

Unit -IV

WSC

- WSC is the foundation of Internet services many people use every day
 - searching, social networking, online maps, video sharing, online shopping, email services so on.
- WSC could be considered a type of datacenter architecture and operation are quite different.
- Datacenters
 - buildings where multiple servers and communication gear are located because of their common environmental requirements and physical security needs, and for ease of maintenance.

WSC

- WSCs act as one giant machine and cost on the order of \$150M
 - the building
 - the electrical
 - cooling infrastructure
 - the servers
 - networking equipment that connects and houses 50,000 to 100,000 servers.

WSC

- WSCs represent a much larger share of the IT market.
- Rapid growth of cloud computing makes WSCs available to anyone with a credit card
- ***WSC architects share many goals and requirements with server architects:***

1. Cost-performance

- Work done per dollar is critical in part because of the scale. Reducing the capital cost of a WSC by 10% could save \$15M.

WSC

2. Energy efficiency

- Power distribution costs are functionally related to power consumption
- you need sufficient power distribution before you can consume power
- need to get out the heat that you put in.
- peak power and consumed power drive both the cost of power distribution and the cost of cooling systems.
- energy efficiency is an important part of environmental stewardship.
- work done per joule is critical for both WSCs and servers.

WSC

3. Dependability via redundancy

- hardware and software in a WSC must collectively provide at least 99.99% of availability
- it must be down less than 1 hour per year
- Redundancy is the key to dependability for both WSCs and servers

WSC

4. Network I/O

- WSC architects must provide a good network interface to the external world
- Networking is needed to keep data consistent between multiple WSCs as well as to interface to the public.

WSC

5. Interactive and batch processing workloads

- highly interactive workloads for services like search and social networking with millions of users
- run massively parallel batch programs to calculate metadata useful to such services.

WSC

- chars not shared with server architecture

1. Ample parallelism

- A concern for a server architect is whether the applications in the targeted marketplace have enough parallelism to justify the amount of parallel hardware to exploit this parallelism.
- A WSC architect has no such concern.
- Batch applications benefit from the large number of independent datasets that require independent processing.
- software as a service (SaaS), can benefit from millions of independent users of interactive Internet services.

WSC

2. Operational costs count

- Server architects usually design their systems for peak performance within a cost budget
- ignore operational costs of a server
- WSCs have longer lifetimes-the building and electrical and cooling infrastructure are often amortized over 10 or more years
- the operational costs for Energy, power distribution, and cooling represent more than 30% of the costs of a WSC in 10 years.

WSC

3. Scale and the opportunities/problems associated with scale

- extreme computers are extremely expensive because they require custom hardware.
- when you purchase 50,000 servers to construct a single WSC, you do get volume discounts.
- WSCs are so massive internally you get economy of scale.

Programming Models WSC

- most popular framework for batch processing in a WSC is Map-Reduce
- its open-source twin Hadoop
- Facebook runs Hadoop on 2000 batch-processing servers of the 60,000 servers it is estimated to have in 2011

Programming Models WSC

- **Map** first applies a programmer-supplied function to each logical input record.
- **Map** runs on thousands of computers to produce an intermediate result of key-value pairs
- **Reduce** collects the output of those distributed tasks and collapses them using another programmer-defined function

Programming Models WSC

- Map and Reduce functions that operate on key-value pairs
- Each map function operates on a collection of records a record is a webpage or a facebook user profile
- The records are in the file system and scattered across several servers thousands of map functions are spawned to work on all records in parallel
- The Reduce function aggregates and sorts the results produced by the Mappers, also performed in parallel

Programming Models WSC

- **MapReduce** program calculates the number of occurrences of every English word in a large collection of documents
- Below is a simplified version of that program,

Programming Models WSC

map(String key, String value):

// key: document name

// value: document contents for each word w in value:

EmitIntermediate(w, "1"); // Produce list of all words

reduce(String key, Iterator values):

//key: a word

// values: a list of counts

int result=0;

for each v in values:

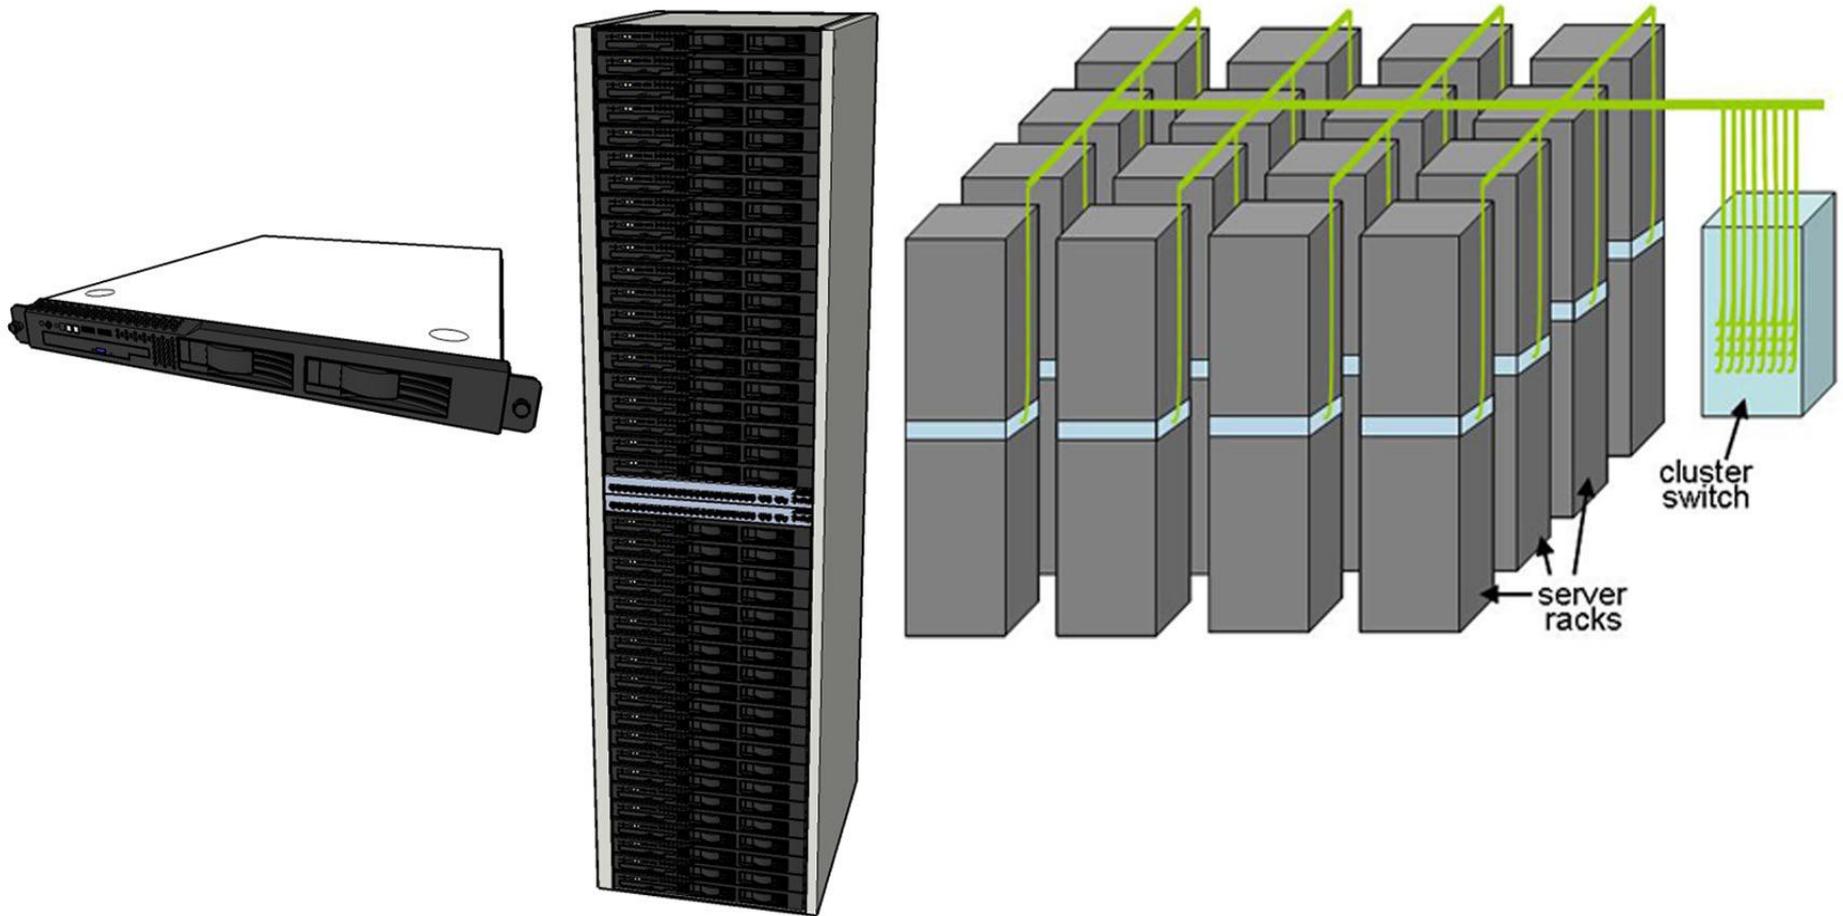
result+= ParseInt(v); // get integer from key-value pair

Emit(AsString(result));

Programming Models WSC

- ***EmitIntermediate*** used in the Map function emits each word in the document and the value one
- ***Reduce*** function sums all the values per word for each document using `ParseInt()` to get the number of occurrences per word in all documents.
- MapReduce can be thought of as a generalization of the SIMD operation

Architecture of Warehouse-Scale Computers



Elements in warehouse-scale systems: 1U server (left), 7' rack with Ethernet switch (middle), and diagram of a small cluster with a cluster-level Ethernet switch/router (right).

Architecture of Warehouse-Scale Computers

- A set of low-end servers, typically in a 1U or blade enclosure format, are mounted within a rack
- interconnected using a local Ethernet switch.
- These rack-level switches, can use 1- or 10-Gbps links, have a number of uplink connections to one or more cluster-level (or datacenter-level) Ethernet switches.
- This second-level switching domain can span more than ten thousand individual servers

Architecture of Warehouse-Scale Computers

1. Storage

- include disks inside the server, and rely on Ethernet connectivity for access to information on the disks of remote servers.
- The alternative would be to use network attached storage (NAS) like Infiniband.
- NAS solution is generally more expensive per terabyte of storage
- It provides many features, including RAID techniques to improve dependability of the storage.
- WSCs generally rely on local disks and provide storage software that handles connectivity and dependability.

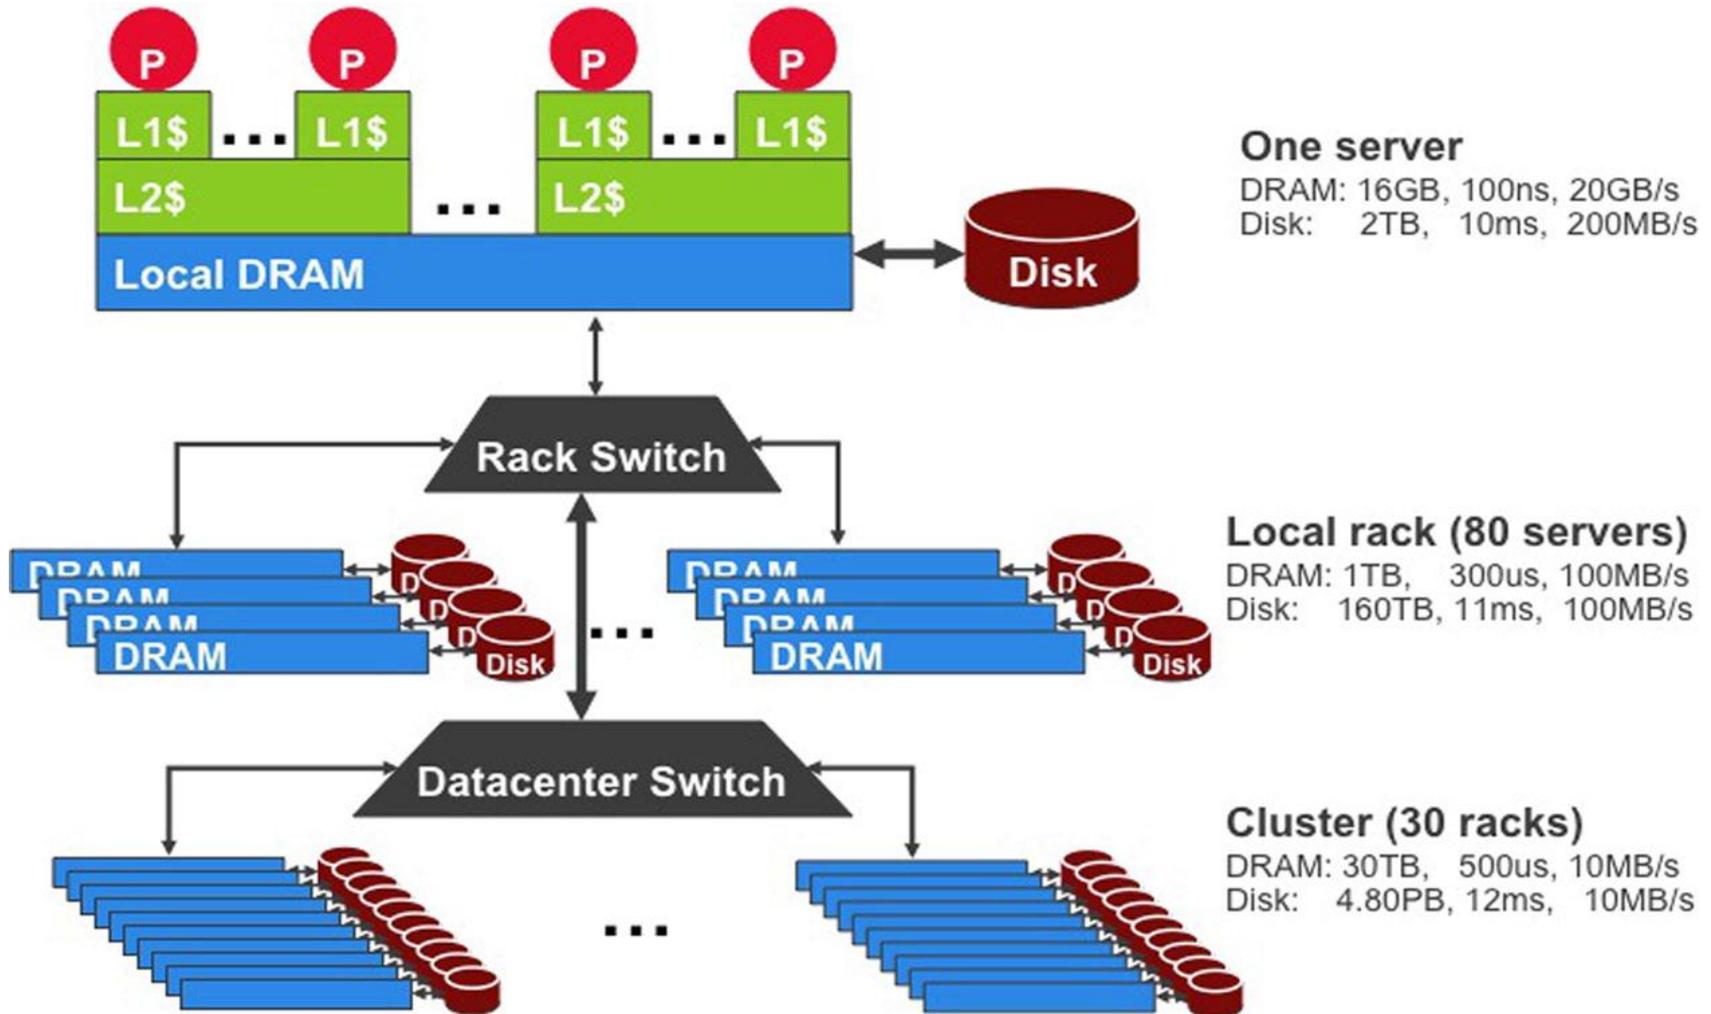
Architecture of Warehouse-Scale Computers

2. Array Switch

- It connects an array of racks is more expensive than the 48-port commodity Ethernet switch
- Network switches with high port counts, which are needed to tie together WSC clusters, have ten times more expensive (per 1-Gbps port) than commodity switches.
- A switch that has 10 times the bi-section bandwidth costs about 100 times as much.
- One reason is that the cost of switch bandwidth for n ports can grow as n^2 .

Architecture of Warehouse-Scale Computers

3. WSC Memory Hierarchy



Physical Infrastructure & Costs of WSC

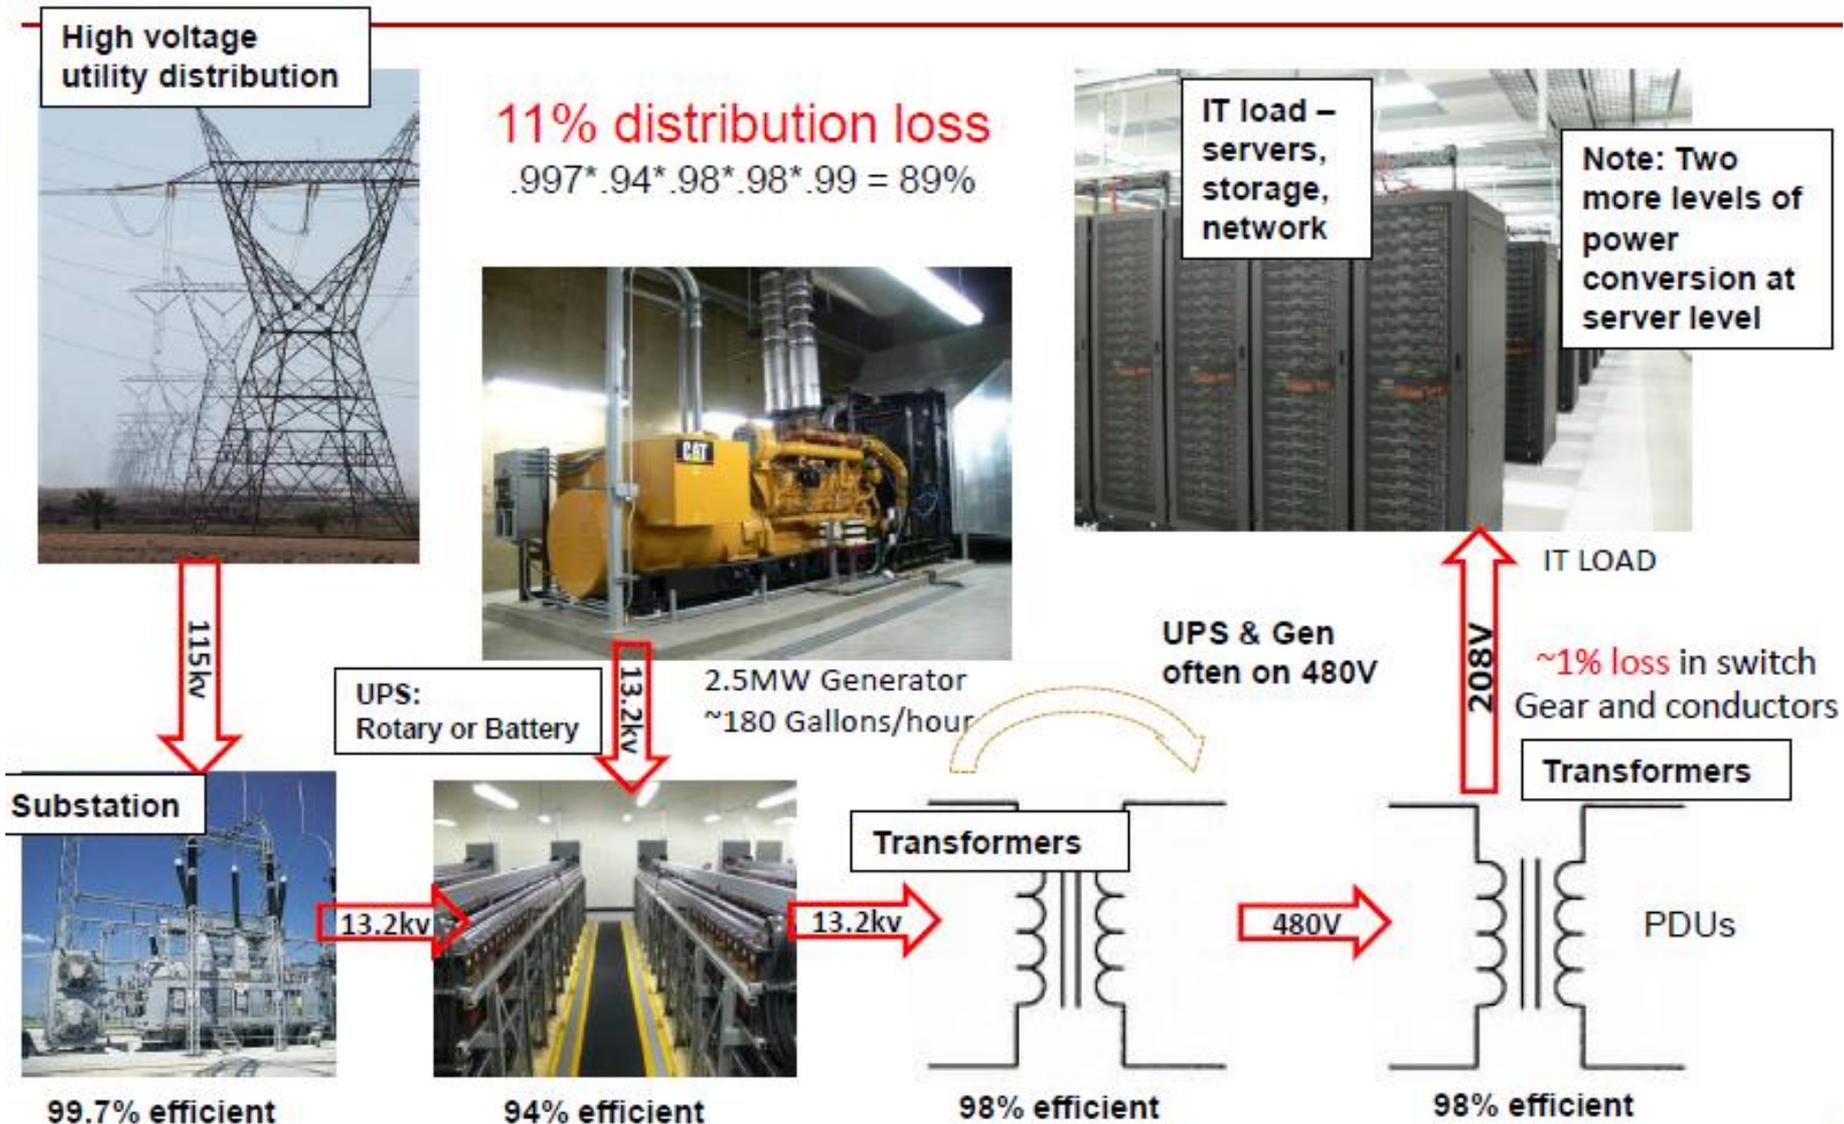


Fig: Power distribution and where losses occur

Physical Infrastructure & Costs of WSC

1. The substation switches from 115,000 volts to medium-voltage lines of 13,200 volts, with an efficiency of 99.7%.
2. To prevent the whole WSC from going offline if power is lost, a WSC has an UPS.
 - it involves large diesel engines that can take over from the utility company

Physical Infrastructure & Costs of WSC

2. UPS plays three roles:

- **power conditioning** (maintain proper voltage levels)
- **holding the electrical load** while the **generators start** and come on line,
- **holding the electrical load** when **switching back from the generators** to the electrical utility.
- Efficiency of this large UPS is **94%**, the facility loses **6%** of the power by having a UPS.
- The WSC UPS can account for **7% to 12%** of the cost of all the IT equipment.

Physical Infrastructure & Costs of WSC

3. Power Distribution Unit (PDU):

- It converts to low voltage, internal, three-phase power at 480 volts.
- The conversion efficiency is 98%.
- A typical PDU handles 75 to 225 kilowatts of load, or about 10 racks.

Physical Infrastructure & Costs of WSC

4. There is yet another down step to two-phase power at 208 volts that servers can use, once again at 98% efficiency.
5. The connectors, breakers, and electrical wiring to the server have a collective efficiency of 99%.

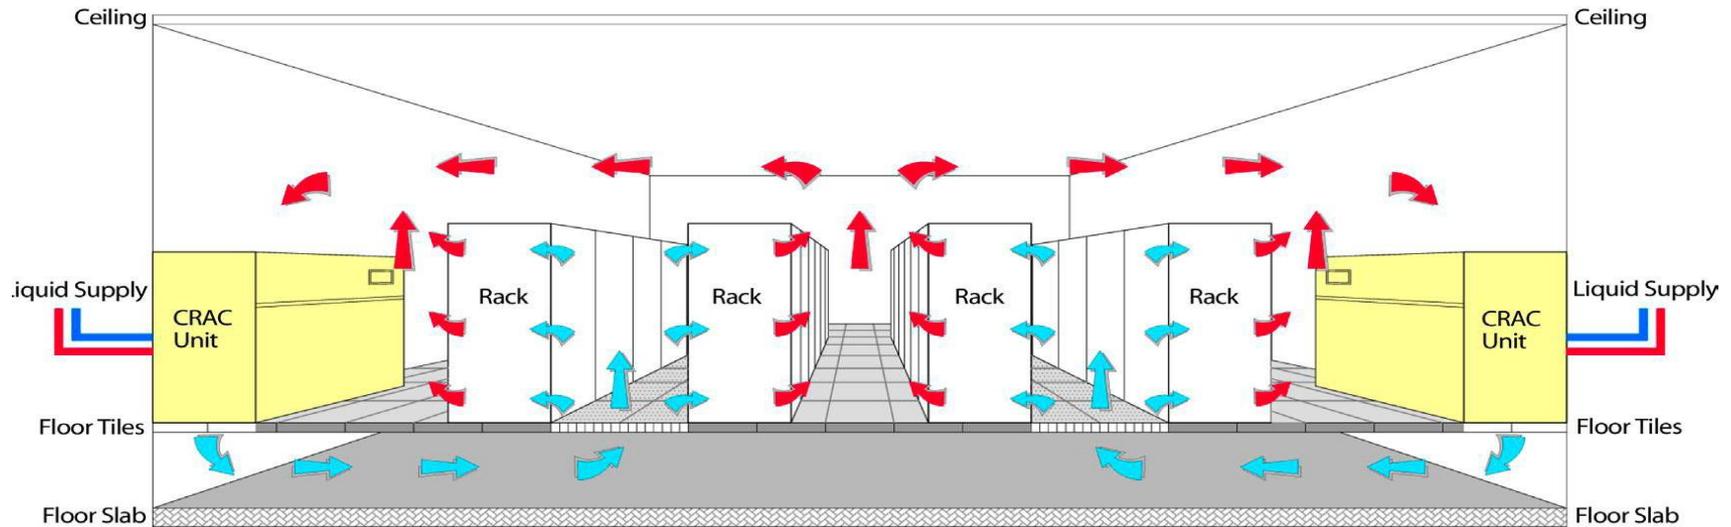
Physical Infrastructure & Costs of WSC

- the efficiency of turning 115,000-volt power from the utility into 208-volt power that servers can use is 89%:

$$99.7\% \times 94\% \times 98\% \times 98\% \times 99\% = 89\%$$

- WSCs outside North America use different conversion values, but the overall design is similar.

Physical Infrastructure & Costs of WSC



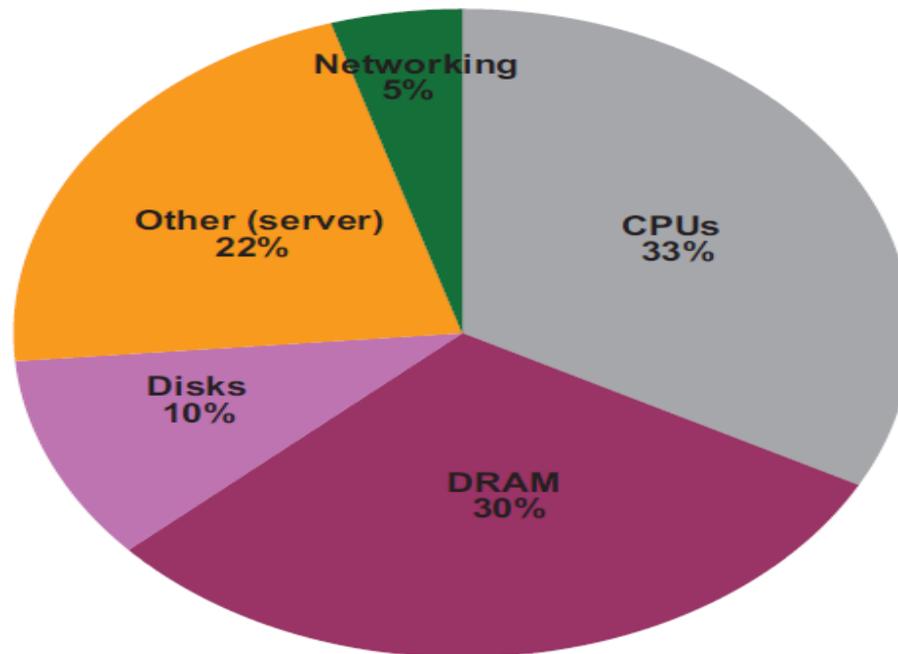
- CRAC = computer room air conditioning
- Cold air goes through servers and exits in hot aisle
- Cold aisles ~18-22C, hot aisles ~35C
- CRAC units consume significant amount of energy!

Physical Infrastructure & Costs of WSC

- The relative power costs of cooling equipment to IT equipment in a typical datacenter [Barroso and Holzle 2009] are as follows:
 - **Chillers account for 30% to 50% of the IT equipment power.**
 - **CRAC accounts for 10% to 20% of the IT equipment power, due to fans.**

Physical Infrastructure & Costs of WSC

- power usage inside the IT equipment itself, Barroso and Holzle [2009] reported the following for a Google WSC deployed in 2007:



Efficiency of a WSC

- Simple metric to evaluate the efficiency of a datacenter or a WSC is called power utilization effectiveness (PUE)
- $PUE = (\text{Total facility power}) / (\text{IT equipment power})$
- Thus PUE must be greater than or equal to 1
- Ranges from 1.33 to 3.03, median of 1.69
- Bigger the PUE the less efficient the WSC.

Cost of a WSC

- WSCs costs includes :
- **Capital expenditure(CAPEK):**
 - Infrastructure costs for the building, power delivery, cooling, and servers
- **Operational expenditure (OPEK):**
 - the monthly bill for energy, failures, personnel, etc
- CapEx can be amortized into a monthly estimate by assuming that
 - the facilities will last 10 years
 - server parts will last 3 years
 - networking parts will last 4 years

Cost of a WSC Case Study

- 8 MW facility : facility cost: \$88M, server/networking cost: \$79M
- Monthly expense: \$3.8M. Breakdown:
 - Servers 53% (amortized CapEx)
 - Networking 8% (amortized CapEx)
 - Power/cooling infrastructure 20% (amortized CapEx)
 - Other infrastructure 4% (amortized CapEx)
 - Monthly power bill 13% (true OpEx)
 - Monthly personnel salaries 2% (true OpEx)

Cloud Computing

- **Cloud Computing** is a general term used to describe a new class of network based computing that takes place over the Internet,
 - basically a step on from Utility Computing
 - a collection/group of integrated and networked hardware, software and Internet infrastructure (called a platform).
 - Using the Internet for communication and transport provides hardware, software and networking services to clients
- These platforms hide the complexity and details of the underlying infrastructure from users and applications by providing very simple graphical interface or API (Applications Programming Interface).

Cloud Computing

- In addition, the platform provides on demand services, that are always on, anywhere, anytime and any place.
- Pay for use and as needed, elastic
 - scale up and down in capacity and functionalities
- The hardware and software services are available to
 - general public, enterprises, corporations and businesses markets

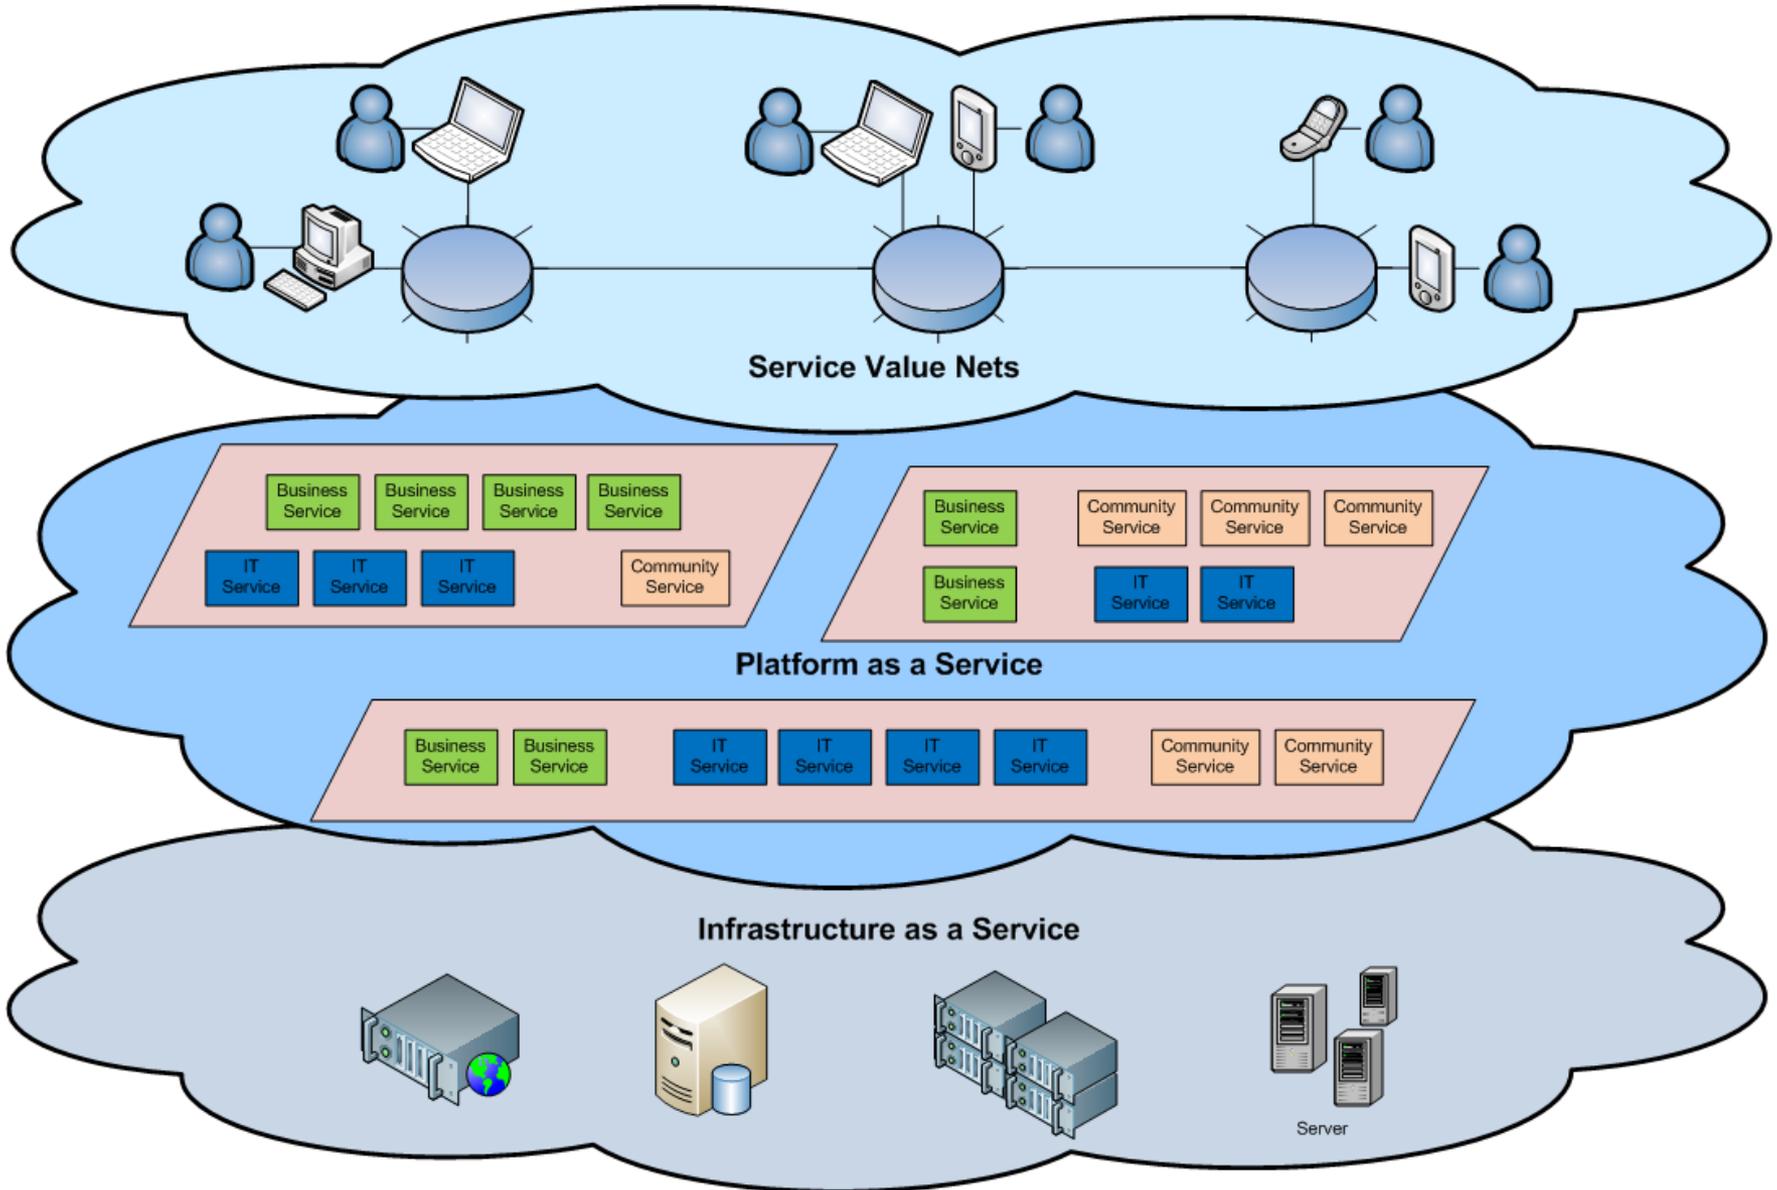
Cloud Computing

- Cloud computing is an umbrella term used to refer to Internet based development and services
- A number of characteristics define cloud data, applications services and infrastructure:
 - **Remotely hosted:** Services or data are hosted on remote infrastructure.
 - **Ubiquitous:** Services or data are available from anywhere.
 - **Commodified:** The result is a utility computing model similar to traditional that of traditional utilities, like gas and electricity - you pay for what you would want!

What is the purpose and benefits?

- Cloud computing enables companies and applications, which are system infrastructure dependent, to be infrastructure-less.
- By using the Cloud infrastructure on “pay as used and on demand”, all of us can save in capital and operational investment!
- Clients can:
 - Put their data on the platform instead of on their own desktop PCs and/or on their own servers.
 - They can put their applications on the cloud and use the servers within the cloud to do processing and data manipulations etc.

Cloud Architecture



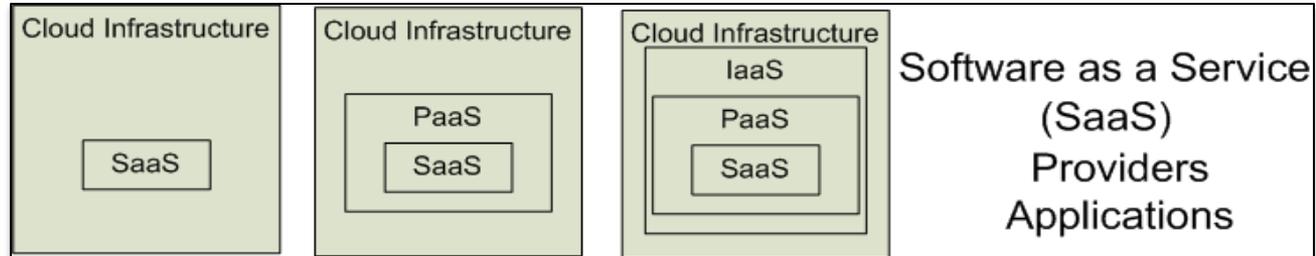
Cloud Service Models

Software as a Service (SaaS)

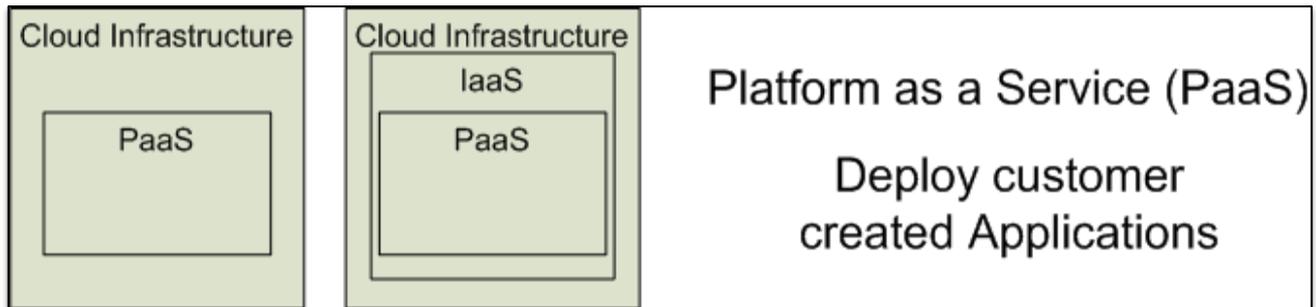
Platform as a Service (PaaS)

Infrastructure as a Service (IaaS)

SalesForce CRM
LotusLive

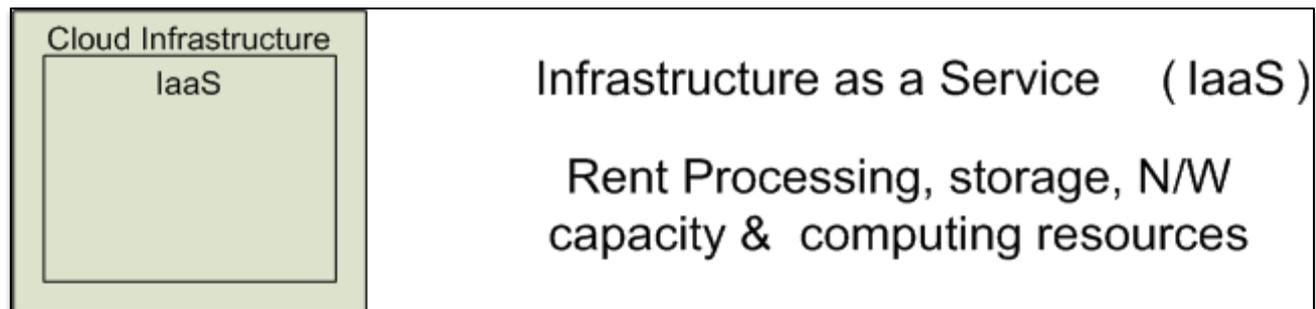


Google App Engine
Windows Azure
The Future Made Familiar



amazon web services™

rackspace®
HOSTING



Cloud Computing Service Layers

	Services	Description
Application Focused	Services	Services – Complete business services such as PayPal, OpenID, OAuth, Google Maps, Alexa
	Application	Application – Cloud based software that eliminates the need for local installation such as Google Apps, Microsoft Online
	Development	Development – Software development platforms used to build custom cloud based applications (PAAS & SAAS) such as SalesForce
Infrastructure Focused	Platform	Platform – Cloud based platforms, typically provided using virtualization, such as Amazon ECC, Sun Grid
	Storage	Storage – Data storage or cloud based NAS such as CTERA, iDisk, CloudNAS
	Hosting	Hosting – Physical data centers such as those run by IBM, HP, NaviSite, etc.

Basic Cloud Characteristics

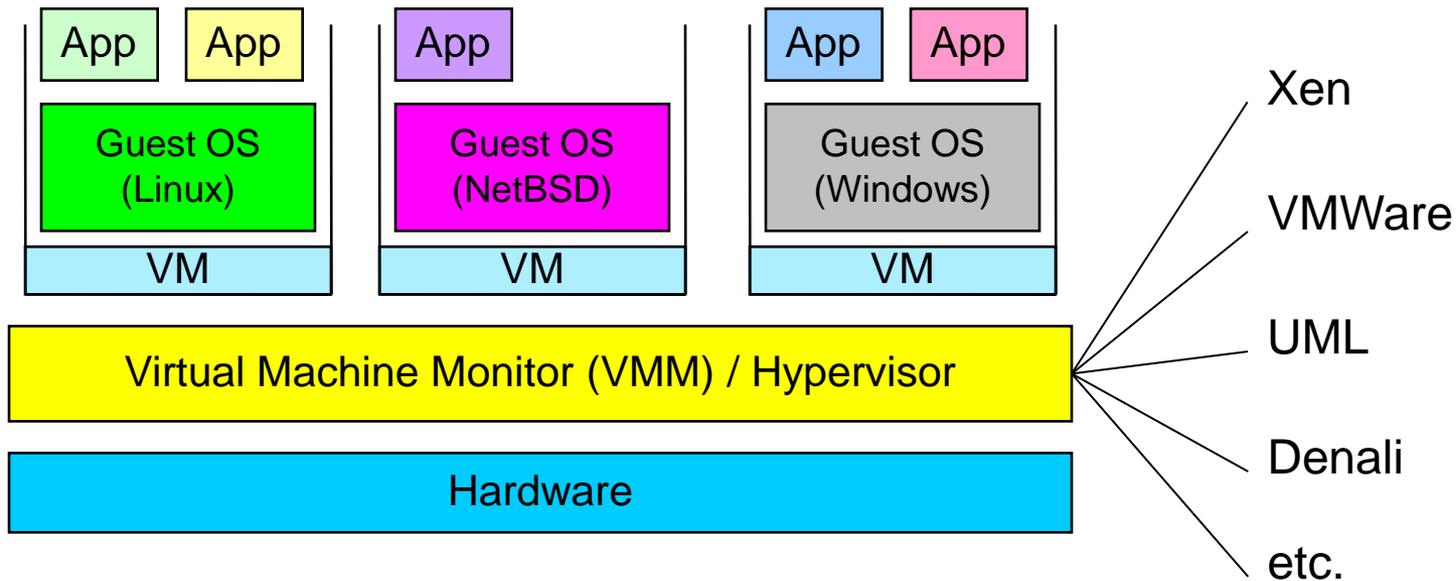
- The “**no-need-to-know**” in terms of the underlying details of infrastructure, applications interface with the infrastructure via the APIs.
- The “**flexibility and elasticity**” allows these systems to scale up and down at will
 - utilising the resources of all kinds
 - CPU, storage, server capacity, load balancing, and databases
- The “**pay as much as used and needed**” type of utility computing and the “**always on!, anywhere and any place**” type of network-based computing

Basic Cloud Characteristics

- Cloud are transparent to users and applications, they can be built in multiple ways
 - branded products, proprietary open source, hardware or software, or just off-the-shelf PCs.
- In general, they are built on clusters of PC servers and off-the-shelf components plus Open Source software combined with in-house applications and/or system software.

Virtual Machines

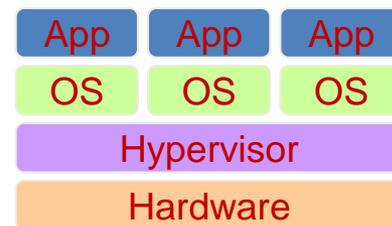
- VM technology allows multiple virtual machines to run on a single physical machine.



Performance: Para-virtualization (e.g. Xen) is very close to raw physical performance!

Virtual Machines

- VM technology allows multiple virtual machines to run on a single physical machine.
- Virtual workspaces:
 - An abstraction of an execution environment that can be made dynamically available to authorized clients by using well-defined protocols,
 - Resource quota (e.g. CPU, memory share),
 - Software configuration (e.g. O/S, provided services).
- Implement on Virtual Machines (VMs):
 - Abstraction of a physical host machine,
 - Hypervisor intercepts and emulates instructions from VMs, and allows management of VMs,
 - VM Ware, Xen, etc.
- Provide infrastructure API:
 - Plug-ins to hardware/support structures



Virtualized Stack

Virtual Machines

- Advantages of virtual machines:
 - Run operating systems where the physical hardware is unavailable,
 - Easier to create new machines, backup machines, etc.,
 - Software testing using “clean” installs of operating systems and software,
 - Emulate more machines than are physically available,
 - Timeshare lightly loaded systems on one host,
 - Debug problems (suspend and resume the problem machine),
 - Easy migration of virtual machines (shutdown needed or not).
 - Run legacy systems!

Cloud Storage

- Several large Web companies are now exploiting the fact that they have data storage capacity that can be hired out to others.
 - allows data stored remotely to be temporarily cached on desktop computers, mobile phones or other Internet-linked devices.
- Amazon's Elastic Compute Cloud (EC2) and Simple Storage Solution (S3) are well known examples
 - Mechanical Turk

Amazon Simple Storage Service (S3)

- Unlimited Storage.
- Pay for what you use:
 - \$0.20 per GByte of data transferred,
 - \$0.15 per GByte-Month for storage used,
 - Second Life Update:
 - 1TBytes, 40,000 downloads in 24 hours - \$200

Opportunities and Challenges

- The use of the cloud provides a number of opportunities:
 - It enables services to be used without any understanding of their infrastructure.
 - Cloud computing works using economies of scale:
 - It potentially lowers the outlay expense for start up companies, as they would no longer need to buy their own software or servers.
 - Cost would be by on-demand pricing.
 - Vendors and Service providers claim costs by establishing an ongoing revenue stream.
 - Data and services are stored remotely but accessible from “anywhere”.

Opportunities and Challenges

- In parallel there has been backlash against cloud computing:
 - Use of cloud computing means dependence on others and that could possibly limit flexibility and innovation:
 - The others are likely become the bigger Internet companies like Google and IBM, who may monopolise the market.
 - Some argue that this use of supercomputers is a return to the time of mainframe computing that the PC was a reaction against.
 - Security could prove to be a big issue:
 - It is still unclear how safe out-sourced data is and when using these services ownership of data is not always clear.
 - There are also issues relating to policy and access:
 - If your data is stored abroad whose policy do you adhere to?
 - What happens if the remote server goes down?
 - How will you then access files?
 - There have been cases of users being locked out of accounts and losing access to data.