MAP REDUCE FEST
TACKLING THE MULTI-HEADED DRAGON

GENOVEVA VARGAS SOLAR

FRENCH COUNCIL OF SCIENTIFIC RESEARCH, LIG-LAFMIA, FRANCE

Genoveva.Vargas@imag.fr

http://www.vargas-solar.com/teaching/map-reduce-fest/

Modern information societies are defined by vast repositories of data, both public and private

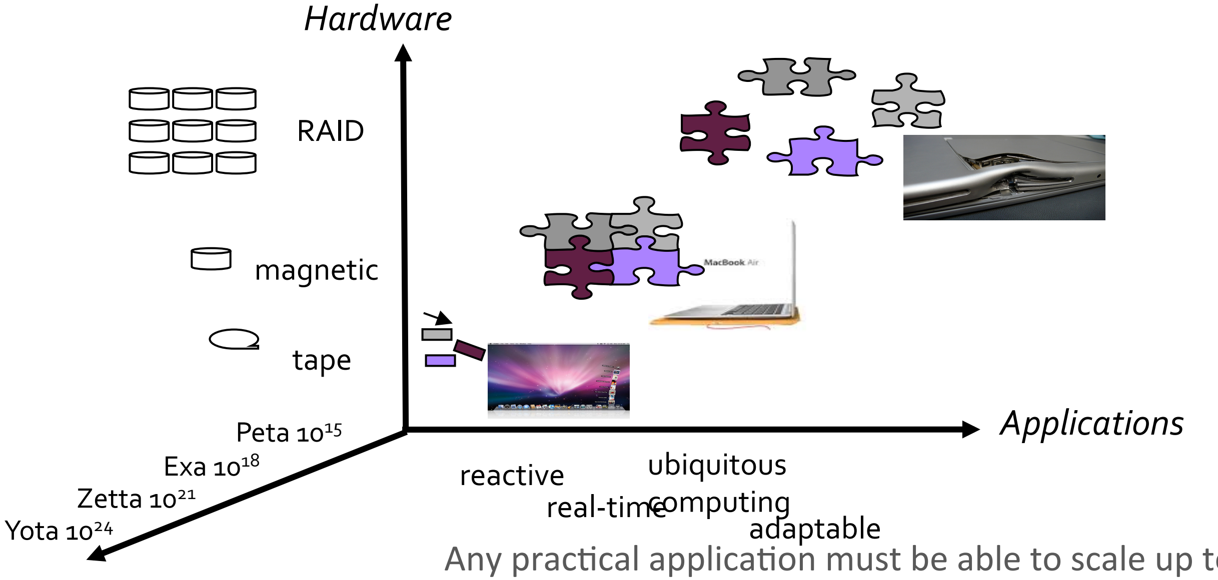

Any practical application must be able to scale up to datasets of interest

Use of memory and computing capacities of all computers and servers distributed in the world communicated by a network (e.g. Internet)

Megabyte(10^6) Gigabyte (10^9) Terabytes (10^{12}), petabytes (10^{15}), exabytes (10^{18}) and zettabytes (10^{21})

DIGITAL INFORMATION SCALE

Unit	Size	Meaning
Bit (b)	1 or 0	Short for binary digit, after the binary code (1 or 0) computers use to store and process data
Byte (B)	8 bits	Enough information to create an English letter or number in computer code. It is the basic unit of computing
Kilobyte (KB)	2^{10} bytes	From "thousand' in Greek. One page of typed text is 2KB
Megabyte (MB)	2^{20} bytes	From "large" in Greek. The complete works of Shakespeare total 5 MB. A typical pop song is 4 MB.
Gigabyte (GB)	2^{30} bytes	From "giant" in Greek. A two-hour film ca be compressed into 1-2 GB.
Terabyte (TB)	2^{40} bytes	From "monster" in Greek. All the catalogued books in America's Library of Congress total 15TB.
Petabyte (PB)	2^{50} bytes	All letters delivered in America's postal service this year will amount ca. 5PB. Google processes 1PB per hour.
Exabyte (EB)	2^{60} bytes	Equivalent to 10 billion copies of The Economist
Zettabyte (ZB)	2^{70} bytes	The total amount of information in existence this year is forecast to be around 1,27ZB

MASSIVE DATA

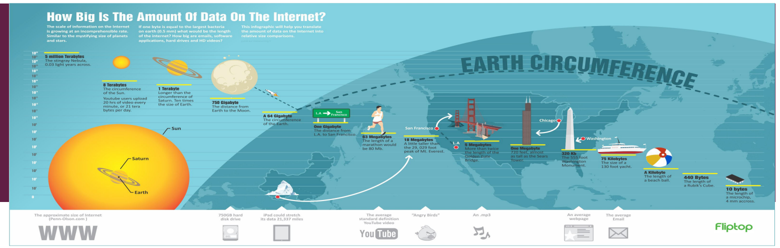

Data sources

- Information-sensing mobile devices, aerial sensory technologies (remote sensing)
- Software logs, posts to social media sites
- Telescopes, cameras, microphones, digital pictures and videos posted online
- Transaction records of online purchases
- RFID readers, wireless sensor networks (number of sensors increasing by 30% a year)
- Cell phone GPS signals (increasing 20% a year)

Massive data

- Sloan Digital Sky Survey (2000-) its archive contains 140 terabytes. Its successor: Large Synoptic Survey Telescope (2016-), will acquire that quantity of data every five days !
- Facebook hosts 140 billion photos, and will add 70 billion this year (ca. 1 petabyte). Every 2 minutes today we snap as many photos as the whole of humanity took in the 1800s !
- Wal-Mart, handles more than 1 million customer transactions every hour, feeding databases estimated at more than 2.5 petabytes (10^{15})
- The Large Hadron Collider (LHC): nearly 15 million billion bytes per year -15 petabytes (10^{15}). These data require 70,000 processors to be processed!

http://blog.websourcing.fr/infographie-la-vrai-taille-dinternet/

WHO PRODUCES DATA ? *DIGITAL SHADOW*

75% of the information is generated by individuals — writing documents, taking pictures, downloading music, etc. — but is far less than the amount of information being created about them in the digital universe

- 3,146 billion mail addresses, of which 360 million Hotmail
- 95.5 million domains.com
- 2.1 billion Internet users, including 922 million in Asia and 485 million in China
- 2.4 billion accounts on social networks
- 1000 billion videos viewed on YouTube, about 140 average per capita on the planet!
- 60 images posted on Instagram every second
- 250 million tweets per day in October

BIG DATA

- Collection of data sets so large and complex that it becomes difficult to process using on-hand database management tools or traditional data processing applications.
- Challenges include capture, curation, storage, search, sharing, analysis, and visualization **within a tolerable elapsed time**
- Data growth challenges and opportunities are three-dimensional (**3Vs model**)
  - increasing volume (amount of data)
  - velocity (speed of data in and out)
  - variety (range of data types and sources)

"Big data are high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization."

DATA MANAGEMENT WITH RESOURCES CONSTRAINTS

STORAGE SUPPORT

RAM

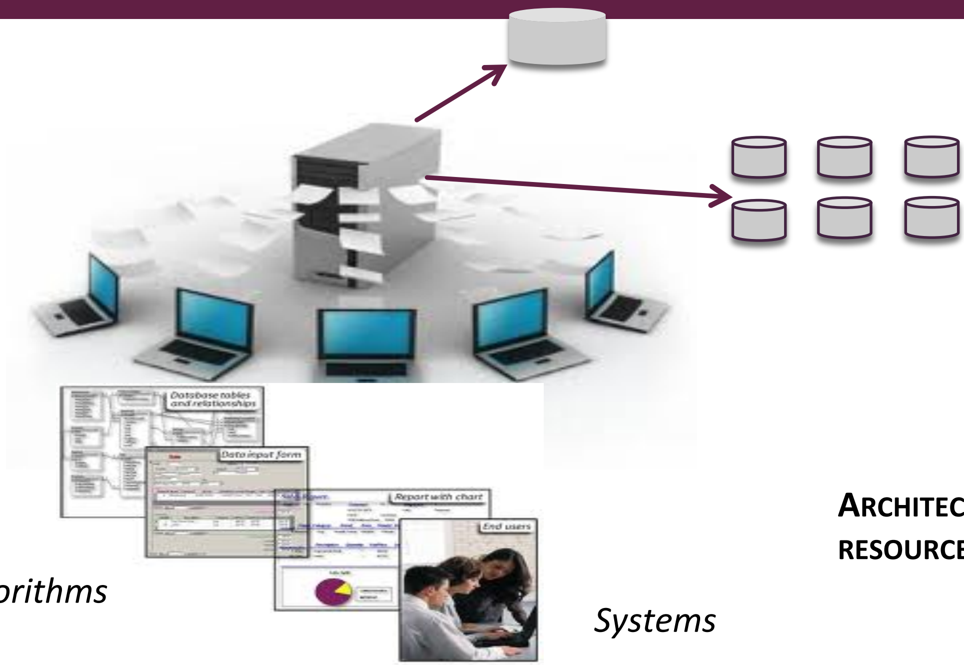

ARCHITECTURE & RESOURCES AWARE

Algorithms

Systems

Efficiently manage and exploit data sets according to given specific storage, memory and computation resources

DATA MANAGEMENT WITHOUT RESOURCES CONSTRAINTS

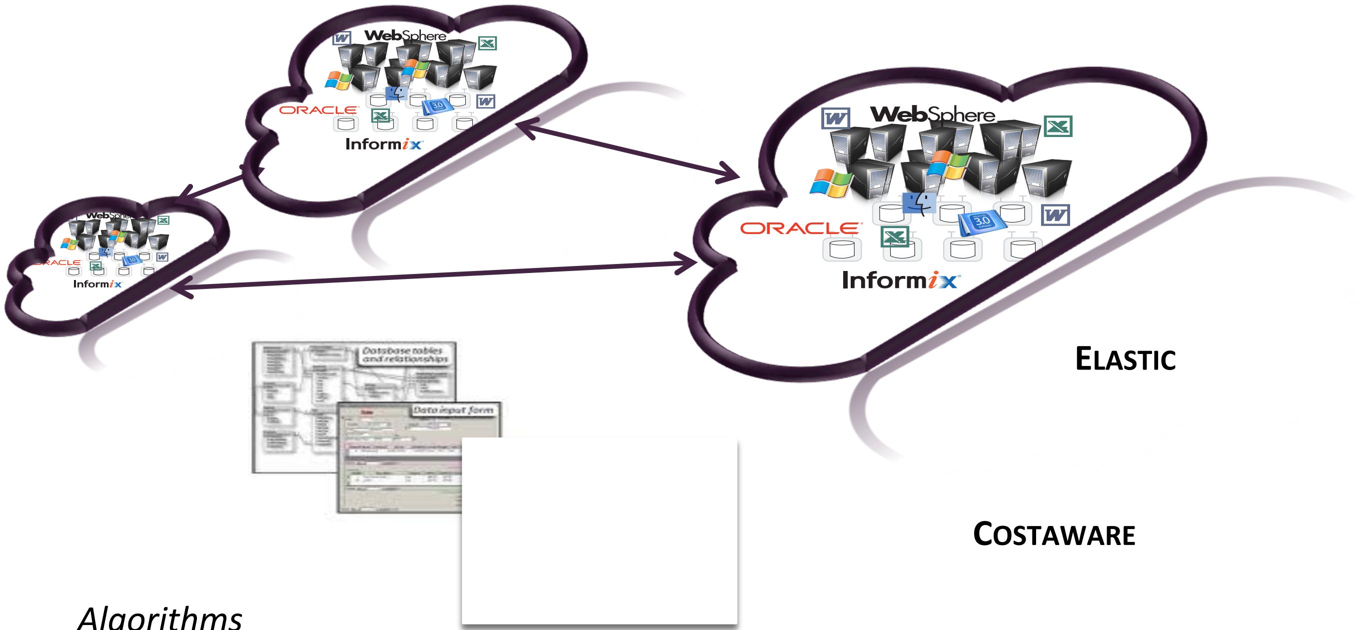

ELASTIC

COSTAWARE

Algorithms

Systems

Costly manage and exploit data sets according to unlimited storage, memory and computation resources

CLOUD IN A NUTSHELL

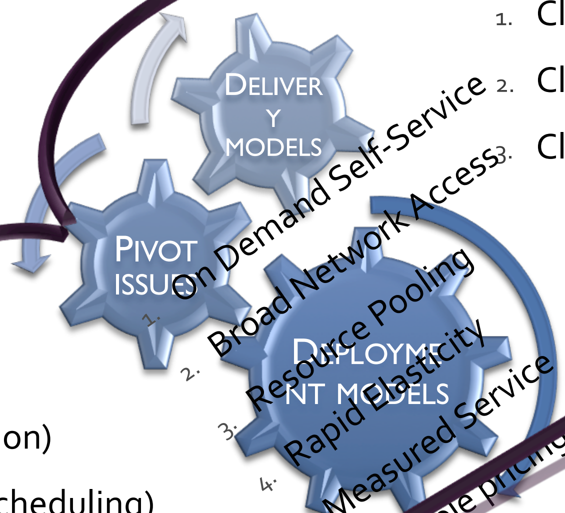

1. Cloud Software as a Service (SaaS)
2. Cloud Platform as a Service (PaaS)
3. Cloud Infrastructure as a Service (IaaS)

1. On Demand Self-Service
2. Broad Network Access
3. Resource Pooling
4. Rapid Elasticity
5. Measured Service
6. Scalable pricing

1. Private cloud
2. Public cloud
3. Hybrid cloud
4. Community Cloud

1. Virtualization
2. Autonomics (automation)
3. Grid Computing (job scheduling)

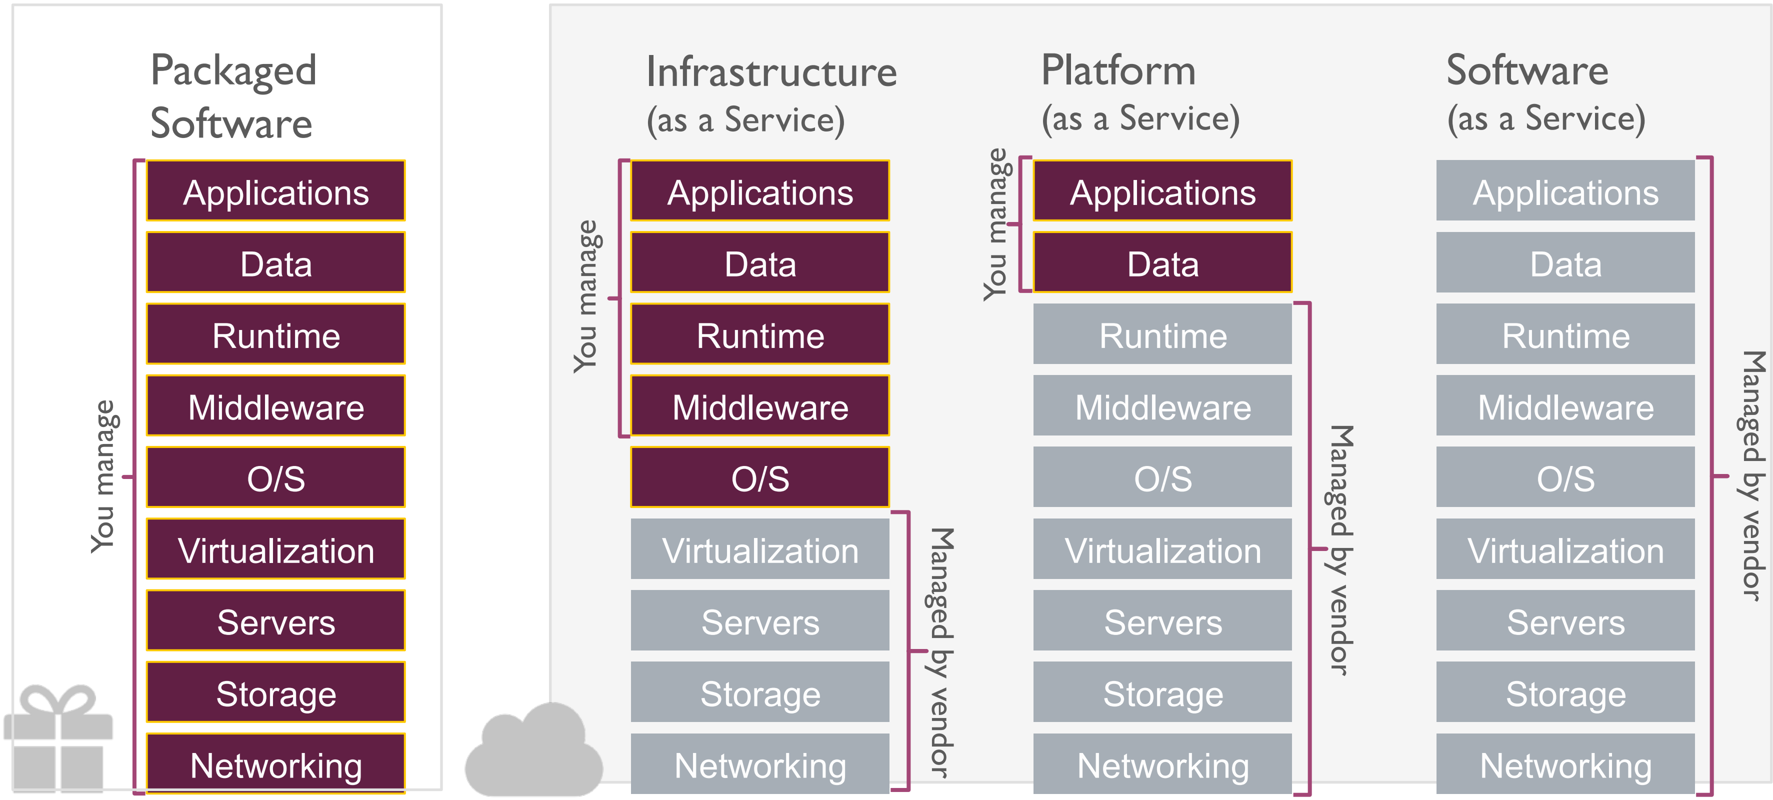
Packaged Software
You manage
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Infrastructure (as a Service)
You manage
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Managed by vendor
Platform (as a Service)
You manage
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Managed by vendor
Software (as a Service)
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Managed by vendor

CLOUD COMPUTING CHARACTERISTICS

- On-demand self-service
  - Available computing capabilities without human interaction
- Broad network access
  - Accessibility from heterogeneous clients (mobile phones, laptops, etc.)
- Rapid elasticity
  - Possibility to automatically and rapidly increase or decrease capabilities

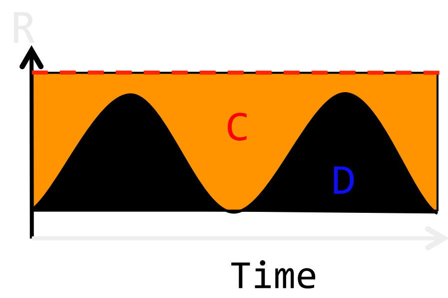

- Measured service
  - Automatic control and optimization of resources
- Resource pooling (virtualization, multi-location)
  - Multi-tenant resources (CPU, disk, memory, network) with a sense of location independence

STORING DATA

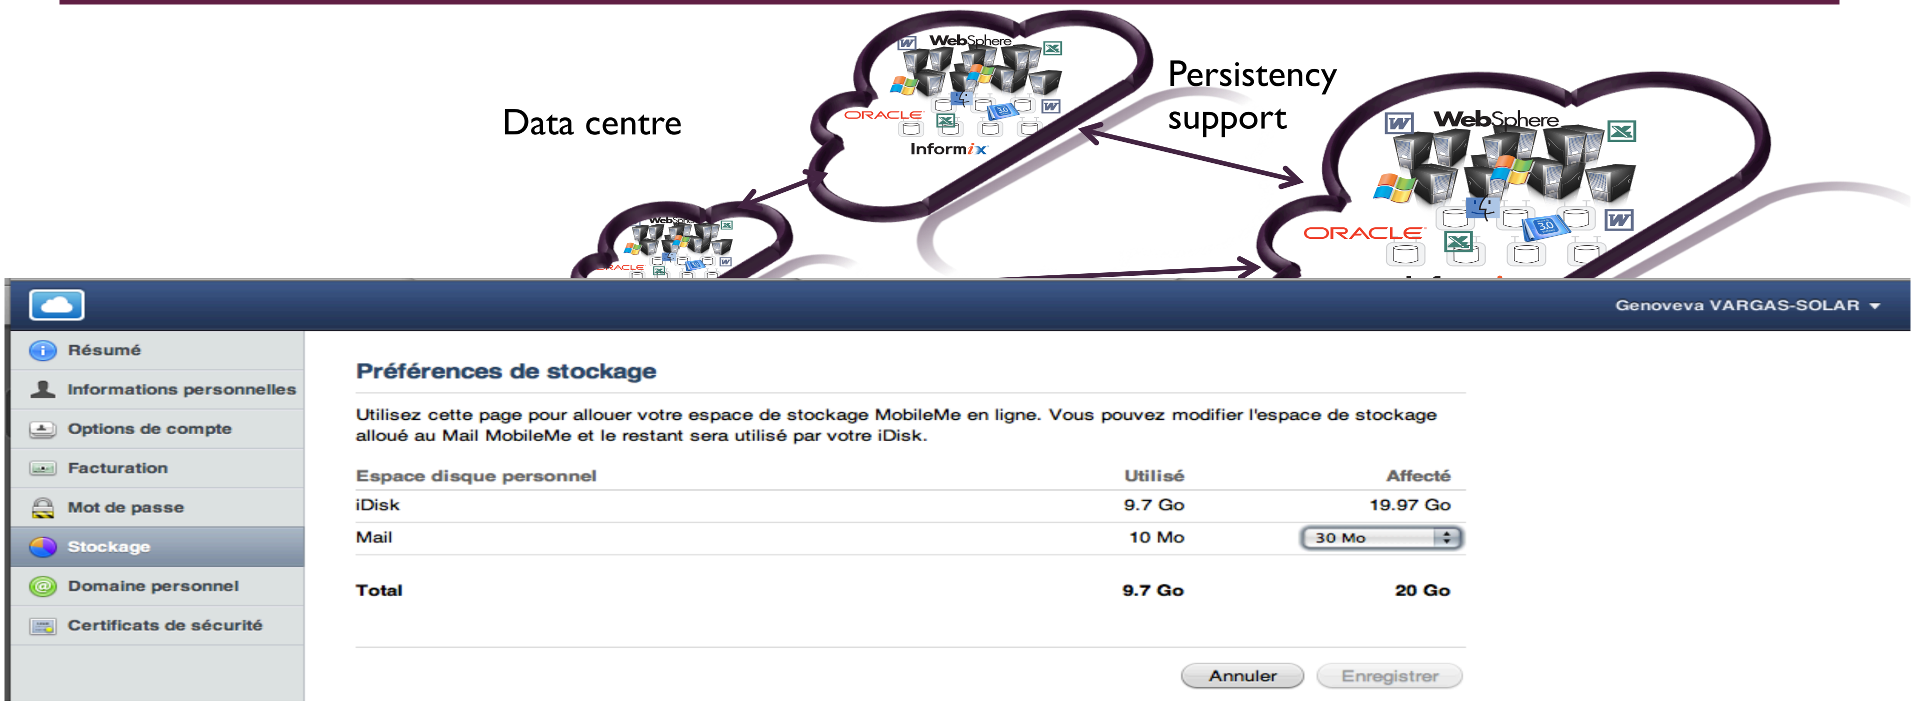

… WITH RESOURCES CONSTRAINTS

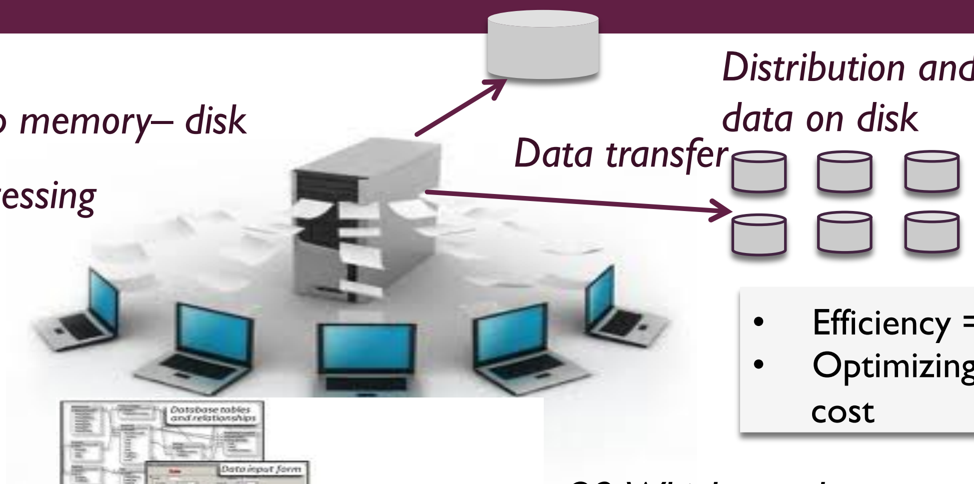

- Efficiency => time cost
- Optimizing memory and computing cost

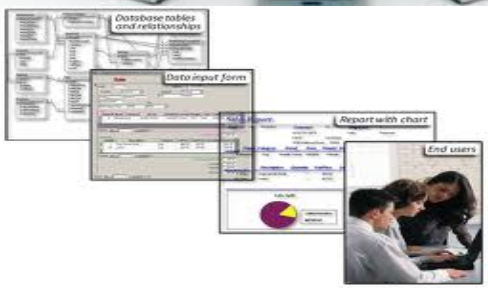

Q2:Which are the consumption rules of Starbucks clients ?

Efficiently manage and exploit data sets according to given specific storage, memory and computation resources

WITHOUT RESOURCES CONSTRAINTS …

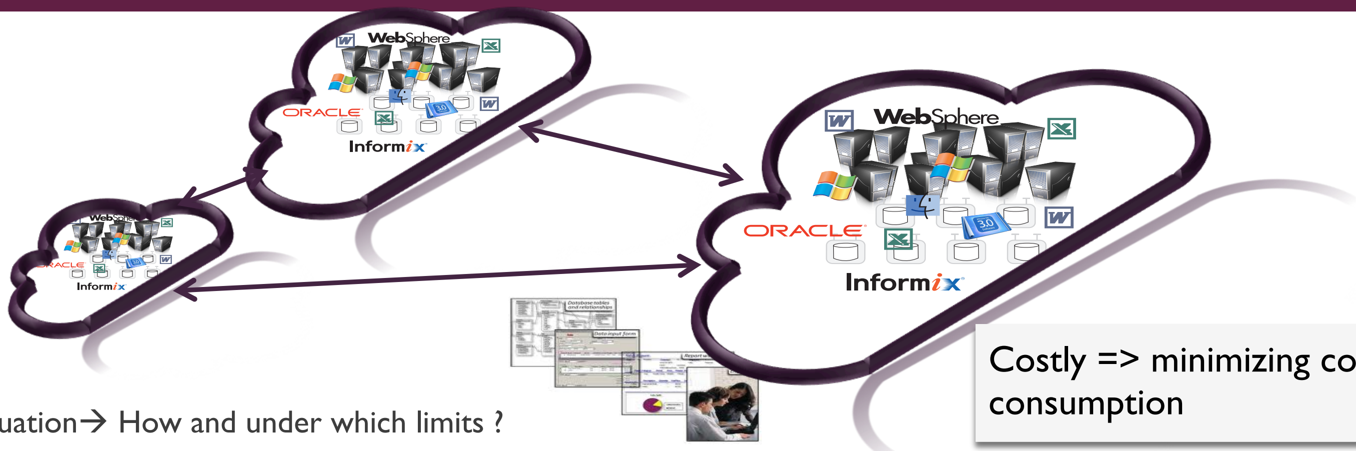

Costly => minimizing cost, energy consumption

- Query evaluation→ How and under which limits ?
  - is not longer completely constraint by resources availability: computing, RAM, storage, network services
  - Decision making process determined by resources consumption and consumer requirements
- Data involved in the query, particularly in the result can have different costs: top 5 gratis and the rest available in return to a credit card number
- Results storage and exploitation demands mor resources

Costly manage and exploit data sets according to unlimited storage, memory and computation resources

EXECUTION MODEL: MAP-REDUCE

- Programming model for expressing distributed computations on massive amounts of data and an execution framework for large-scale data processing on clusters of commodity servers
- Open-source implementation called Hadoop, whose development was led by Yahoo (now an Apache project)
- Divide and conquer principle: partition a large problem into smaller subproblems
  - To the extent that the sub-problems are independent, they can be tackled in parallel by different workers
  - intermediate results from each individual worker are then combined to yield the final output
- Large-data processing requires bringing data and code together for computation to occur:

→ no small feat for datasets that are terabytes and perhaps petabytes in size!

MAP-REDUCE ELEMENTS

$$\text{map: } (k_1, v_1) \longrightarrow [(k_2, v_2)]$$
$$\text{reduce: } (k_2, [v_2]) \longrightarrow [(k_3, v_3)]$$

- **Stage 1**: Apply a user-specified computation over all input records in a dataset.
  - These operations occur in parallel and yield intermediate output (key-value couples)
- **Stage 2**: Aggregate intermediate output by another user-specified computation
  - Recursively applies a function on every pair of the list
- Execution framework coordinates the actual processing
- Implementation of the programming model and the execution framework

MAP REDUCE EXAMPLE

Count the number of occurrences of every word in a text collection

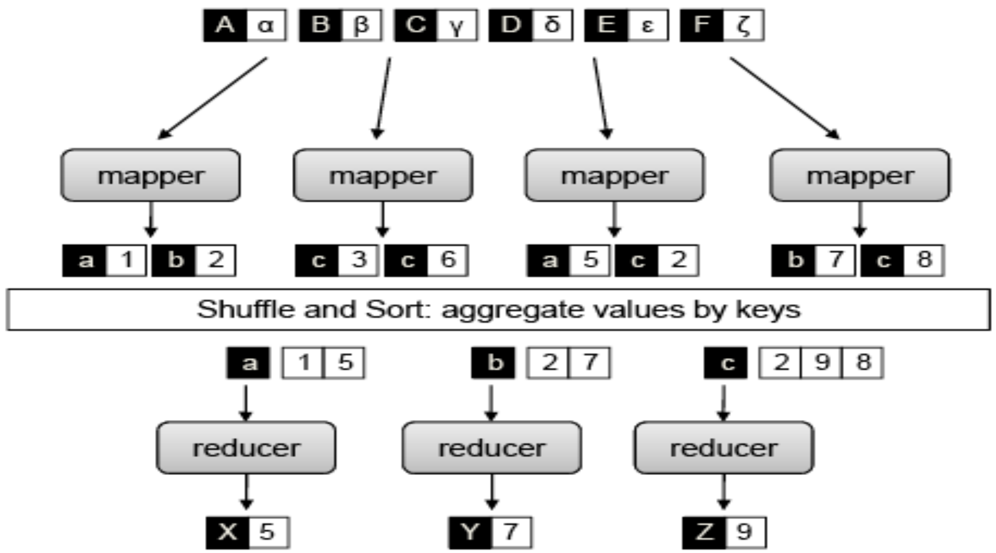

EXECUTION FRAMEWORK

- Important idea behind MapReduce is separating the what of distributed processing from the how
- A MapReduce program (job) consists of
  - code for mappers and reducers packaged together with
  - configuration parameters (such as where the input lies and where the output should be stored)
- Execution framework responsibilities: scheduling
  - Each MapReduce job is divided into smaller units called tasks
  - In large jobs, the total number of tasks may exceed the number of tasks that can be run on the cluster concurrently → manage tasks queues
  - Coordination among tasks belonging to different jobs

MAP-REDUCE ADDITIONAL ELEMENTS

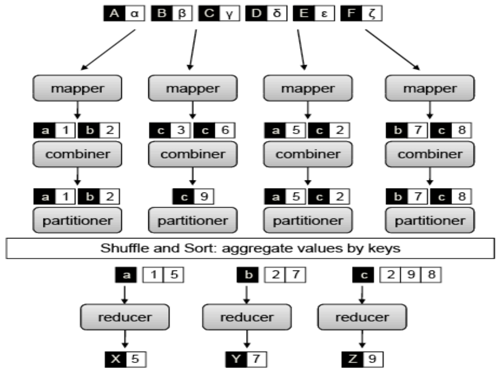

- Partitioners are responsible for dividing up the intermediate key space and assigning intermediate key-value pairs to reducers
  - the partitioner species the task to which an intermediate key-value pair must be copied
- Combiners are an optimization in MapReduce that allow for local aggregation before the shuffle and sort phase

BIG VALUE FROM BIG DATA

Big data is not a "thing" but instead a dynamic/activity that crosses many IT borders

- Big data is not only about the original content stored or being consumed but also about the information around its consumption.
  - Gigabytes of stored content can generate a petabyte or more of transient data that we typically don't store digital TV signals we watch but don't record voice calls that are made digital in the network backbone during the duration of a call
- Big data technologies describe a **new** generation of technologies and architectures, designed to **economically** extract value from very large volumes of a wide variety of data, by enabling **high-velocity** capture, discovery, and/or analysis

http://www.emc.com/collateral/analyst-reports/idc-extracting-value-from-chaos-ar.pdf

WHAT ARE THE FORCES BEHIND THE EXPLOSIVE GROWTH OF THE DIGITAL UNIVERSE?

- Technology has helped by driving the cost of creating, capturing, managing, and storing information
- … Prime mover is financial
- Since 2005, the investment by enterprises in the digital universe has increased 50% — to $4 trillion
  - spent on hardware, software, services, and staff to create, manage, and store — and derive revenues from the digital universe
  - the trick is to generate value by extracting the right information from the digital universe

BIG DATA TECH ECONOMY

- Data-driven world guided by a rapid ROD (**Return on Dat**a)

→ reducing cost, complexity, risk and increasing the value of your holdings thanks to a mastery of technologies

- The key is how quickly data can be turned in to currency by:
  - Analysing patterns and spotting relationships/trends that enable decisions to be made faster with more precision and confidence
  - Identifying actions and bits of information that are out of compliance with company policies can avoid millions in fines
  - Proactively reducing the amount of data you pay ($18,750/gigabyte to review in eDiscovery) by identifying only the relevant pieces of information
  - Optimizing storage by deleting or offloading non-critical assets to cheaper cloud storage thus saving millions in archive solutions

CONTENT

Part I: First and commun steps

- Introduction
  - 1.1.Principle and objective of map - reduce
  - 1.2.Preparing huge data collections
- Overview of the Hadoop infrastructure

Part II: programming patterns

- A first step to map reduce: counting words
- Summarization pattern
  - Processing "Les Miserables": counting words with an in-combiner pattern
  - Dealing with the municipal accounts of an imaginary city: numerical summarization
- A first step to Social Networks' posts processing: inverted index
- Optimizing NoSQL queries: bloom filter pattern

CONTENT

Part III: Putting MapReduce in perspective

- From parallel programming to parallel DBMS who is who?

METHODOLOGY: HACKATON STYLE

- Collaborative knowledge construction
- Synchronized "theory" content presentations
- Free hours: intensive programming exercises

Thanks Merci

Gracias

Contact: Genoveva Vargas-Solar, CNRS, LIG-LAFMIA

Genoveva.Vargas@imag.fr

http://vargas-solar.imag.fr

Open source polyglot persistence tools

http://code.google.com/p/exschema/

http://code.google.com/p/model2roo/

Want to put cloud data management in practice ?

cf. http://vargas-solar.imag.fr/academika/cloud-data-management/

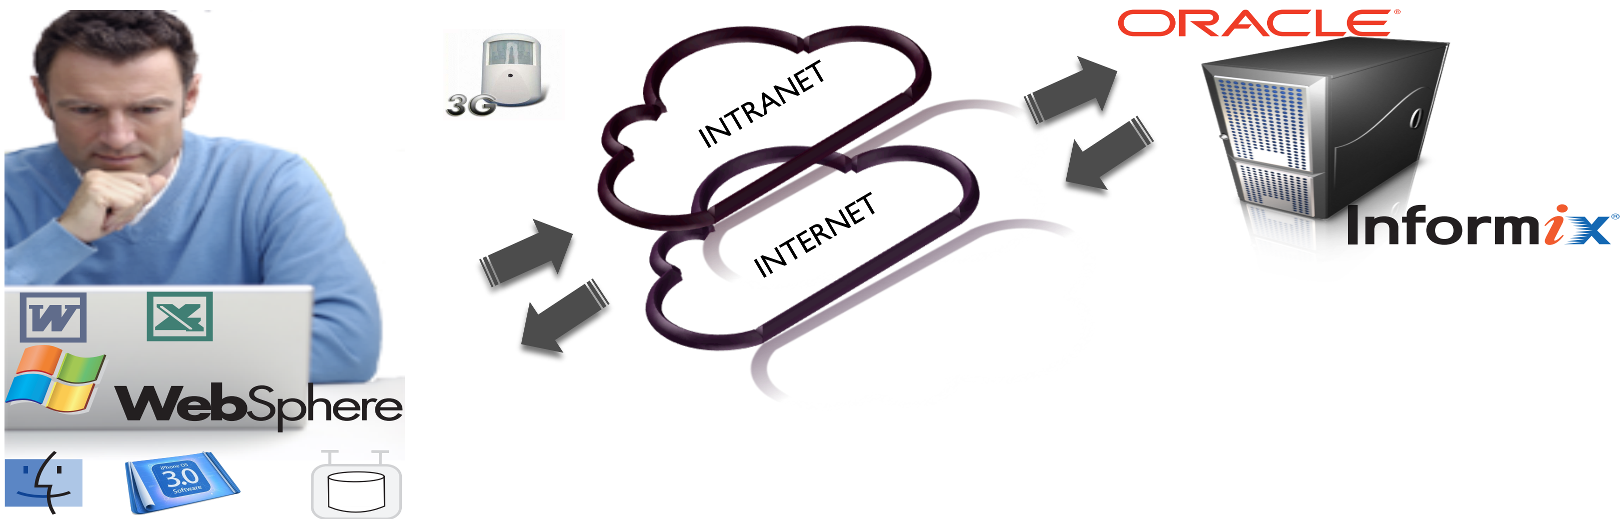
WebSphere
3G
INTRANET
INTERNET
ORACLE
Informix

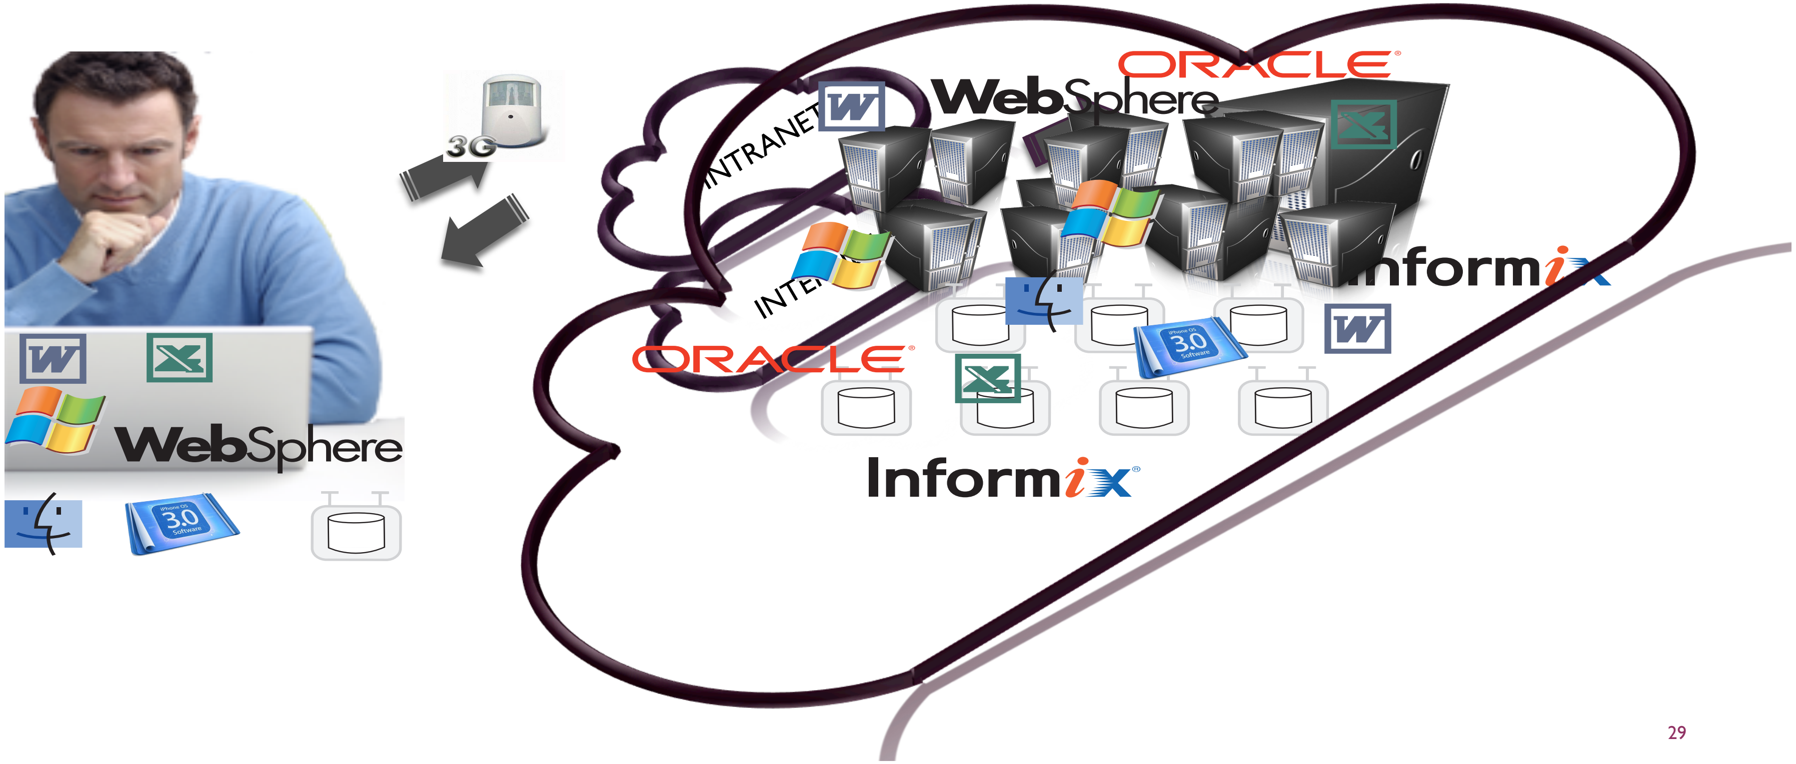
3G
WebSphere
ORACLE
WebSphere
INTRANET
INTE
ORACLE
Informix
Informix
3.0

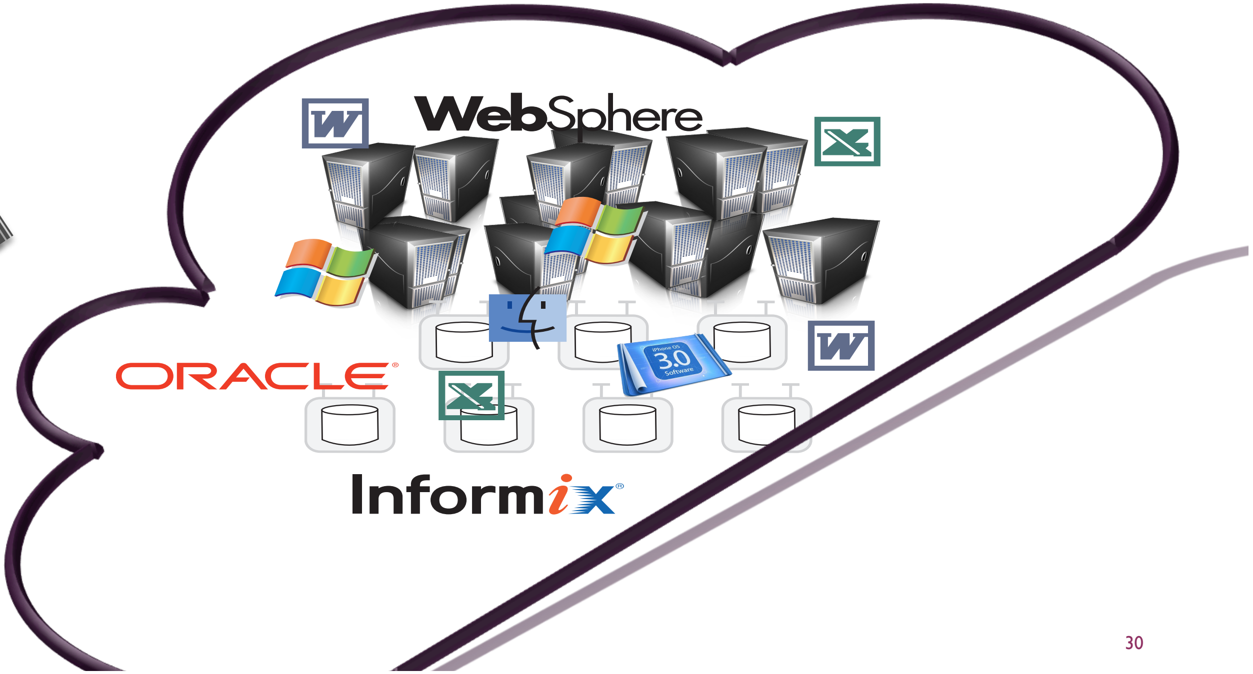
WebSphere
ORACLE
Informix
3.0

DEFINITION

Technological vision

- Cloud computing is a model for enabling on-demand network access to a shared pool of virtualized computing resources
  - e.g., networks, servers, storage, applications, devices/ mobiles and services)
- Rapidly provisioned and released with minimal management effort or service provider interaction (self-service model through API or web portals)

Market vision

- Pay-per-use (or pay-as-you-go) billing models
- An application may exist to
  - run a job for a few minutes or hours
  - provide services to costumers on a long-term basis
- Billing is based on resource consumption: CPU hours, volumes of data moved, or gigabytes of data stored

AMBITION

- Provides a platform for new software applications that run across a large collection of physically separate computers and free computation in front of a user
- Heralds a revolution comparable to PC [J. Larus, MS]
  - Supply on-demand internet computer resources on a vast scale and at low price
  - Exists beyond services offered by Amazons' AWS, Microsoft Azure or Google AppEngine

DELIVERY MODELS

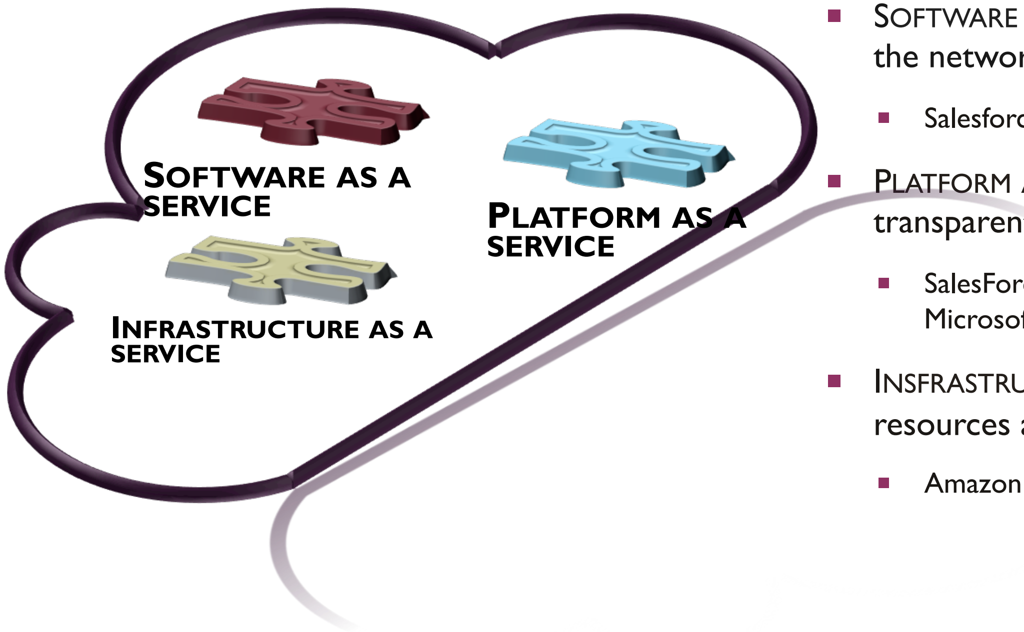

- SOFTWARE AS A SERVICE: applications accessible through the network (Web services, REST/SOAP)
  - Salesforce.com (CRM) and Google (Gmail, Google Apps)
- PLATFORM AS A SERVICE: provide services for transparently managing hardware resources
  - SalesForce.com (Force.com), Google (Google App Engine), Microsoft (Windows Azure), Facebook (Facebook Platform)
- INSFRASTRUCTURE AS A SERVICE: provide Data centers resources and others like CPU, storage and memory
  - Amazon (EC2/S3) and IBM (Bluehouse)

DEPLOYMENT MODELS

Public cloud

- Run by third parties and applications from different clients mixed together on the clouds' servers, storage systems and networks

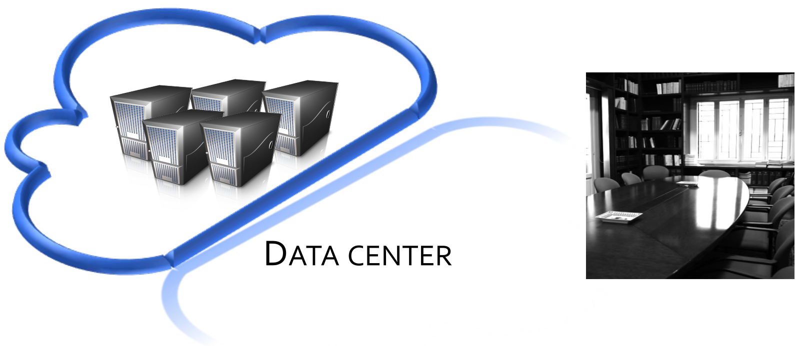

Private cloud

- Used by an exclusive client providing utmost control over data, security and QoS
- The company owns and controls the infrastructure

DEPLOYMENT MODELS

Hybrid cloud

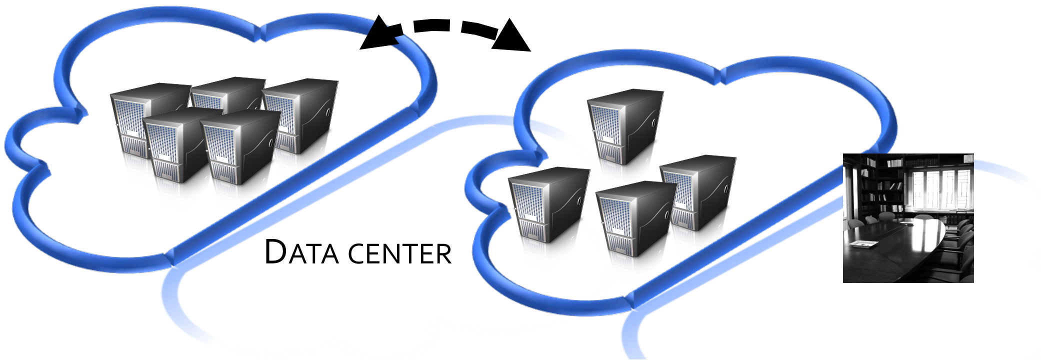

Community cloud

- My data and services anywhere (online, on home devices, on mobile devices) from anywhere
- Extension of Personal Cloud with sharing
- Support for the Internet of Things, M2M platforms

VIRTUALIZATION

- Virtual machines and virtual appliances become standard deployment object
- Abstract the hardware to the point where software stacks can be deployed and redeployed without being tied to a specific physical server
  - Servers provide a pool of resources that are harnessed as needed
  - The relationship of applications to compute, storage, and network resources changes dynamically to meet workload and business demands
  - Applications can be deployed and scaled rapidly without having to produce physical servers

VIRTUALIZATION

- Full virtualization is a technique in which a complete installation of one machine is run on another
  - A system where all software running on the server is within a virtual machine
  - Applications and operating systems
  - Means of accessing services on the cloud
- A compute cloud is a self-service proposition where a credit card can purchase compute cycles, and a Web interface or API is used to create virtual machines and establish network relationships between them

PROGRAMMABLE INFRASTRUCTURE

- Cloud provider API
  - to create an application initial composition onto virtual machines
  - To define how it should scale and evolve to accommodate workload changes
    - Self monitoring and self expanding applications
- Applications must be assembled by assembling and configuring appliances and software
- Cloud services must be composable so they can be consumed easily

PROGRAMMABLE INFRASTRUCTURE