情報科学入門

#13 データセンター・エネルギー問題

Yutaka Yasuda

Electricity usage of Data Center, US

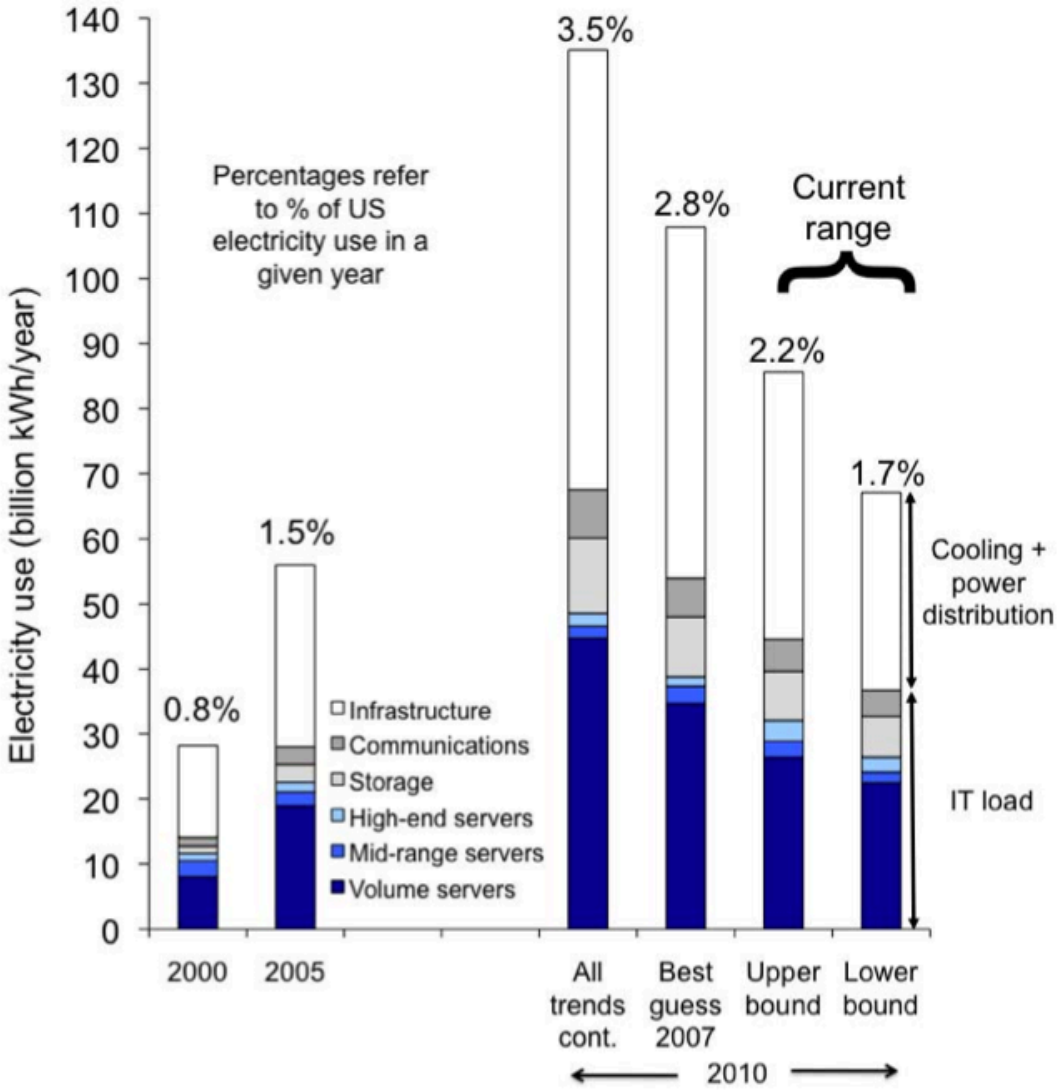

データセンターは 2011 年に米国全体の **2%** の電力を使っている

Growth in Data Center Electricity Use 2005 TO 2010
Jonathan G. Koomey, Ph.D. (2011)

ENIAC, the first computer (1945)

ENIAC vacuum tubes, at computer history museum, 2009

The Evolution of Computing

IBM 1401 (1959)

proof of the law, in real world

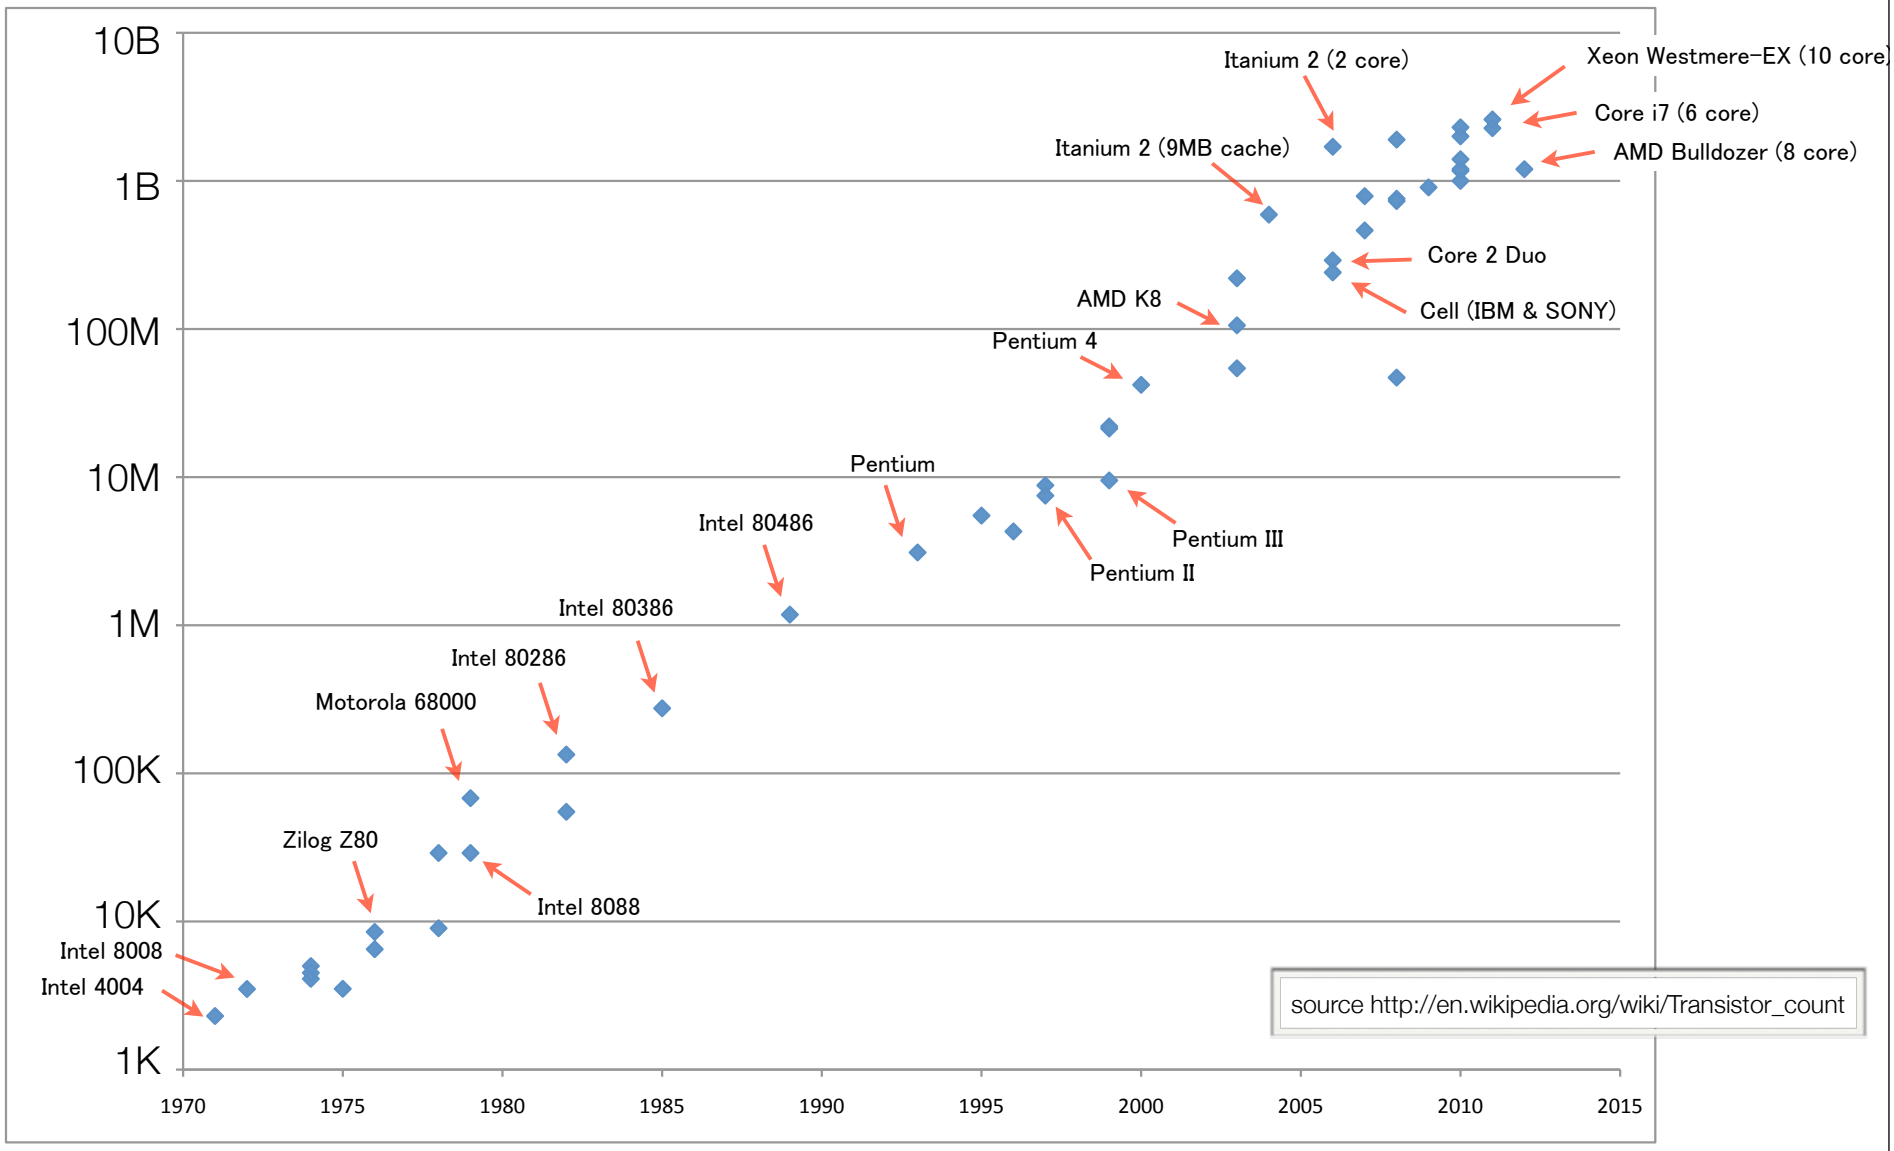

proof of the law, in real world

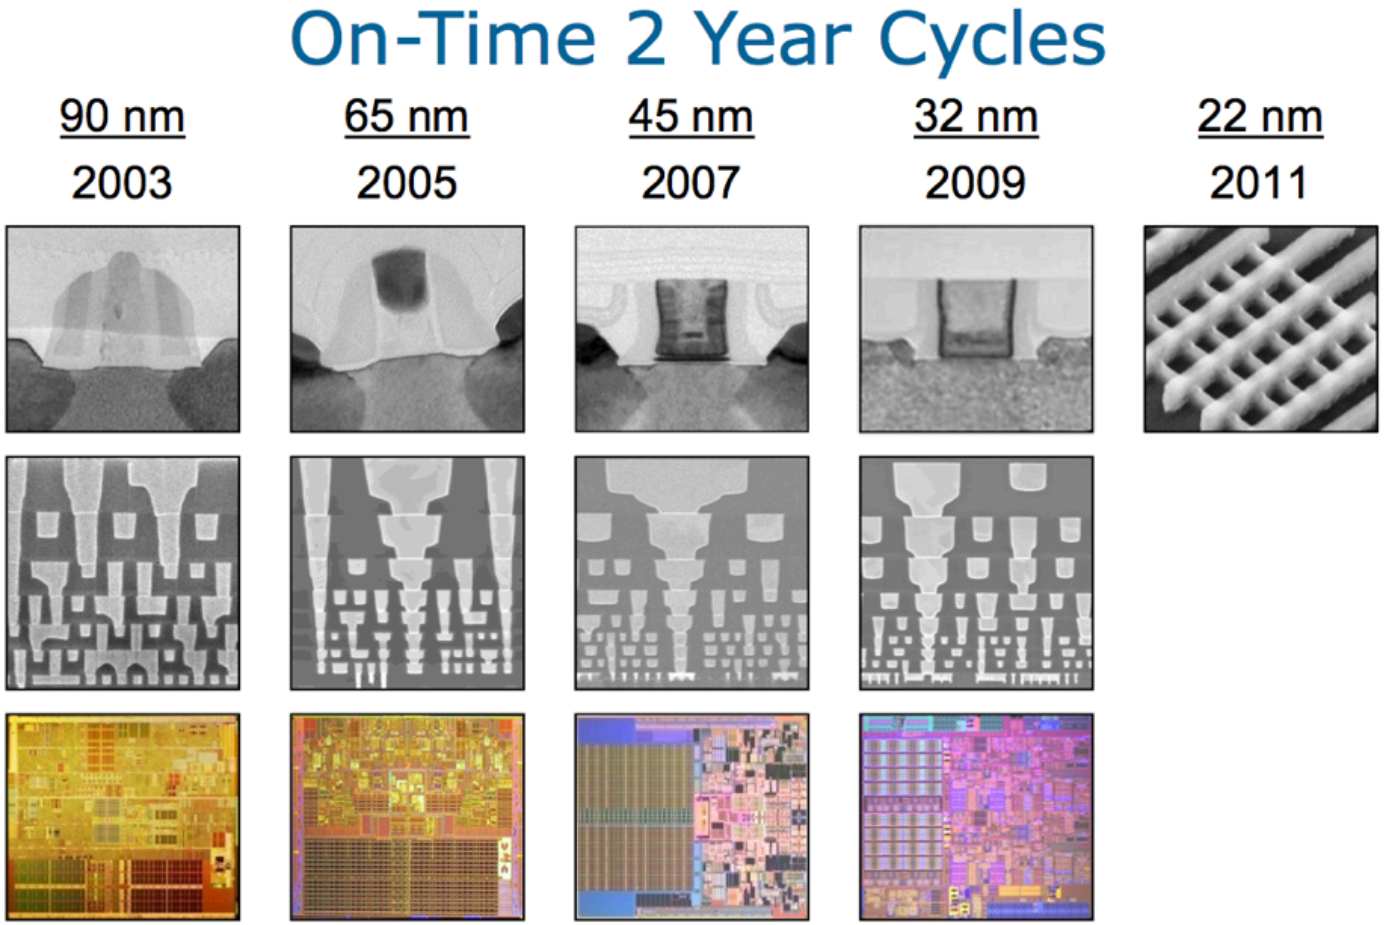

Intel continues to successfully introduce leading edge process + products on a 2 year cadence

Intel's Revolutionary 22 nm Transistor Technology, Mark Bohr, Kaizad Mistry, Intel (2011)

Fabs of Intel

Intel's Revolutionary 22 nm Transistor Technology, Mark Bohr, Kaizad Mistry, Intel (2011)

Solution : Mass Production

350 M processors per year

RESULT : cheapest, mass produced, for consumer product has the highest performance

Internet Service Model

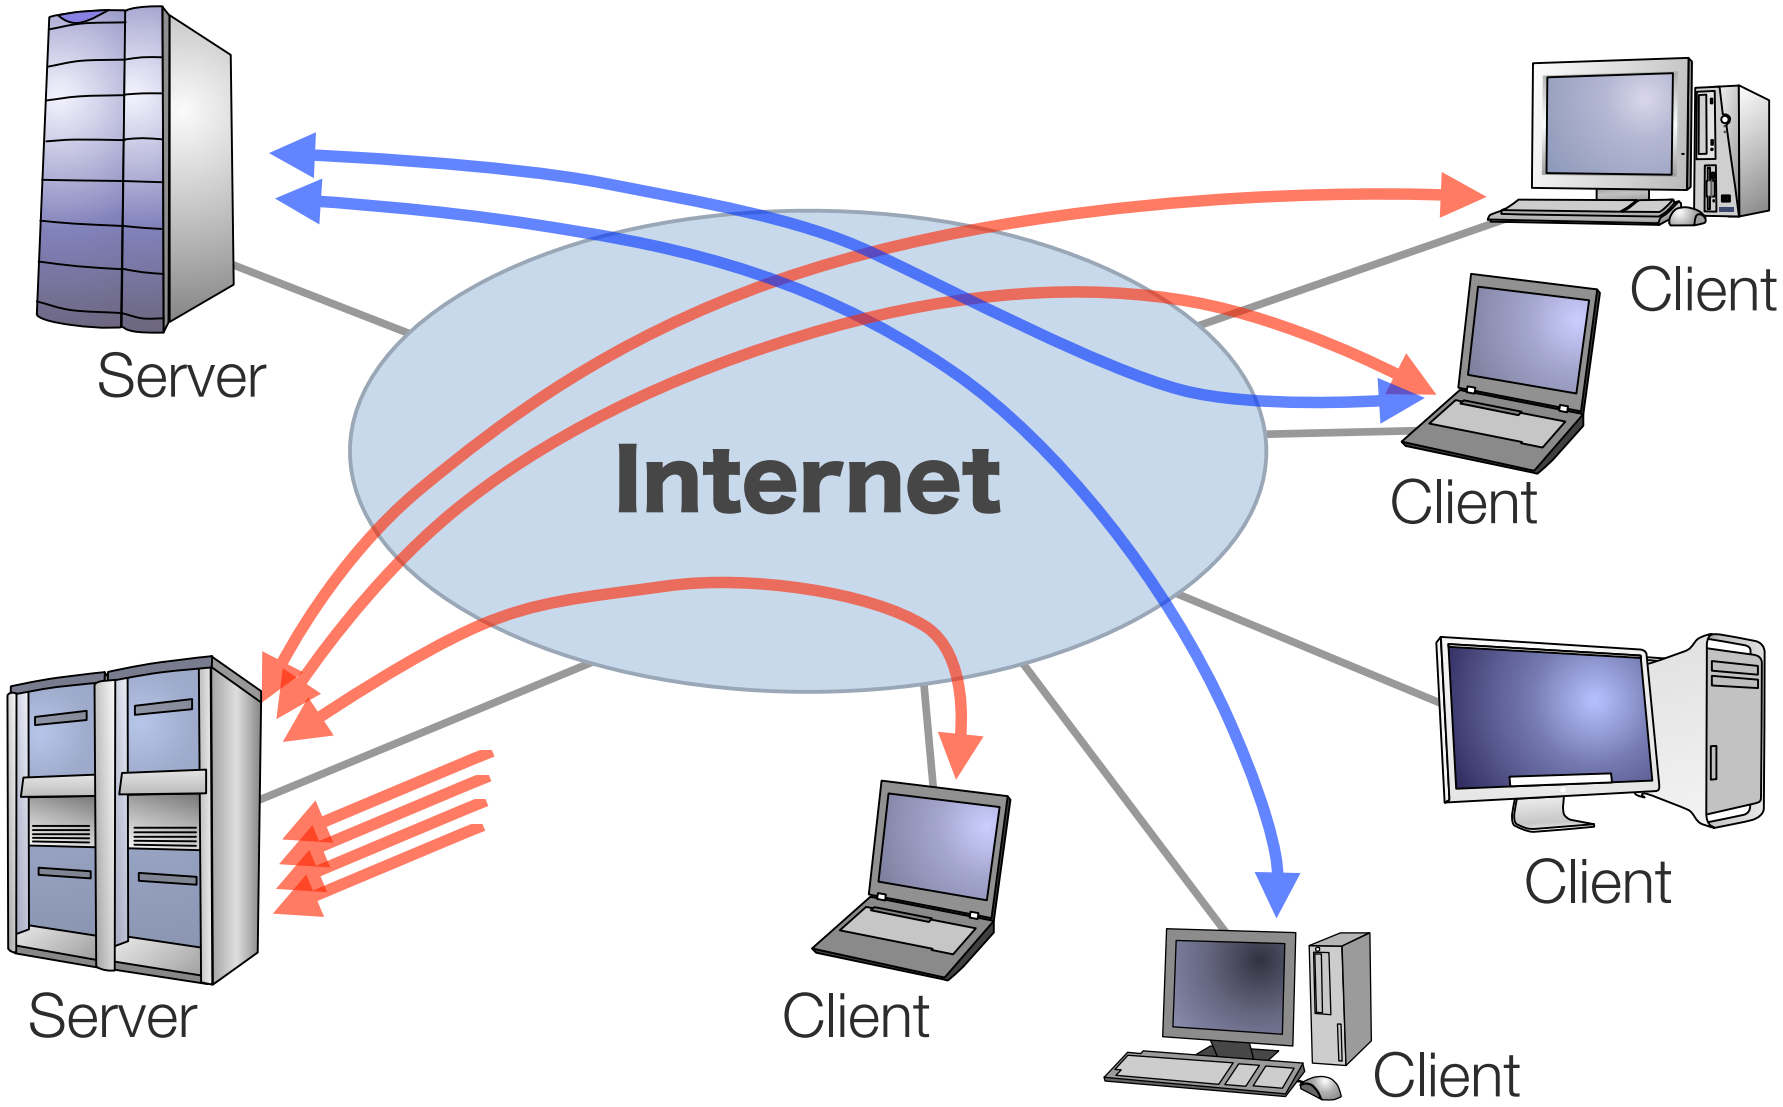

Server Farm : parallel and distributed processing at centralized facility

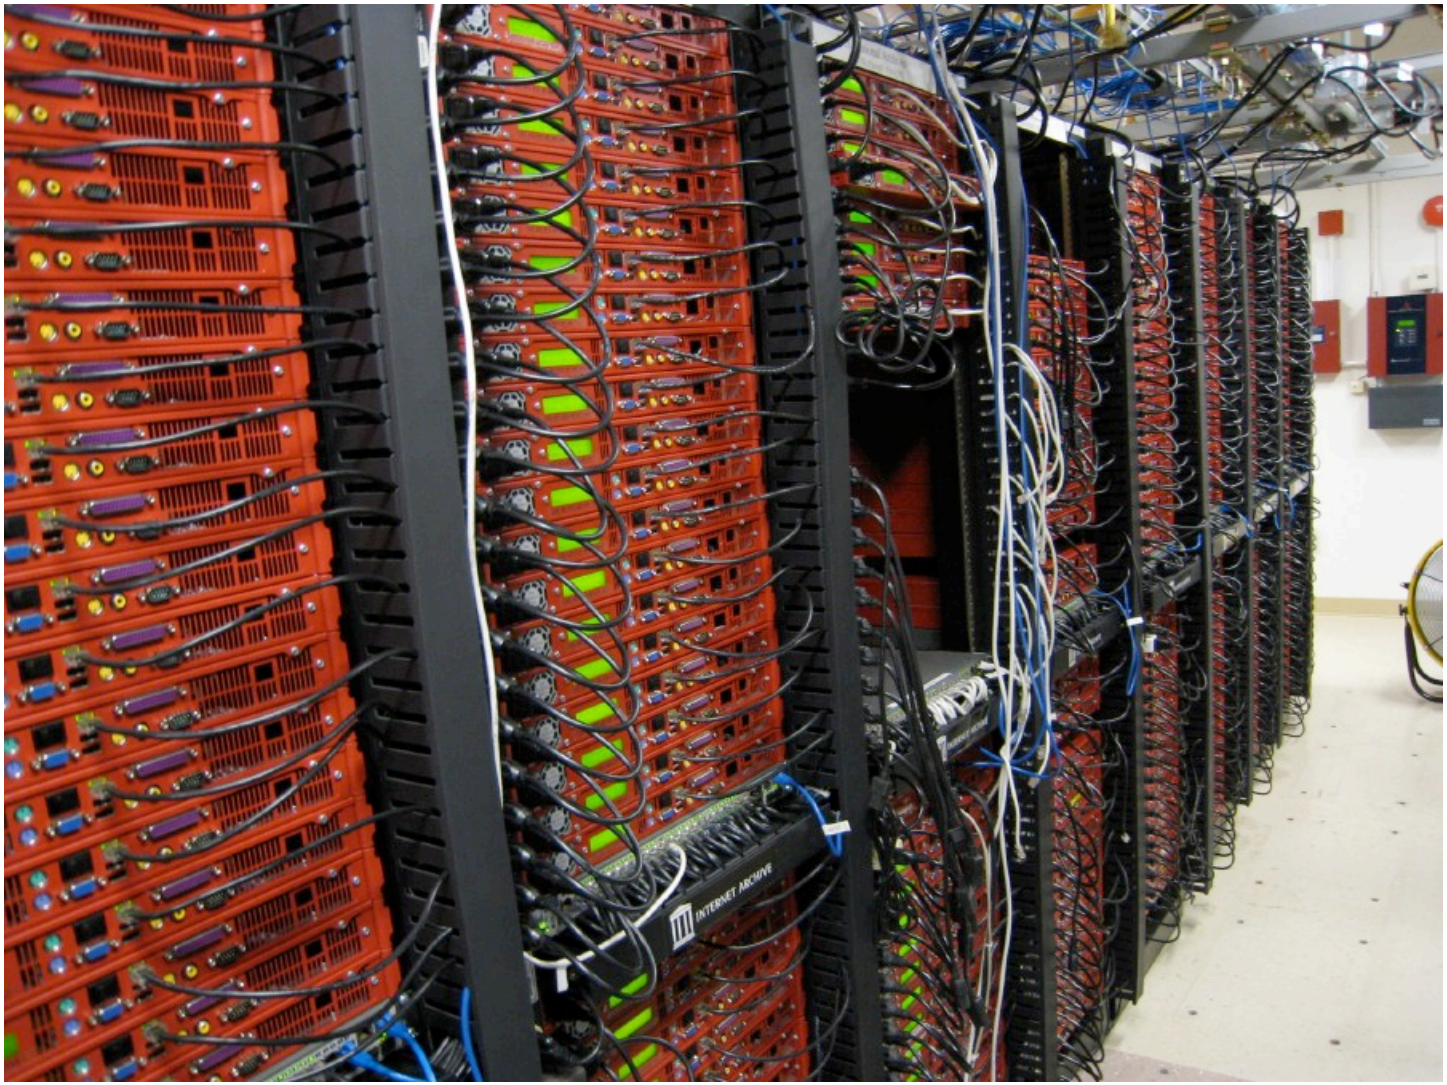

Internet Archive, 2008, San Francisco, US

technology review
Published by MIT

Inside Facebook's Not-So-Secret New Data Center

http://www.technologyreview.com/computing/37295/page8/#photo

technology review
Published by MIT

Inside Facebook's Not-So-Secret New Data Center

http://www.technologyreview.com/computing/37295/page8/#photo

technology review
Published by MIT

Inside Facebook's Not-So-Secret New Data Center

http://www.technologyreview.com/computing/37295/page8/#photo

technology review
Published by MIT

Inside Facebook's Not-So-Secret New Data Center

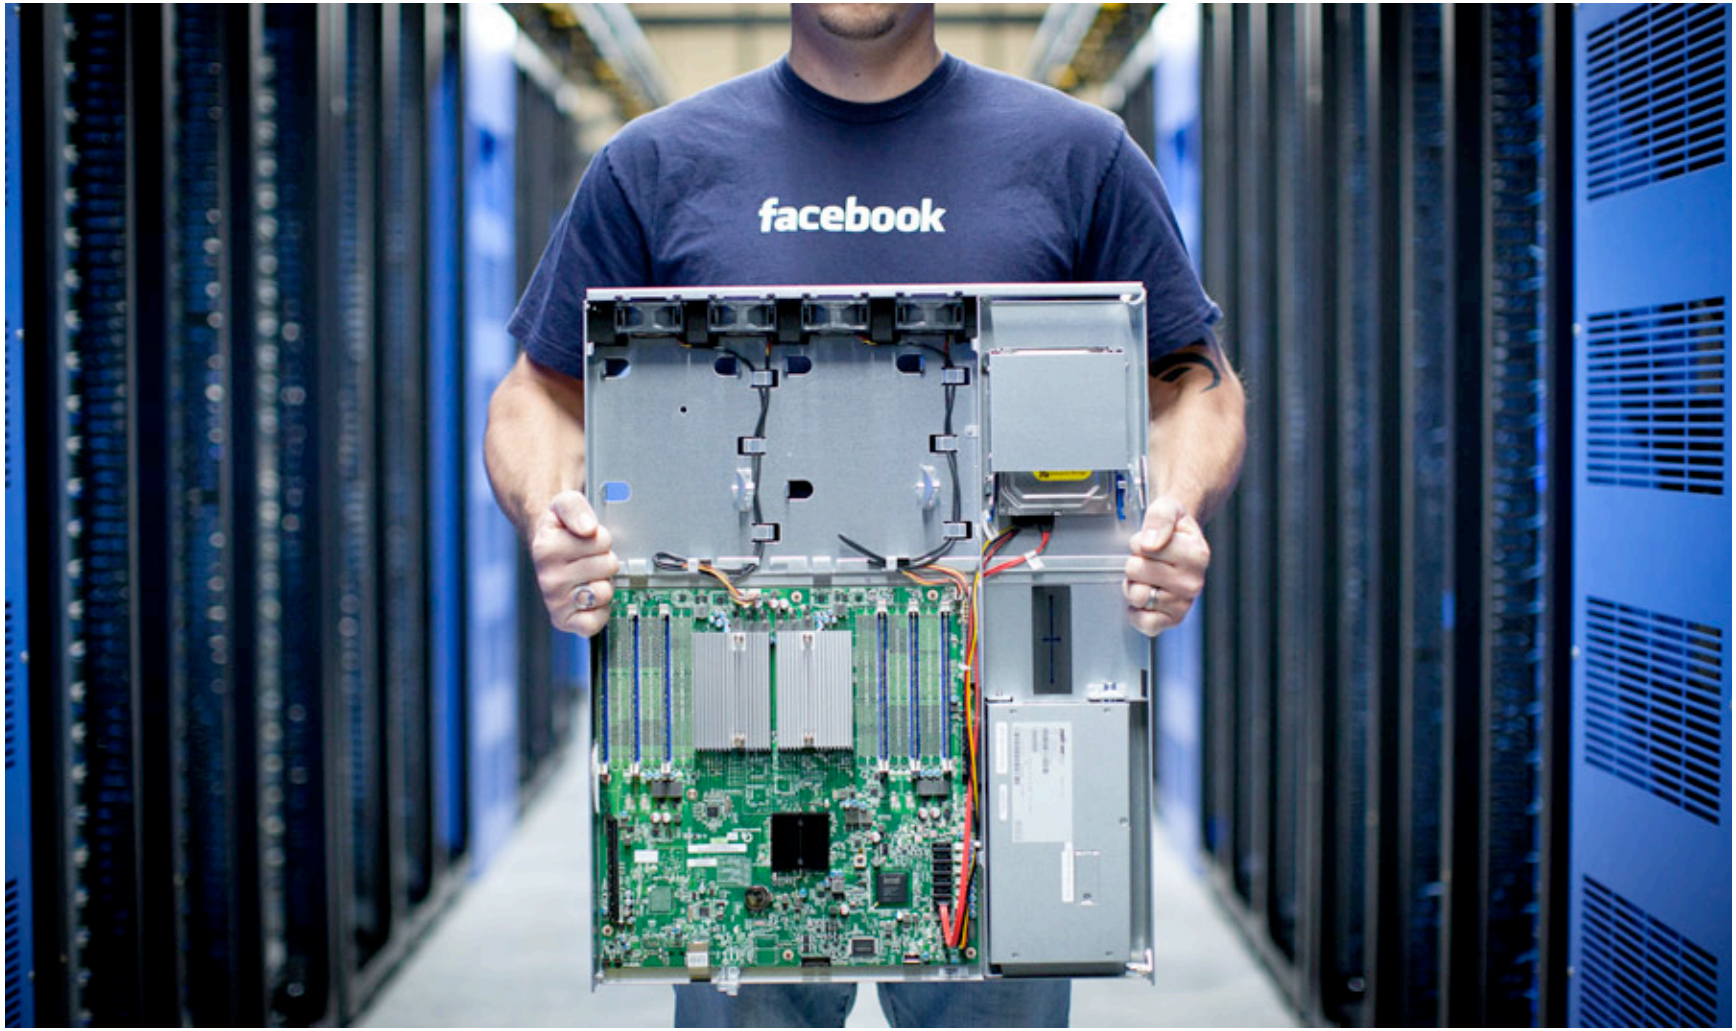

http://www.technologyreview.com/computing/37295/page8/#photo

Giant Services on Cloud Computing

Cloud Computing

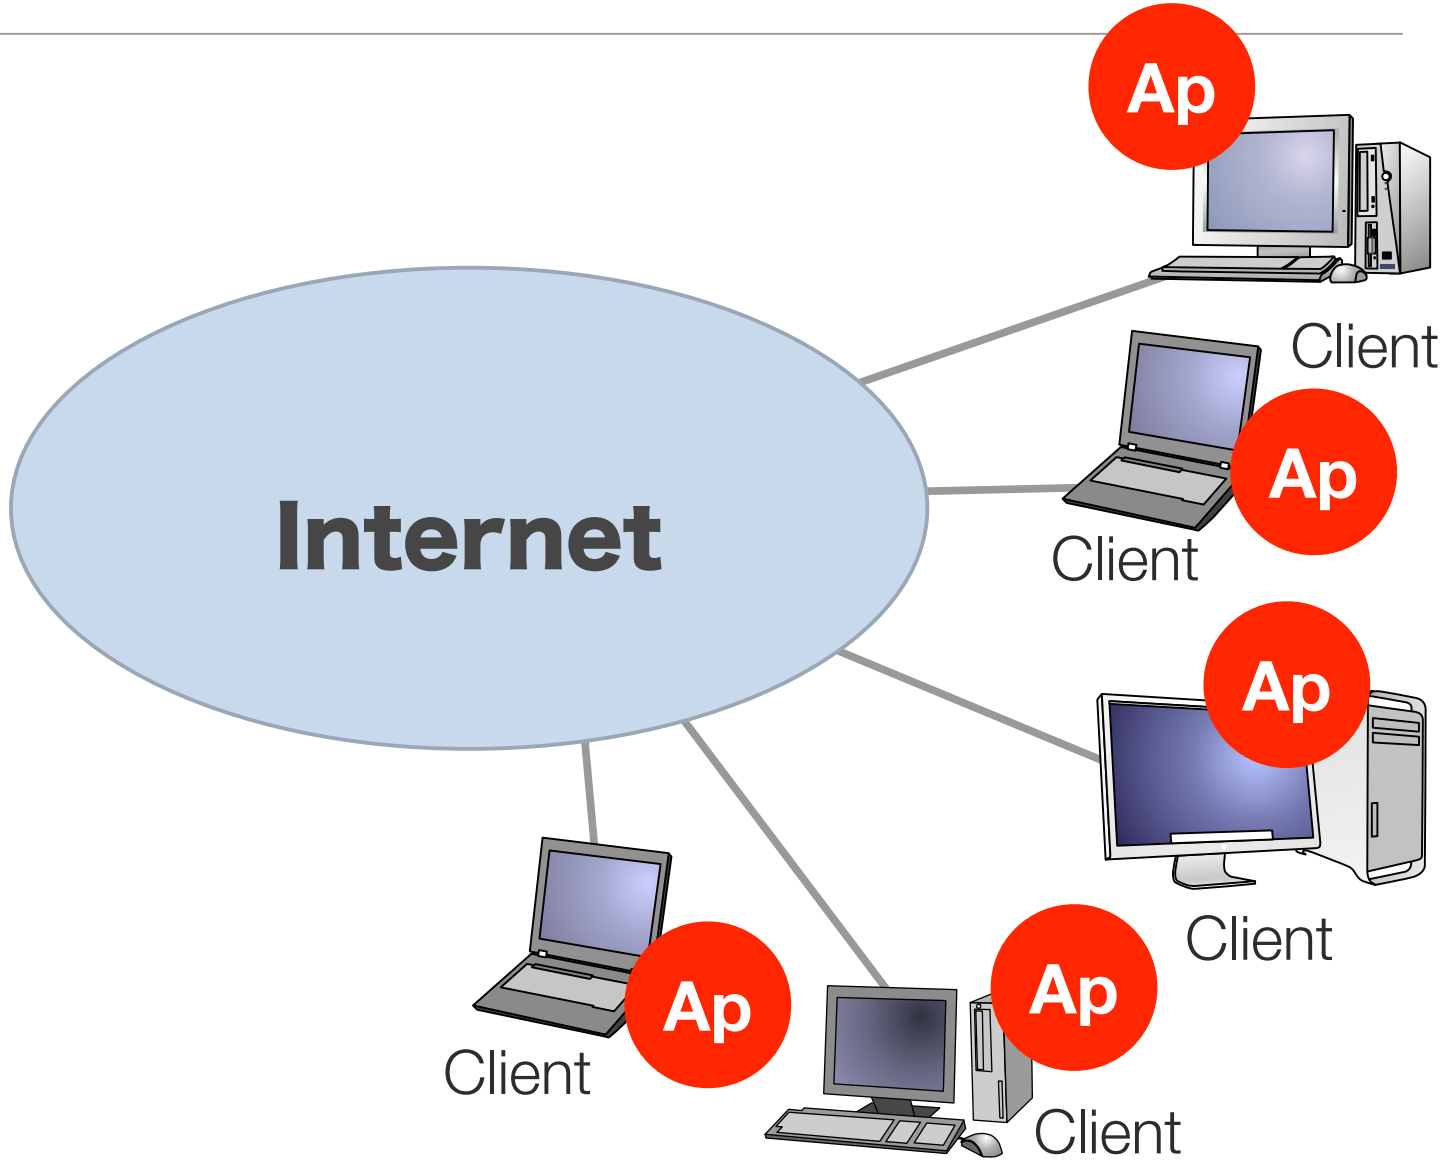

Cloud Computing

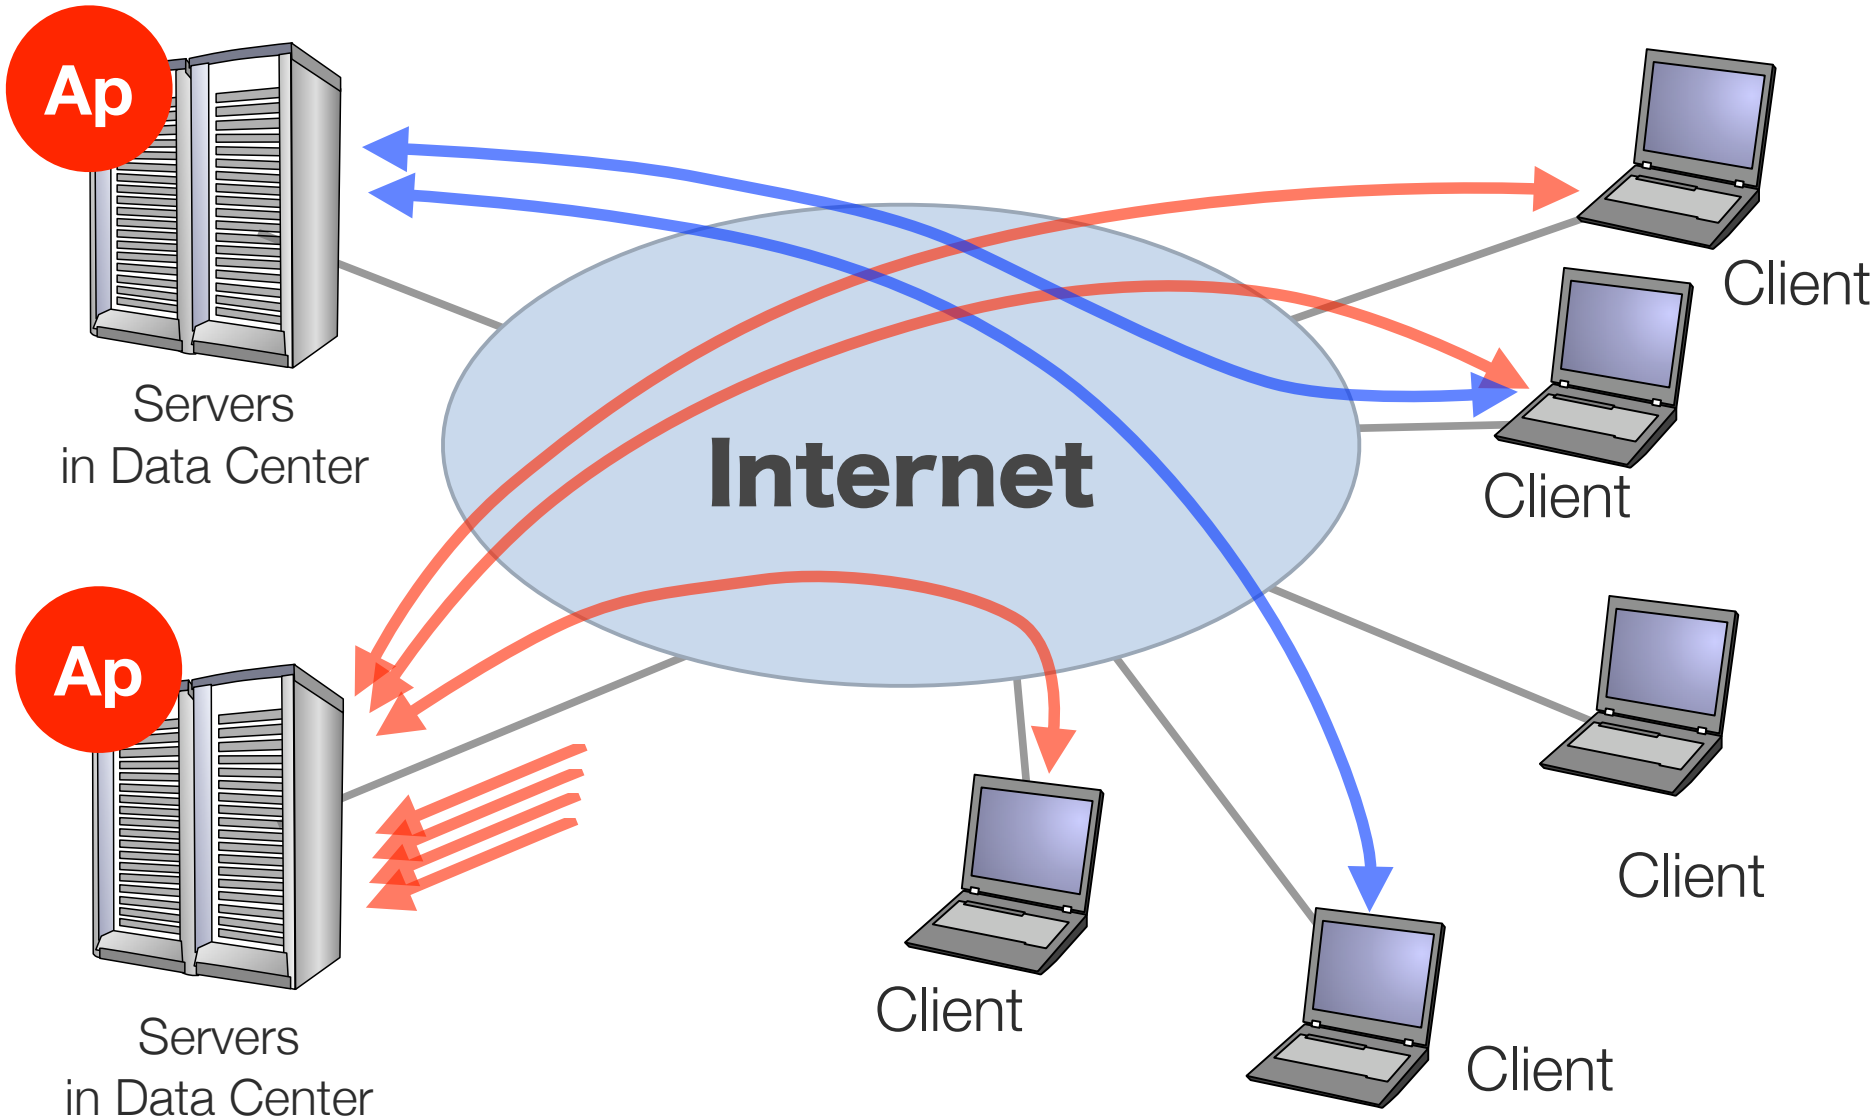

Cloud Computing

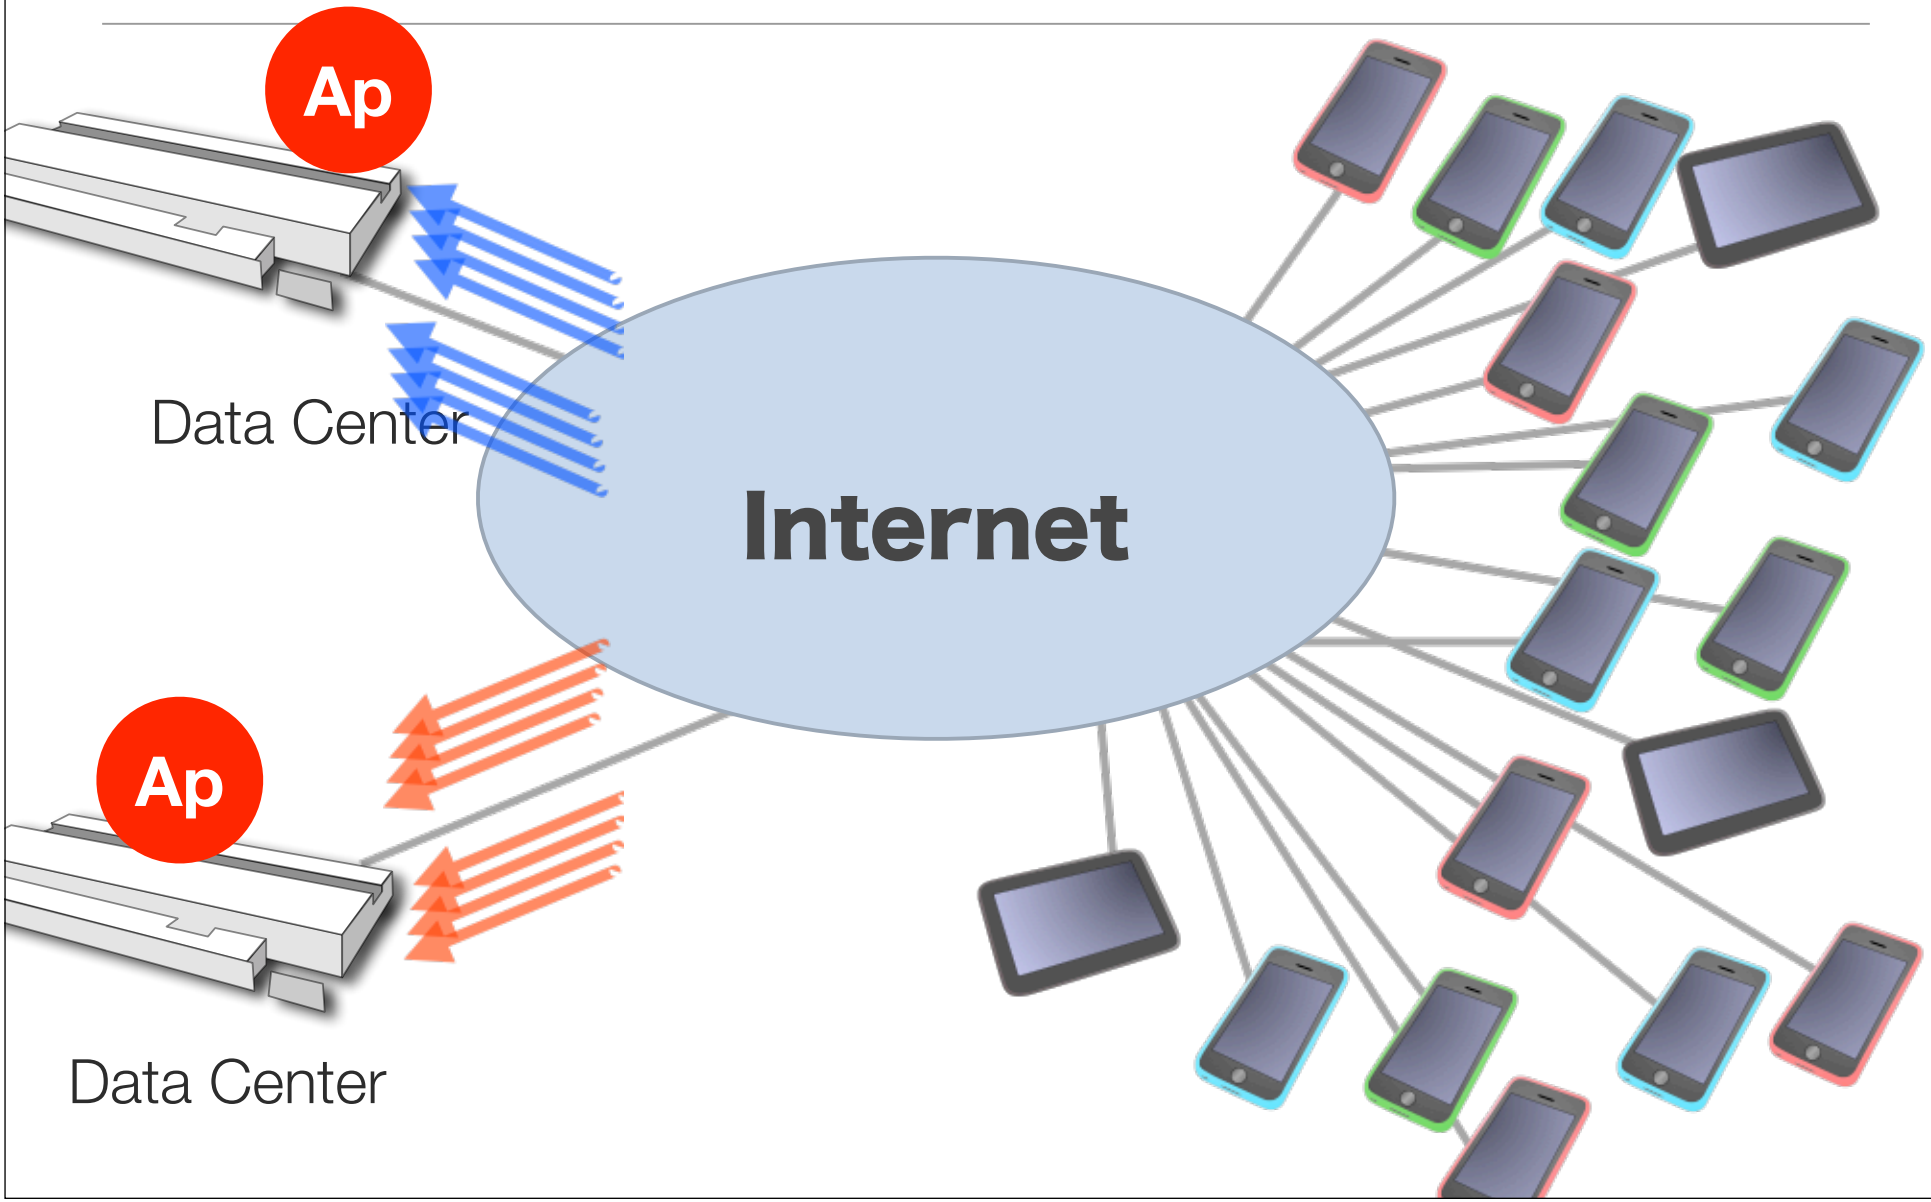

The light is not fast enough

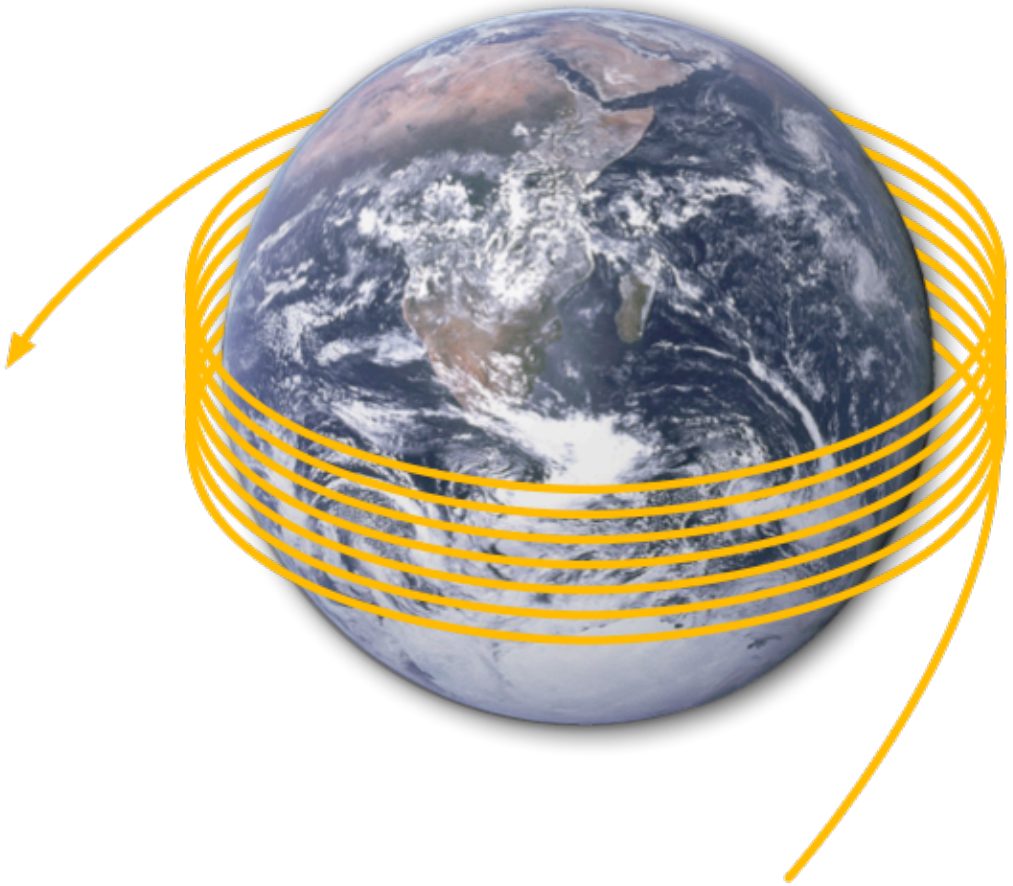

earth image: NASA, http://commons.wikimedia.org/wiki/File:The_Earth_seen_from_Apollo_17_with_white_background.jpg

Google Data Centers

Americas

- Berkeley County, South Carolina
- Council Bluffs, Iowa
- Douglas County, Georgia
- Mayes County, Oklahoma
- Lenoir, North Carolina
- The Dalles, Oregon
- Quilicura, Chile

EU

- Hamina, Finland
- St Ghislain, Belgium
- Dublin, Ireland

Asia

- Singapore
- Changhua County, Taiwan

Google Data Centers http://www.google.com/about/datacenters/locations/index.html

EPA's prediction in 2007

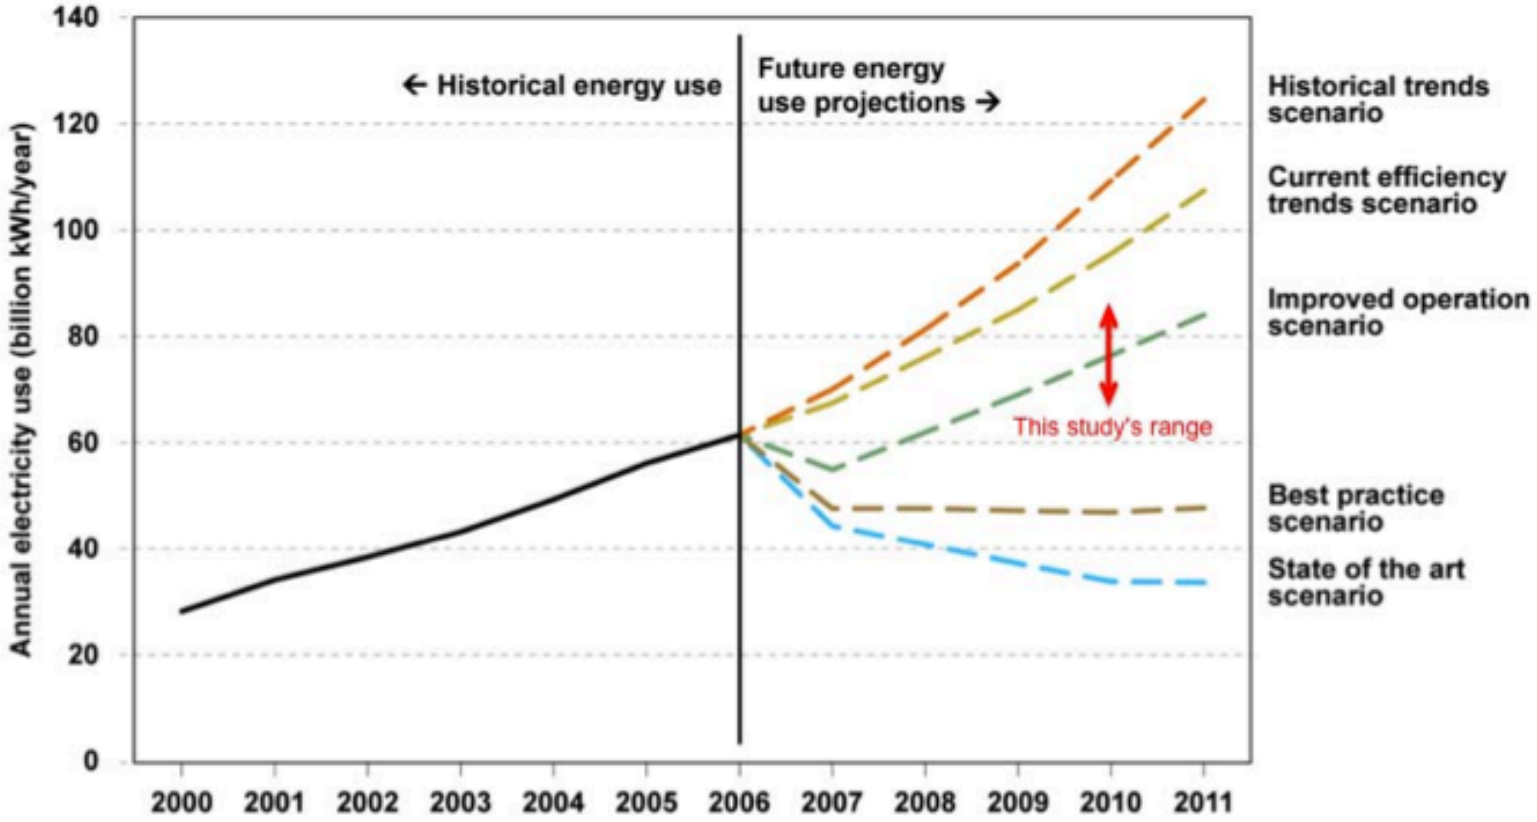

Growth in Data Center Electricity Use 2005 TO 2010
Jonathan G. Koomey, Ph.D. (2011)

Top Ten Traffic Providers 2007 - 2010

2007

Rank	Name	%
1	ISP A	5.77
2	ISP B	4.55
3	ISP C	3.35
4	ISP D	3.20
5	ISP E	2.60
6	ISP F	2.77
7	ISP G	2.24
8	ISP H	1.82
9	ISP I	1.35
10	ISP J	1.35

2009

Rank	Name	%
1	ISP A	9.41
2	ISP B	5.70
3	Google	5.20
4	ISP F	5.00
5	ISP H	3.22
6	Comcast	3.12
7	ISP D	3.08
8	ISP E	2.32
9	ISP C	2.05
10	ISP G	1.89

2010

Rank	Name	%
1	ISP A	9.09
2	Google	7.00
3	ISP B	4.70
4	ISP F	3.00
5	ISP H	2.96
6	ISP K	2.89
7	ISP L	2.82
8	ISP M	2.60
9	ISP E	2.30
10	Comcast	2.07

total 30% **total 41%** **total 40%**

Top ten account contribution grows from 30-40%

Need to Be Cost Conscious

How much is your search?

Urs Hoelzle, a Google senior vice president, said context mattered in energy statistics.

0.3Wh for average search query

"Google Details, and Defends, Its Use of Electricity", The New York Times, 2011/9/9
http://www.nytimes.com/2011/09/09/technology/google-details-and-defends-its-use-of-electricity.html

Designing devices by themselves

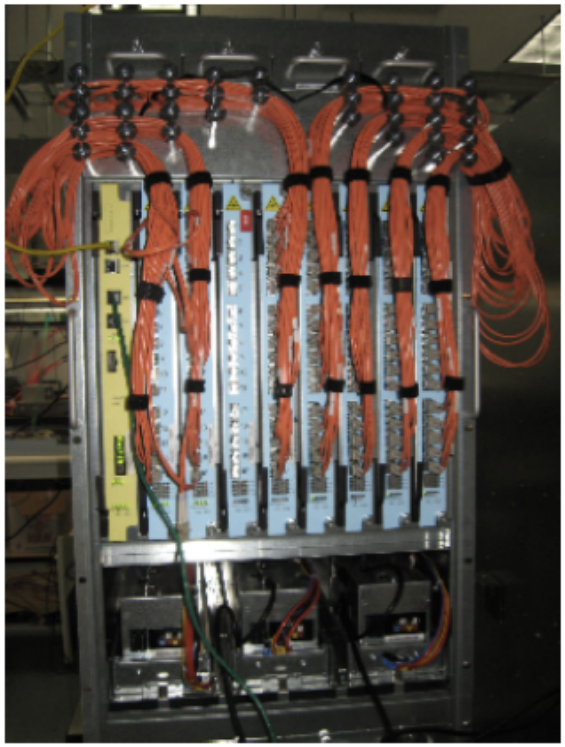

Google's network switch
G-Scale Network Hardware
10G Ethernet x 128 ports switch
OpenFlow support

"OpenFlow @ Google" by Urs Hoelzle, SVP Google, 2012
http://opennetsummit.org/talks/ONS2012/hoelzle-tue-openflow.pdf

Facebook's storage system
Knox storage system
15 harddrives in one tray

Inside Facebook's lab: A mission to make hardware open source, ZDnet, 2012 Sep 24.
http://www.zdnet.com/inside-facebooks-lab-a-mission-to-make-hardware-open-source-7000004557/

Designing their own servers too

http://www.google.com/about/datacenters/gallery/

Designing their own servers too

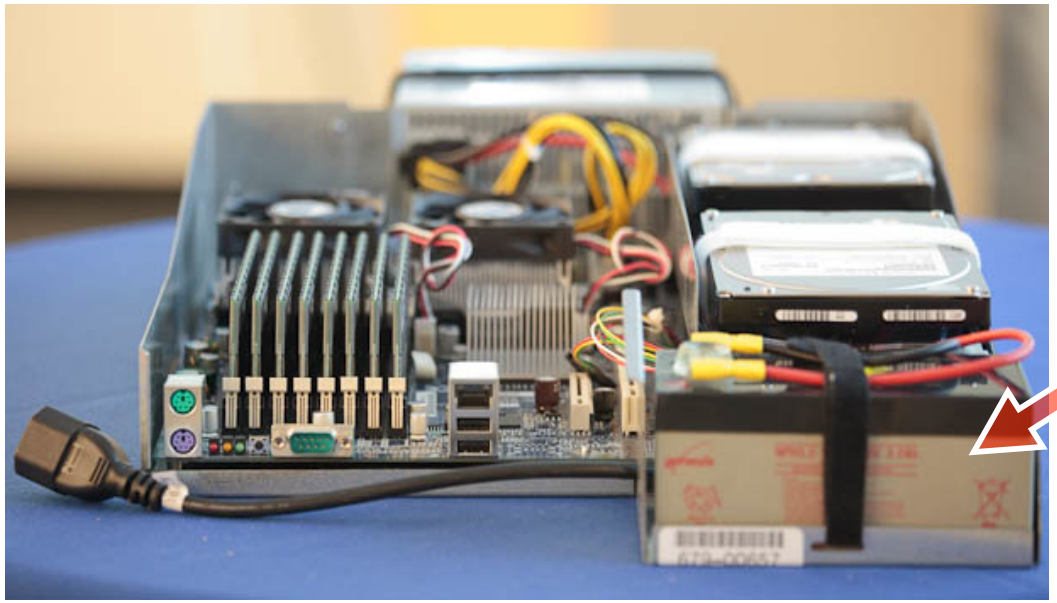

battery!

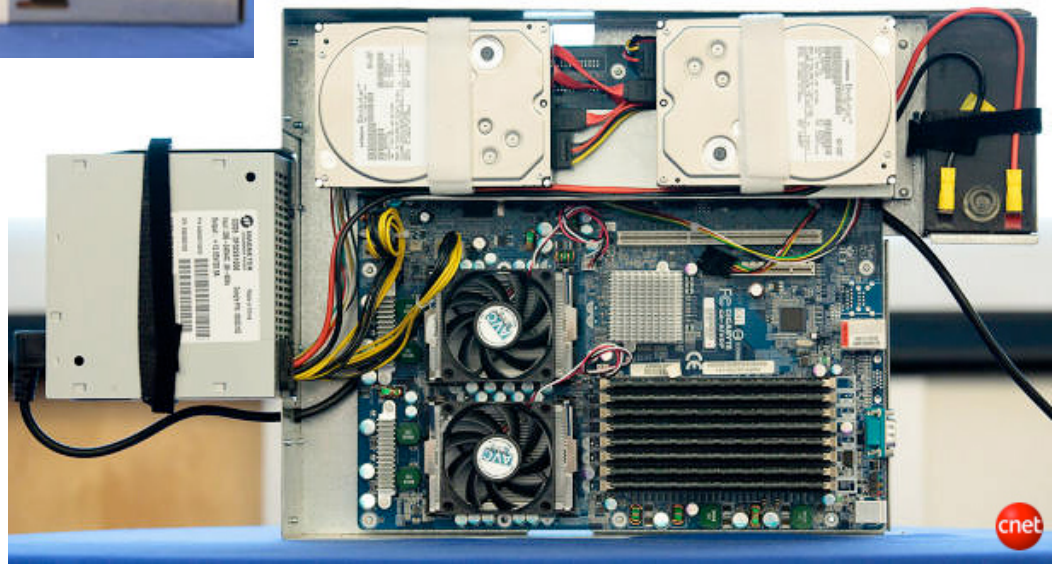

Google uncloaks once-secret server, Cnet, 2009
http://news.cnet.com/
8301-1001_3-10209580-92.html

PUE, Power Usage Effectiveness

$$\mathbf{PUE} = \frac{\textbf{Total facility power}}{\textbf{IT equipment power}}$$

Facility energy efficiency

A Google data center in Hamina, Finland, is housed in a former paper mill.

http://www.nytimes.com/2011/09/09/technology/google-details-and-defends-its-use-of-electricity.html

Facility energy efficiency

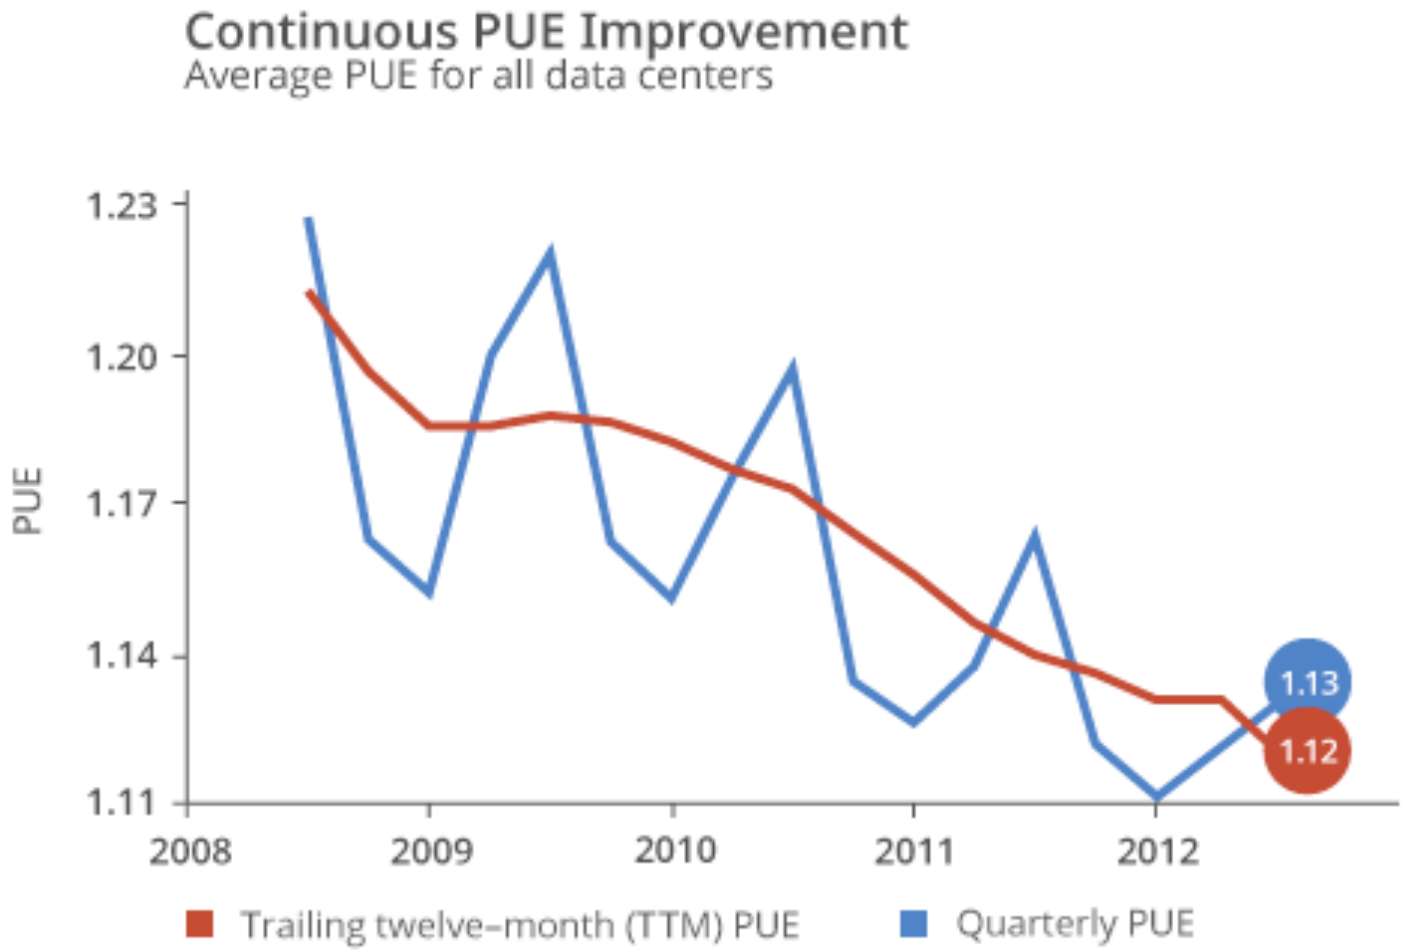

Efficiency: How we do it, Google
http://www.google.com/about/datacenters/efficiency/internal/

Set the thermostat

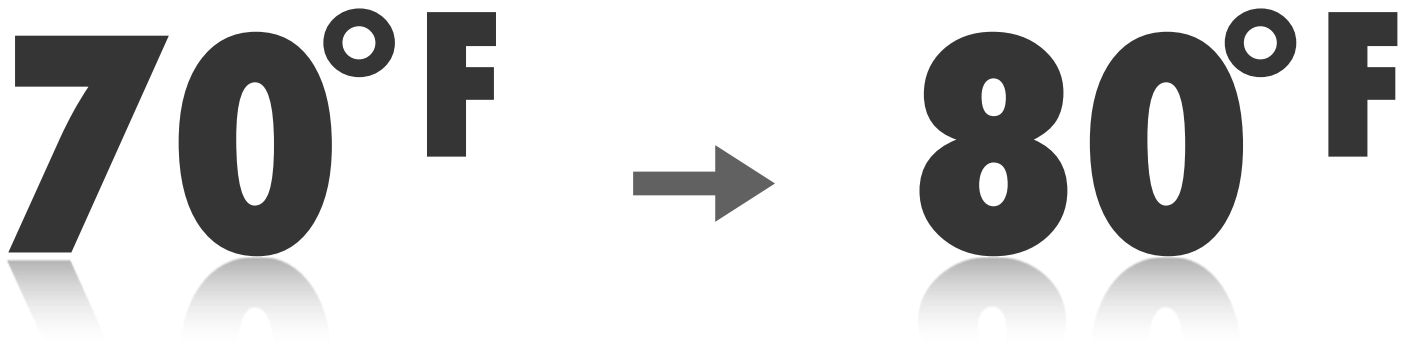

"70°F is a myth"

"Raise it to 80°F (or higher)

(70°F = 21°C, 80°F = 27°C)

technology review
Published by MIT

Inside Facebook's Not-So-Secret New Data Center

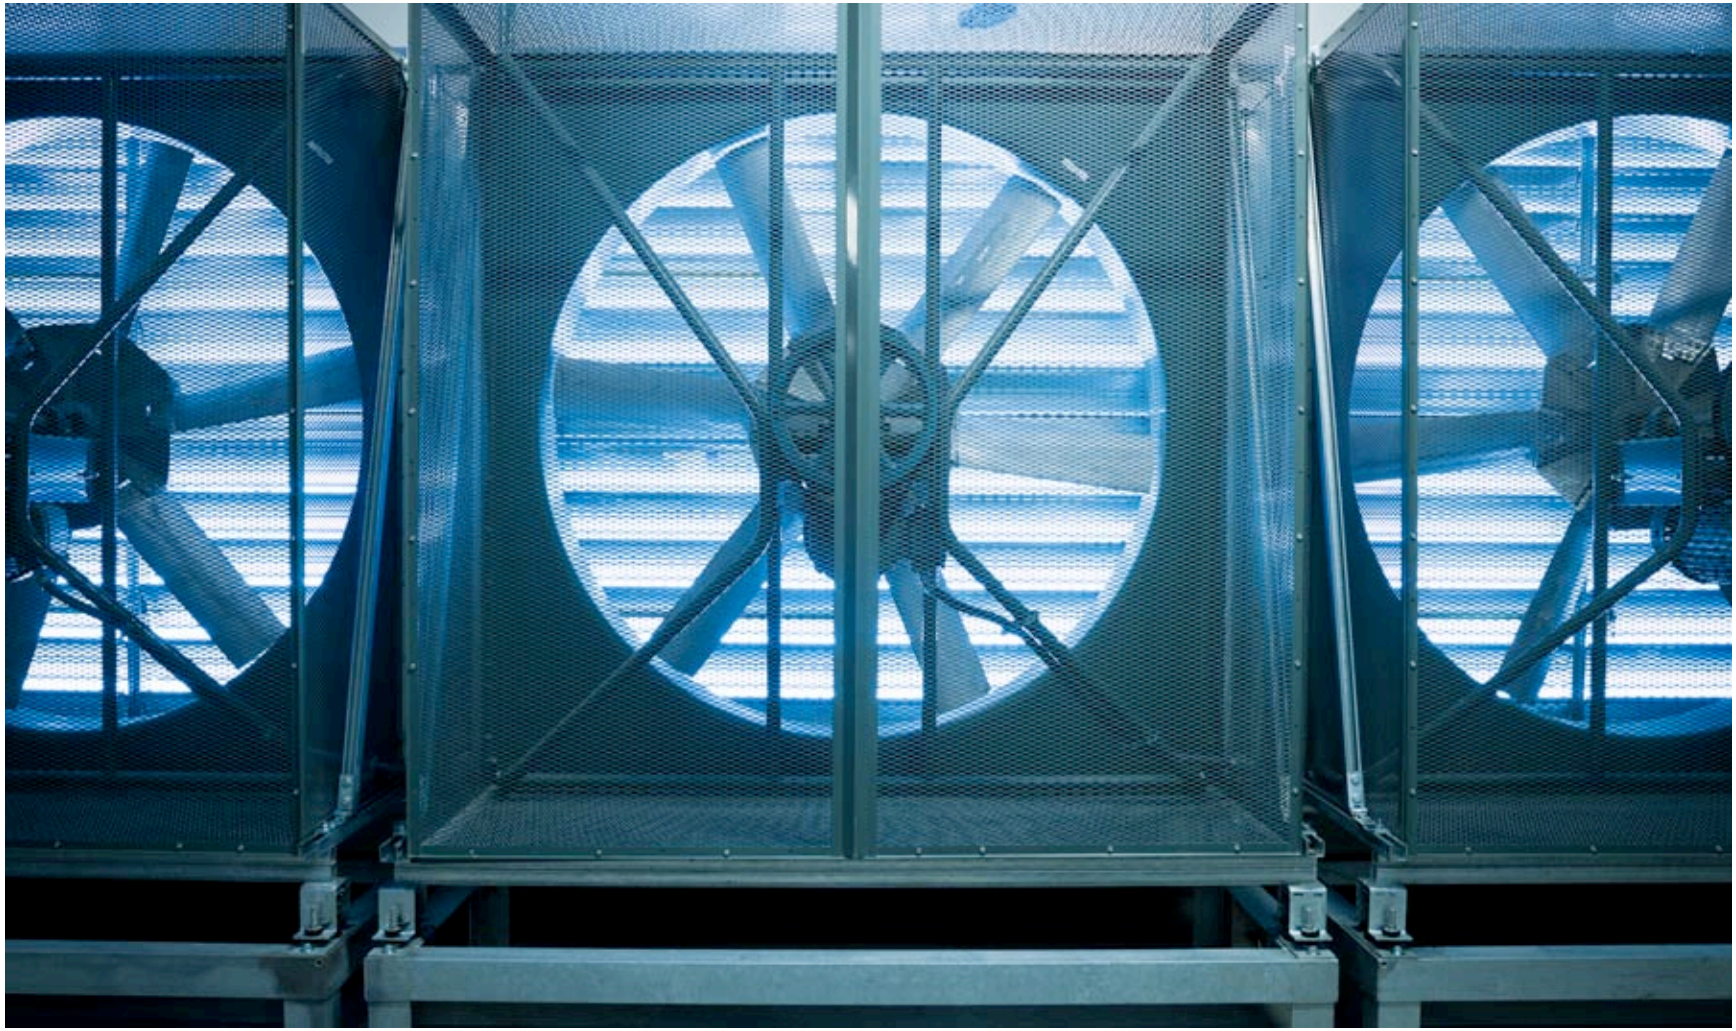

http://www.technologyreview.com/computing/37295/page8/#photo

technology review
Published by MIT

Inside Facebook's Not-So-Secret New Data Center

http://www.technologyreview.com/computing/37295/page8/#photo

Facebook Datacenter Specification

Condition H:
Unacceptable OA Conditions (Smoke or Dust)

Economizer at minimum OA (recirculation mode). **We acknowledge** that this is a condition that can cause **potential shutdown**. We already have filtration installed and will run evaporative cooling at full capacity to reduce smoke and particulates in the event of a fire or contamination. **Then, depending on intensity, we can utilize time for orderly shutdown**, or else run for a prolonged period of time at minimum OA. We have a provision for a closed-loop system that uses indirect cooling; see section 7.

http://opencompute.org

Electricity price (TEPCO, Japan)

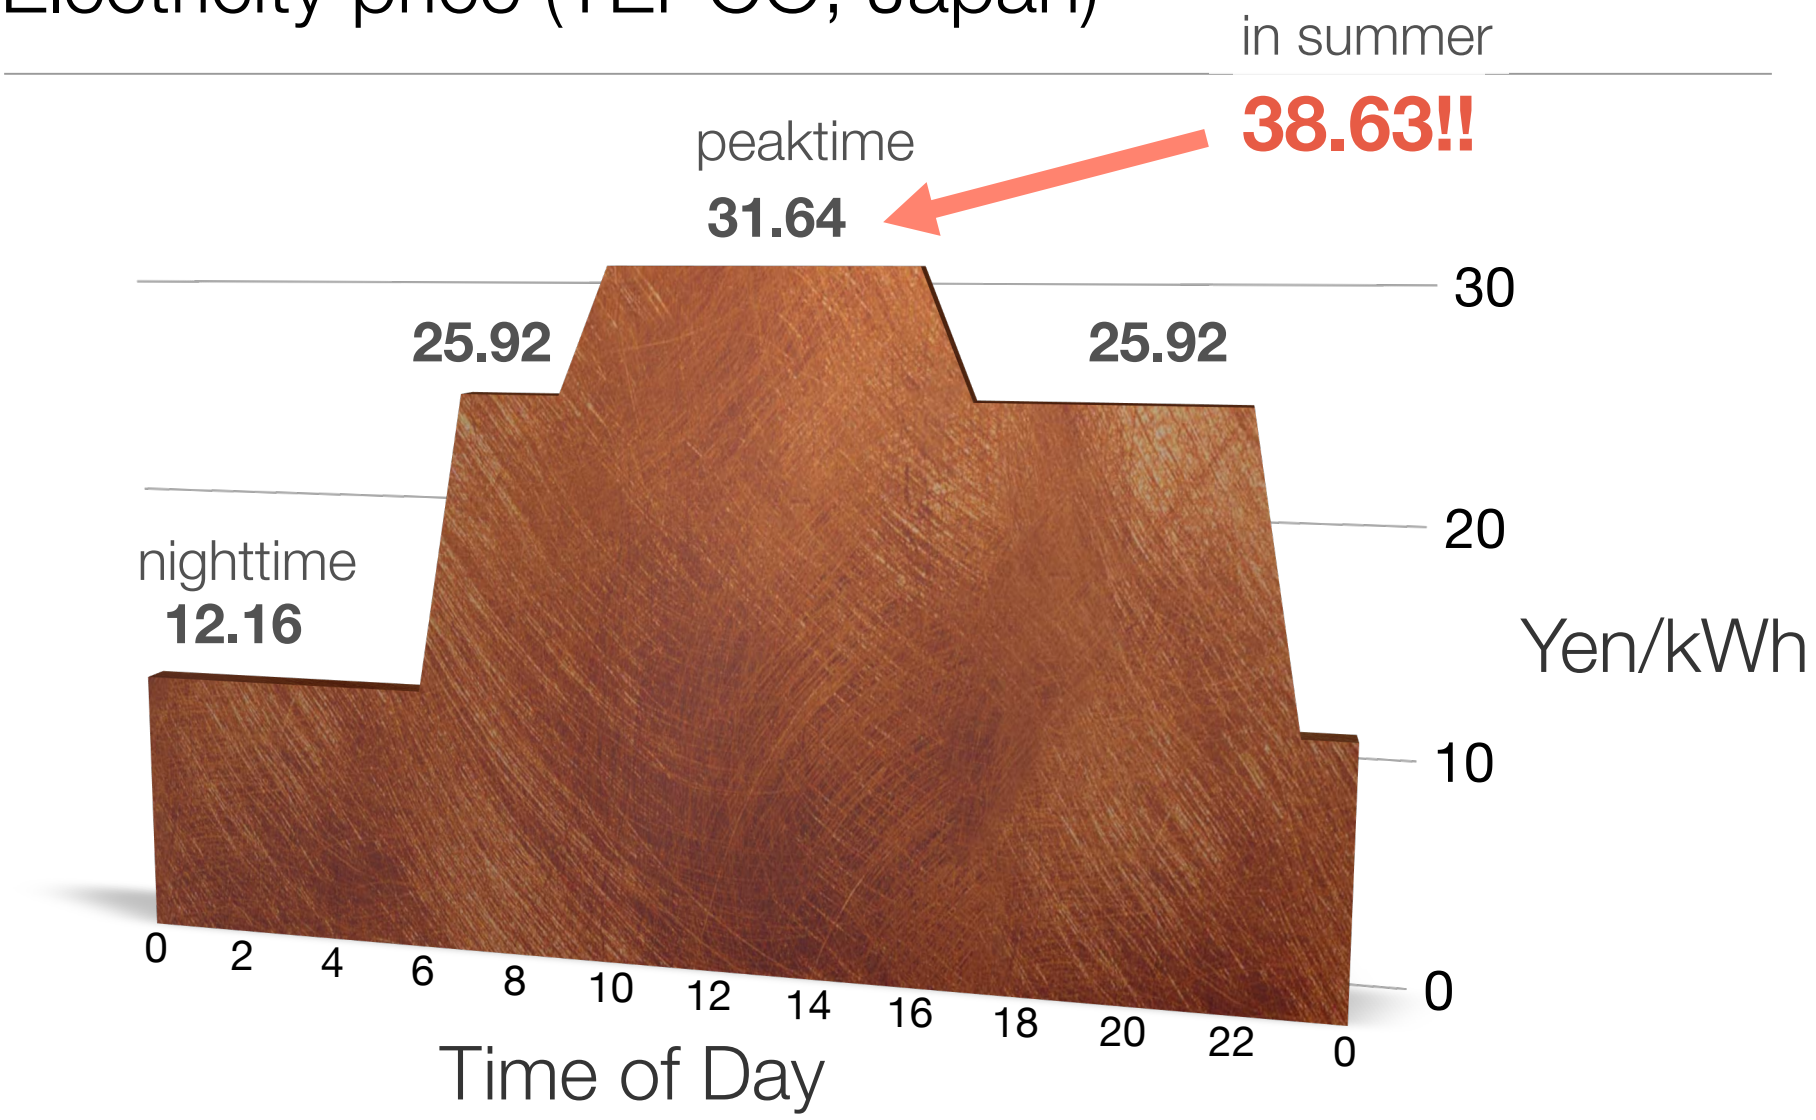

data source : 東京電力 電化上手
http://www.tepco.co.jp/e-rates/individual/menu/home/home01-j.html

Transformation of Distributed Computing

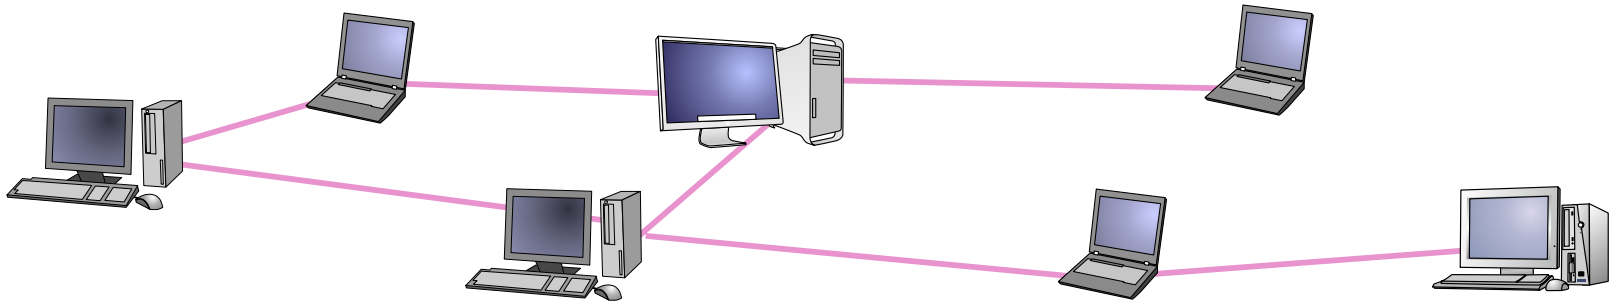

Transformation of Distributed Computing

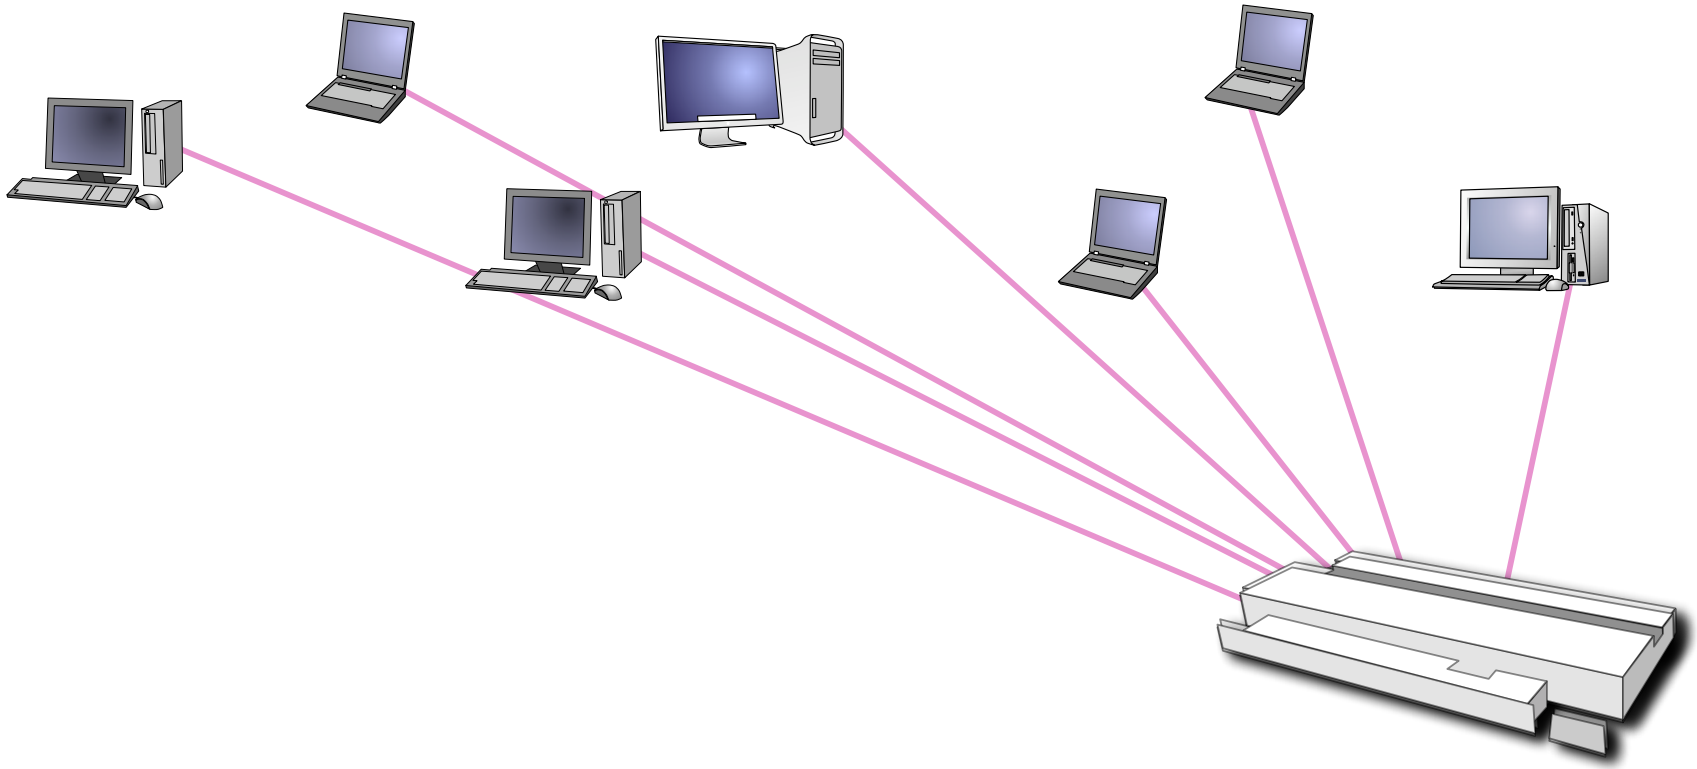

Data Center

Transformation of Distributed Computing

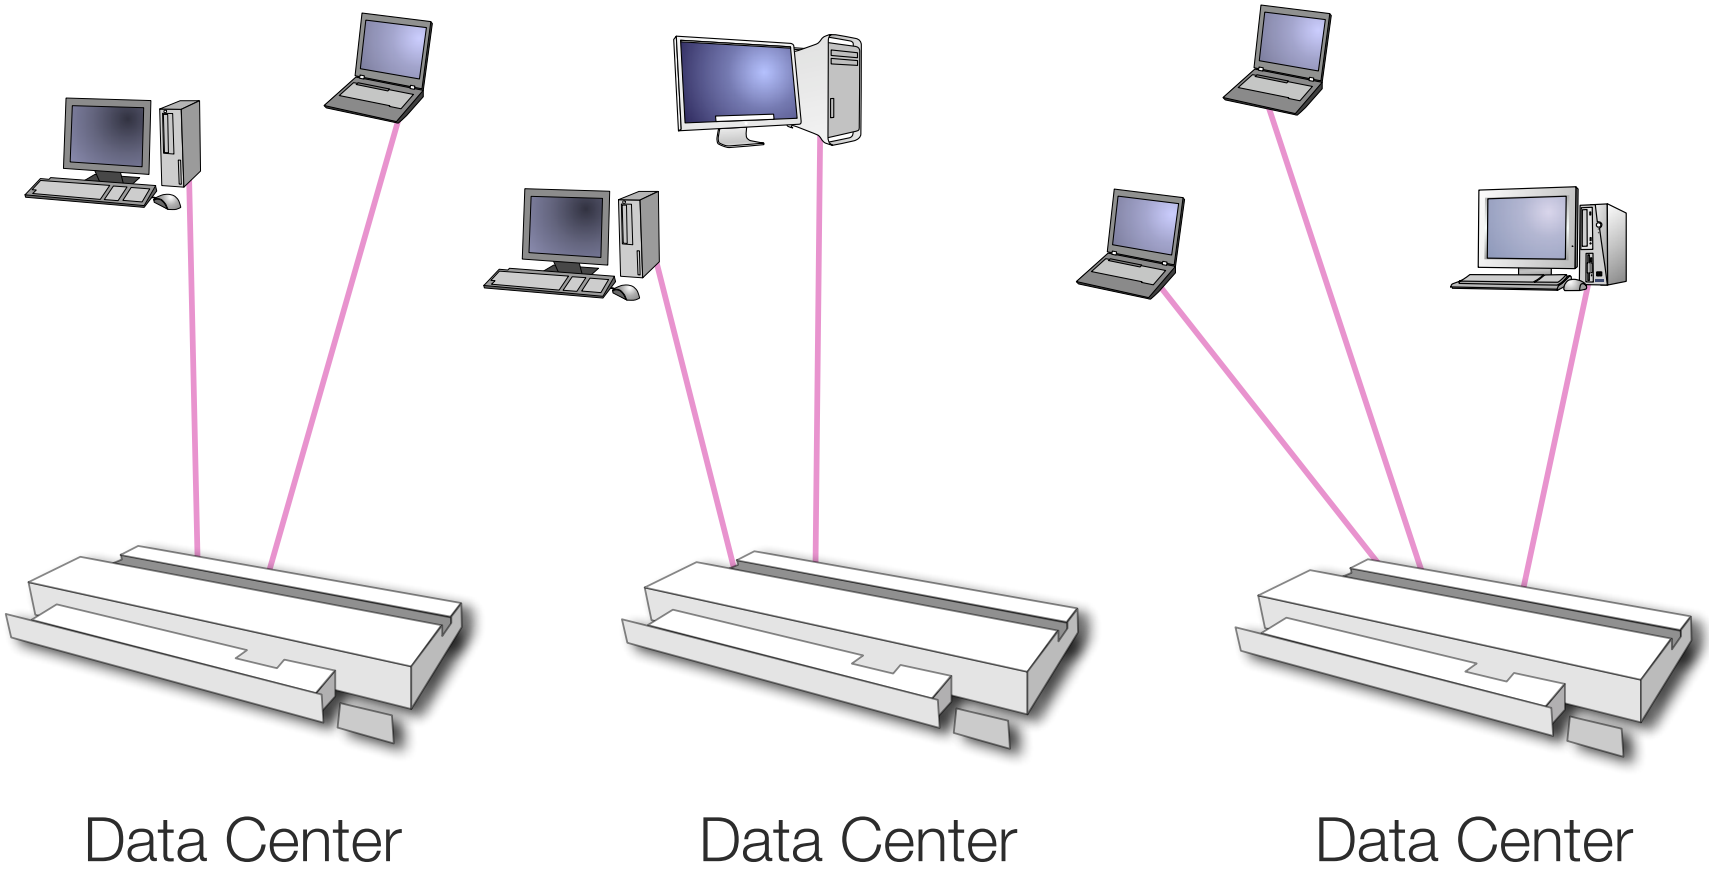

FORTE : Flow Optimization based framework for request-Routing and Traffic Engineering

"FORTE dynamically controls the fraction of user traffic directed to each datacenter in response to changes in both request workload and carbon footprint."

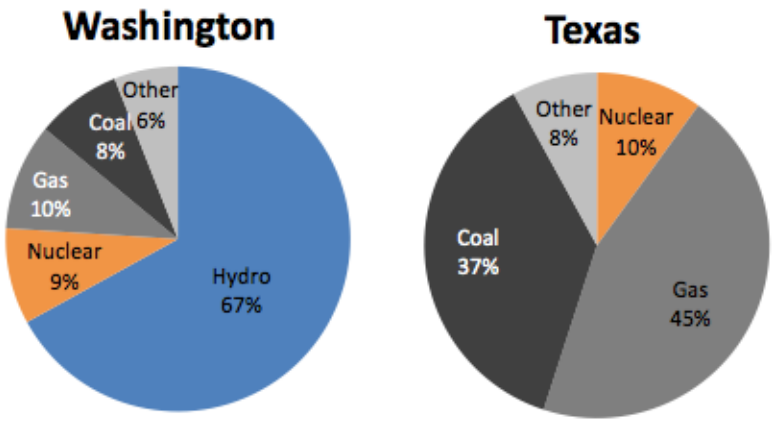

Figure 2: Generator fuel type in Washington and Texas

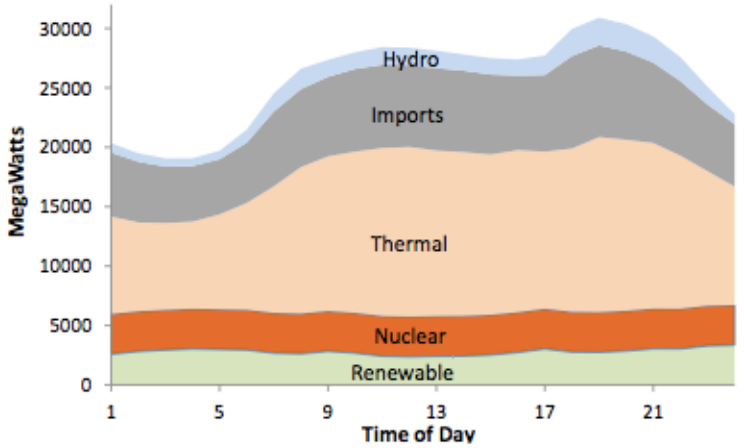

Figure 3: Hourly breakdown of total production by resource type, California, Jan 23rd 2012 [5]

"It's Not Easy Being Green", Peter Xiang Gao et. al., in Proceedings of ACM SIGCOMM 2012

3 way tradeoffs of access latency, carbon footprint, and electricity costs

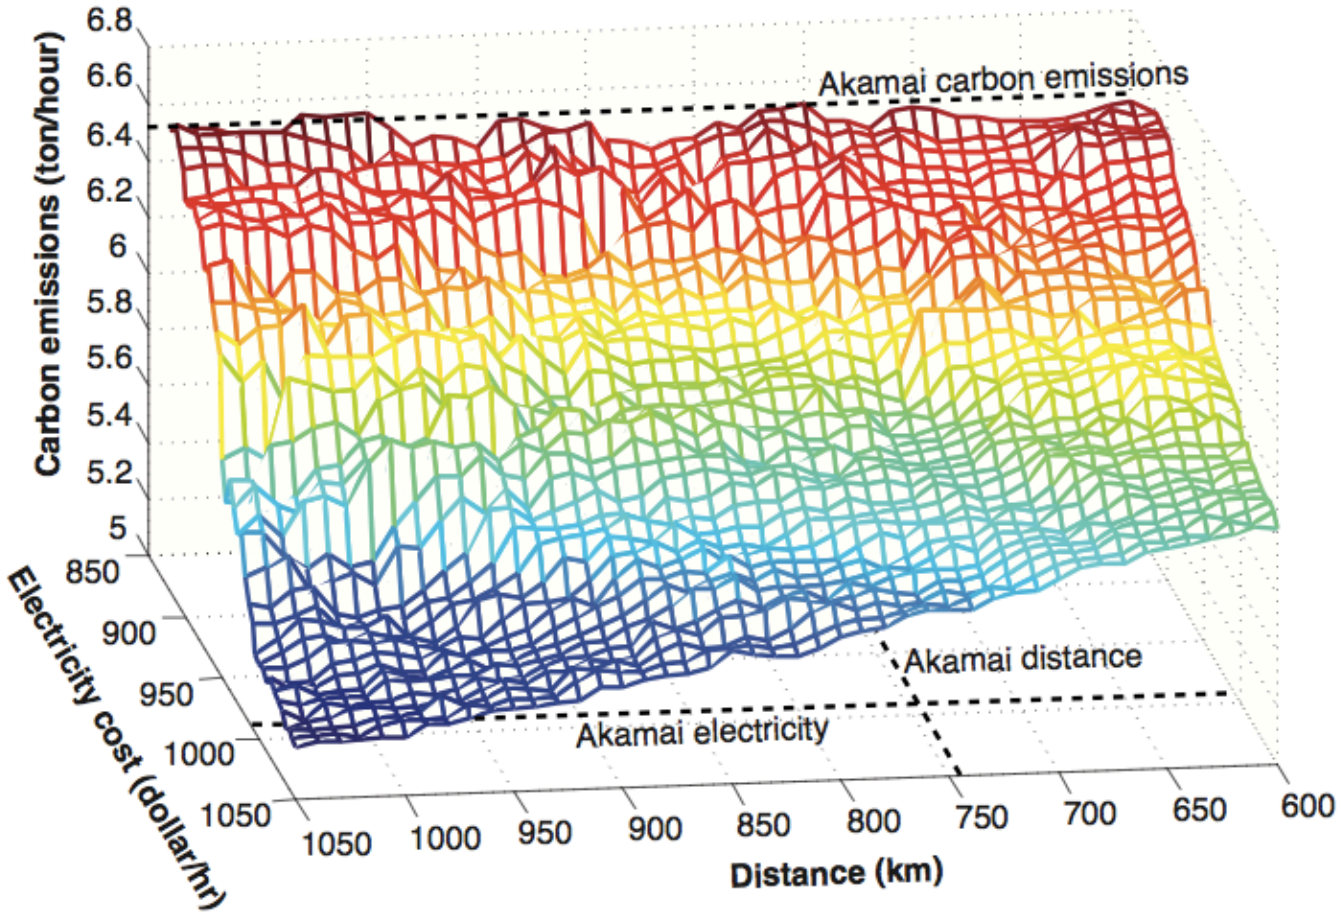

Figure 12: Tradeoff between carbon emissions, average distance, and electricity costs.

"It's Not Easy Being Green", Peter Xiang Gao et. al., in Proceedings of ACM SIGCOMM 2012