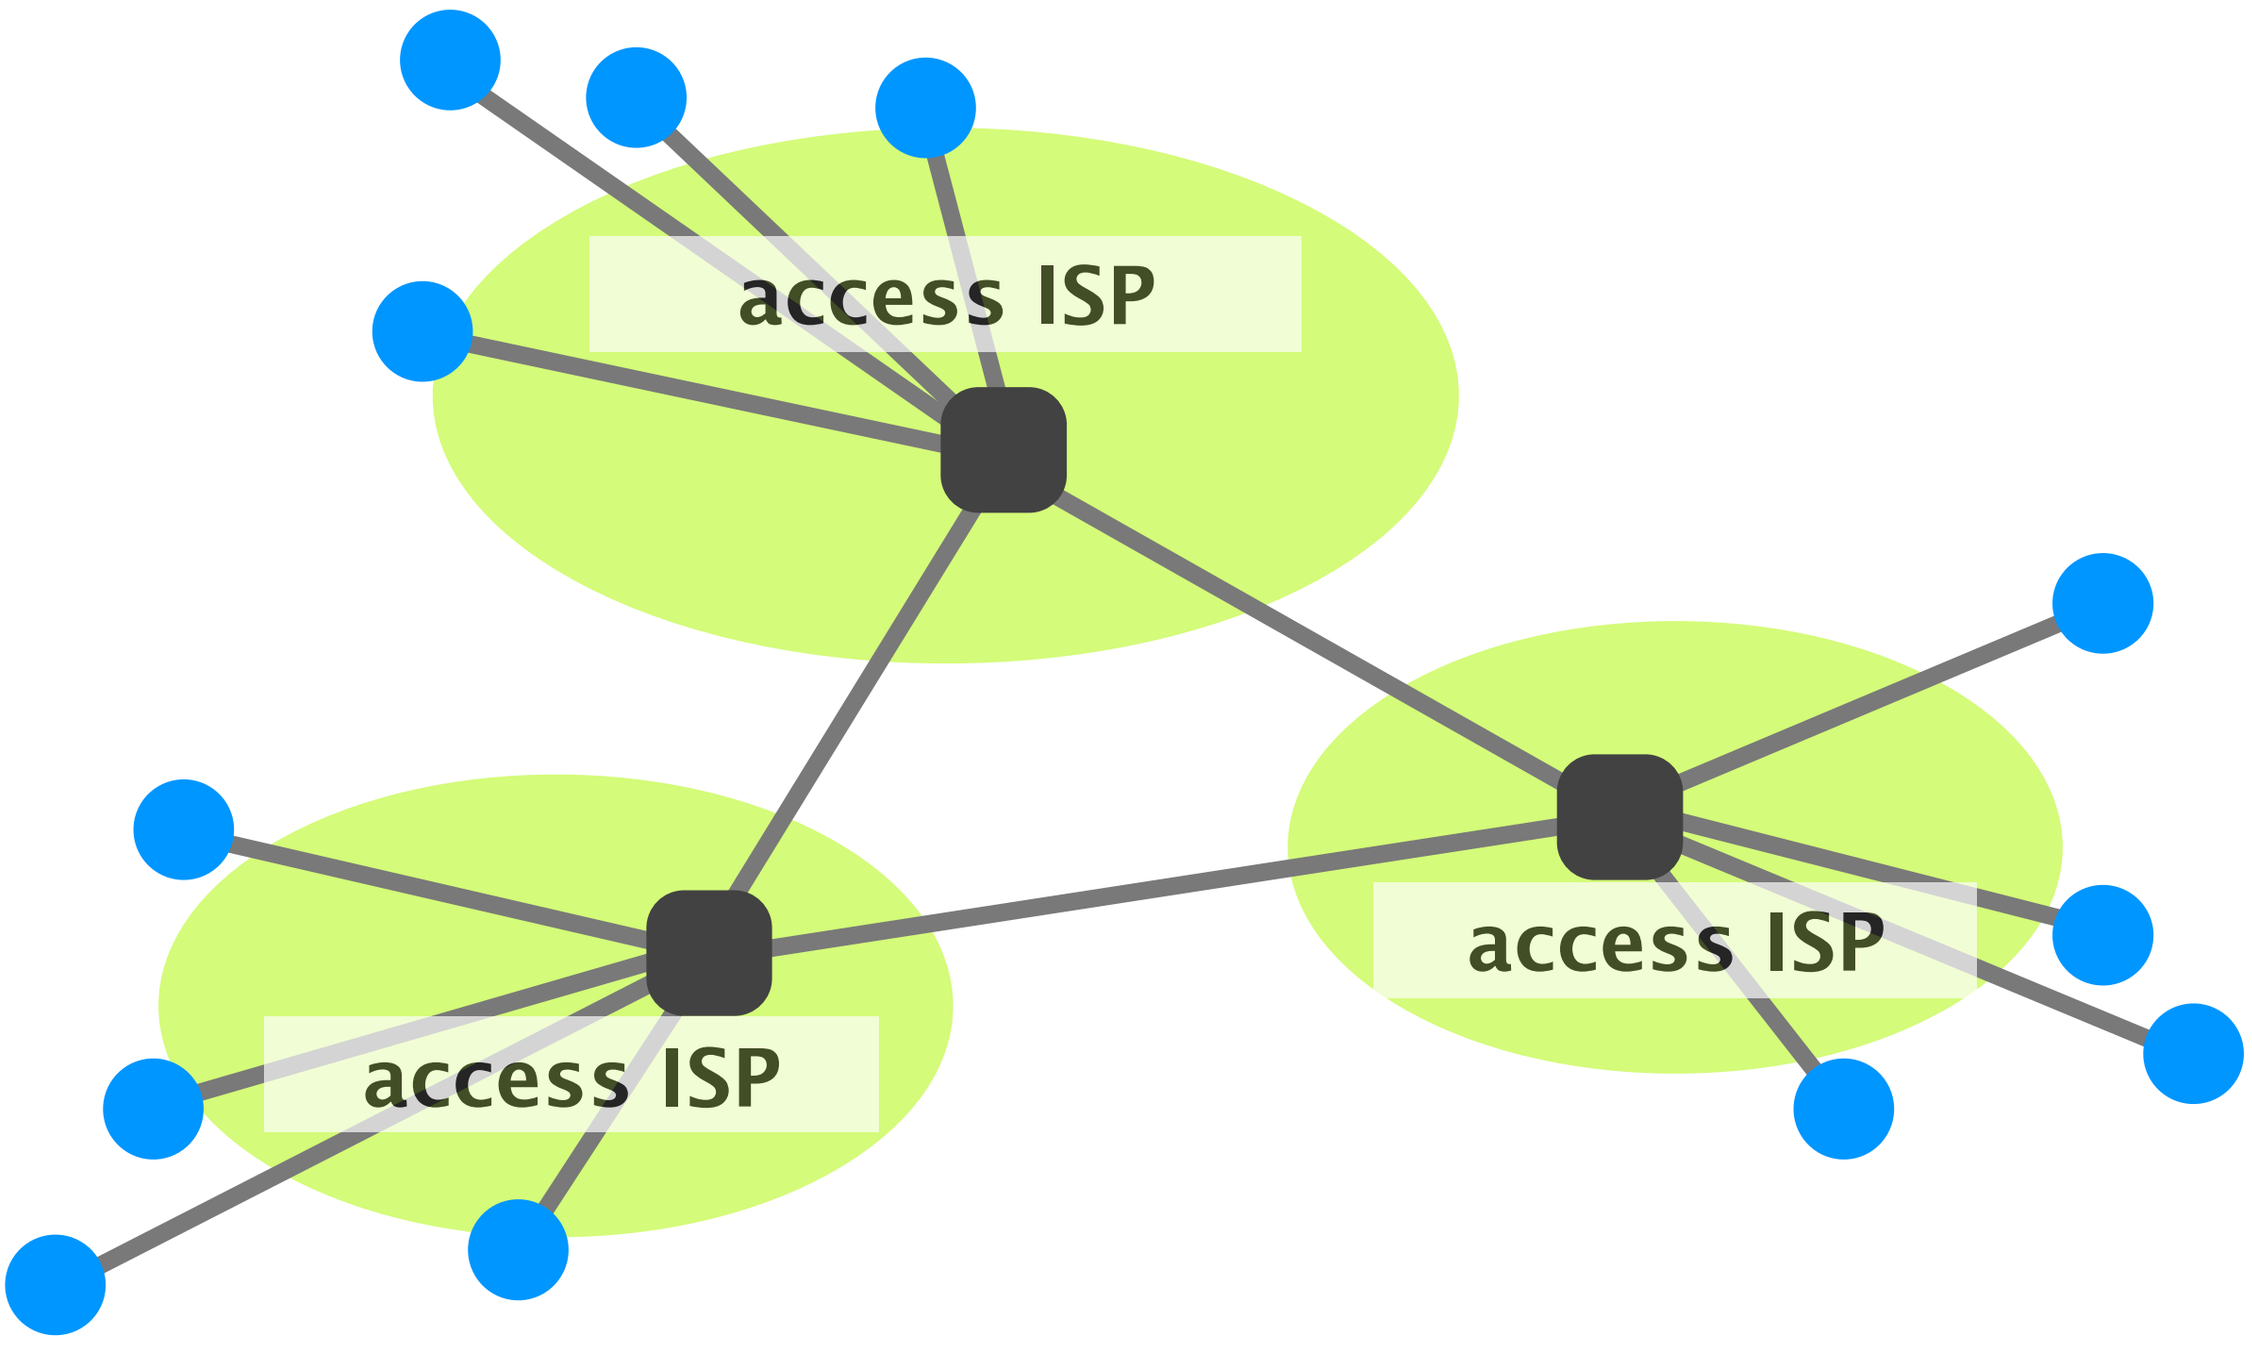
access ISP
access ISP
access ISP

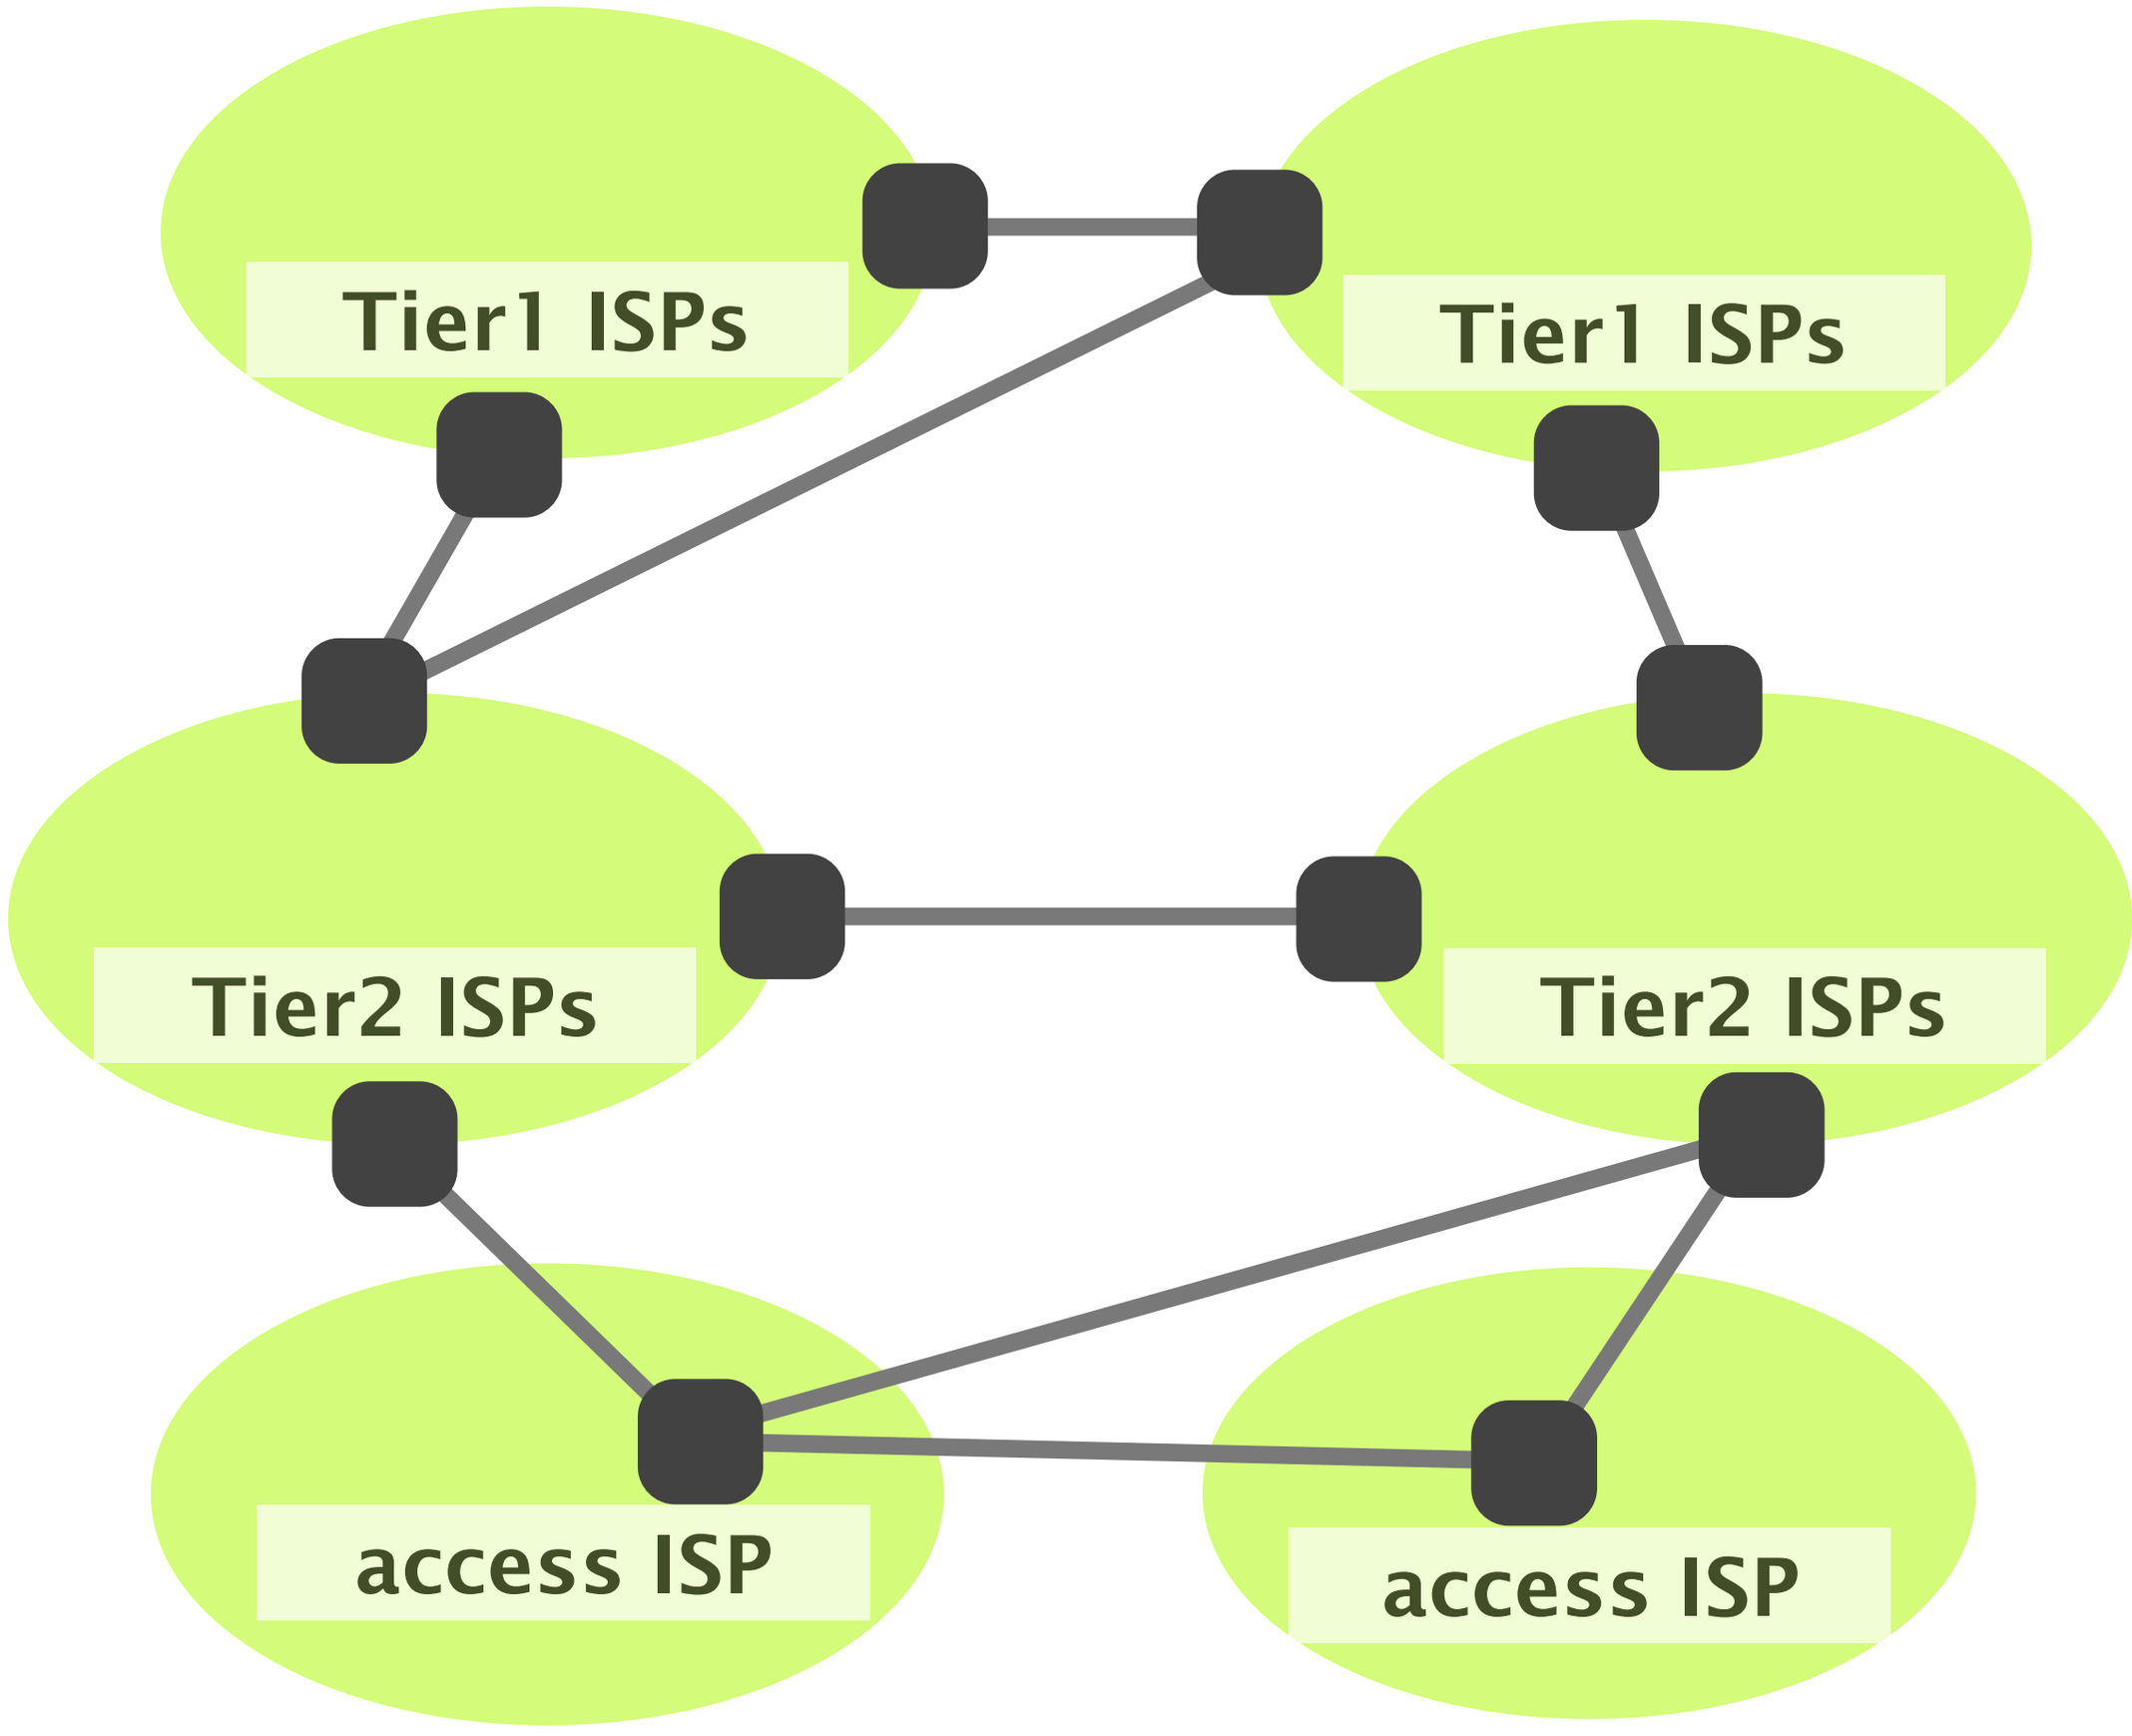
Tier1 ISPs
Tier1 ISPs
Tier2 ISPs
Tier2 ISPs
access ISP
access ISP

The Internet has a hierarchical structure

Tier-1 international	have no provider
Tier-2 national	provide transit to Tier-3s have at least one provider
Tier-3 local	do not provide any transit have at least one provider

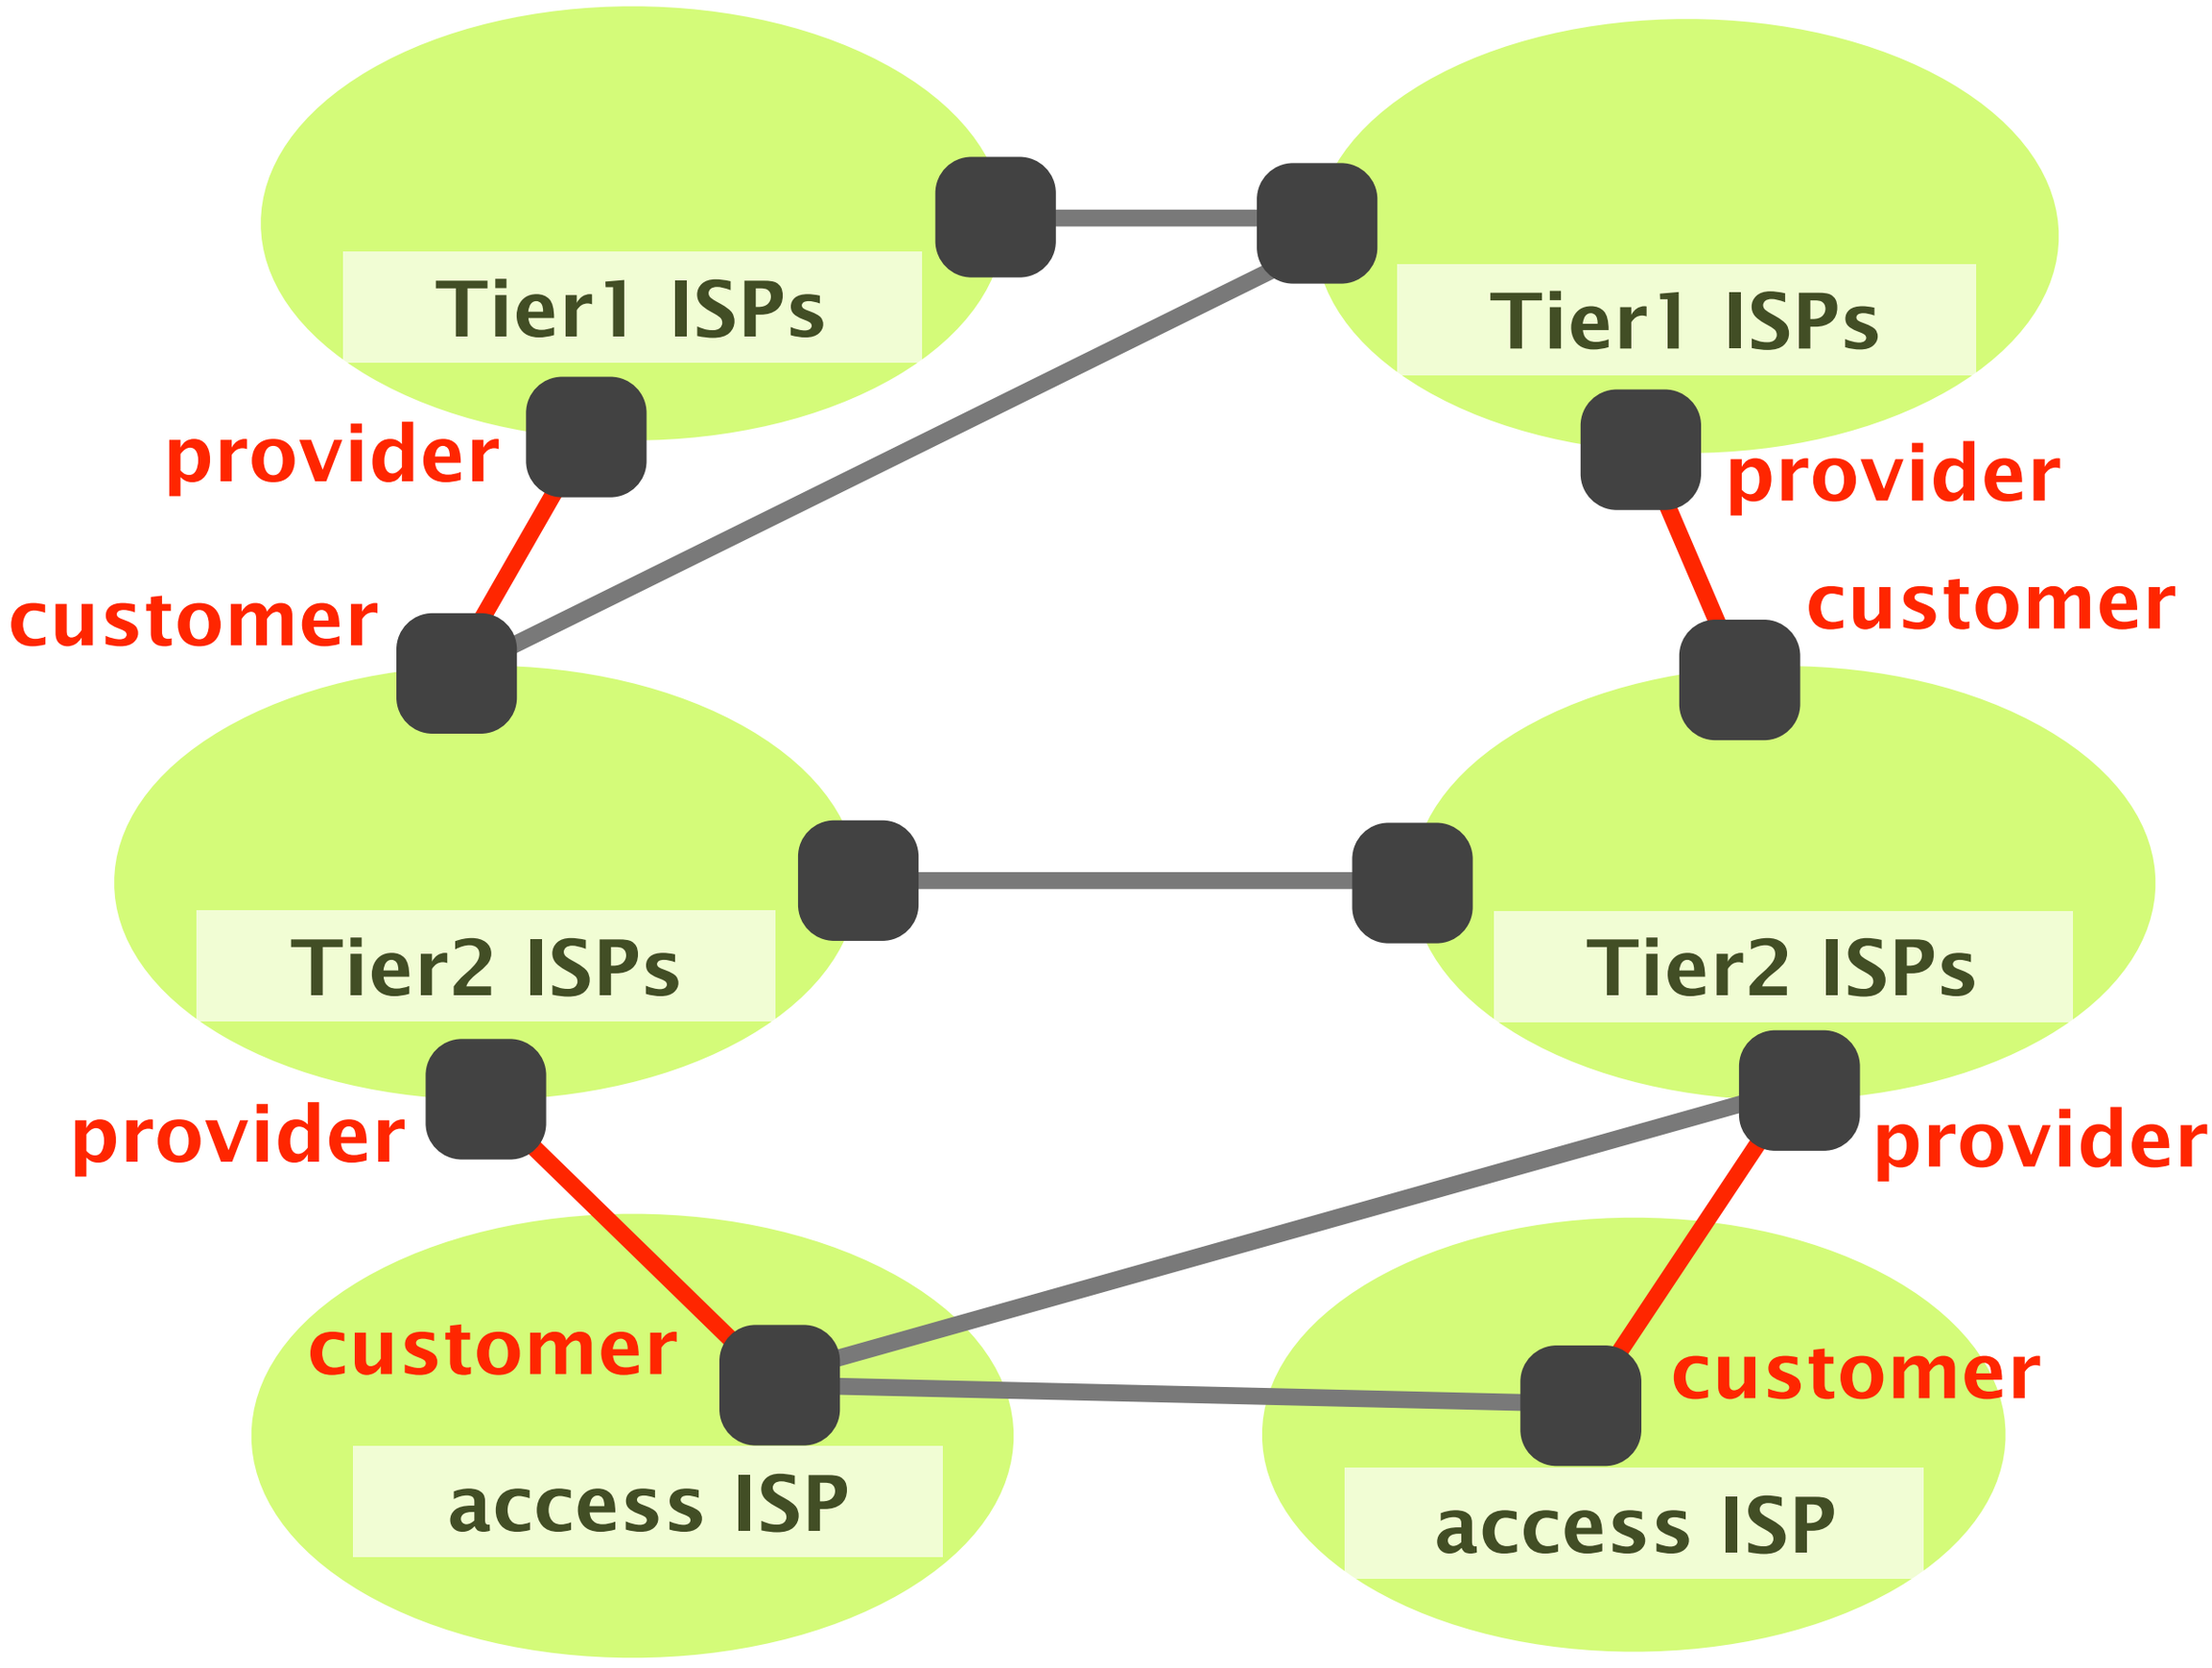
Tier1 ISPs
Tier1 ISPs
provider
customer
provider
customer
Tier2 ISPs
Tier2 ISPs
provider
provider
customer
customer
access ISP
access ISP

The distribution of networks in Tiers is extremely skewed towards Tier-3s

	total	~50,000 networks
Tier-1 international	have no provider	~12
Tier-2 national	provide transit to Tier-3s have at least one provider	1,000s
Tier-3 local	do not provide any transit have at least one provider	85-90%

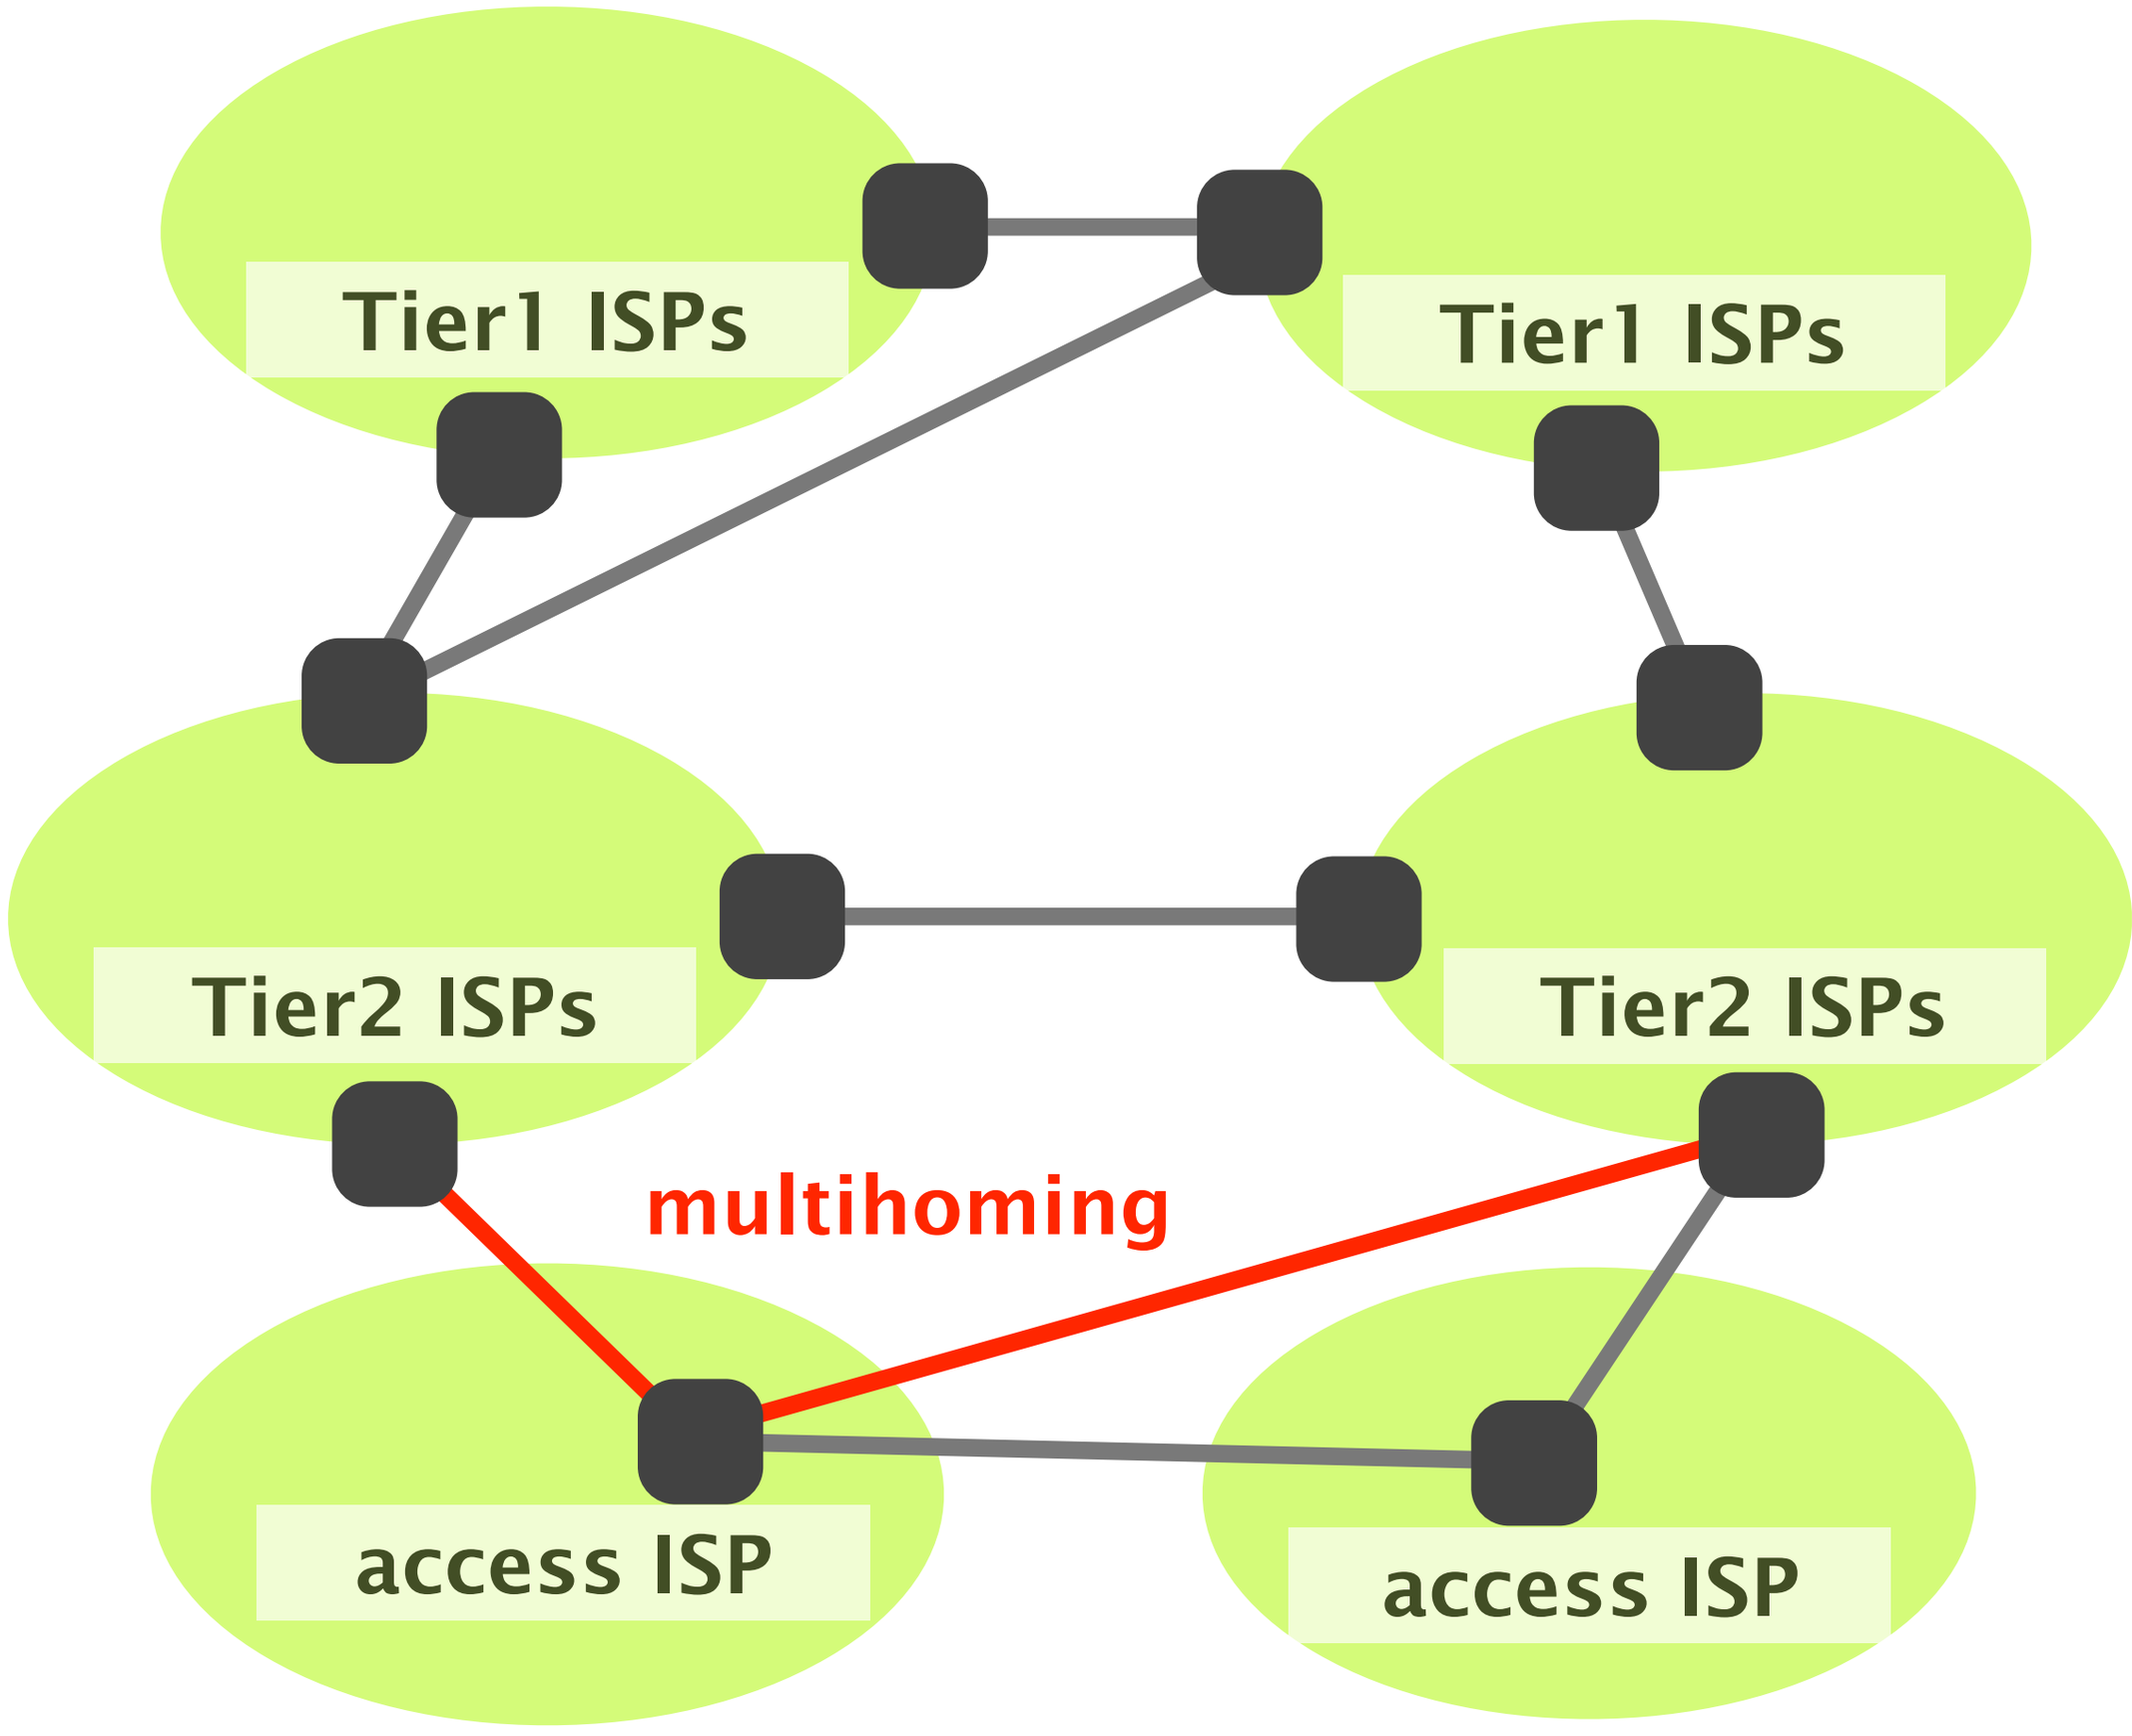
Tier1 ISPs
Tier1 ISPs
Tier2 ISPs
Tier2 ISPs
multihoming
access ISP
access ISP

Some networks have an incentive to connect directly, to reduce their bill wrt their own provider

This is known as “peering”

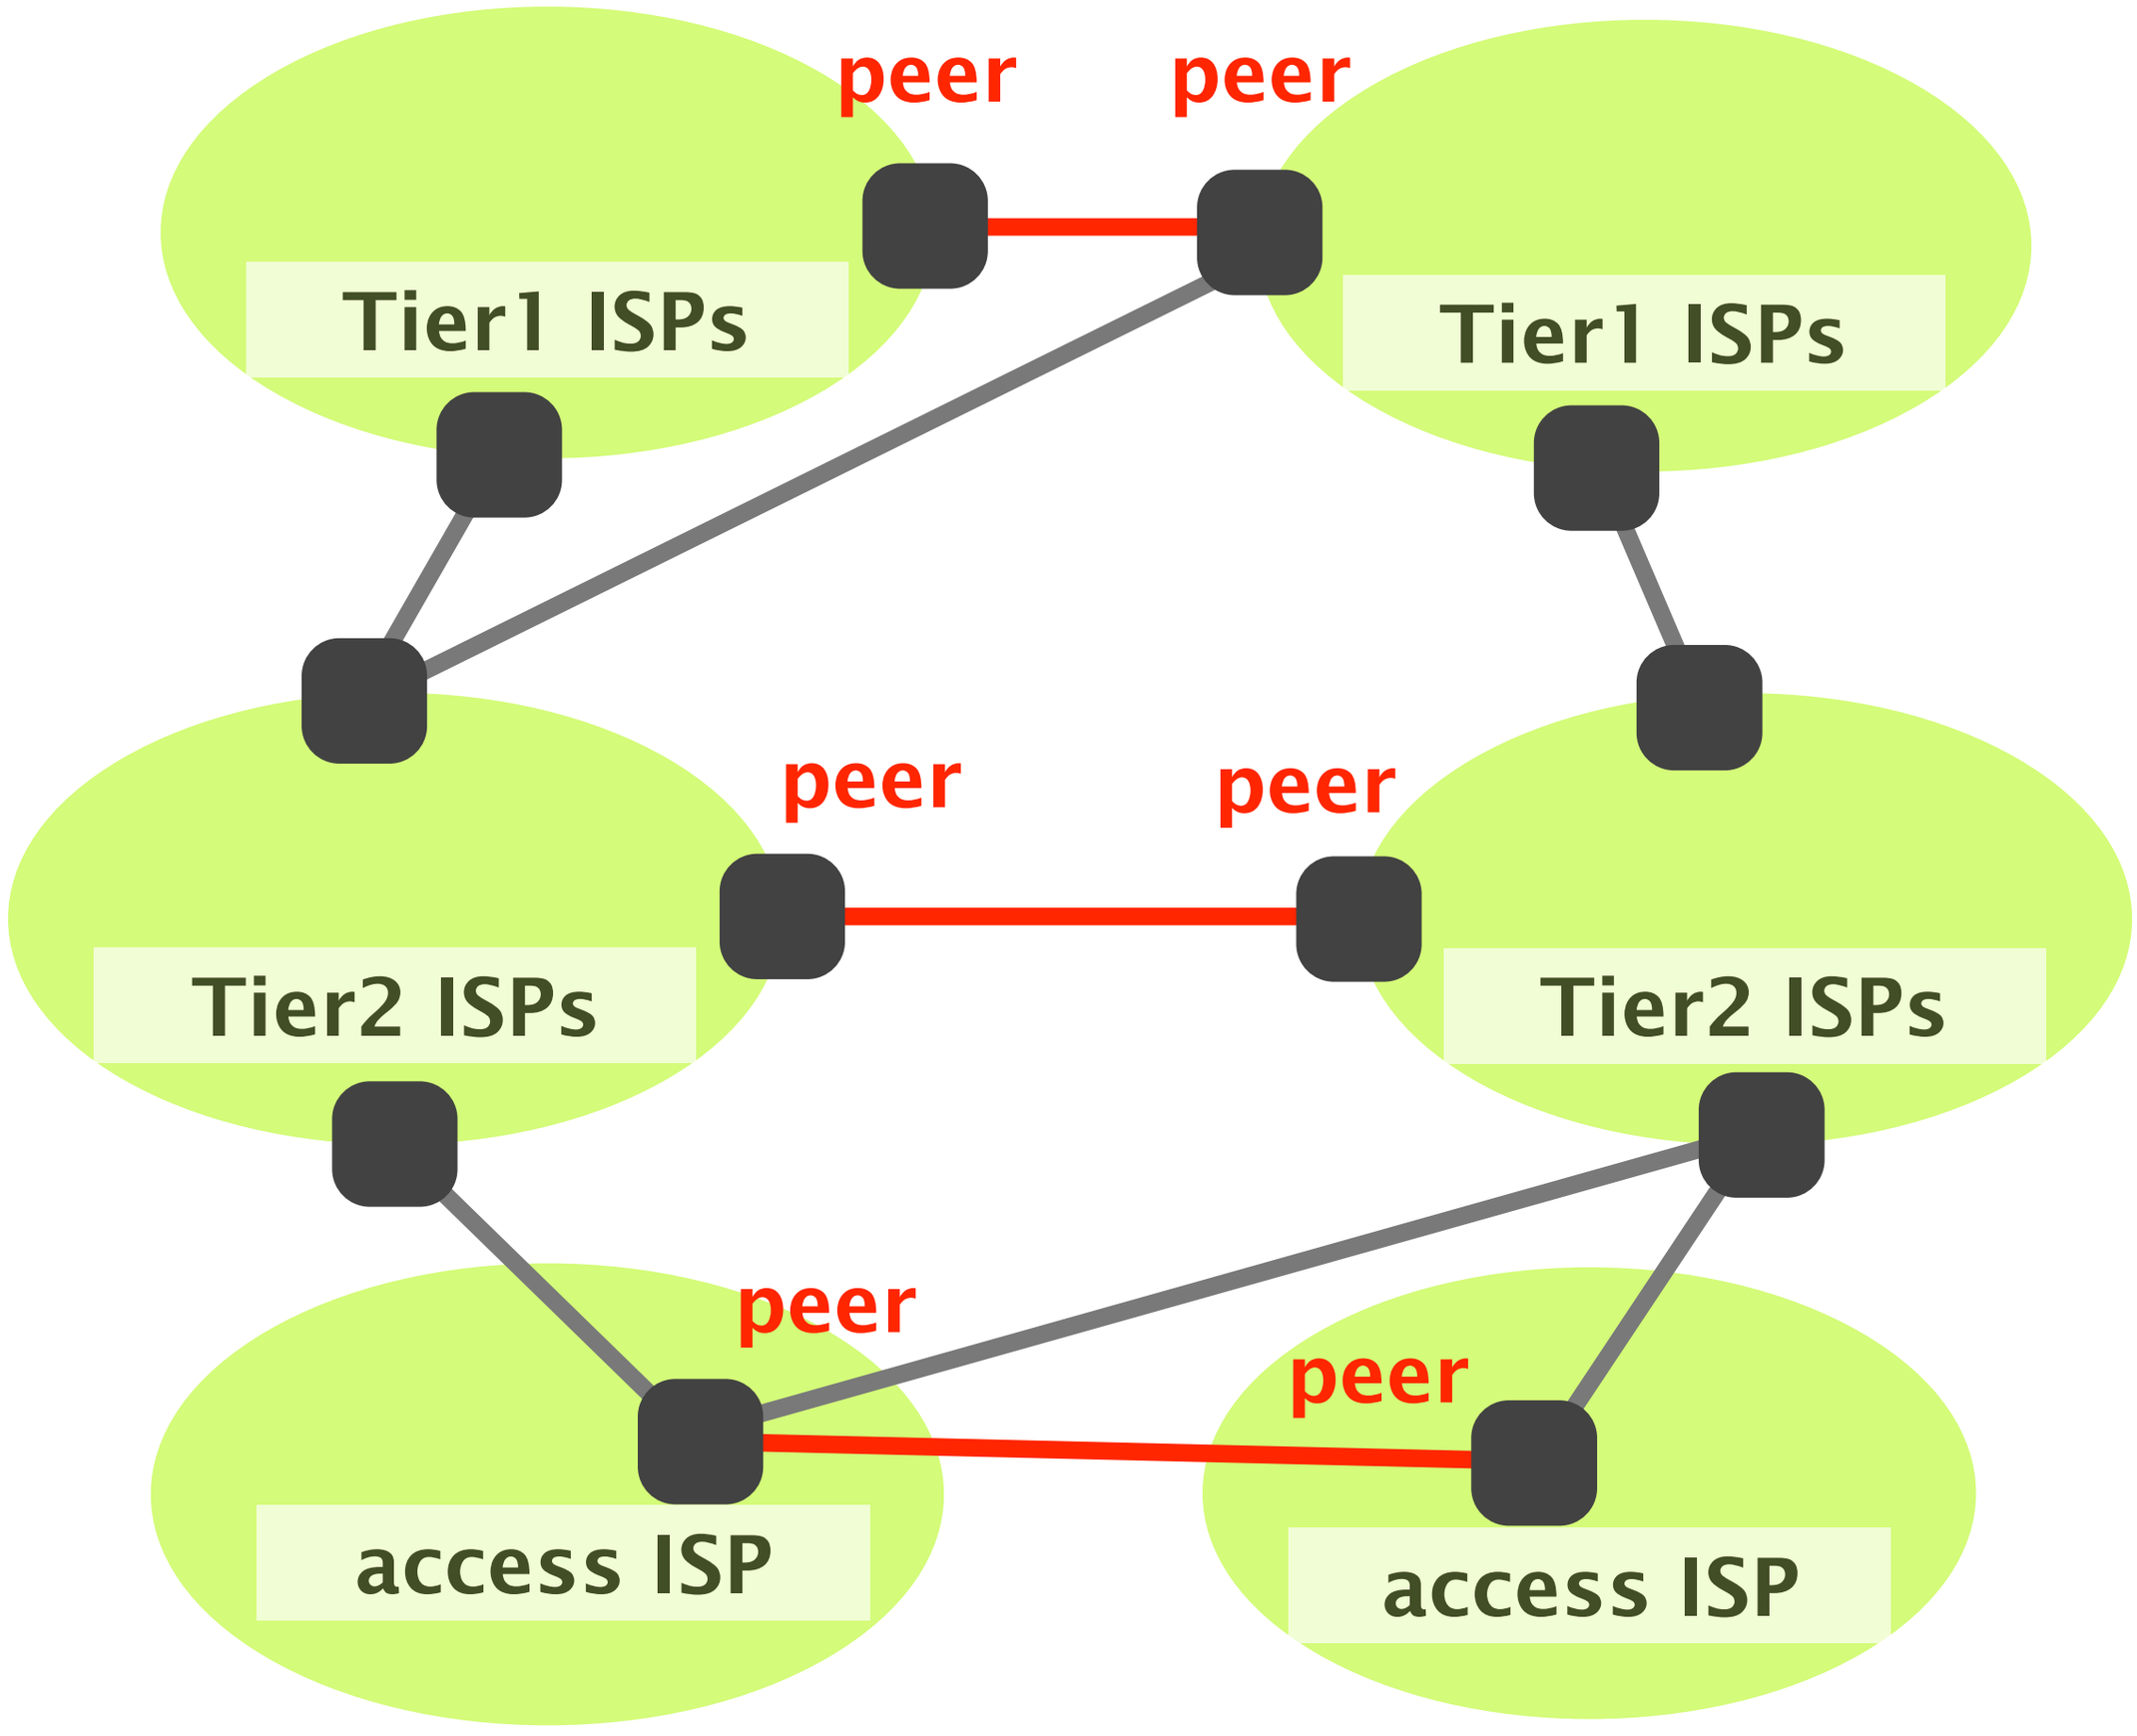
peer
peer
Tier1 ISPs
Tier1 ISPs
peer
peer
Tier2 ISPs
Tier2 ISPs
peer
peer
access ISP
access ISP

Interconnecting each network to its neighbors one-by-one is not cost effective

Physical costs

of provisioning or renting physical links

Bandwidth costs

a lot of links are not necessarily fully utilized

Human costs

to manage each connection individually

Internet eXchange Points (IXPs) solve these problems by letting *many* networks connect in one location

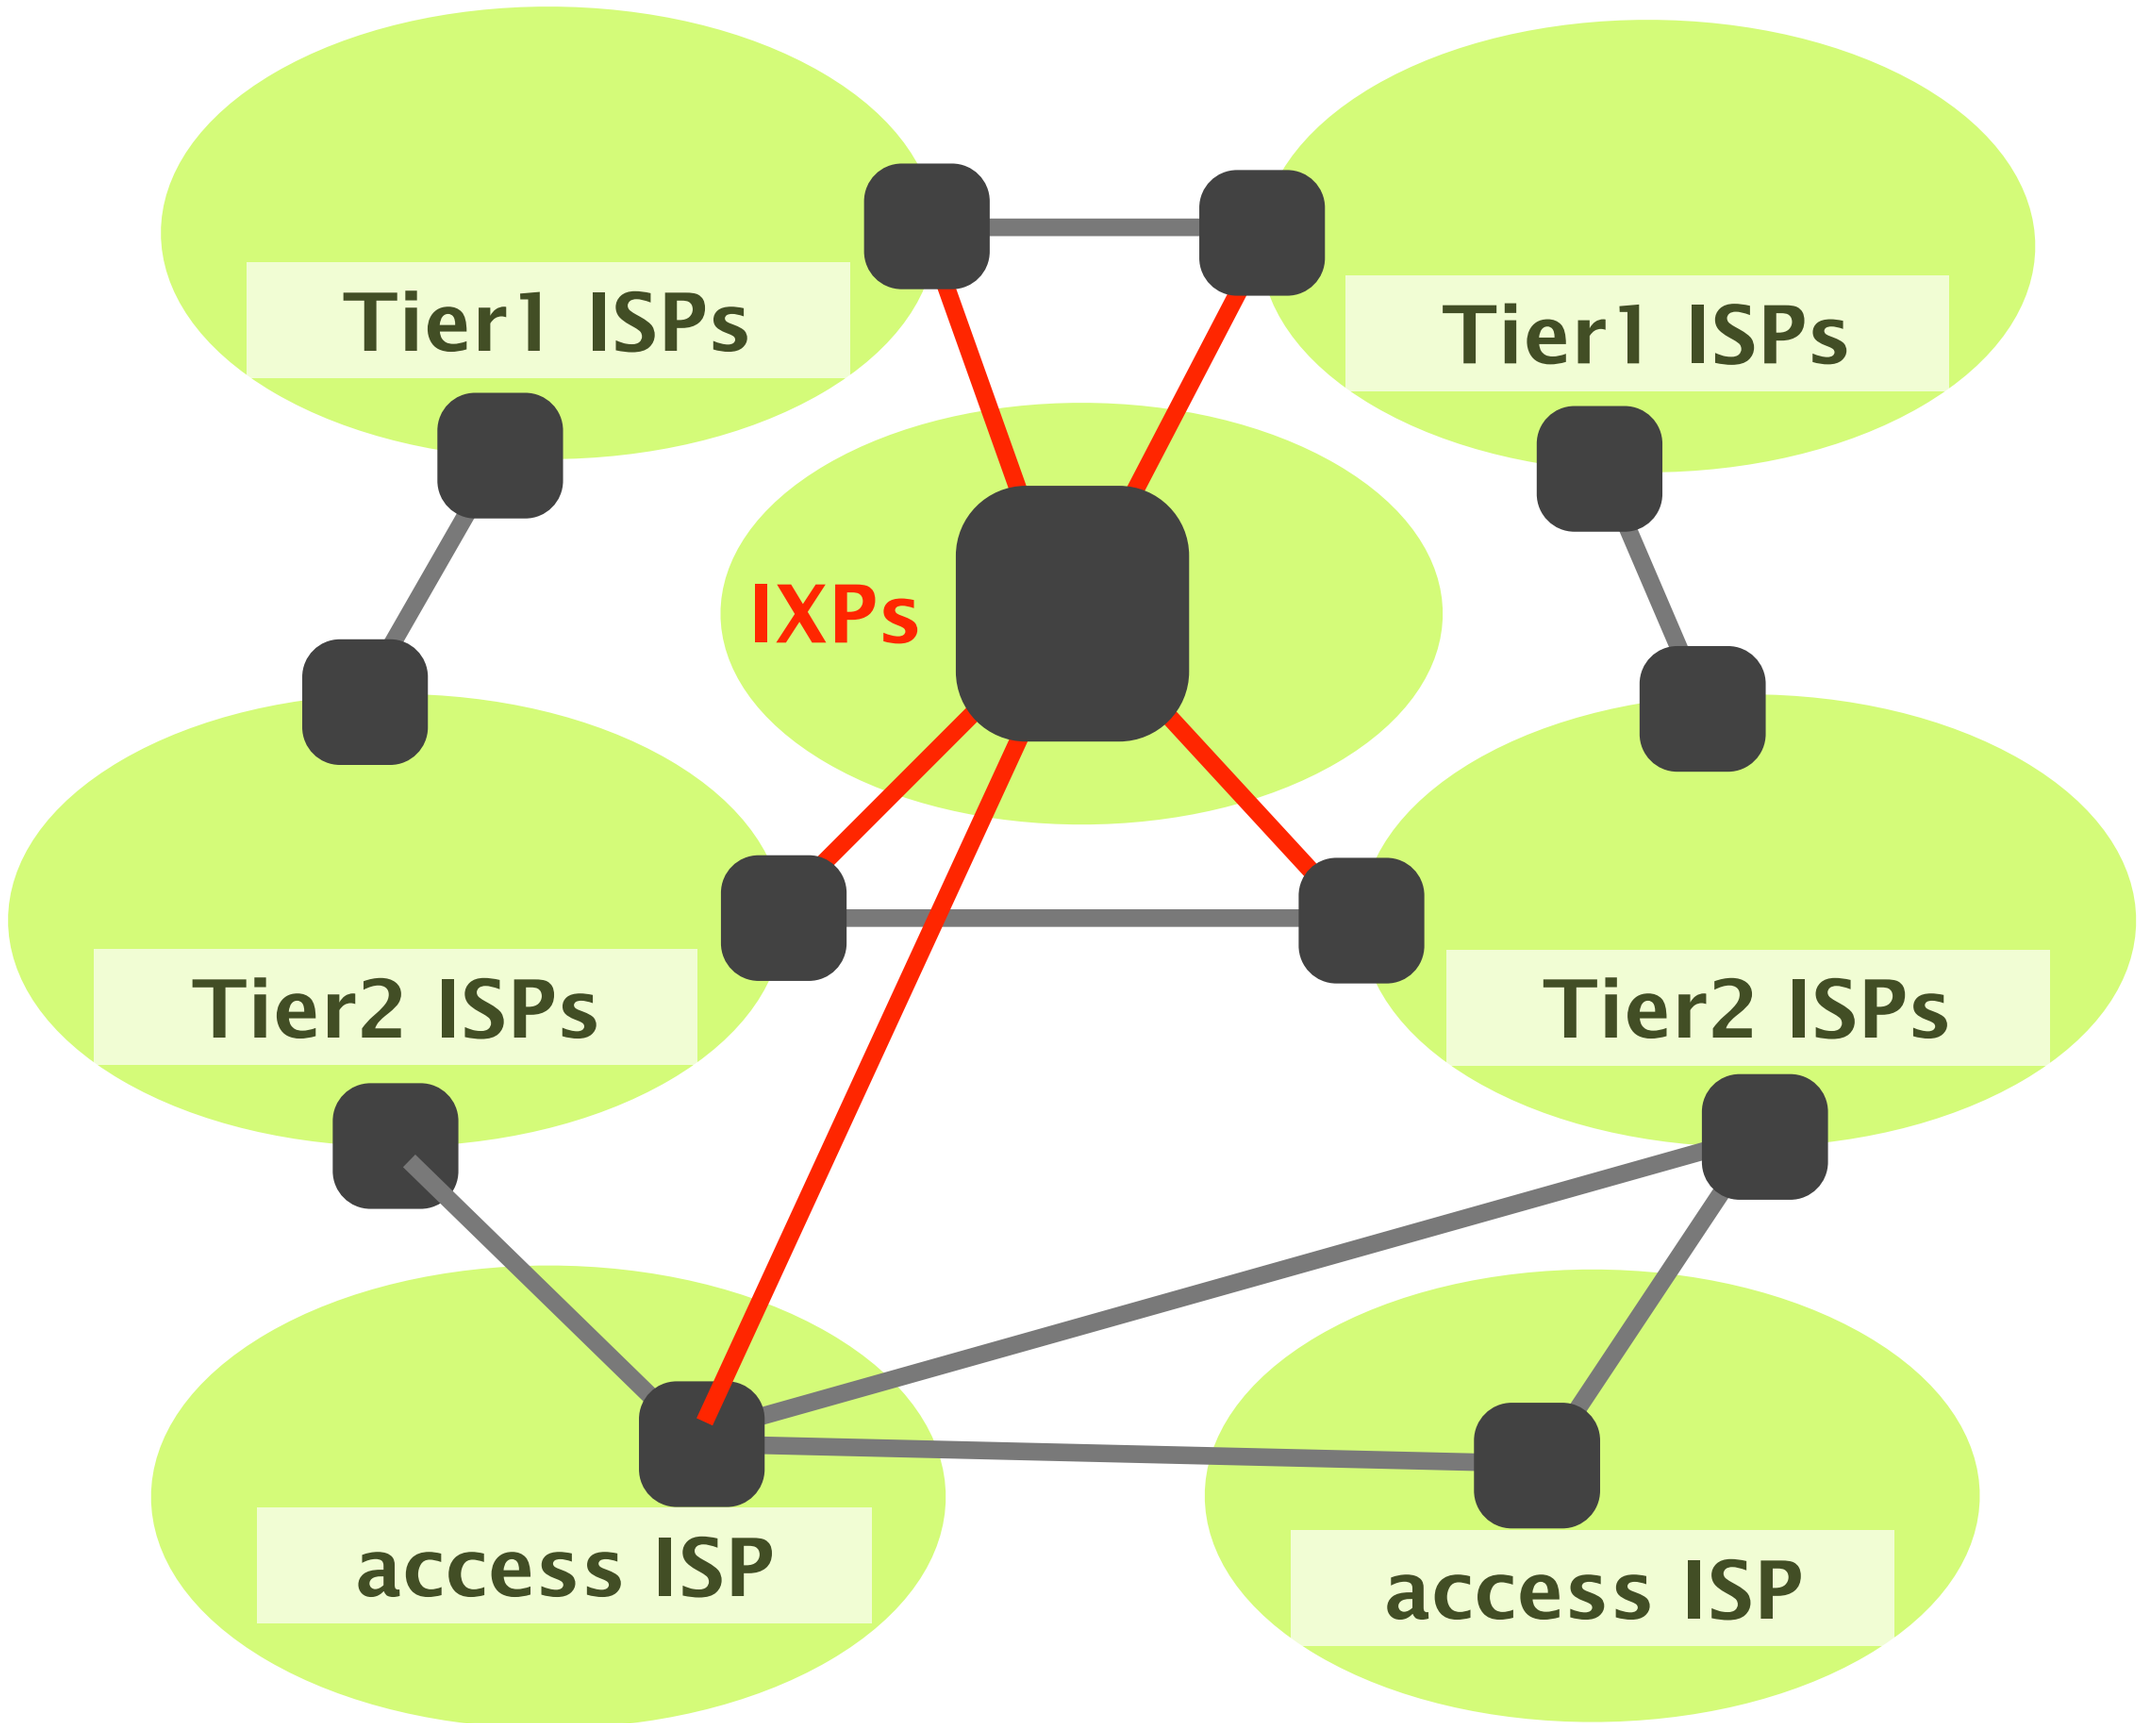
Tier1 ISPs
Tier1 ISPs
IXPs
Tier2 ISPs
Tier2 ISPs
access ISP
access ISP

Communication Networks

Part 1: Overview

#1 What is a network made of?

#2 How is it shared?

#3 How is it organized?

#4 How does communication happen?

#5 How do we characterize it?