To Cloud or Not To.

An exploration of the economics of clouds and cyber-security.

Radu Sion

C3 Cloud Computing Center

National Science Foundation
WHERE DISCOVERIES BEGIN

Center of Excellence
Wireless and Information Technology

Stony Brook University

ver 1.9

Feynman Moment

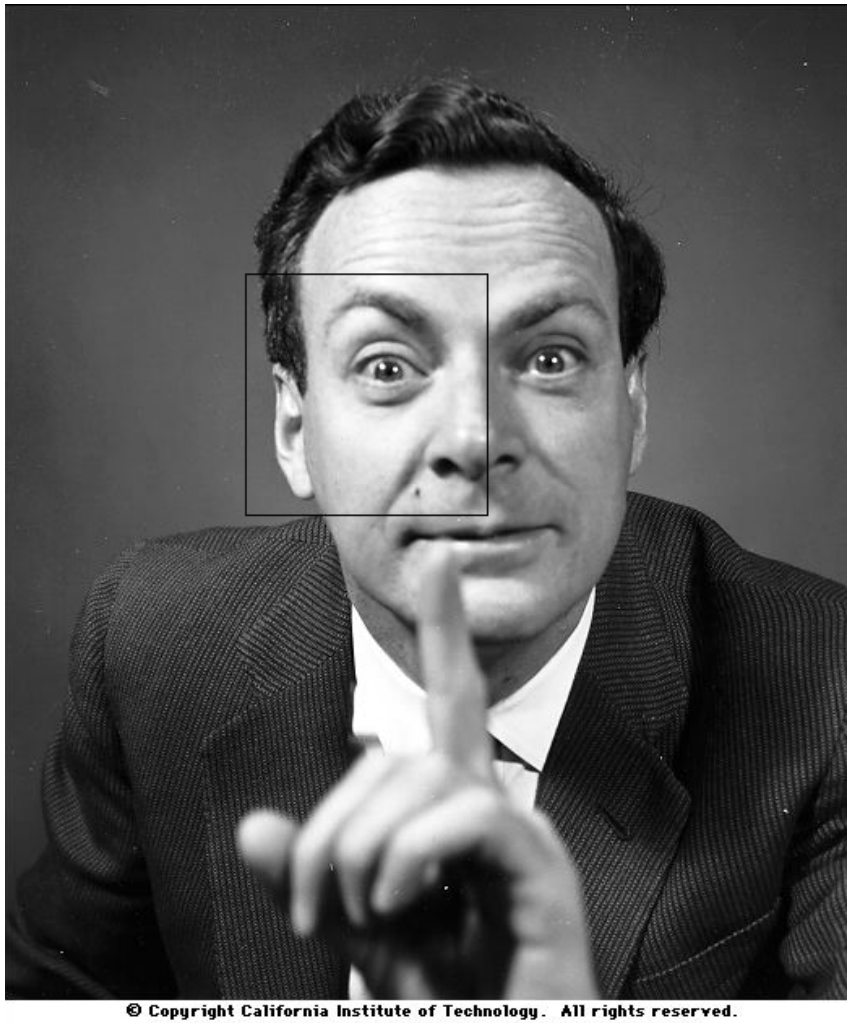

"I have experience only in teaching graduate students […] and as a result […] I know that I don't know how to teach."

so: please interrupt and engage

Gordon Moore

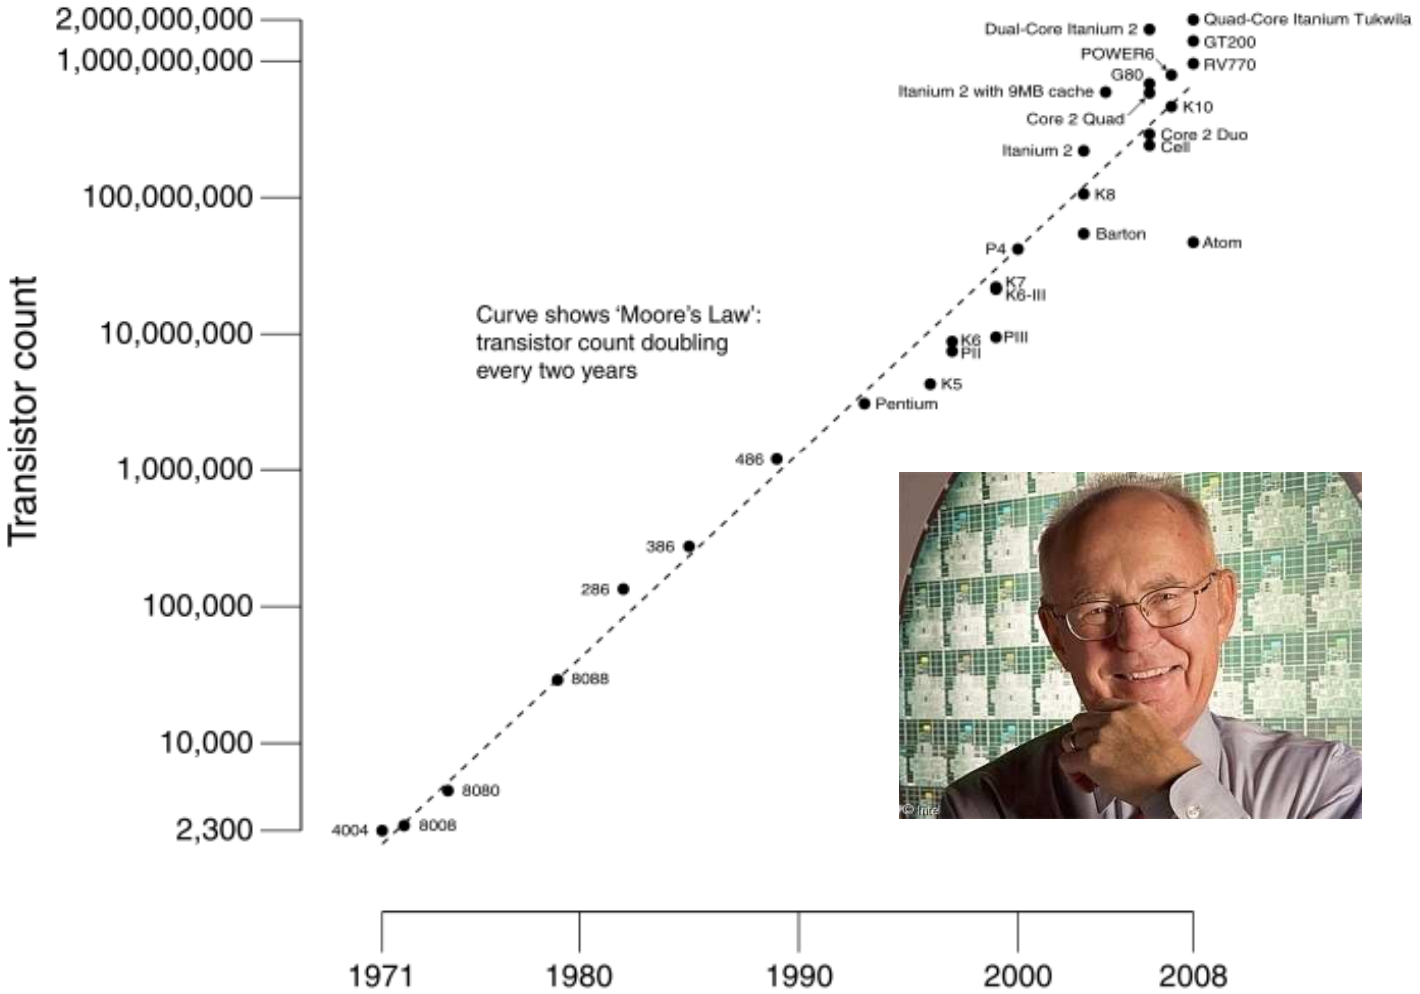

CPU Speeds Follow Moore

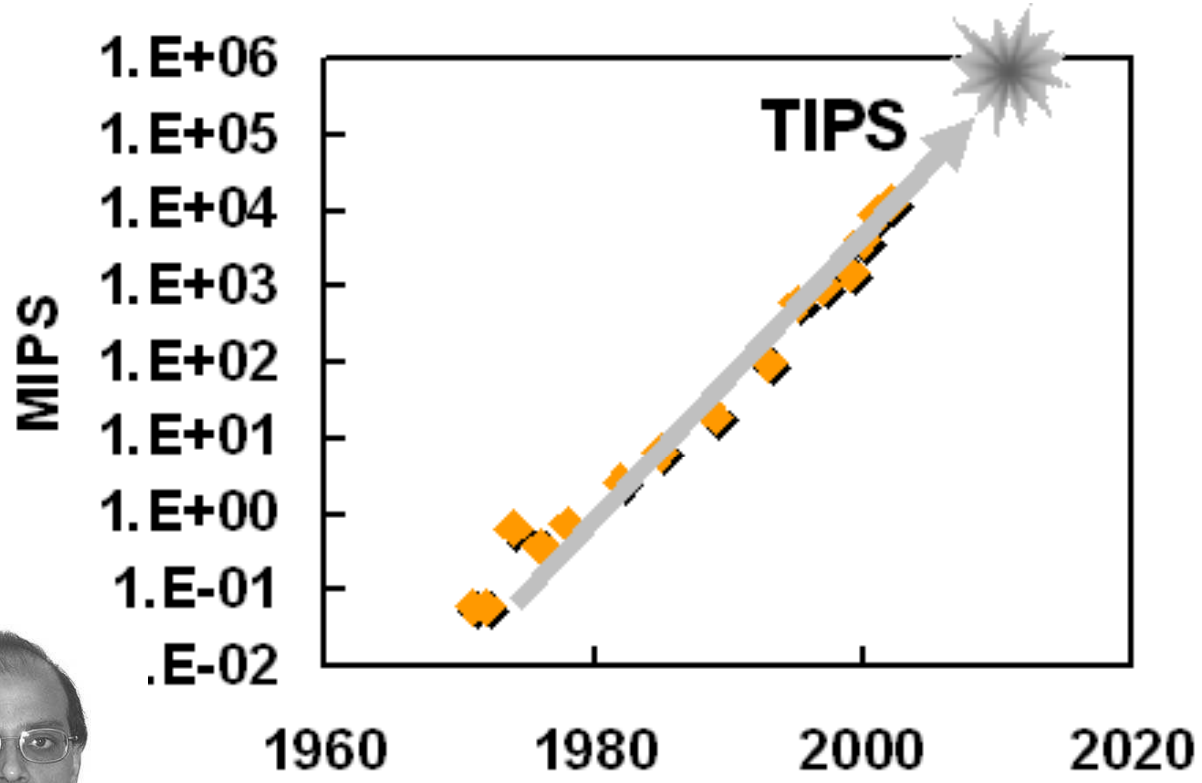

Source: "Gigascale Integration-Challenges and Opportunities", Shekhar Borkar, Director, Circuit Research, Intel Corp.

Jakob Nielsen

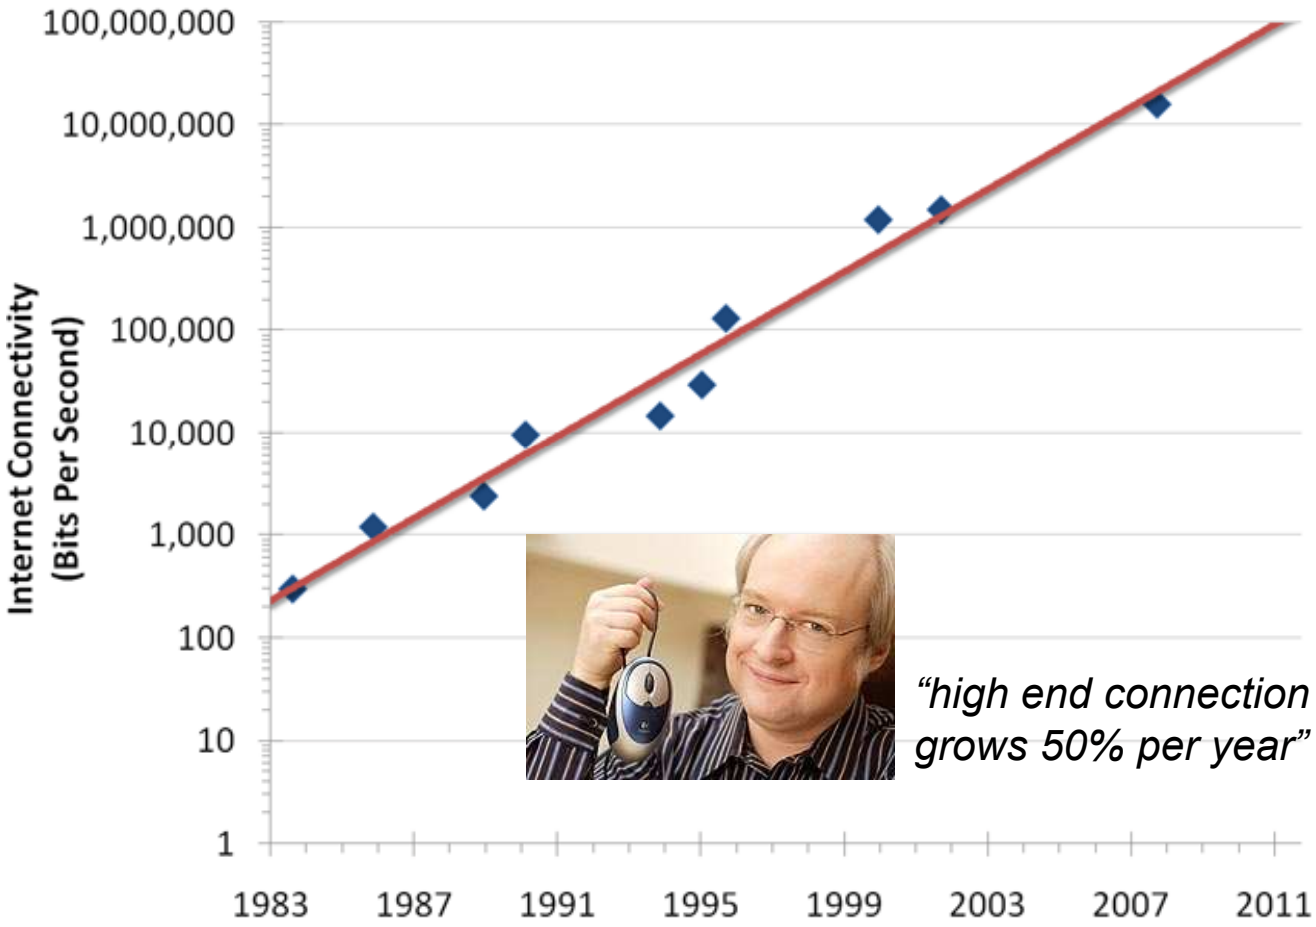

"high end connection speed grows 50% per year"

Thus The Cloud

government
corporations
healthcare

consumers

mobile

"Utility" Computing
Transparency
Availability
Extremely cheap
Efficient
On-Demand

Online Services
Software
Platforms
Infrastructure
Knowledge

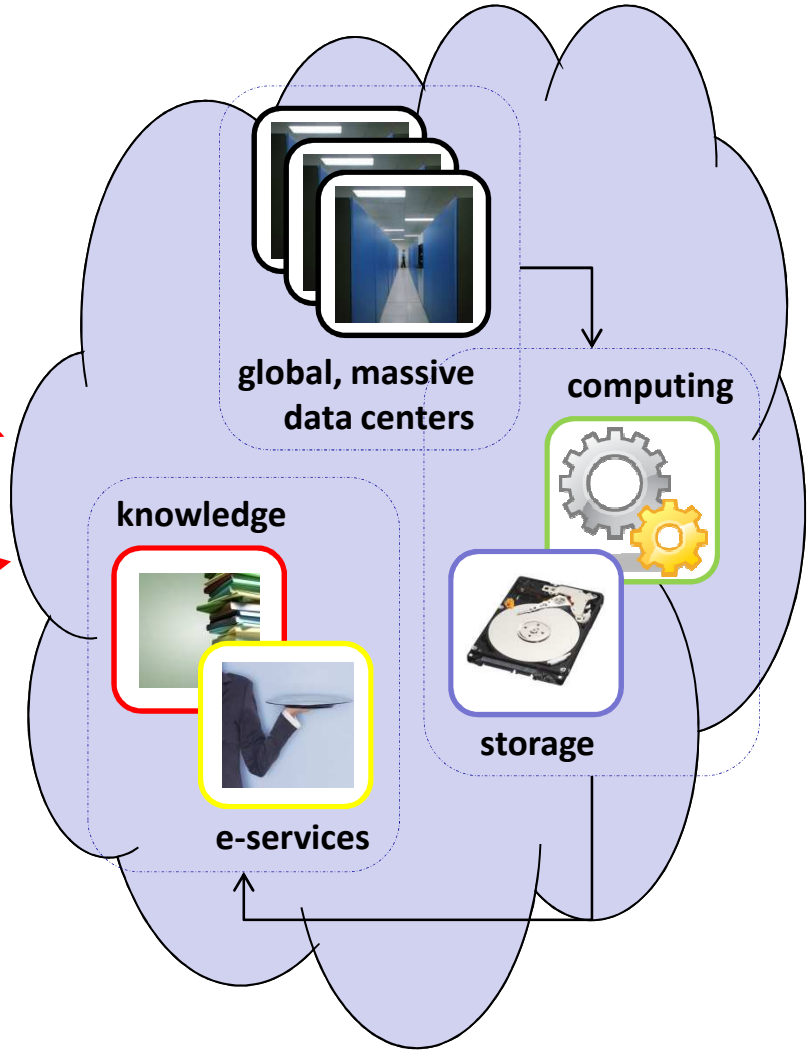

Clouds v. Grids v. Clusters v. …

+ Illusion of “Unlimited”
+ No up-front commitment (“pay as you go”)
+ On-demand
+ (Very) Short-term allocation
+ Close to 100% Transparency
+ Increased Platform Independence
+ It is actually here and happening!

Cloud Flavors

Traditional Outsourcing [(Semi)Private Clouds]
ACME Corp. manages servers for XYZ Financials

Clouds
Amazon EC2, Google Apps, MS Azure

Managed servers

Un-managed hardware

But: is it worth it?

costs vs. benefits

costs
technology costs
cost of security
etc.

clients

benefits
availability
opportunity
consolidation
etc.

the "cloud"

“Core costs of computing”

+ Storage ($/MByte/year)
+ Computing ($/CPU Cycles)
+ Networking ($/bit)

Reality is more mundane

Hardware

servers, disks, **network**, racks, power, cooling

Energy

power, cooling, infrastructure

People/Service

maintenance, development

Space

Size matters

Home Users (1-10 CPUs)

"no" rent/cooling/administration

Small Enterprises (up to 1k)

no custom hardware, low utilization

Mid-size Enterprises (up to 10k)

better network service, better utilization

Large/Clouds (10k+)

Clouds

+ Custom hardware
+ Efficient cooling
+ Cross-timezone load shifting
+ High CPU utilization
+ Preferential network deals
+ High Power Usage Efficiency (PUE)

Finding out the cost of a CPU Cycle

Parameters	H	S	M	L
CPU utilization	5-8%	10-12%	15-20%	40-56%
server:admin ratio	N.A.	100-140	140-200	800-1000
Space (sqft/month)	N.A.	$0.5	$0.5	$0.25
PUE	N.A.	2-2.5	1.6-2	1.2-1.5

$$\frac{\lambda_s \cdot N_s/\tau_s + (w_p \cdot \mu + w_i \cdot (1-\mu)) \cdot PUE \cdot \lambda_e + \frac{N_s}{\alpha} \cdot \lambda_p + \lambda_w \cdot N_w/\tau_w + \lambda_f \cdot \frac{(w_p \cdot \mu + w_i \cdot (1-\mu)) \cdot PUE}{\beta}}{\mu \cdot \nu \cdot N_s}$$

CPU Cycle Cost

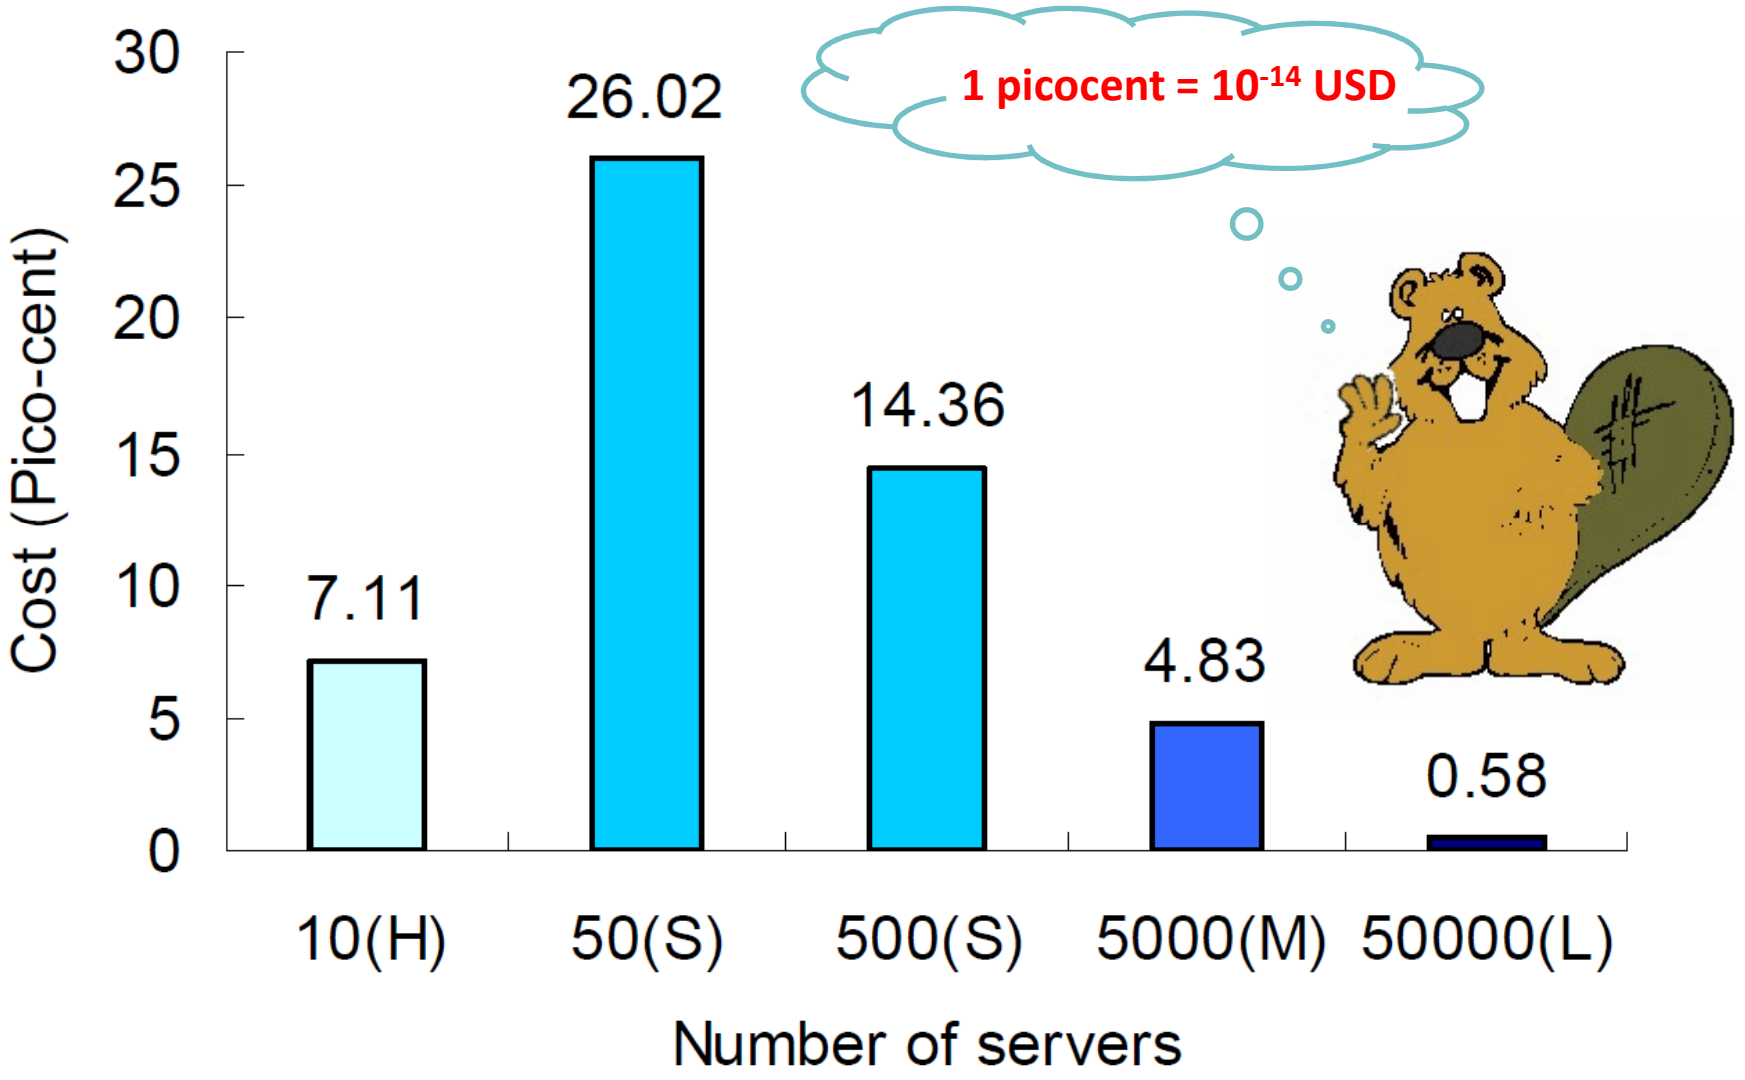

Today's consumer clouds

Provider	Picocents
[redacted]	0.93 - 2.36
	up to 2.31
	up to 1.96

Cost Breakdown

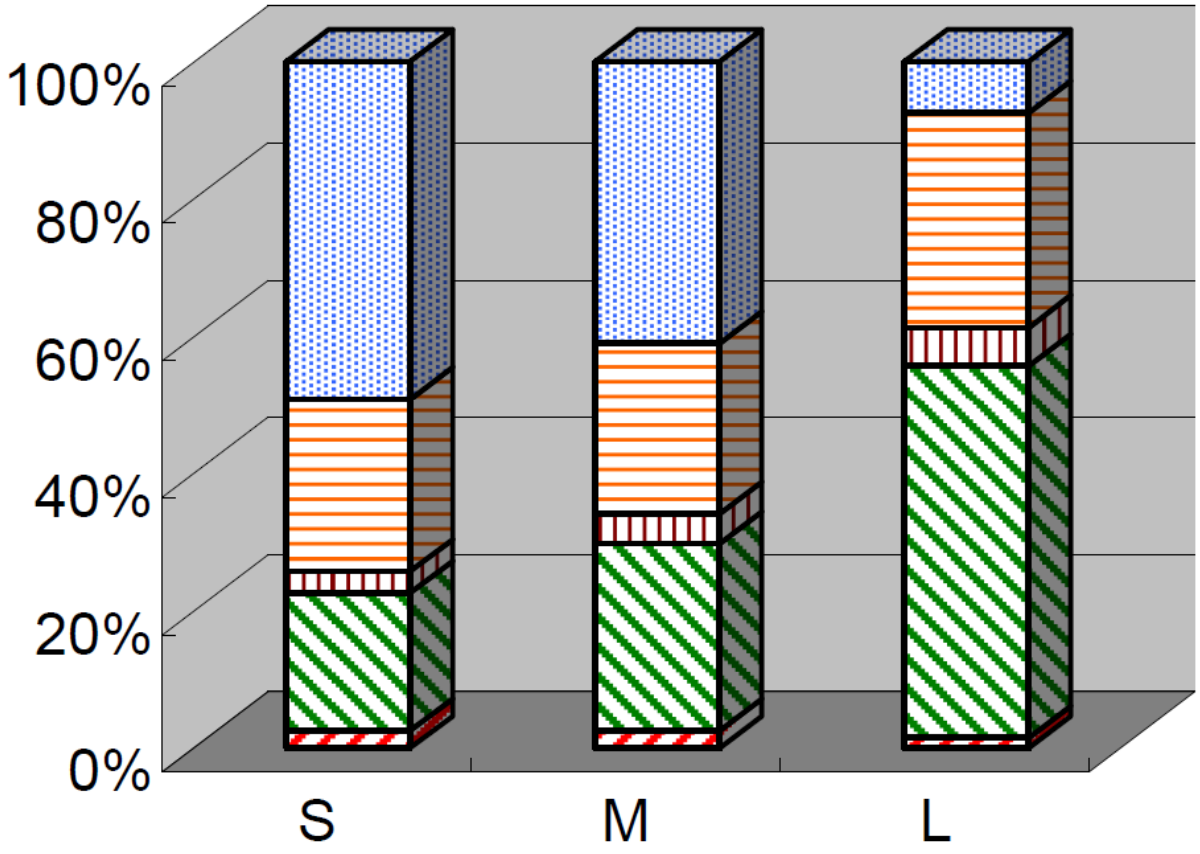

So: is it worth it?

Mostly yes ...

Why ?

1 client cycle
6-27 US picocents

clients

1 cloud cycle
0.58 picocents

But we are far!

provider	monthly	bandwidth (d/u)	picocent/bit
[redacted]	$29.95	15 Mbps /5 Mbps	77/231
	$44.9	30 Mbps /5 Mbps	58/346
	>$1000	5-1000 Mbps	5000 (est.)
	$19.99	1 Mbps/384 Kbps	771/2008
	$29.99	3 Mbps/768 Kbps	386/1506
	$42.99	7.1 Mbps/768 Kbps	233/2160
Mid-size	$95 (est.)	1 Mbps (dedicated)	3665 (est.)
Large/cloud	$13 (est.)	1 Mbps (dedicated)	500 (est.)

Per bit transfer cost	
H → cloud	800
S → cloud	6,000
M → cloud	4,500

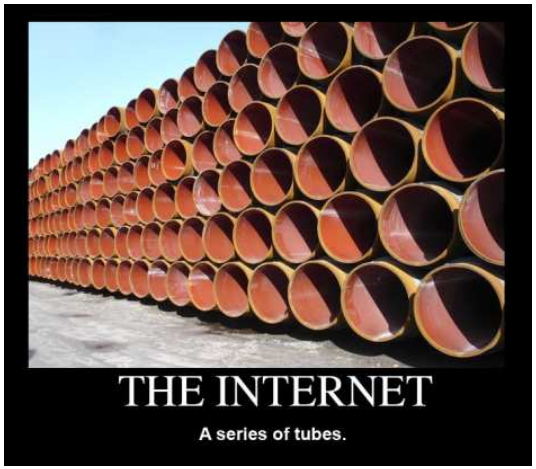

Storage?

Disk	cap.	price	Adj. MTBF	amort. acq.	power	power2	power3	power cost	total cost	acq. %	avg. seek	avg. seek4	power5	read cost
	(GB)	(USD)	(mil.hrs)	(pcent/bit/yr)	seek (W)	idle (W)	(W)	(pcent/bit/yr)	(pcent/bit/yr)		time (ms)	cost (pcents)	read (W)	(pcent/bit)
Maxtor Diamond Max	500	53	0.35	32.89	13.6	8.10	10.85	237.62	270.50	12.16	9.00	377542	11.16	0.03
Hitachi Deskstar 7k500	500	67	0.29	49.89	15	9.60	12.30	269.37	319.26	15.63	8.50	407953		
Hitachi Ultrastar A7K1000	1024	153	0.35	46.36	14	9.00	11.50	122.97	169.33	27.38	8.20	417631		
WD Caviar GP Low Power	1024	103	0.29	37.45	7.5	4.00	5.75	61.49	98.93	37.85	8.90	271994	7.40	0.02
Seagate Barracuda 7200.10	750	63	0.35	26.06	12.6	9.30	10.95	159.87	185.93	14.02	9.25	369615	13.00	0.06
WD Caviar SE16	500	62	N/A		8.77	8.40	8.59	188.01			9.90		8.77	0.04
Samsung SSD	32	269	0.29	3129.65	1	1.00	1.00	342.19	3471.83	90.14	1.70	47912	0.5	0.0017
Intel SSD X18-M	80	389	0.35	1508.59	0.15	0.06	0.11	14.37	1522.96	99.06			0.15	0.0002
Intel SSD X25-M	160	765	0.35	1483.38	0.15	0.06	0.11	7.19	1490.57	99.52			0.15	0.0002

Up to 350 for 3 year lifetime!

Storage Capacity over Time (retail)

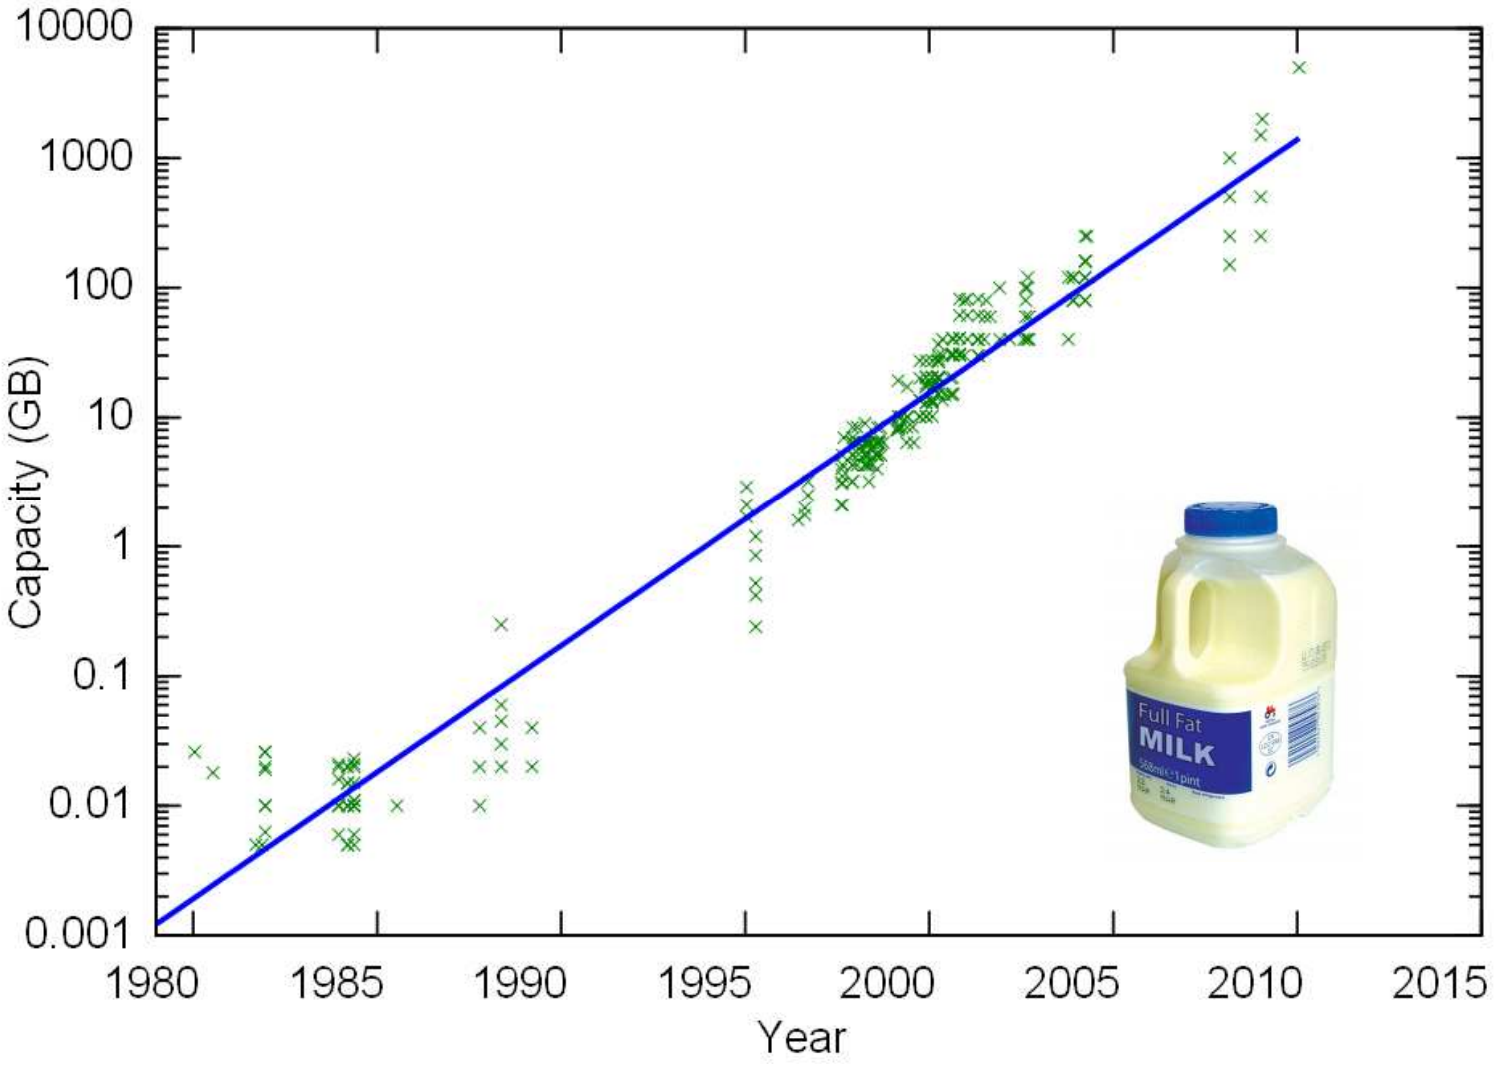

So are clouds worth it ?

… not always.

CPU Cycle
6-27 picocents

1 bit storage/year
6 picocents

CPU Cycle
0.58 picocents

1 bit storage/year
5.3-6 picocents

1 bit network transfer
800-6000 picocents

clients

Application Owner == Sole Client

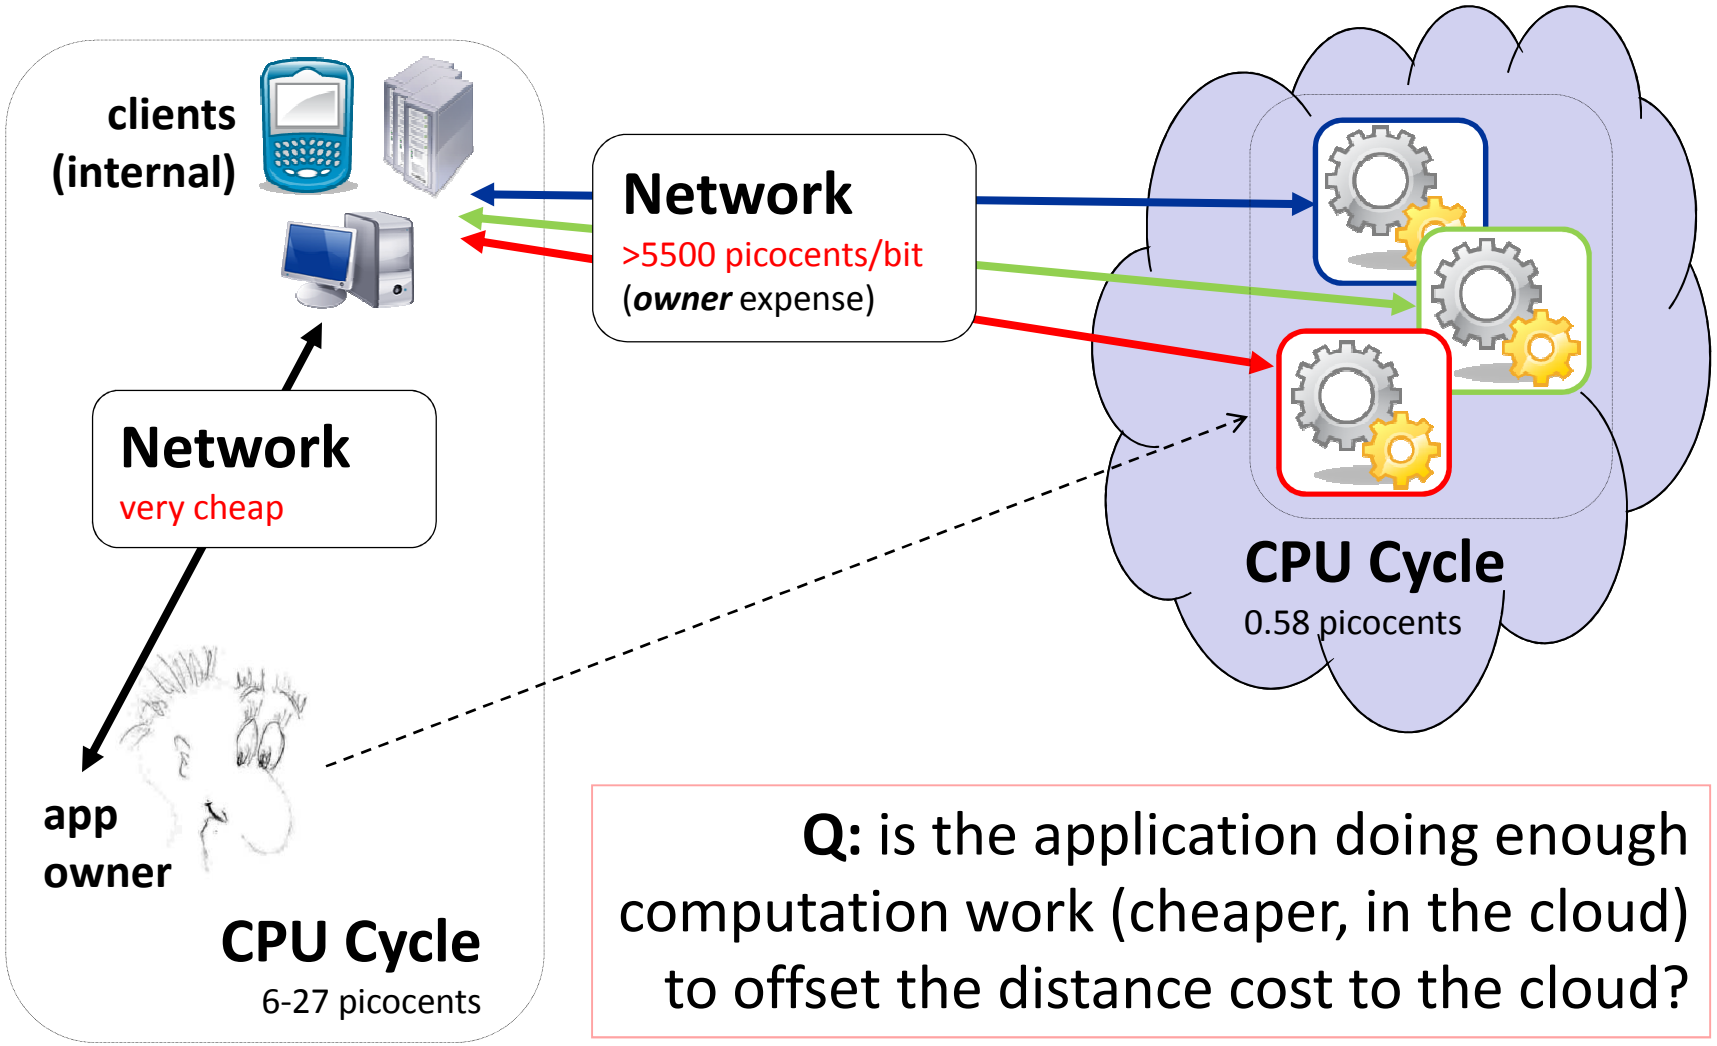

Q: is the application doing enough computation work (cheaper, in the cloud) to offset the distance cost to the cloud?

So when is it clearly worth it ?

Q: is the application doing enough computation work (cheaper) to offset the distance cost to the cloud?

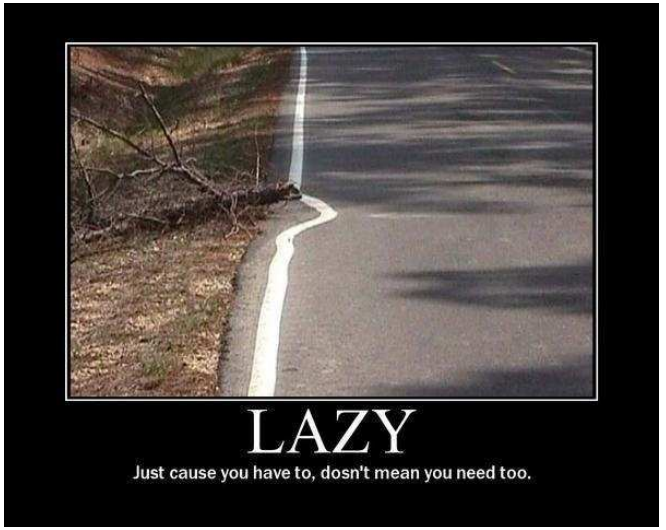

First Principle of Cloud Viability

It is not worth outsourcing any task of less than 4000 CPU cycles per transferred 32-bit input.

But wait … there's more!

… we had only a partial view.

The **actual question** to ask**:** what is the overall application profile (comp+net+storage)

Second Principle of Cloud Viability (paraphrased)
"It is almost always worth outsourcing"

Application Owner != Client(s)

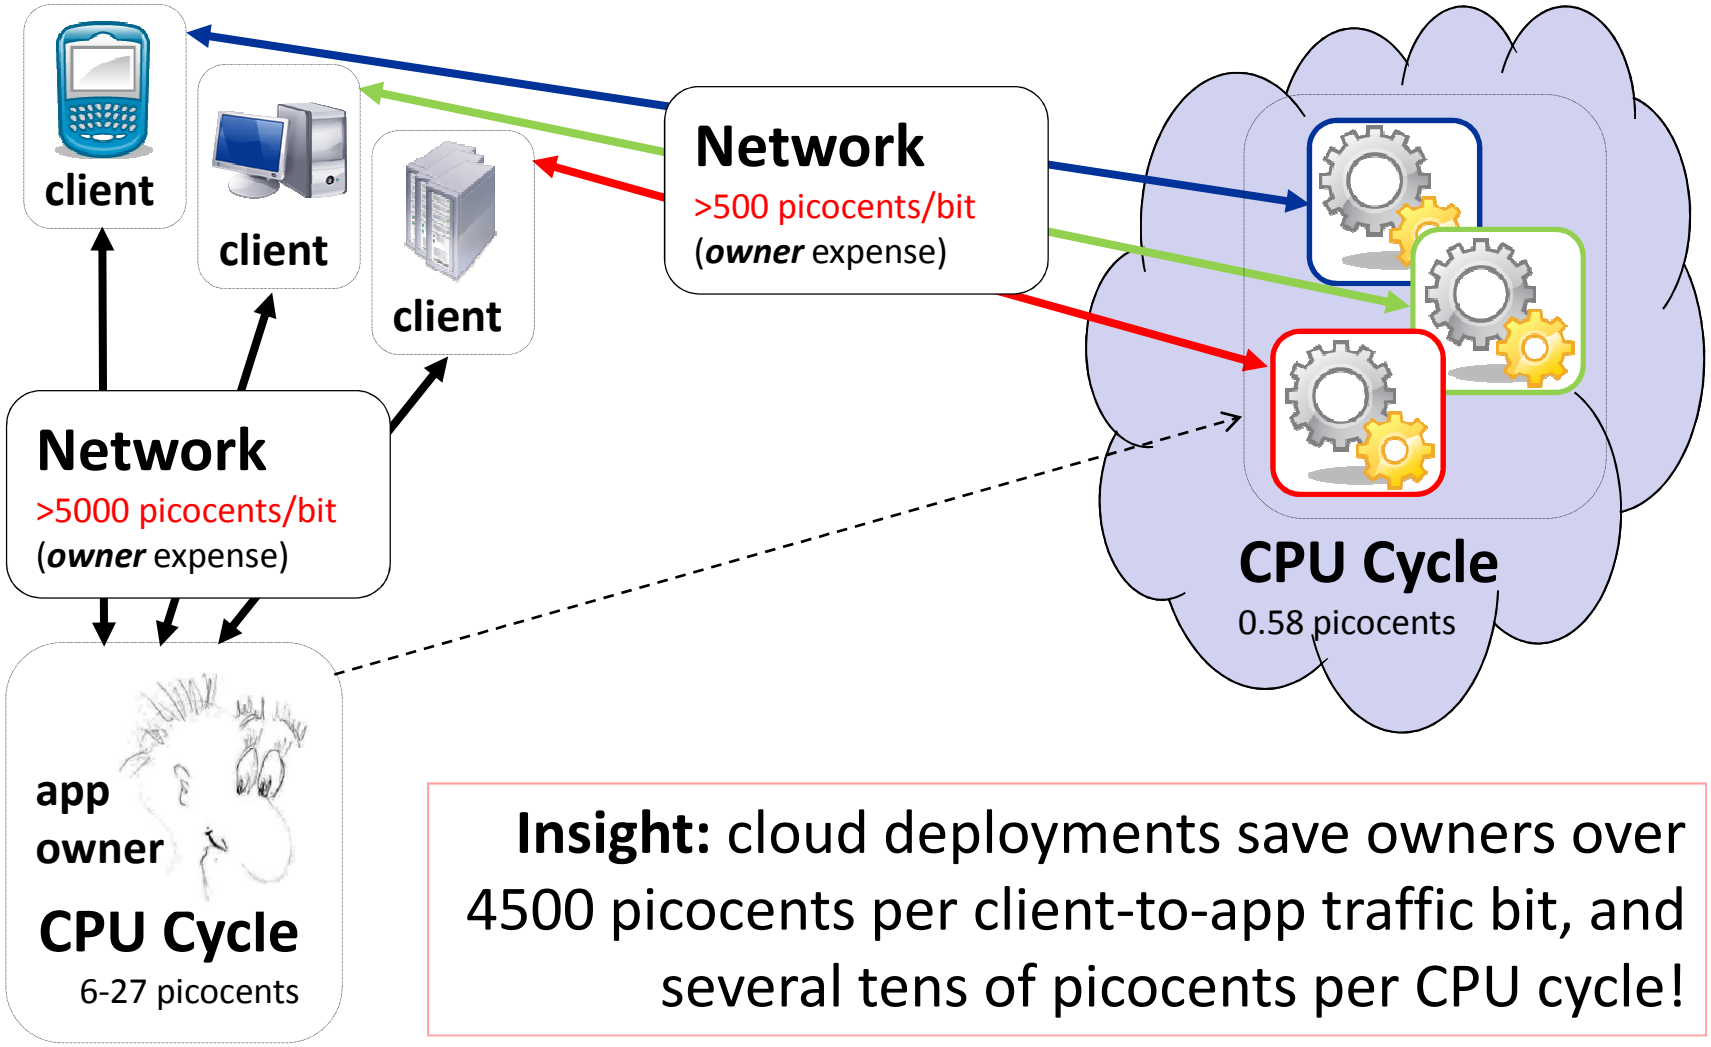

Insight: cloud deployments save owners over 4500 picocents per client-to-app traffic bit, and several tens of picocents per CPU cycle!

Challenges

+ **Interoperability and Standards**
+ Cyber-Security
+ **Privacy and Data Confidentiality**
+ **Shift in Liability**
+ **Regulatory Compliance**
+ Transparent Infrastructure Scalability
+ New Energy Efficient Designs
+ Application Deployment Mechanisms
+ Economic Modeling of new Market
+ **Portability for Legacy IT in Clouds**

take home stuff

Clouds are coming to stay
technology (fast networks, cheap CPUs, storage) and markets are ripe and will soon reach a critical mass

Cloud Computing is extremely cost-feasible
but the savings are a function of application footprints and requirements.

/bin/yes > /dev/null

Sailing Hobiecats on Long Island ☺