InGeoCloudS
Inspired GEOdata CLOUD Services

cip competitiveness and innovation framework programme 2007–2013

Cloud Computing in a nutshell

InGeoCLOUDS Experts Workshop
May 9th-10th, Paris

Claudio Lucchese, Nicola Tonellotto
ISTI. - CNR

How to deal with petabytes of data?

… fluctuating in time…

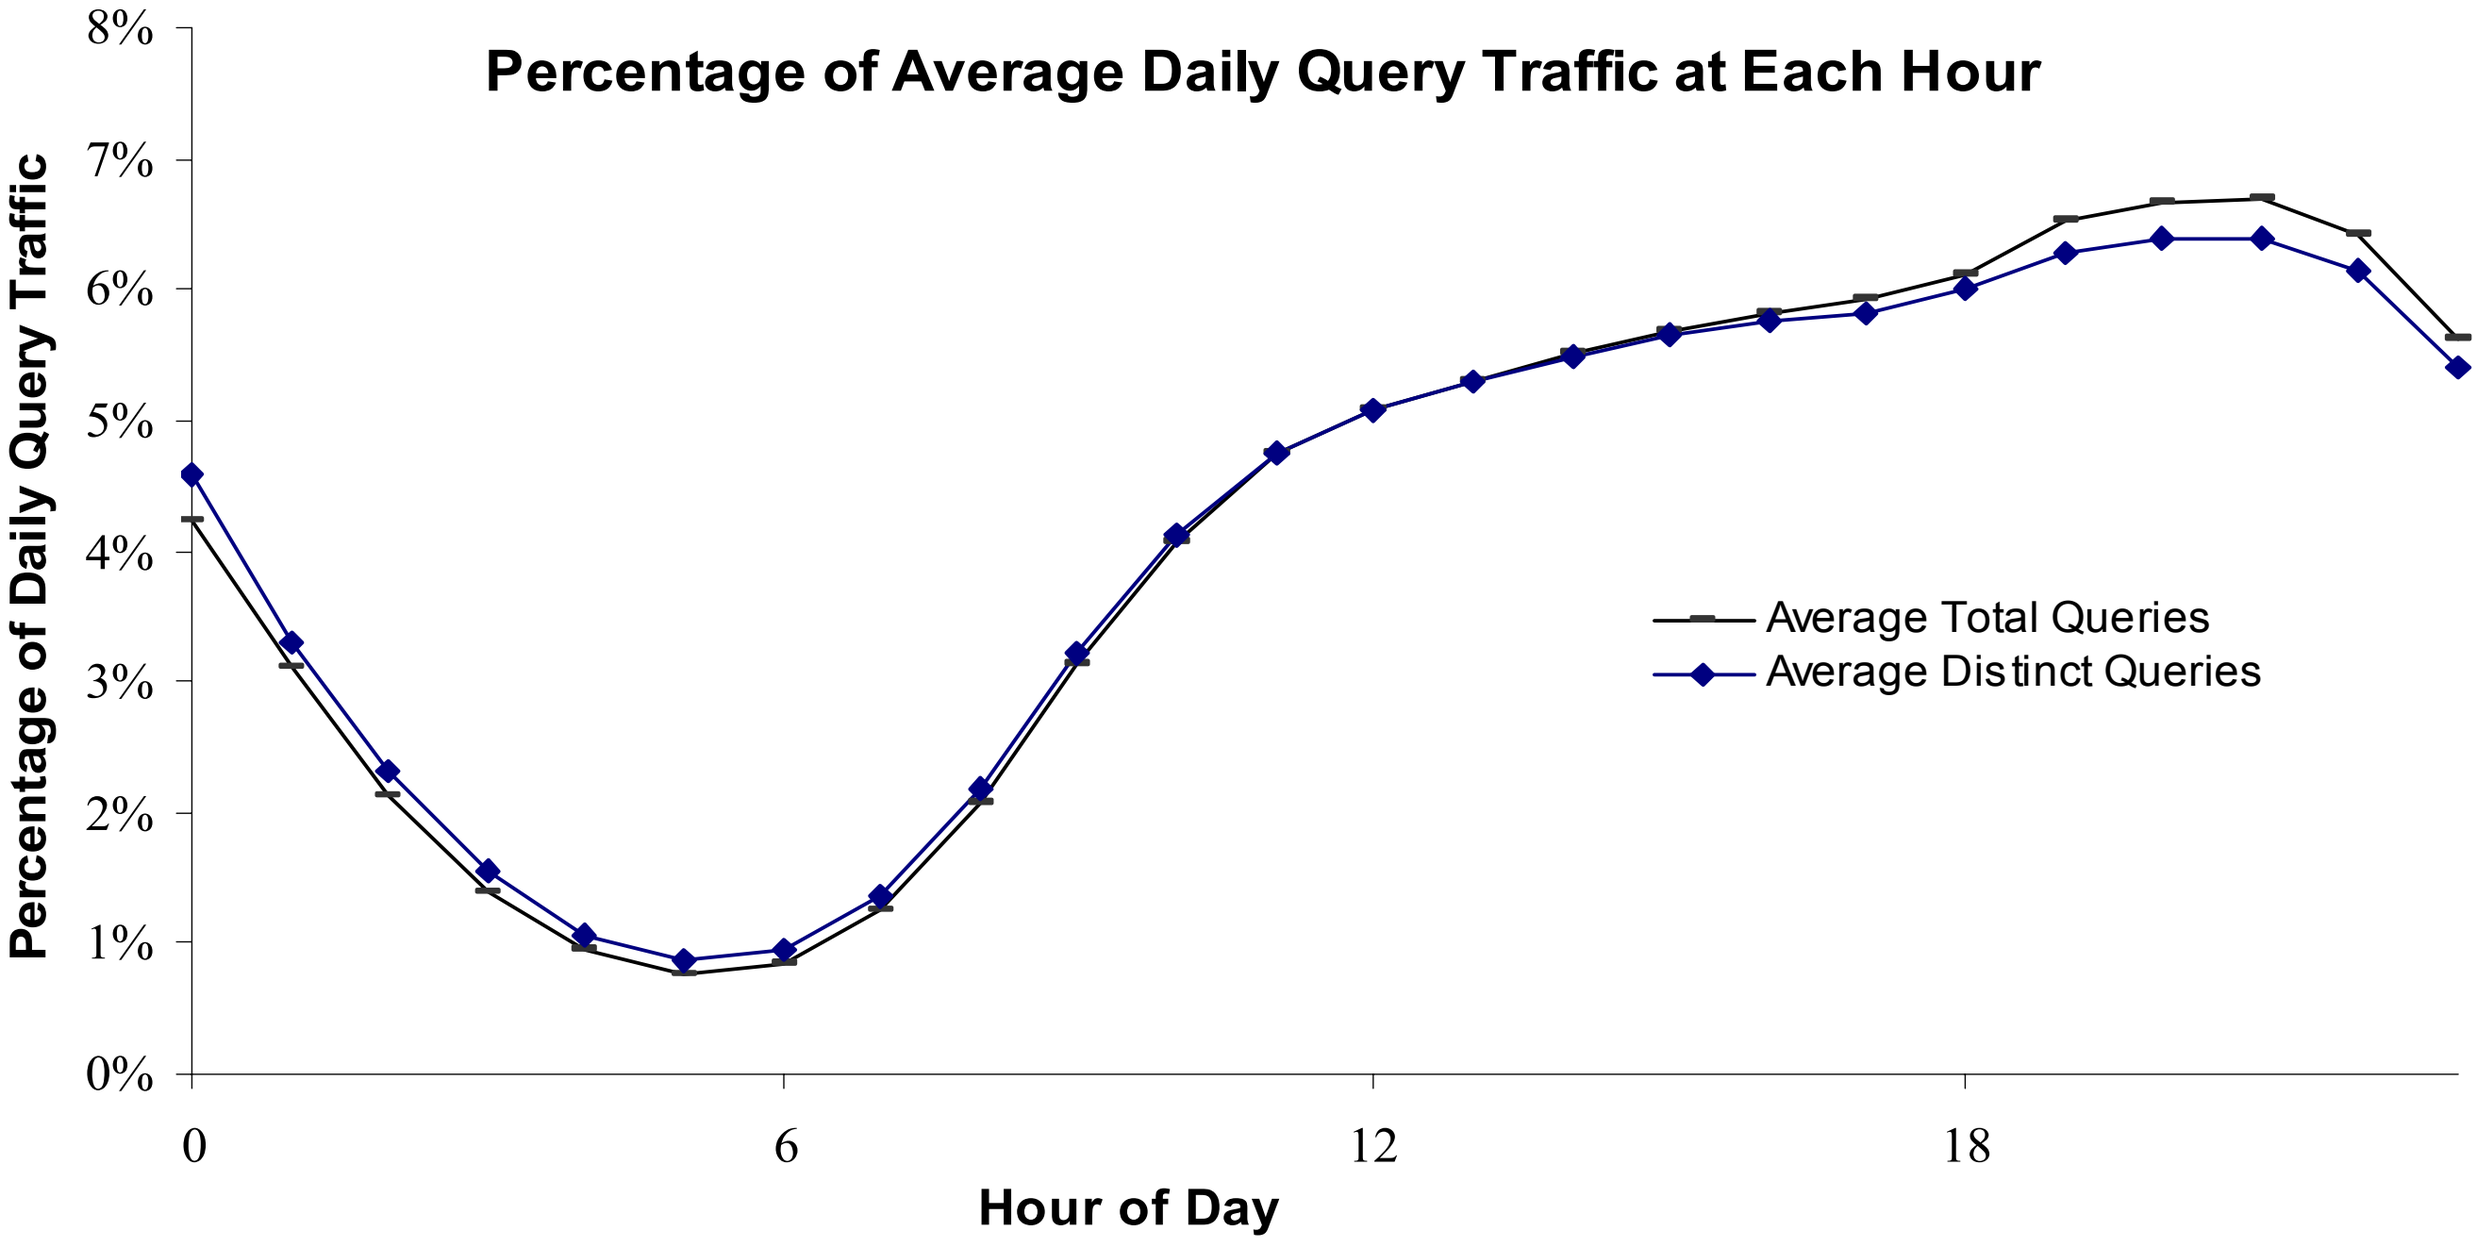

[Steven M. Beitzel, Eric C. Jensen, Abdur Chowdhury, David Grossman, Ophir Frieder. “Hourly Analysis of a Very Large Topically Categorized Web Query Log”. SIGIR 2004]

and how to take advantage of it ?

- "***Credit-card companies*** monitor every purchase and can identify ***fraudulent*** ones with a high degree of accuracy, using rules derived ***by crunching through billions of transactions***."

- "***Mobile-phone operators***, meanwhile, analyse ***subscribers' calling patterns*** to determine, for example, whether most of their frequent contacts are on a rival network."

- "***Retailers***, offline as well as online, are masters of data mining."

- "[…] mankind created 150 exabytes (billion gigabytes) of data in 2005. This year, it will create 1,200 exabytes. Merely keeping up with this flood, and storing the bits that might be useful, is difficult enough.
  Analysing it, to spot patterns and extract useful information, is harder still."

The Economist, Feb 2010

Cloud Computing is the solution !

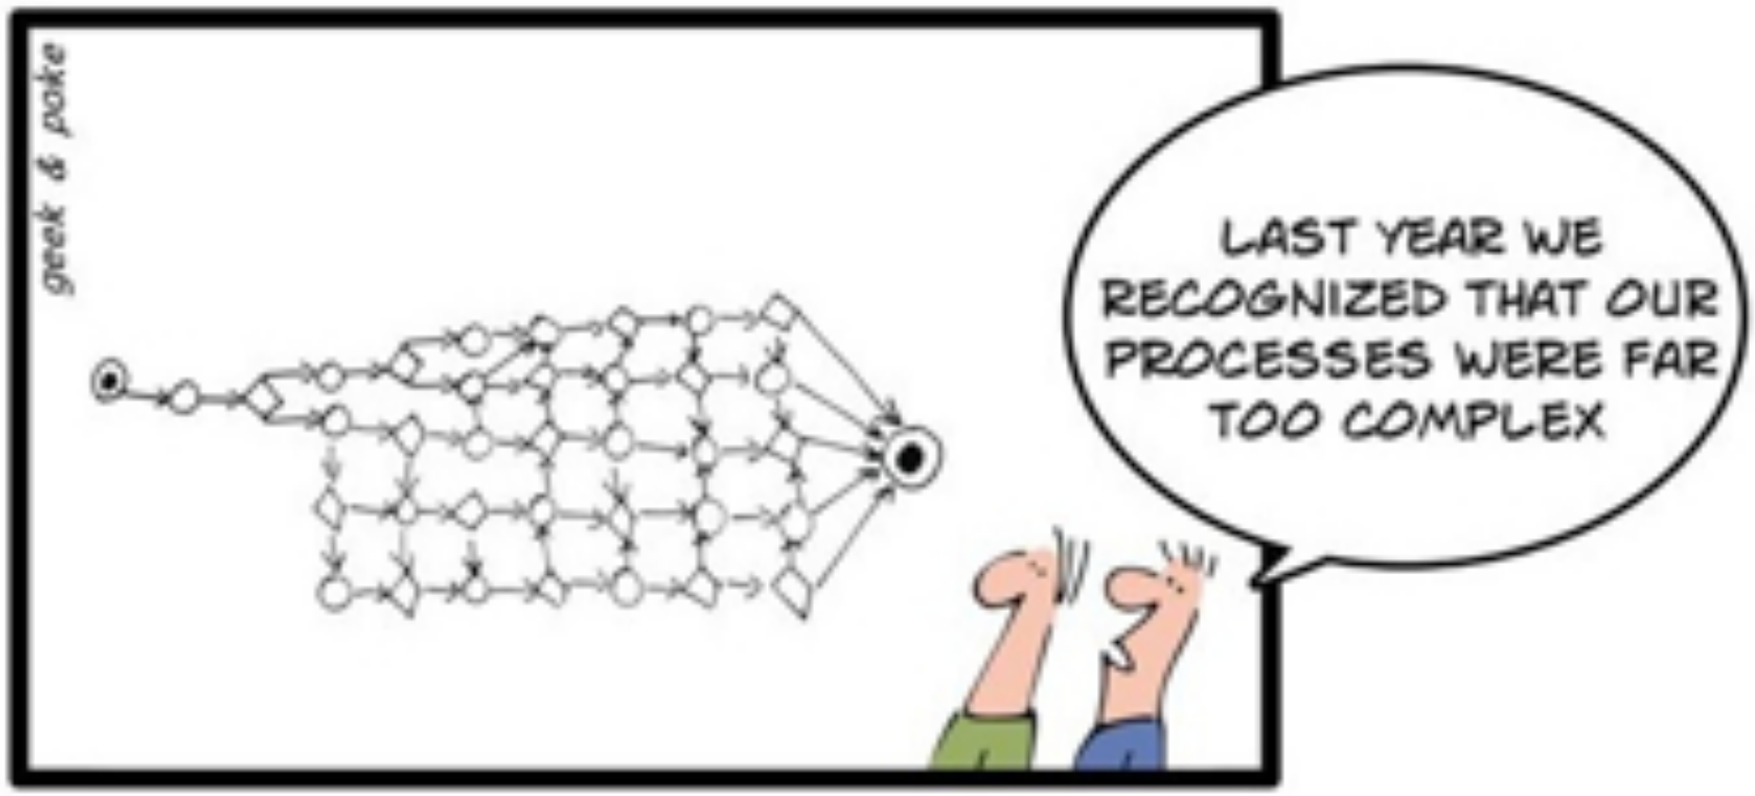

Cloud Computing is the solution !

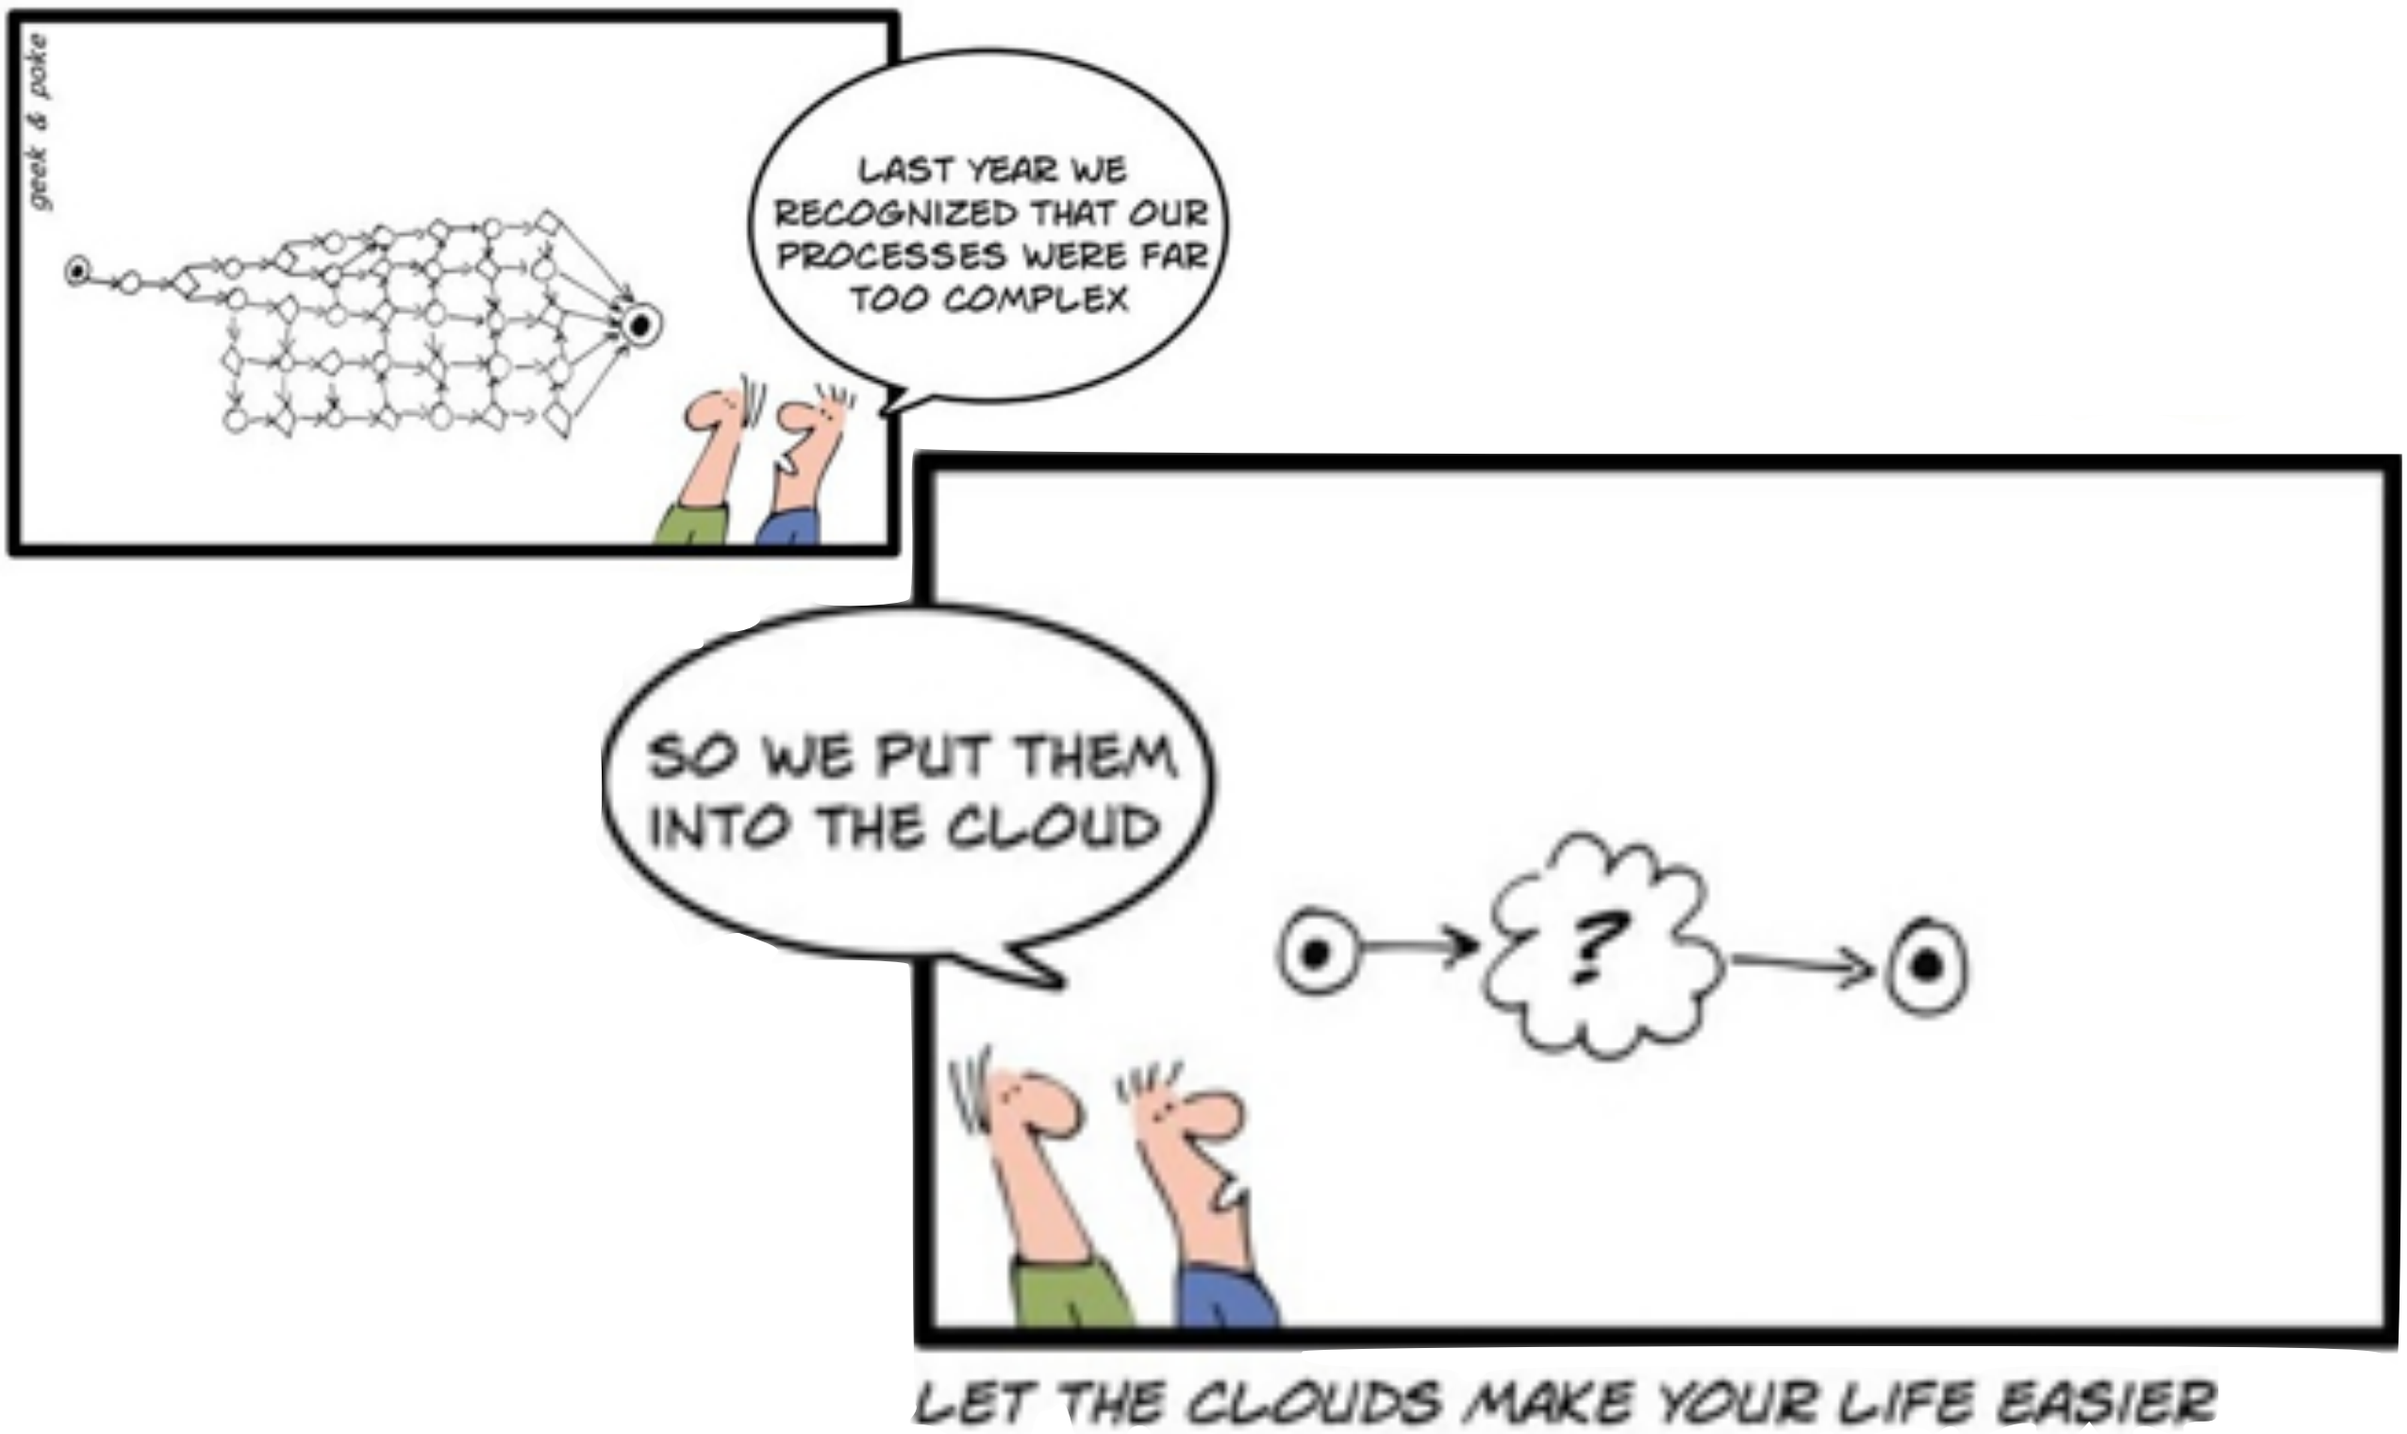

Old-fashioned Data/Computing Center

- glue together as many machines as possible like the following one:

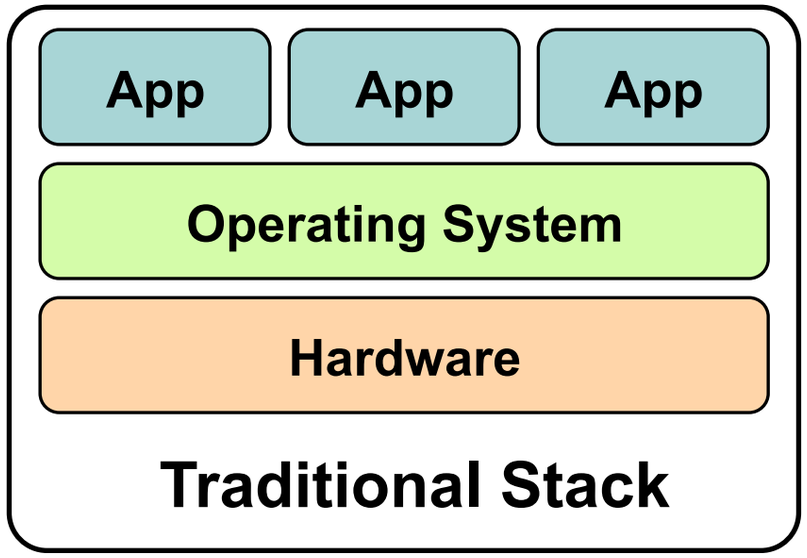

Old-fashioned Data/Computing Center

- glue together as many machines as possible like the following one:

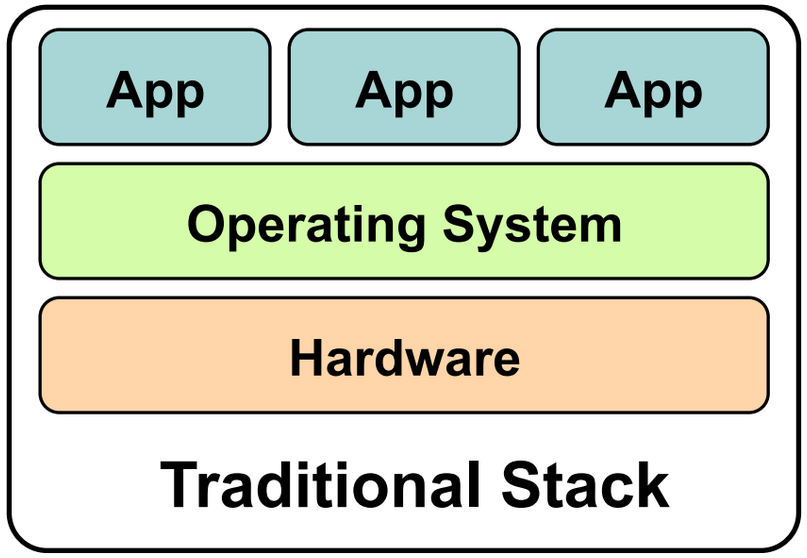

- and get a very large data center:

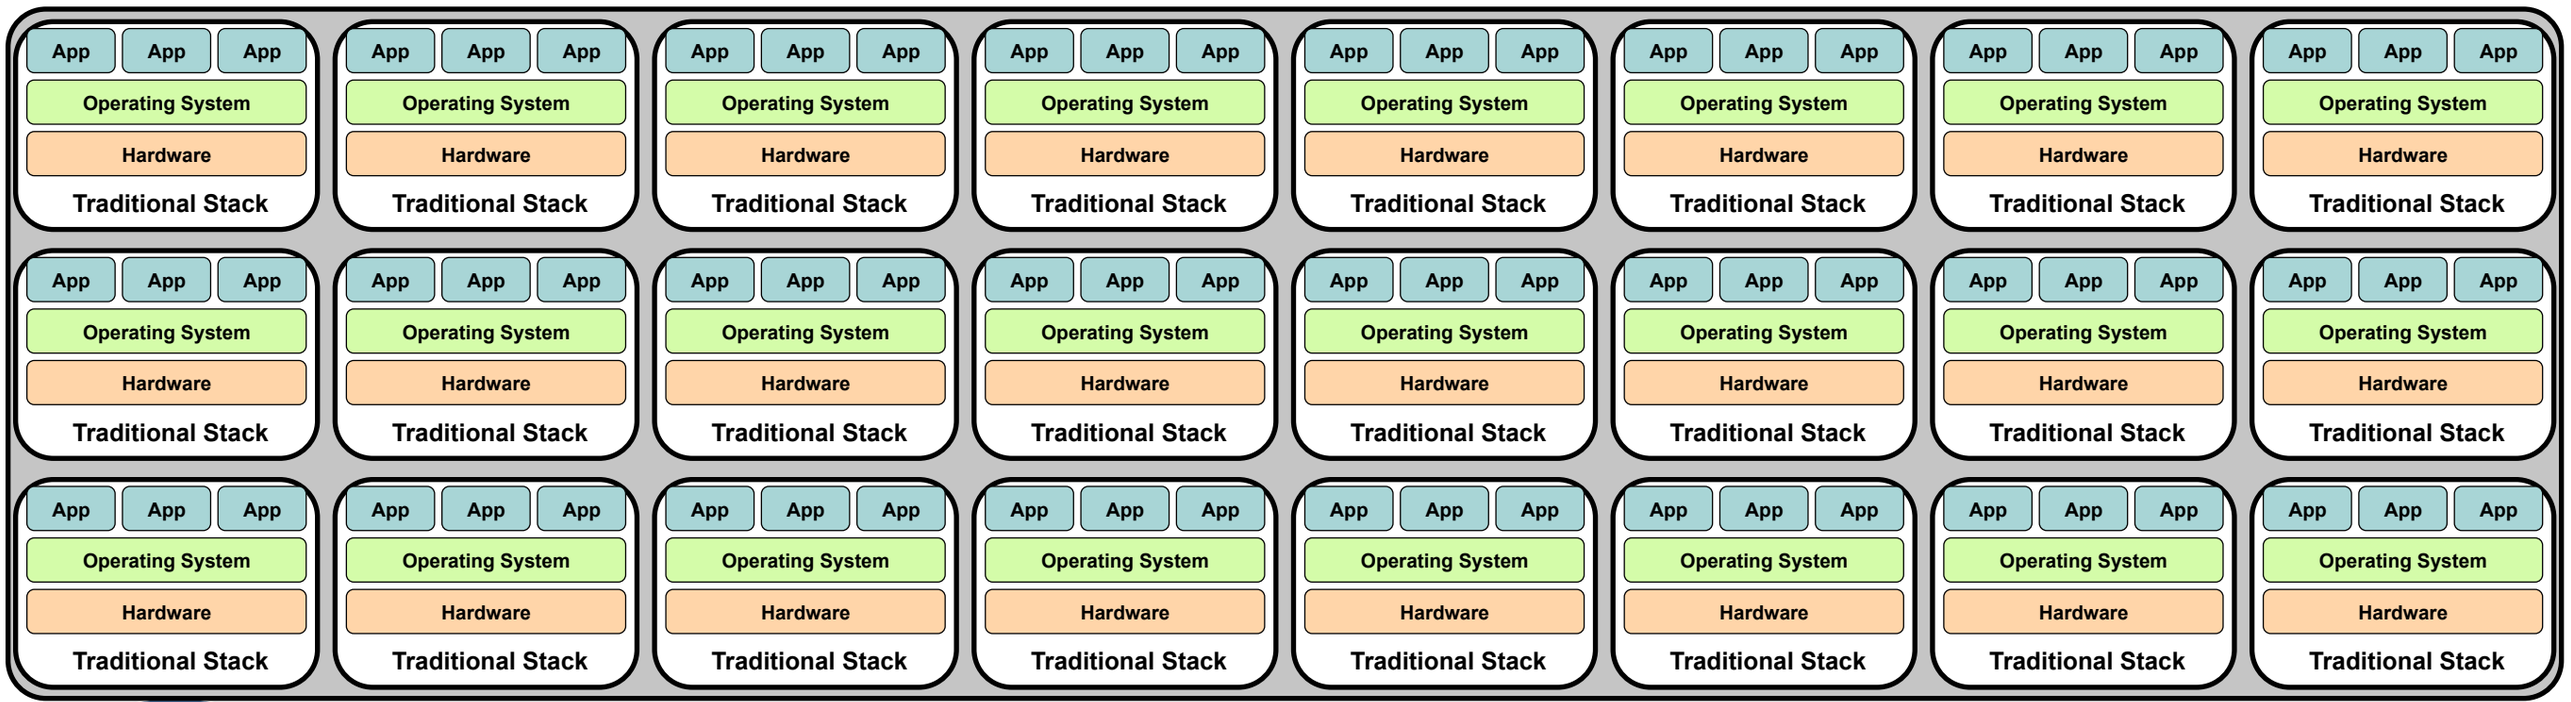

Old-fashioned Data/Computing Center

- The machine/server is the ***building block*** of the system:
  - you must to know ***how many*** servers you need
  - for each server you must specify its ***CPU, Memory, Disks, OS***, etc.
  - each machine as a ***specific role***: Web Server, Database Server, etc.
  - if a server is overloaded you must ***buy a new one*** and "link it"
  - if a server goes broken you must ***replace*** it
    - *what happens in the meanwhile ?*

- An example:
  - *How many machines to buy for your successful brand new application ??*

Capacity

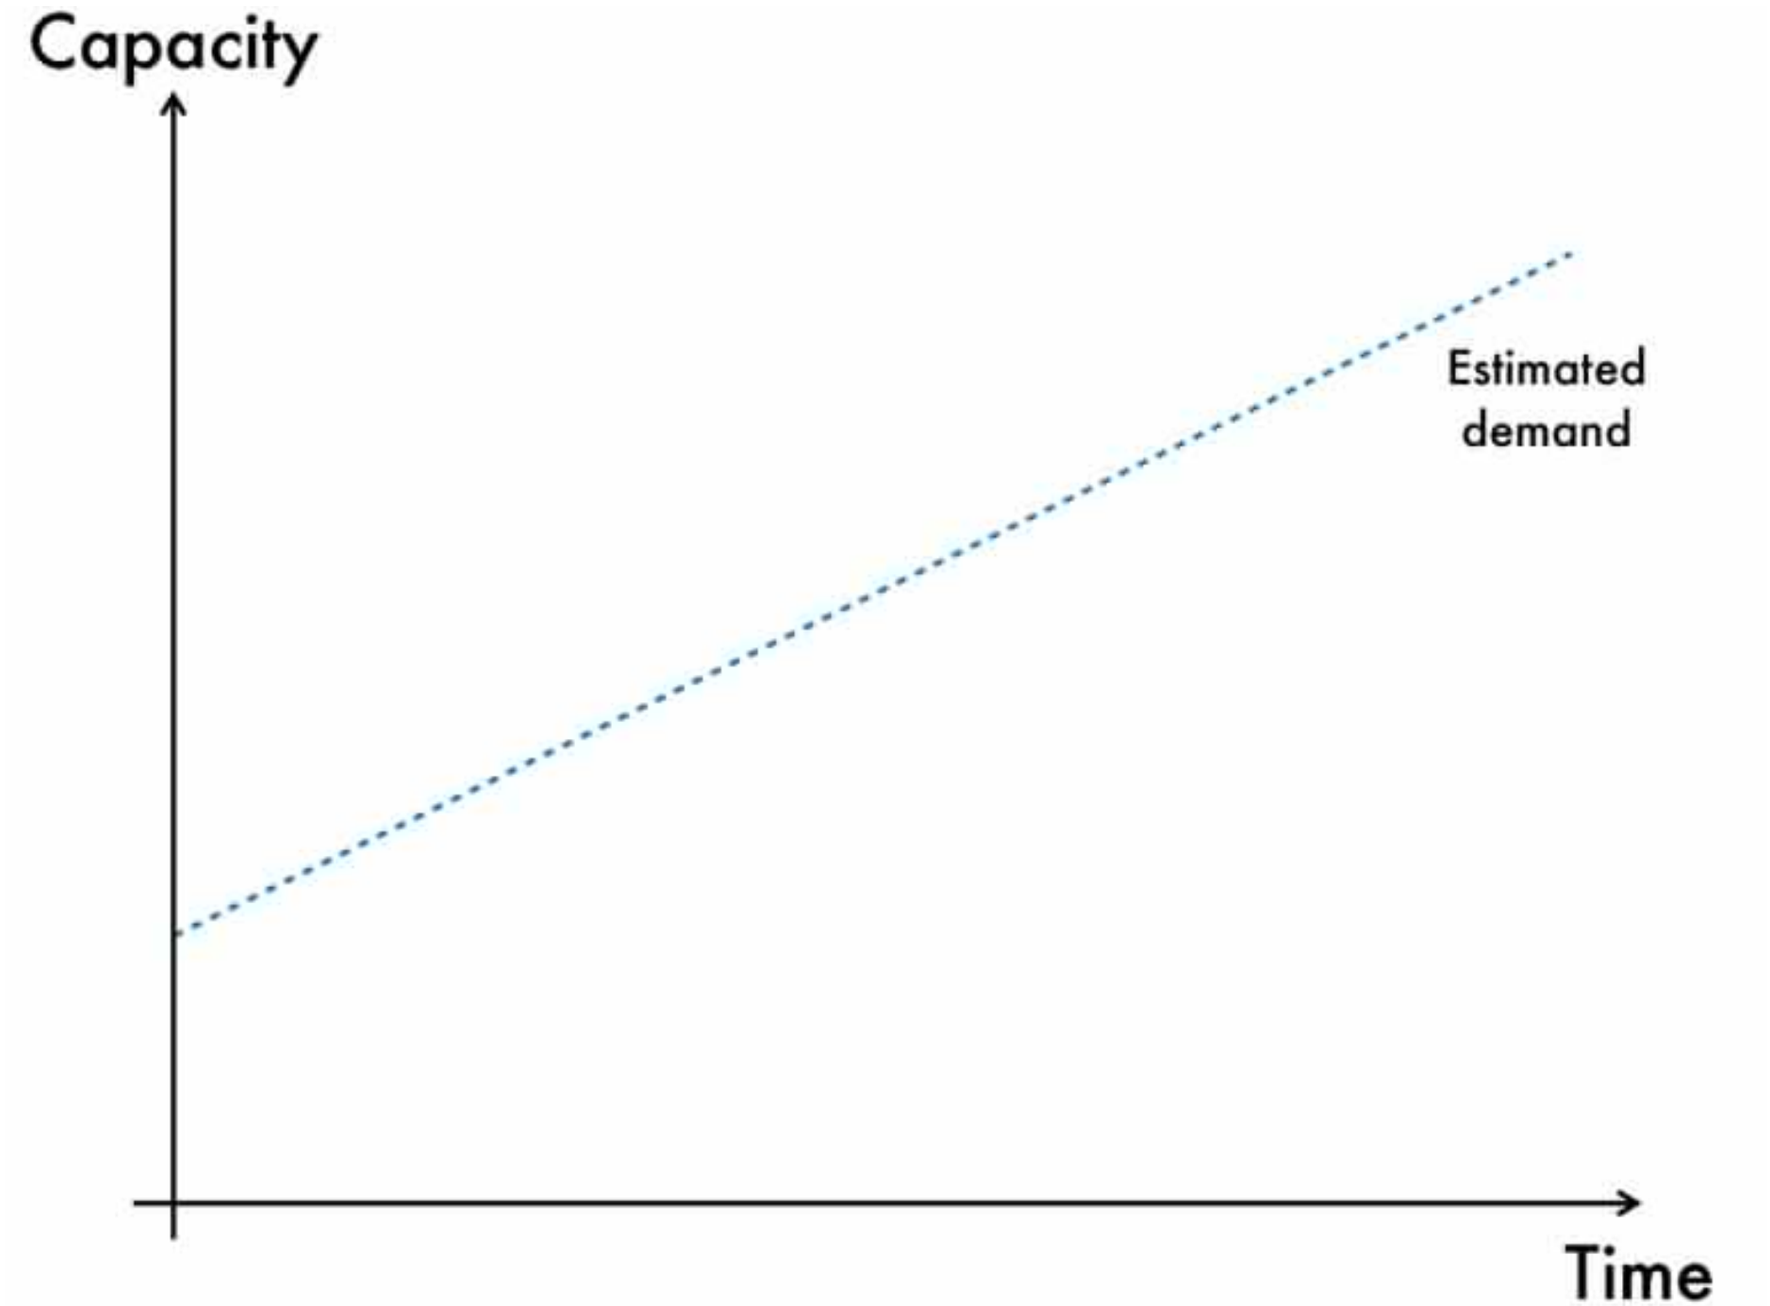

Capacity

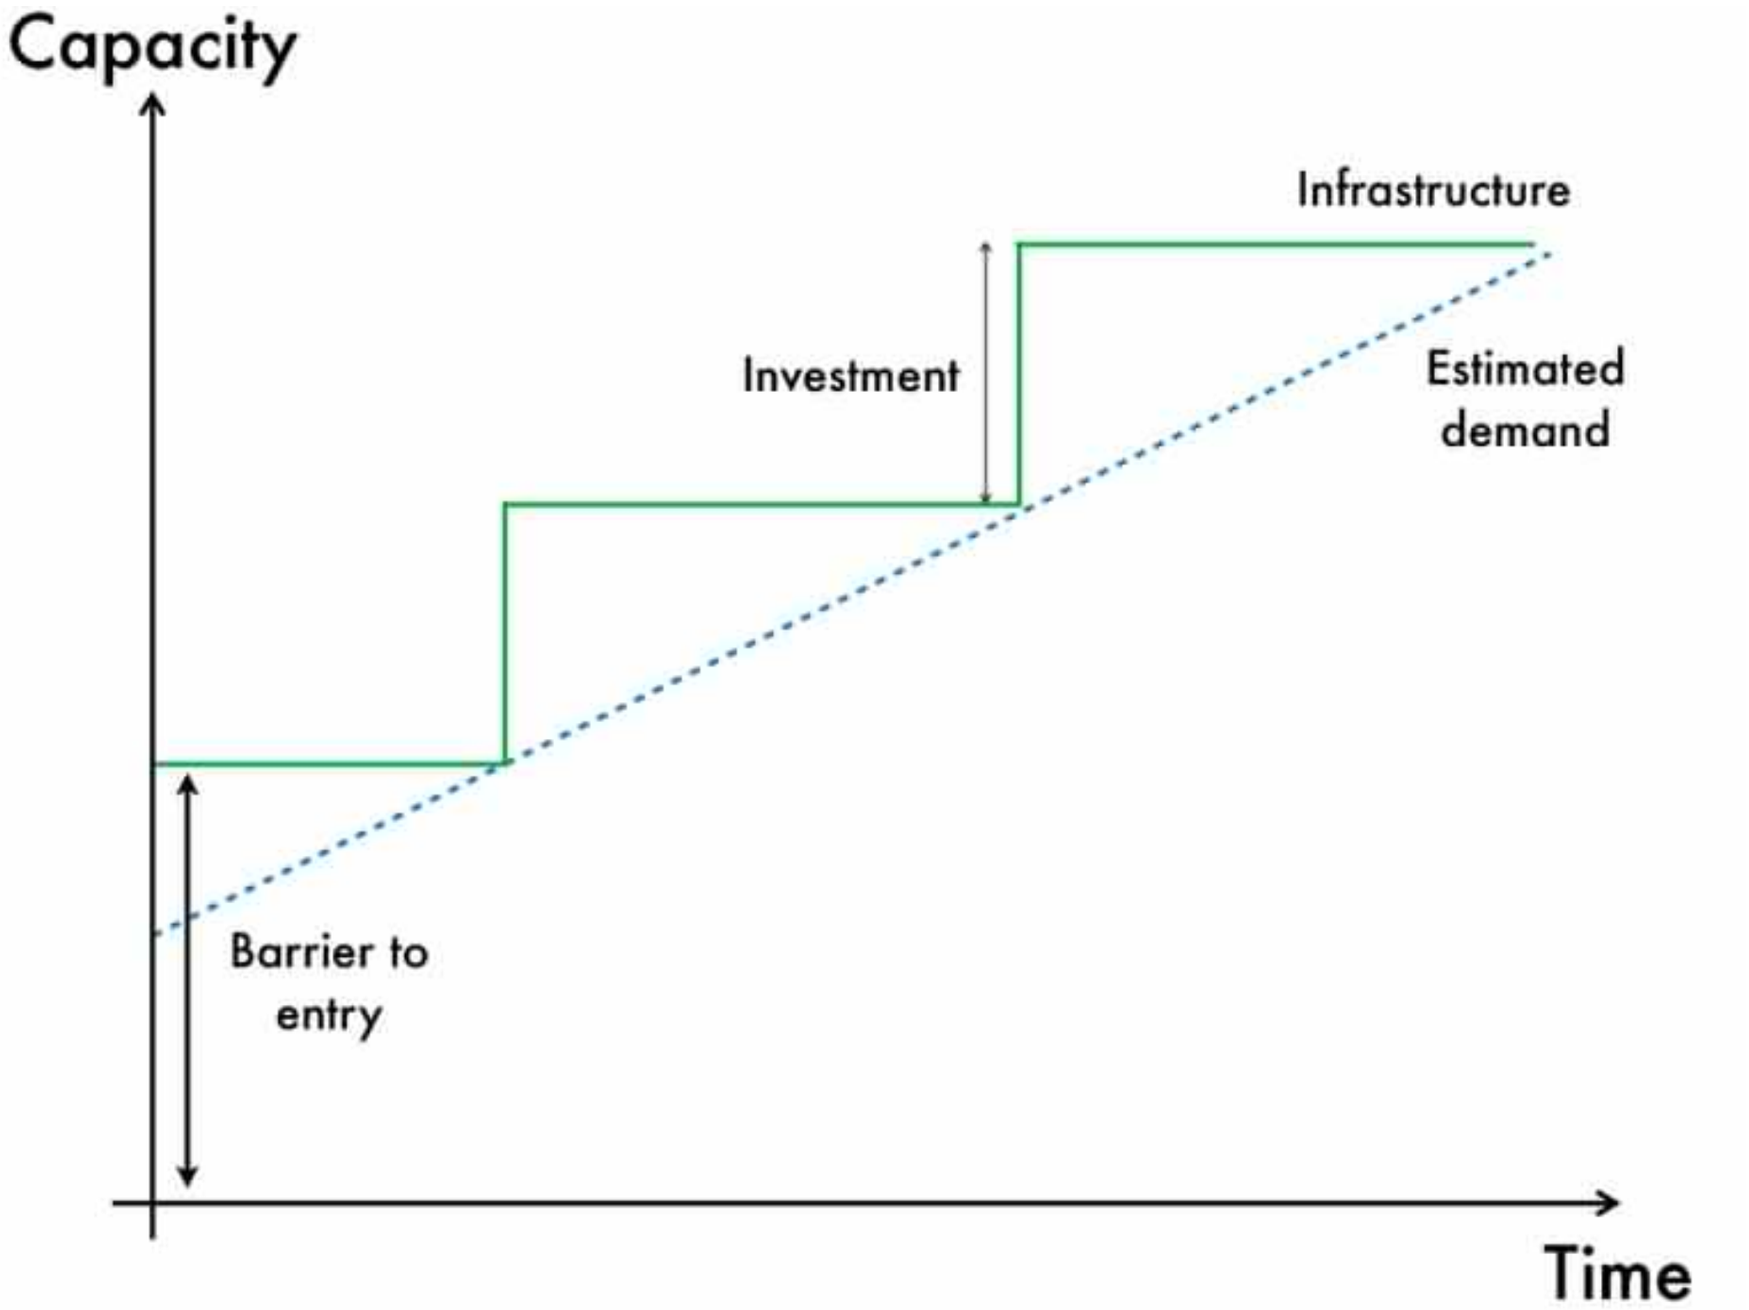
Capacity
Infrastructure
Investment
Estimated
demand
Barrier to
entry
Time

Capacity

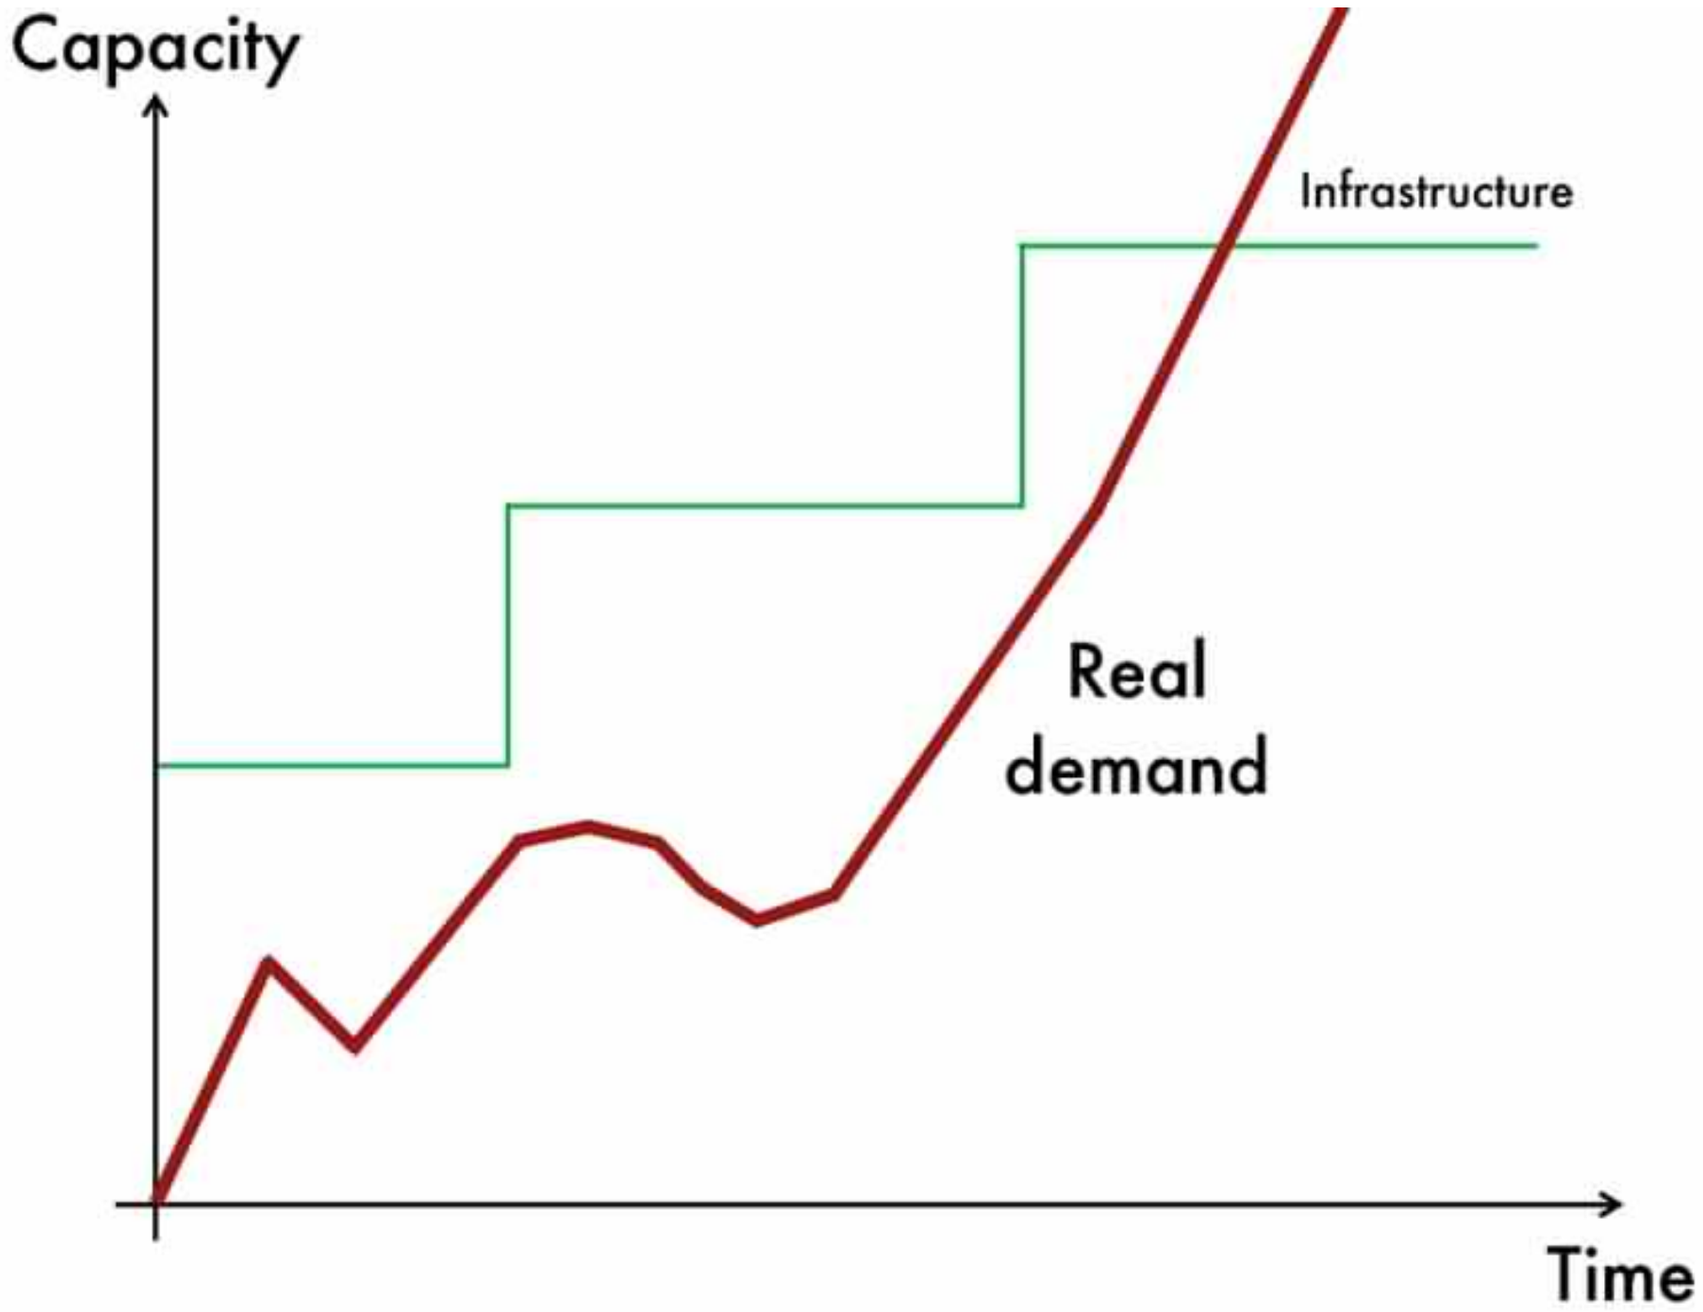

Cloud Computing

- Cloud computing comes from the convergence of:
  - ***service oriented architectures***
    - ... loose coupling of services with operating systems and technologies ...
  - ***parallel computing***
    - large scale data analysis, up to thousands of machines
  - ***virtualization***
    - independence from physical hardware

> Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. (NIST)

http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Virtualization

- **Virtual Machine (VM)**
  - Abstraction of a physical host machine
    - with a specific hardware, OS, etc.
  - Hypervisor
    - de-couples the execution environment of the application from the physical hardware
    - intercepts and emulates instructions from VMs, and allows management of VMs
  - It is a well established technology:
    - Software: VMWare, Xen, etc.
    - Hardware: multicore architectures, Intel VT, AMD-V

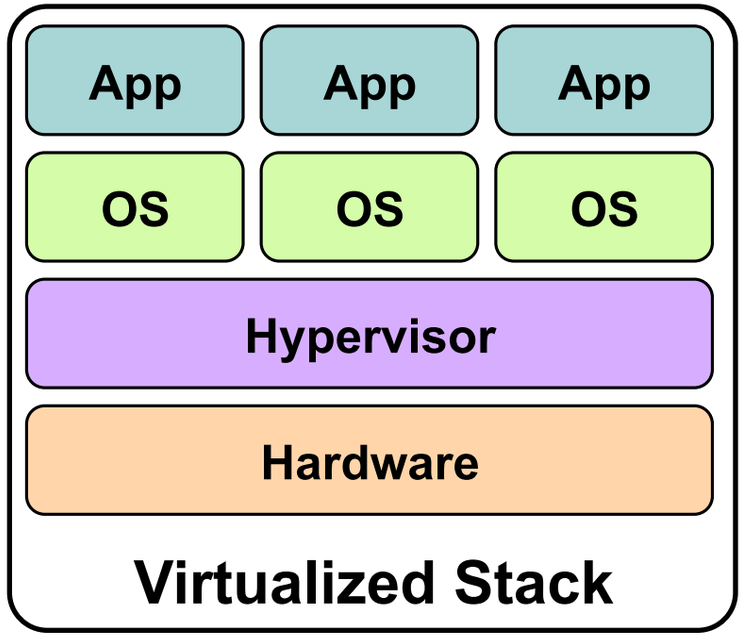

Virtualization

- **Key concepts**:
  - A VM image is stored as a simple file
    - includes: OS, hw configuration, application, etc.
  - ***A VM machine can be executed by any physical hardware***
  - The same VM image can be instantiated multiple times
  - A single machine can host multiple VMs
  - A VM can be seamlessly moved from physical machine to another without interrupting its execution

- **Elasticity:**
  - in case of hardware failure, the hosted VMs can be instantiated from scratch on a different hardware
  - new machines can be added at will at any time
  - when applications are truly designed to run in a distributed environment, they can fully exploit the available resources by tuning their allocation

Should I set up my own cloud ?

- Yes, if you are Google or Microsoft ...

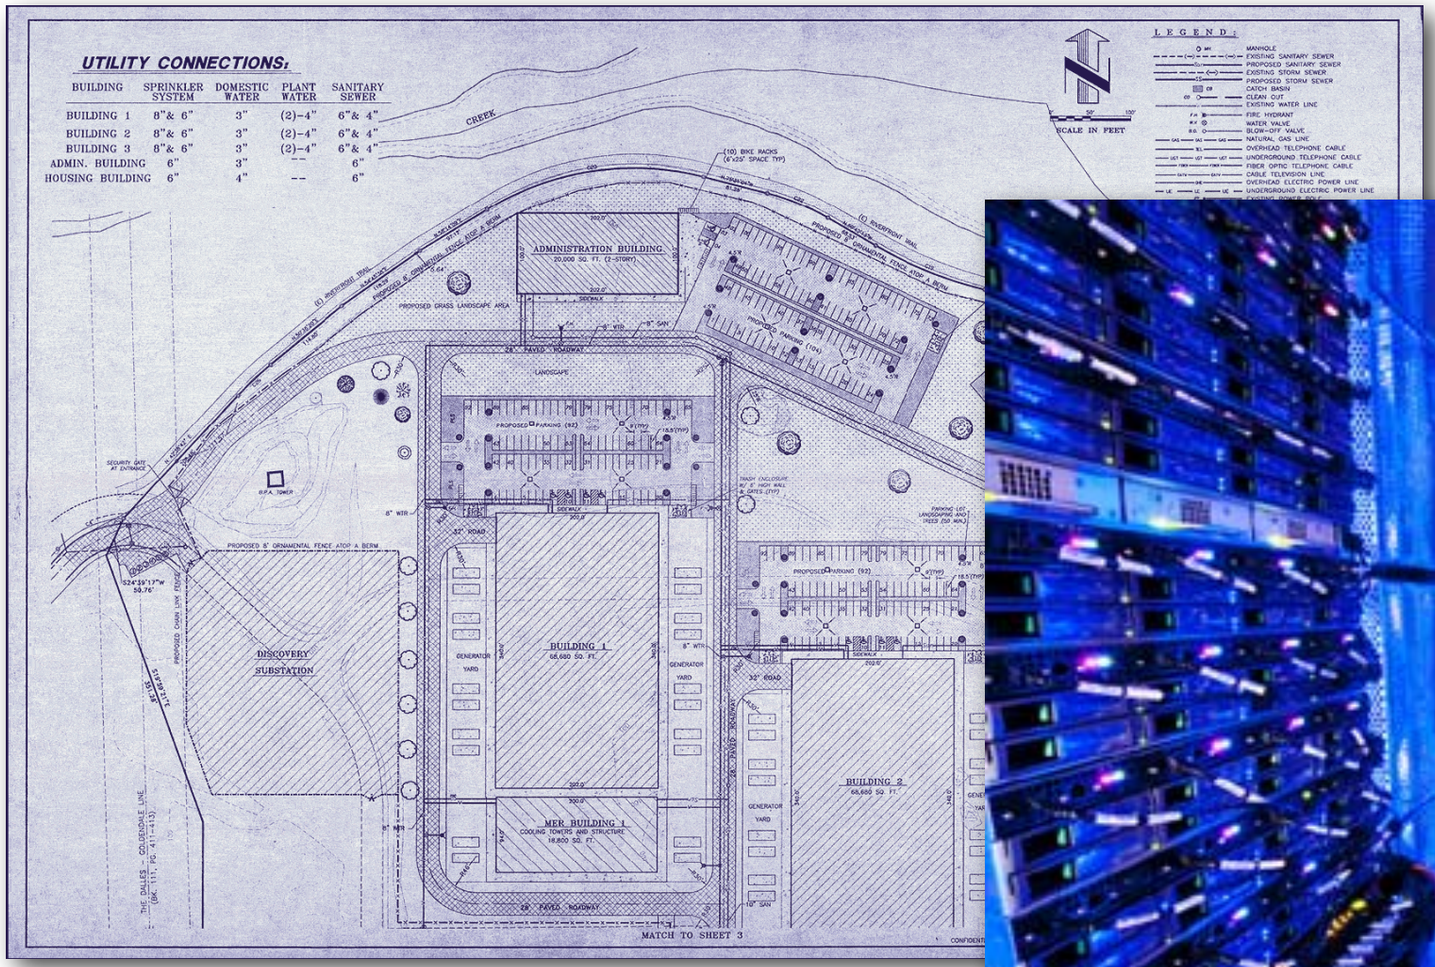

Oregon Google data center blueprint
http://harpers.org/archive/2008/03/0081946

A 40-foot container inside the new Microsoft Chicago data center
http://www.forbes.com/forbes/2009/1116/outfront-ibm-cloud-microsoft-new-cloud-computing.html

Should I set up my own cloud ?

- Or, you can ***pay-as-you-go*** from a cloud provider

- An example of an Amazon EC2 instance:
  - **Large Instance**: 7.5 GB of memory, 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each), 850 GB of local instance storage ..... €0.36 per hour
  - **Data Transfers IN**..... €0.00
  - **Data Transfers OUT**: 1 GB free, ... €0.12 per GB
  - **Elastic Block Store**: ..... €0.11 GB/month
    ..... €0.11 per Mi. I/O req.s
-

Should I set up my own cloud ?

- Or, you can ***pay-as-you-go*** from a cloud provider

- An example of an Amazon EC2 instance:
  - **Large Instance**: 7.5 GB of memory, 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each), 850 GB of local instance storage ..... €0.36 per hour
  - **Data Transfers IN**..... €0.00
  - **Data Transfers OUT**: 1 GB free, ... €0.12 per GB
  - **Elastic Block Store**: ..... €0.11 GB/month
    ..... €0.11 per Mi. I/O req.s

- ***You do <u>not</u> pay for***:
  - bits, bolts, cables, power consumption, physical space, technical staff

- *You can define rules to rent new machines as needed automatically !*

Cloud Elasticity

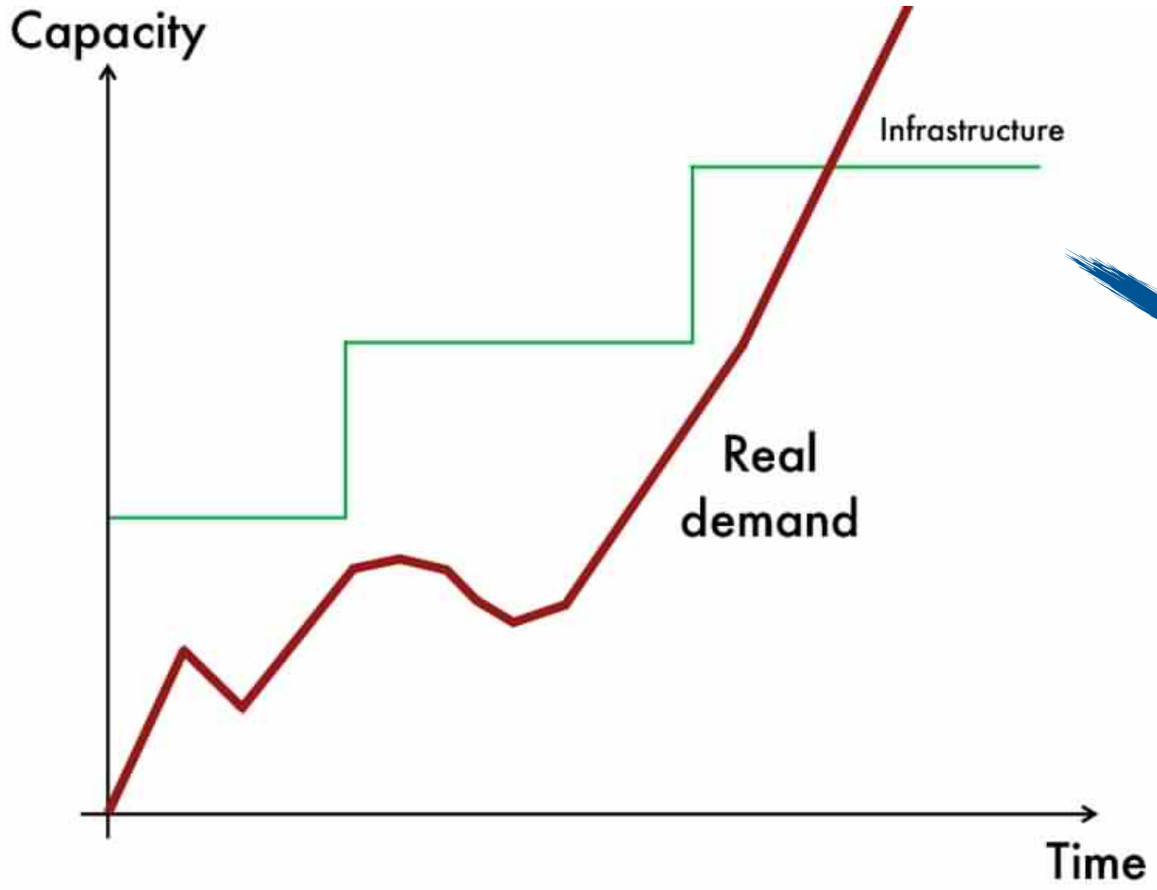

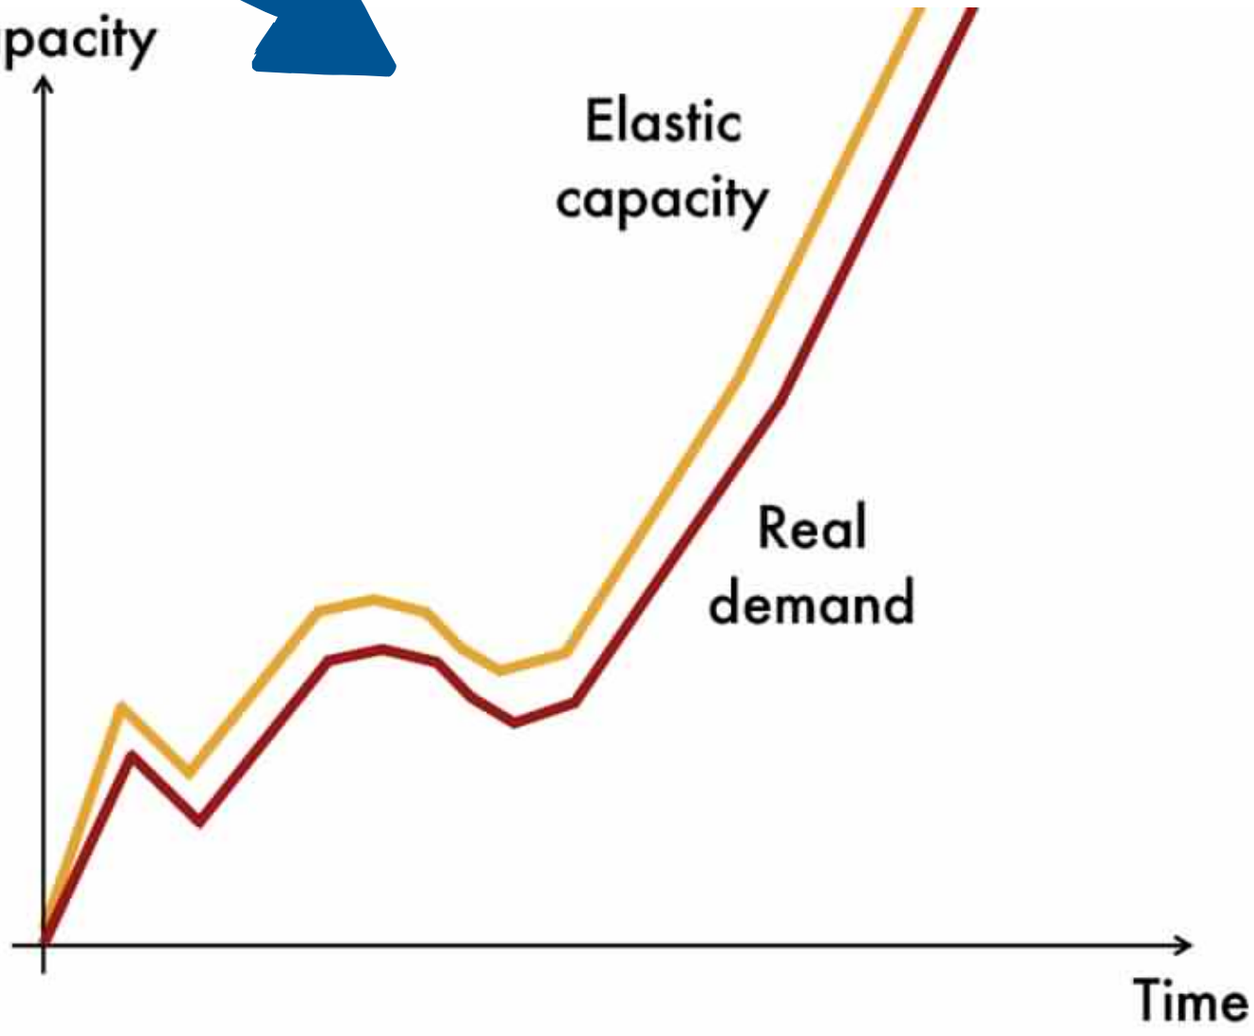

- Little entrance cost
- Adaptive capacity
  - Scale-down is as important as scale-up

The big players

Service Delivery Models (I)

- **Infrastructure as a service (IaaS)**
  - This is the lowest abstraction layer
  - Users ***rent a bunch of virtual machines***
- **Key concepts**:
  - Users are in charge of configuring the VMs, as for physical machines
  - Users do not know where these VMs are running
  - Users do not know how many physical machines will actually be used
    - but some quality of service may be provided
- **Benefits**:
  - Elasticity: users can rent more/less VMs as needed
  - Reliability: hardware is replaced and maintained as needed transparently
  - Pay-as-you-go: users can switch on/off any VM, only up-time is charged
- **IaaS providers**:
  - Amazon, Rackspace, GoGrid, CityCloud, etc.

Service Delivery Models (II)

- **Platform as a service (PaaS)**
  - You rent an ***solution stack***
- **Key concepts:**
  - "Unlimited machine" from the user perspective
  - The user most develop applications according to the cloud provider API
- **Benefits:**
  - Automatic scaling
  - Automatic load balancing
  - Reliable storage storage

- **Storage solutions**:
  - **Amazon DynamoDB**:
    - "Amazon DynamoDB is a fully managed *NoSQL database service that provides fast and predictable performance with seamless scalability*. DynamoDB automatically spreads the data and traffic [...] over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent, fast performance."

Service Delivery Models (II)

- **Computational solutions**:
  - Hadoop MapReduce:
    - "Hadoop MapReduce is a programming model and software framework for writing *applications that rapidly process vast amounts of data in parallel on large clusters of compute nodes*."

- **Application environments**:
  - **Google App Engine**:
    - "Applications run in a secure environment that provides limited access to the underlying operating system. These limitations allow App Engine to *distribute web requests for the application across multiple servers, and start and stop servers to meet traffic demands*. The sandbox isolates your application in its own secure, reliable environment that is independent of the hardware, operating system and physical location of the web server."

  - **Microsoft Azure**:
    - The Windows Azure platform uses a specialized operating system, called Windows Azure, to run its "fabric layer" — a cluster hosted at Microsoft's data centers that *manages computing and storage resources of the computers and provisions the resources* (or a subset of them) to applications running on top of Windows Azure.

Service Delivery Models (III)

- **Software as a service (SaaS)**
  - This is the highest level of abstraction
  - You rent a ***capability***
- **Key Concepts**:
  - A full-fledged service running on a cloud platform is provided
- **Benefits**:
  - No need to install/configure/upgrade the service software stack
  - Seamless availability

- **Google Docs (now Google Drive)**:
  - "[...] is a free, Web-based office suite and data storage service [...]. It allows users to create and edit documents online while collaborating in real-time with other users."

- **Dropbox**:
  - "Dropbox is a Web-based file hosting service [...] that uses networked storage to enable users to store and share files and folders with others across the Internet using file synchronization."

- **Spotify**:
  - "Spotify is a music streaming service offering digitally restricted streaming of selected music from a range of major and independent record labels"

Clouds users

- "The ***New York Times*** a few years ago used cloud computing and Hadoop to convert over ***400,000 scanned images*** from its archives, from 1851 to 1922. By harnessing the power of hundreds of computers, it was able to do the job ***in 36 hours***."

- "***Visa***, a credit-card company, in a recent trial with Hadoop crunched two years of test records, or ***73 billion transactions***, amounting to 36 terabytes of data. The processing time fell from one month with traditional methods to a mere ***13 minutes***."

The Economist, Feb 2010

- ***1/3 of all Internet users will access an Amazon AWS cloud site on average at least once a day***"

http://blog.deepfield.net/

Cloud users

- **Dryad**
  - Written at Microsoft Research, Silicon Valley
  - Deployed since 2006
  - Running 24/7 on >> 104 machines
  - Sifting through > 10Pb data daily
  - Clusters > 3000 machines
  - Jobs with > 105 processes each
  - Platform for rich software ecosystem

- Virtualized 2-D pipeline:

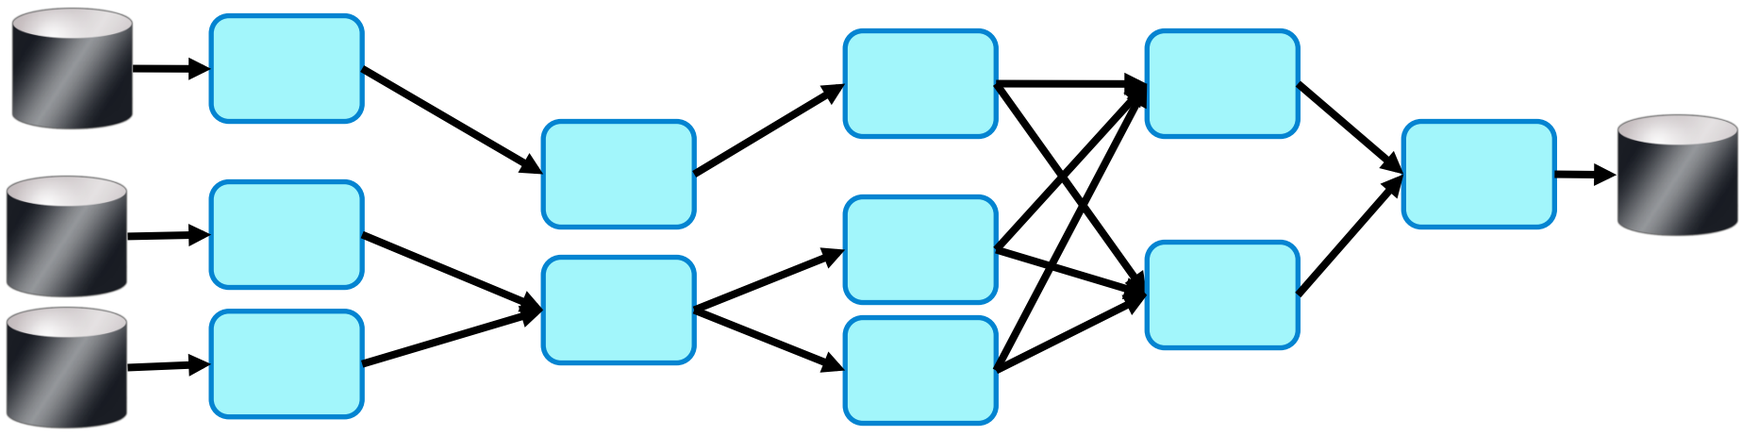

Cloud users

- **Dryad**
  - Written at Microsoft Research, Silicon Valley
  - Deployed since 2006
  - Running 24/7 on >> 104 machines
  - Sifting through > 10Pb data daily
  - Clusters > 3000 machines
  - Jobs with > 105 processes each
  - Platform for rich software ecosystem

- Virtualized 2-D pipeline:

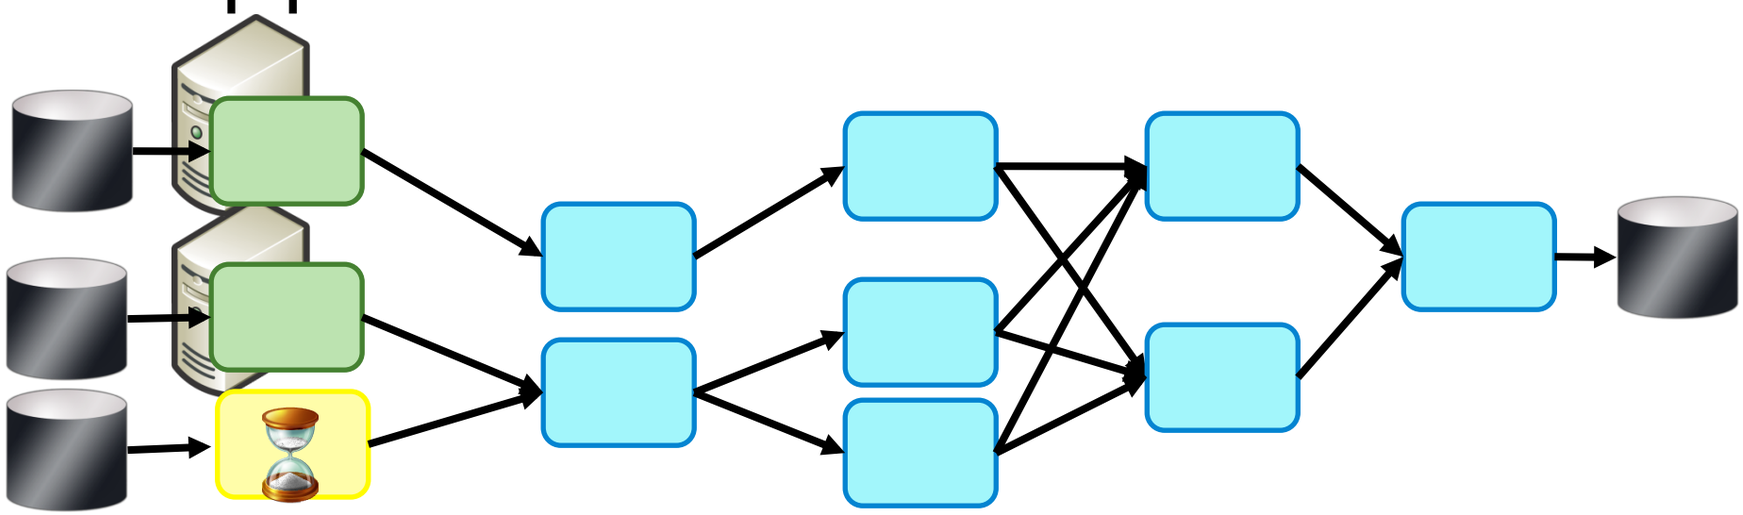

Cloud users

- **Dryad**
  - Written at Microsoft Research, Silicon Valley
  - Deployed since 2006
  - Running 24/7 on >> 104 machines
  - Sifting through > 10Pb data daily
  - Clusters > 3000 machines
  - Jobs with > 105 processes each
  - Platform for rich software ecosystem

- Virtualized 2-D pipeline:

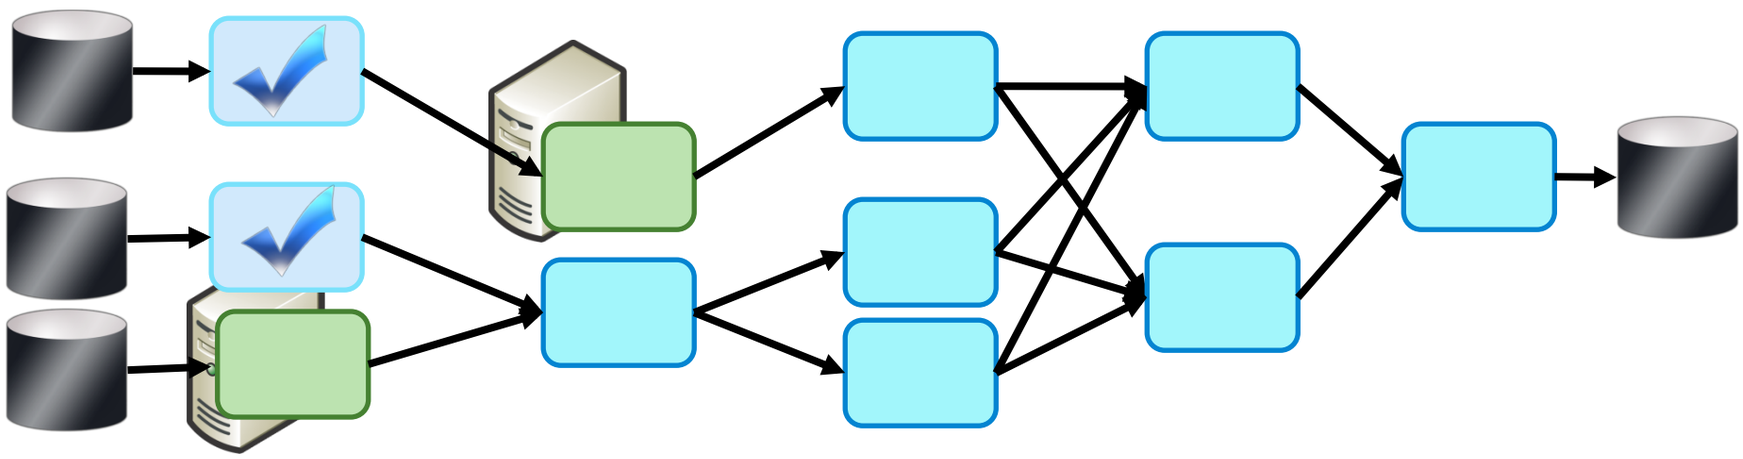

Cloud users

- **Dryad**
  - Written at Microsoft Research, Silicon Valley
  - Deployed since 2006
  - Running 24/7 on >> 104 machines
  - Sifting through > 10Pb data daily
  - Clusters > 3000 machines
  - Jobs with > 105 processes each
  - Platform for rich software ecosystem

- Virtualized 2-D pipeline:

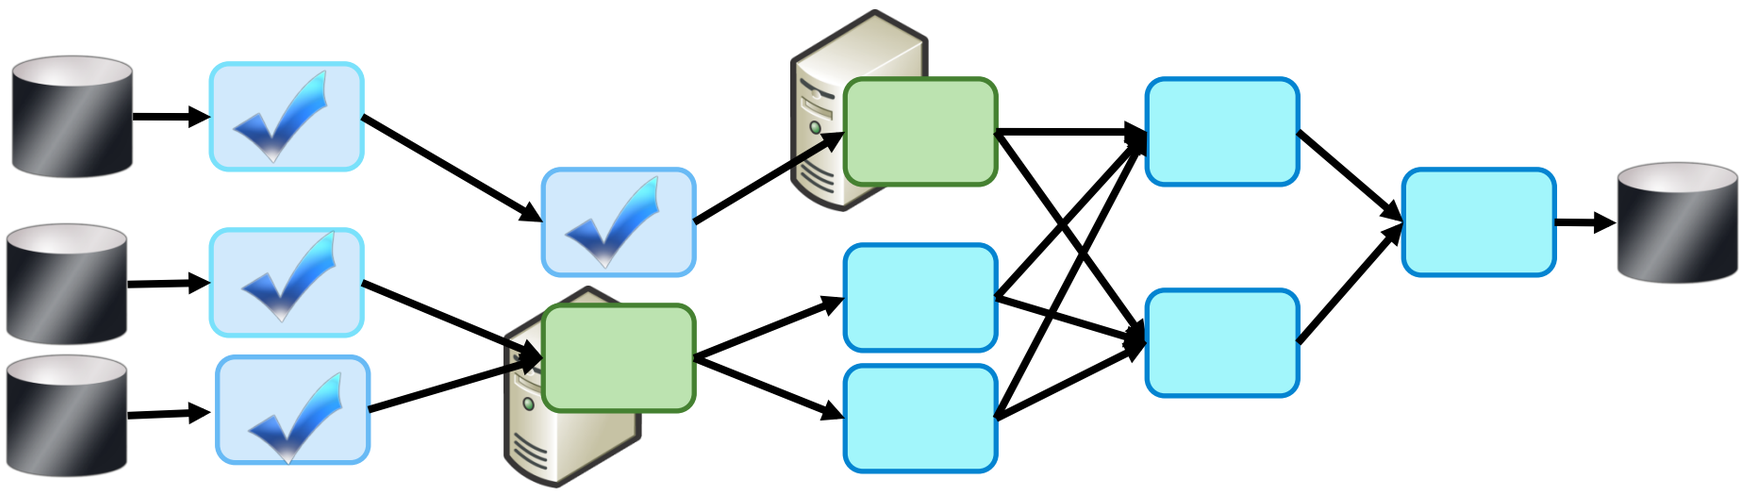

Cloud users

- **Dryad**
  - Written at Microsoft Research, Silicon Valley
  - Deployed since 2006
  - Running 24/7 on >> 104 machines
  - Sifting through > 10Pb data daily
  - Clusters > 3000 machines
  - Jobs with > 105 processes each
  - Platform for rich software ecosystem

- Virtualized 2-D pipeline:

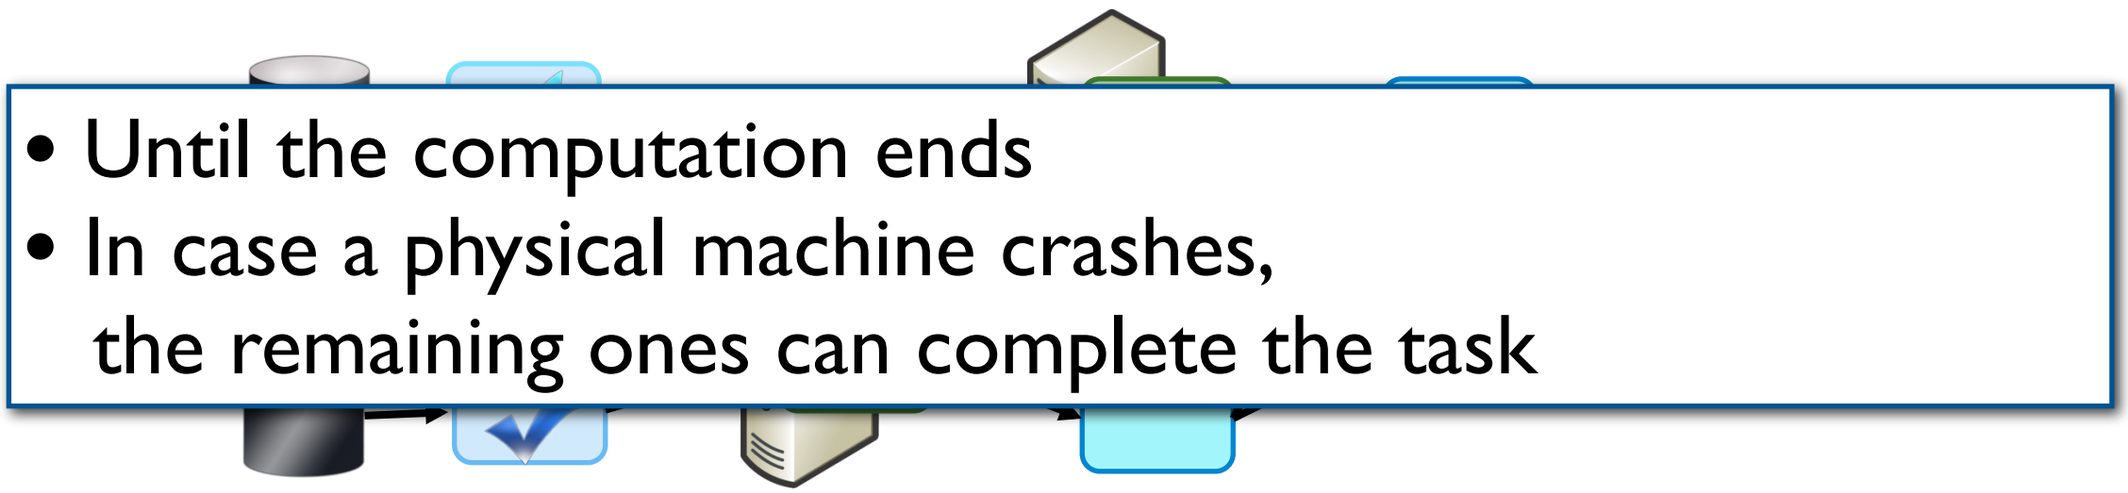

Dryad was used for …

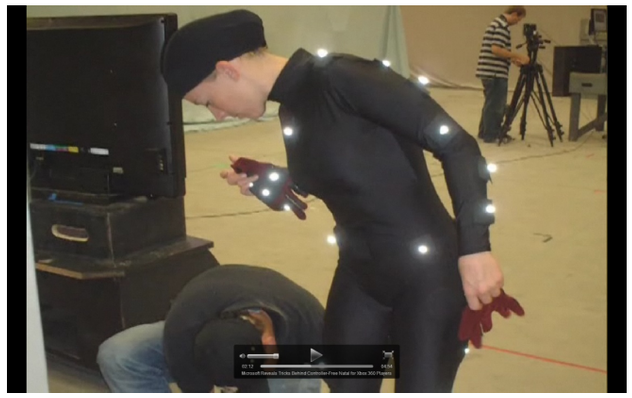

Motion Capture
(ground truth)

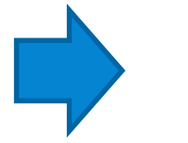

Rasterize

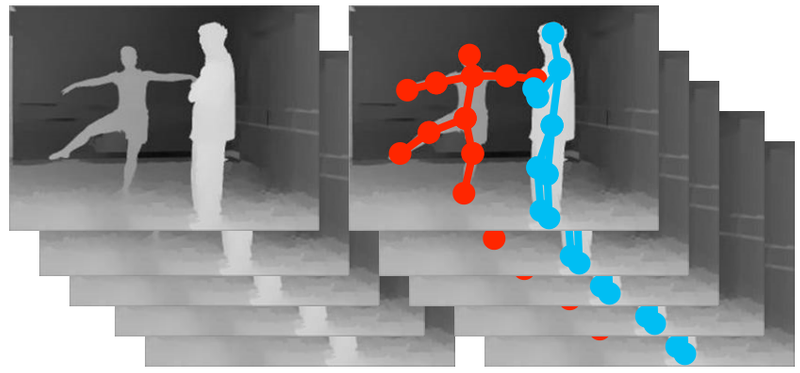

Training examples

Machine
learning

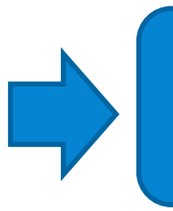

- >10^22 objects
- Sparse, multi-dimensional data structures
- Complex datatypes
  (images, video, matrices, etc.)
- Complex application logic and dataflow
  - >35000 lines of .Net
  - 140 CPU days
  - >10^5 processes
  - 30 TB data analyzed
  - 140 avg parallelism (235 machines)
  - 300% CPU utilization (4 cores/machine)

Dryad was used for …

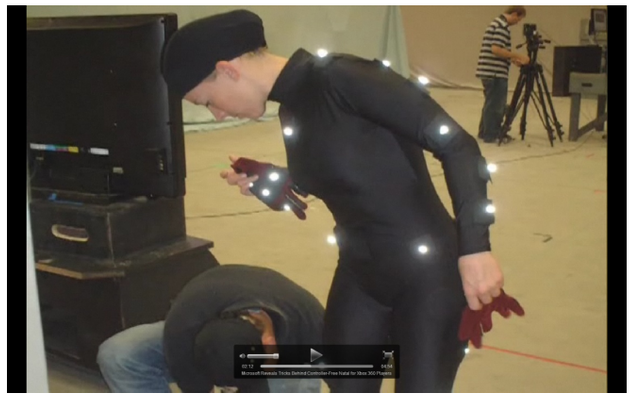

Motion Capture
(ground truth)

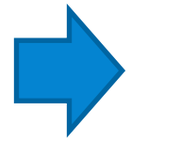

Rasterize

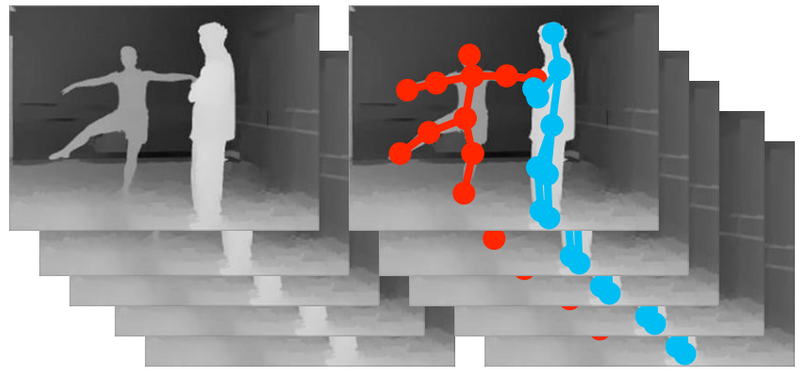

Training examples

Machine learning

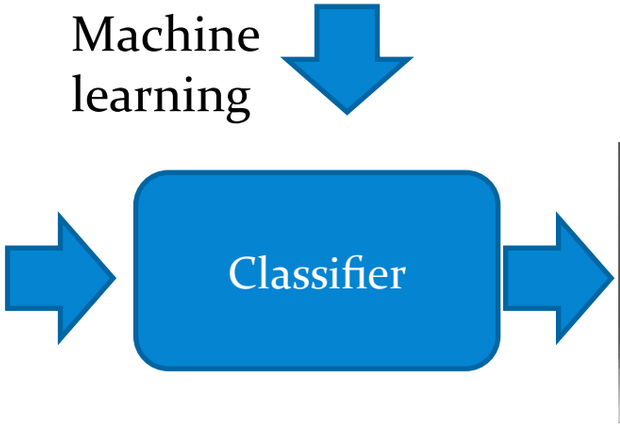

- >10^22 objects
- Sparse, multi-dimensional data structures
- Complex datatypes
  (images, video, matrices, etc.)
- Complex application logic and dataflow
  - >35000 lines of .Net
  - 140 CPU days
  - >10^5 processes
  - 30 TB data analyzed
  - 140 avg parallelism (235 machines)
  - 300% CPU utilization (4 cores/machine)

Wrap up: Key Features

- **On-demand self-service**.
  - No human intervention with providers
  - Users abstracted from implementation
  - Near real-time delivery
- **Broad network access**.
  - Heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).
  - Almost always IP, HTTP and REST
- **Resource pooling**.
  - Resources are drawn from a common pool
  - Location independence
  - Builds economies of scale
  - High efficiency
- **Rapid elasticity**.
  - Resources dynamically allocated
  - Resource dynamically released
  - Fully automated
  - Unlimited resources in real-time from user perspective
- **Measured service**.
  - Service are metered, like gas or power
  - Users pay only for services used
  - Complete control on services used

Wrap up: Main Achievements

- **Goal 1: Cost control**
  - Cloud computing can accommodate different computing demands:
    - Batch processing (e.g. New York Times)
    - Variable loads with peaks (e.g. Web applications)
    - Unknown (e.g. start-ups)

- **Goal 2: Elasticity**
  - Cloud computing can provide additional computing power immediately
  - No need to wait for a server cluster to be purchased, delivered, installed, ...
  - You can give it back if it is not needed anymore

- **Goal 3: Stick to your Business**
  - Users should not take care of the bare metal
  - Immediate and transparent access to computing resources and software services

InGeoClouds Use Case

A Medium-sized InGeoCloudS Instance	
CPU	5
Memory (GB)	16
Storage (GB)	100
Data access (GB/month)	300
Service access (Mill. req.s/month)	24
Network Traffic (GB/month)	2,400

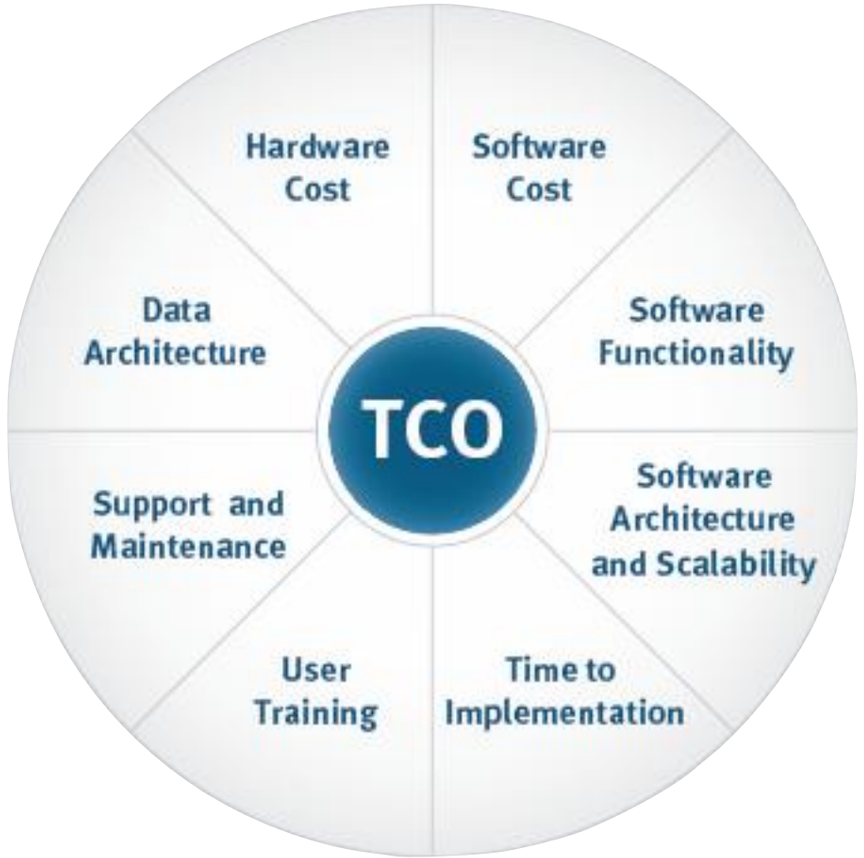

InGeoClouds Use Case

A Medium-sized InGeoCloudS Instance	
CPU	5
Memory (GB)	16
Storage (GB)	100
Data access (GB/month)	300
Service access (Mill. req.s/month)	24
Network Traffic (GB/month)	2,400

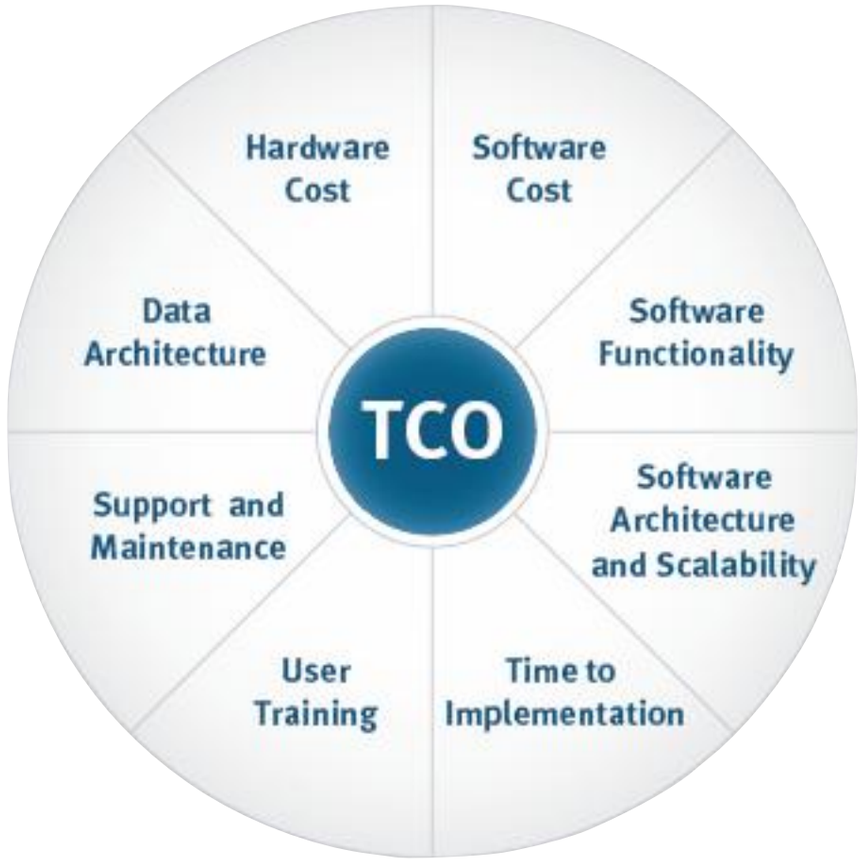

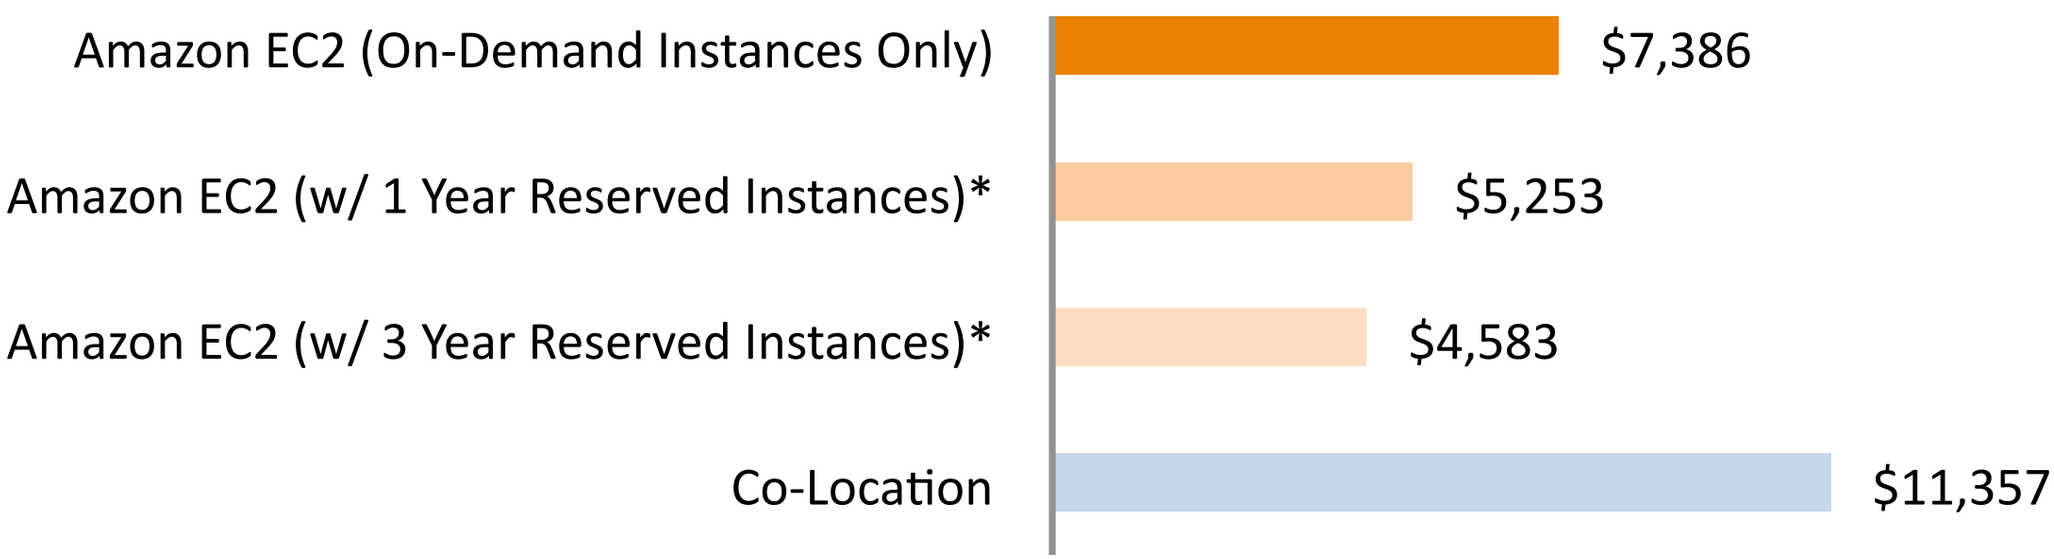

http://d36cz9buwru1tt.cloudfront.net/Amazon_EC2_Cost_Comparison_Calculator_042810.xls

Envisioned advantages for InGeoCloudS

- **Costs**:
  - a single partner can reduce its costs
  - ***economies of scale*** are possible within a large consortium
    - Databases, Web Applications, Maps, can share disks, CPU, RAM

- **It is not only about money !**
  - the cloud platform also provides extremely large computing power!
    - for ***large-scale computer simulations*** other analyses
  - having the data all in one place will:
    - foster ***successful collaborations***
    - push the ***adoption of international standards***
    - enable ***"inter-disciplinary" studies***

- The end.

- You already know my answer to your questions...

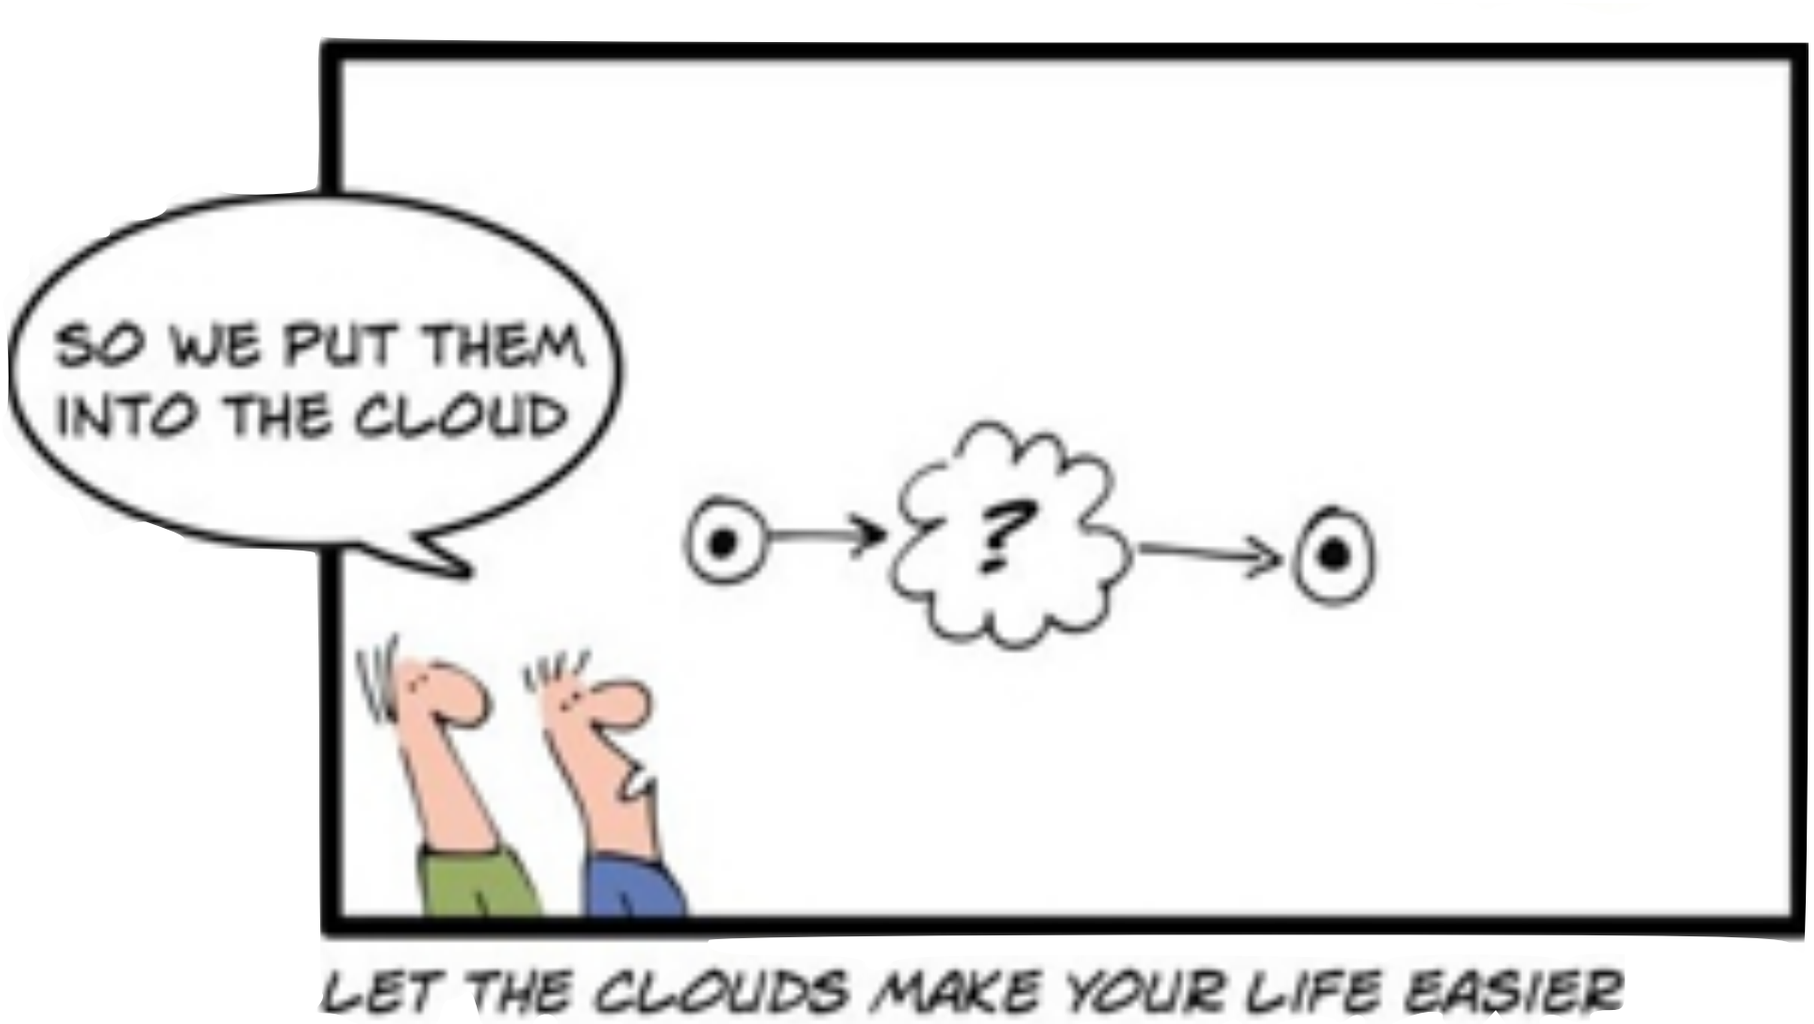

Disclaimer

- All pictures and graphs reported here are property of the respective owners.
- They are used only for didactic purposes.
- Whenever available, sources are cited.

References

- The data deluge. http://www.economist.com/opinion/displayStory.cfm?story_id=15579717. Feb 2010.
- Data, data everywhere. http://www.economist.com/specialreports/displaystory.cfm?story_id=15557443. Feb 2010.
- Jeffrey Dean, Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. OSDI 2004: 137-150.
- Chris Anderson. The Long Tail. Wired, October 2004.
- Benjamin Reed. Pig: A high level language for doing large scale data processing using Hadoop. Workshop on High Performance Analytics - Algorithms, Implementations, and Applications, SDM2010.
- Mihai Budiu. Large-scale Machine Learning using DryadLINQ. Workshop on High Performance Analytics - Algorithms, Implementations, and Applications, SDM2010.
- Large Language Models in Machine Translation Thorsten Brants, Ashok C. Popat, Peng Xu, Franz J. Och, Jeffrey Dean. Google, Inc. Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, pp. 858–867, Prague, June 2007.
- Large Scale Parallel Document Mining for Machine Translation. Jakob Uszkoreit, Jay M. Ponte, Ashok C. Popat, Moshe Dubiner, Google, Inc. Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010), pages 1101–1109, Beijing, August 2010.
- Nature 455, 1. 4 Sept. 2008.
- Predicting consumer behavior with Web search. Sharad Goel,, Jake M. Hofman, Sébastien Lahaie, David M. Pennock, Duncan J. Watts. PNAS 2010.
- Mark Baker. An Introduction and Overview of Cloud Computing.
- N Wilde, T Huber. Virtualization and Cloud Computing.
- Jimmy Lin. What is Cloud Computing? (and an intro to parallel/distributed processing)