Cloud Computing

Overview of Distributed Computing

Agenda

- Trends of computing
  - Utility computing
    - **Utility computing** is a service provisioning model in which a service provider makes **computing** resources and infrastructure management available to the customer as needed, and charges them for specific usage rather than a flat rate.
  - Cluster computing
    - A **computer cluster** consists of a set of loosely or tightly connected **computers** that work together so that, in many respects, they can be viewed as a single system. Unlike grid **computers**, **computer** clusters have each node set to perform the same task, controlled and scheduled by software.
  - Grid computing
    - **Grid computing** is a **computer** network in which each **computer's** resources are shared with every other **computer** in the system. Processing power, memory and data storage are all community resources that authorized users can tap into and leverage for specific tasks.
  - Cloud Computing

- What is the Difference Between Cluster Computing and Grid Computing?
- https://www.youtube.com/watch?v=fBURiN1omLo

- Cloud computing explained: https://www.youtube.com/watch?v=QJncFirhjPg

What's Computing

- Calculate
- Game theory
- Embattle
- Your thinking
- …

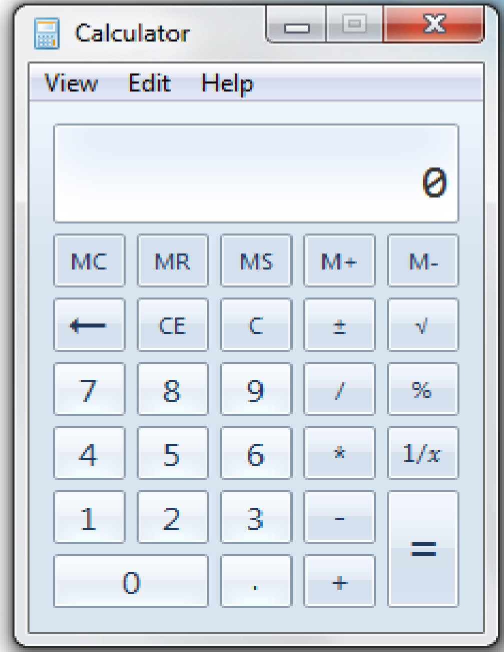

What's Computing

- In computer science
  - In a general way, we can define computing to mean any goal-oriented activity requiring, benefiting from, or creating computers.

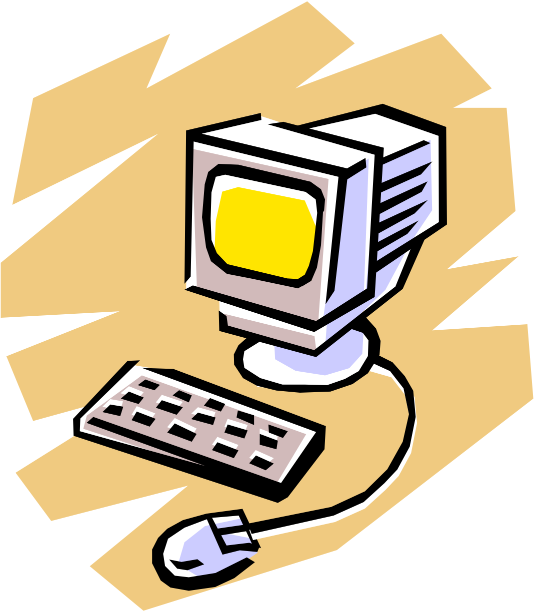

Computing in Computer Science

- Computing includes
  - Designing and building hardware and software systems for a wide range of purposes
  - Processing, structuring, and managing various kinds of information
  - Doing scientific studies using computers
  - Making computer systems behave intelligently
  - Creating and using communications and entertainment media
  - …

- https://www.youtube.com/watch?v=_HCWBS6k8j0
- Google’s AI Beats World’s Best ‘Go’ Player
- (20) Nvidia Finally Reveals The Future Of AI In 2025... – YouTube
- Nvidia Finally Reveals The Future Of AI In 2025

…computing may someday be organized as a public utility just as the telephone system is a public utility… The computer utility could become the basis of a new and important industry.

– John McCarthy (a professor of MIT) 1961.

As of now, computer networks are still in their infancy, but as they grow up and become sophisticated, we will probably see the spread of computer utilities which, like present electric and telephone utilities, will service individual homes and offices across the country.

– L. Kleinrock (one of the chief scientists of the original ARPANET project) 1969.

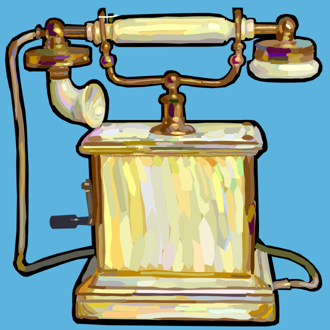

Computing will one day be…

The 5th Utility

The 5th Utility

- Traditional utilities
  - Water
  - Electricity
  - Natural gas
  - Telephone network

- *Computing* is being transformed to a model consisting of services that are commoditized and delivered in a manner similar to traditional utilities

→ Utility Computing

Utility Computing

- Computing as a Utility
  - Provide the basic level of computing service that is considered essential to meet the everyday needs
  - Users access services based on their requirements without regard to where the services are hosted or how they are delivered

- Offering computing resources as a metered service

The Creation of The Internet

- The Internet is the foremost milestone

- Enables individual computers to communicate with any other computers located elsewhere in the world

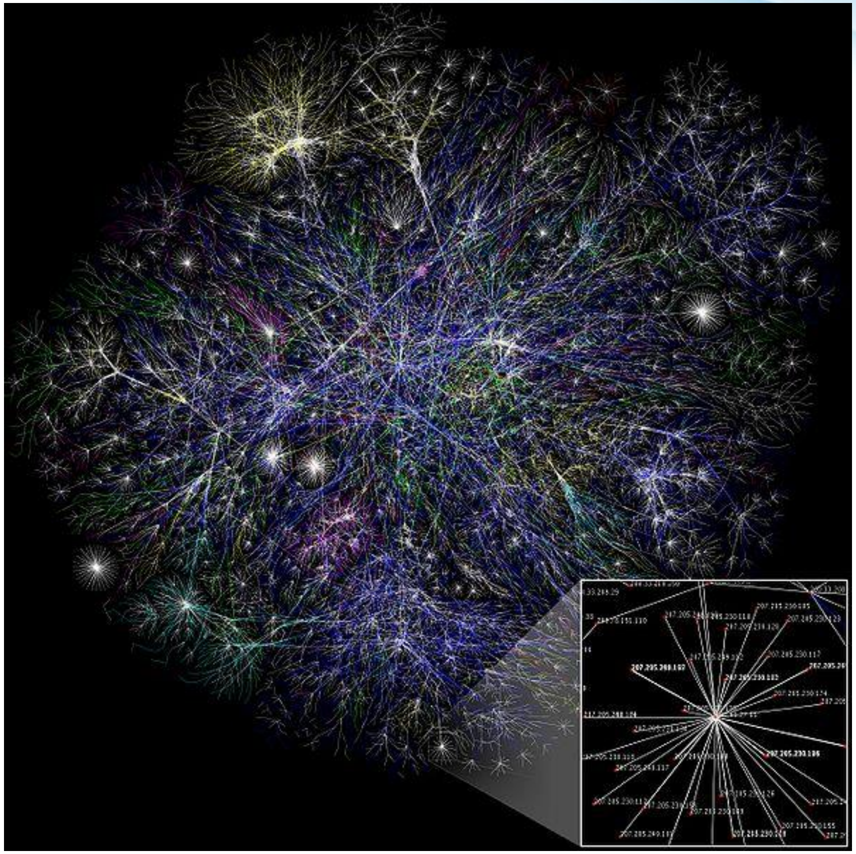

Variety of Paradigms

- New computing paradigms have been proposed and adopted to edge closer toward achieving the vision of *computer utilities*

Paradigms

Utility Computing
Cluster Computing
Grid Computing
Cloud Computing
…

+

Attribute

Accessibility
Manageability
Autonomic
Performance
Scalability
Availability
QoS
SLA
…

→

The fundamental notion is…

Distributed Computing

Distributed Computing

- A computer system in which several *interconnected computers* share the computing tasks assigned to the system

- Paradigms
  - Cluster computing
  - Grid computing
  - Cloud computing

Demand for Computing Power

- Solving grand challenge applications using computer *modeling*, *simulation* and *analysis*

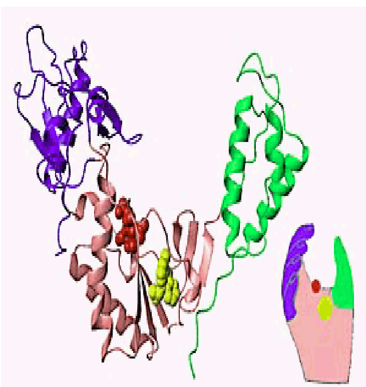

Life Sciences

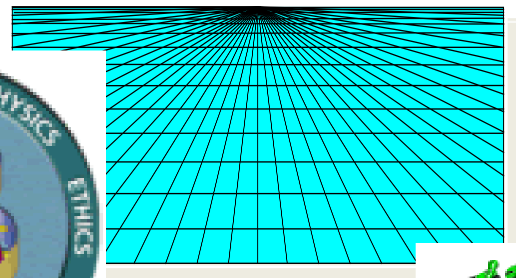

Aerospace

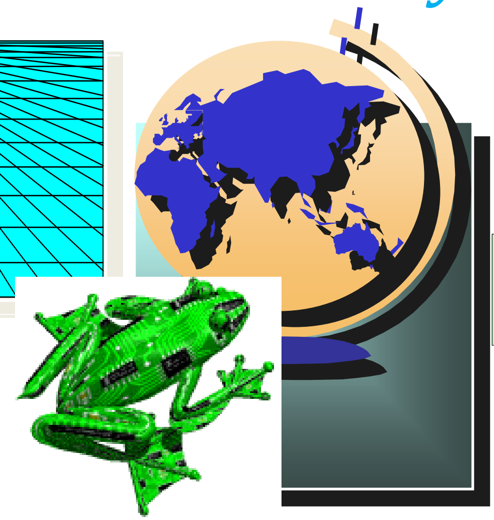

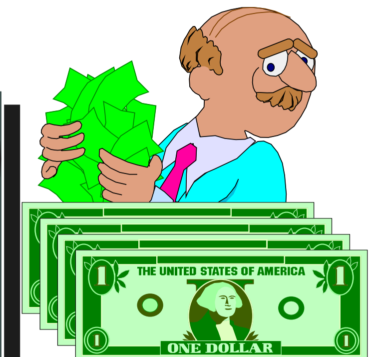

Internet & Ecommerce

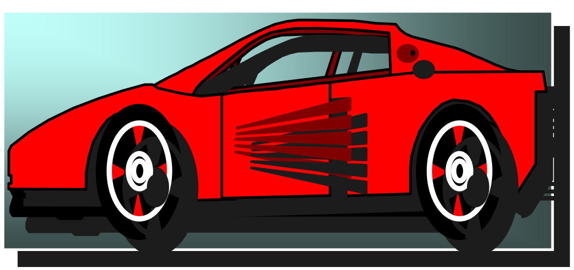

CAD/CAM

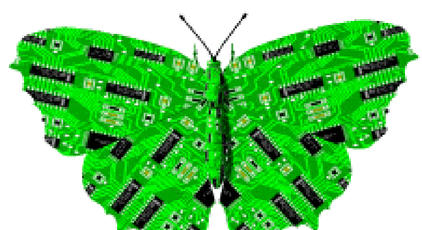

Digital Biology

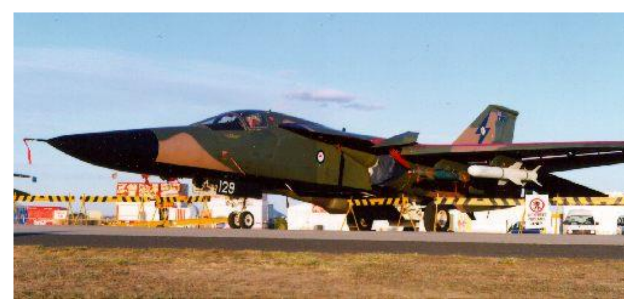

Military Applications

How to Run Applications Faster

- There are 3 ways to improve performance:
  - Work Harder
  - Work Smarter
  - Get Help

- Computer analogy
  - Using faster hardware
  - Using optimized algorithms and techniques to solve computational tasks
  - Using multiple computers to solve a particular task

History

- In the 1980s
  - Computer performance was best improved by creating faster and more efficient processors

- In the early 1990s
  - An increasing trend to move away from expensive and specialized proprietary parallel supercomputers towards networks of workstations

Towards Commodity Computing

- Trend of computing
  - From specialized traditional supercomputing platforms
  - To inexpensive, general purpose systems consisting of loosely coupled components built up from single or multiprocessor PCs or workstations
- Low-cost commodity supercomputing
  - Linking together two or more computers to jointly solve some computational problem
  - Providing high performance computational facilities for large-scale and grand-challenge applications

Traditional Food Chain

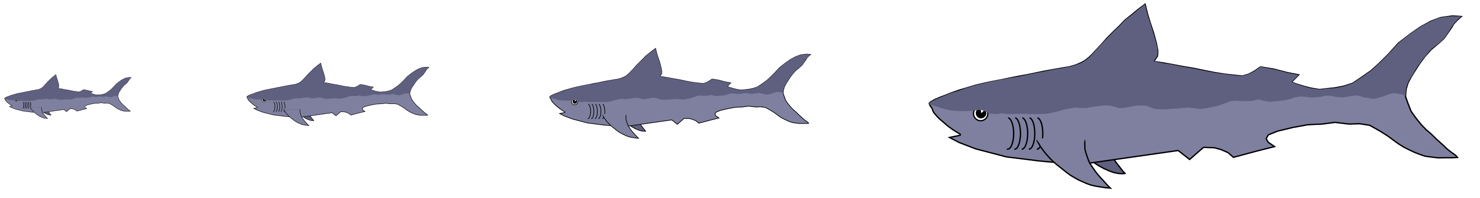

Food Chain of Computer

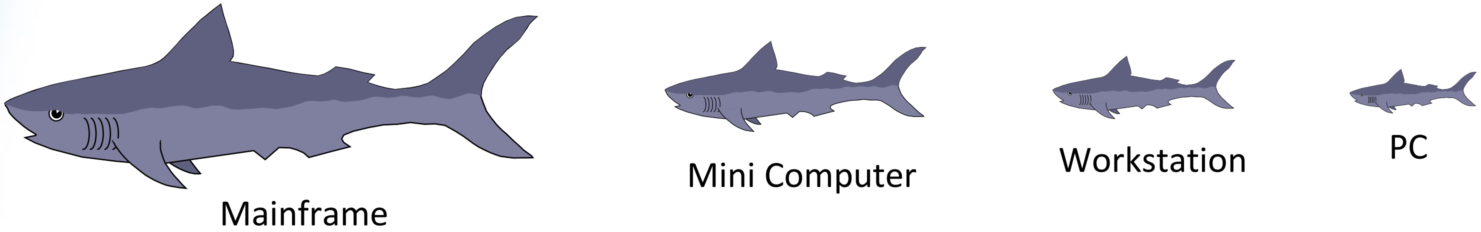

Food Chain of Distributed Computing

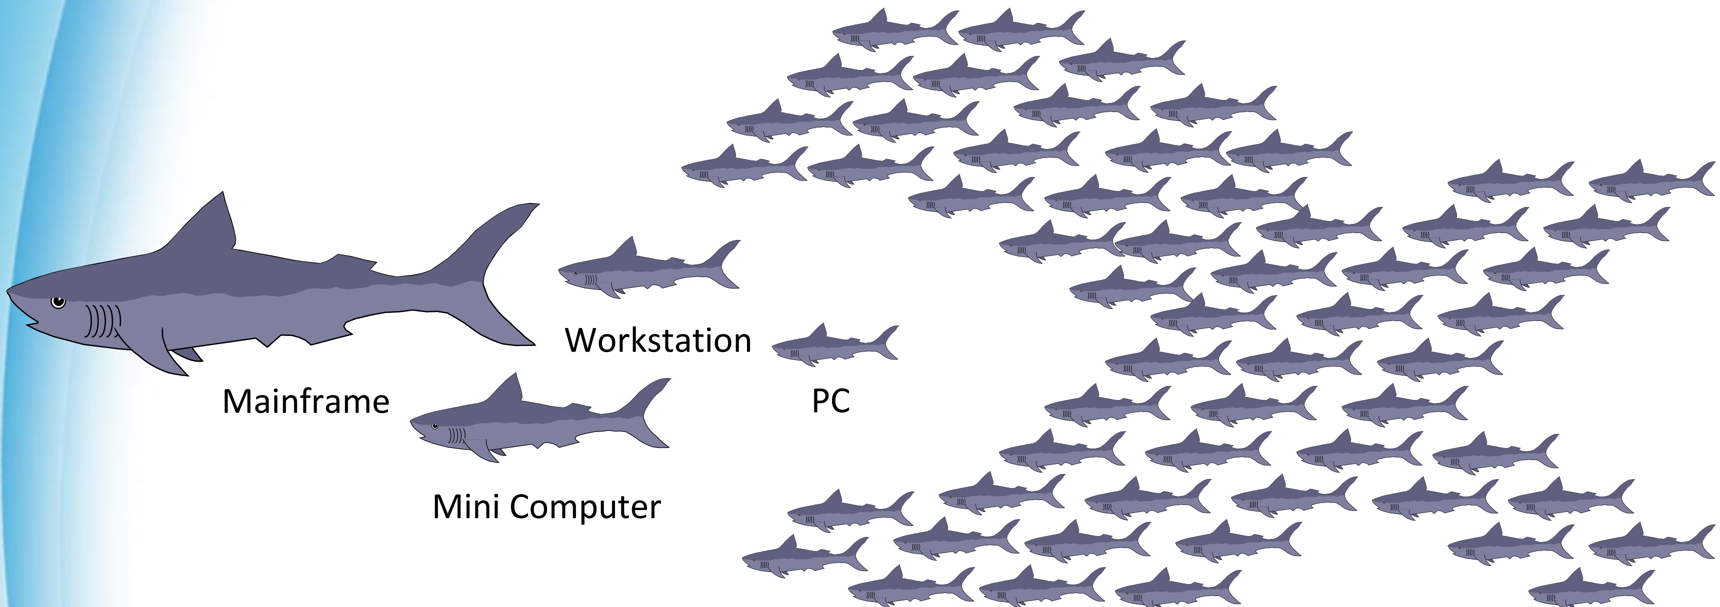

Rise & Fall of Computing Technologies

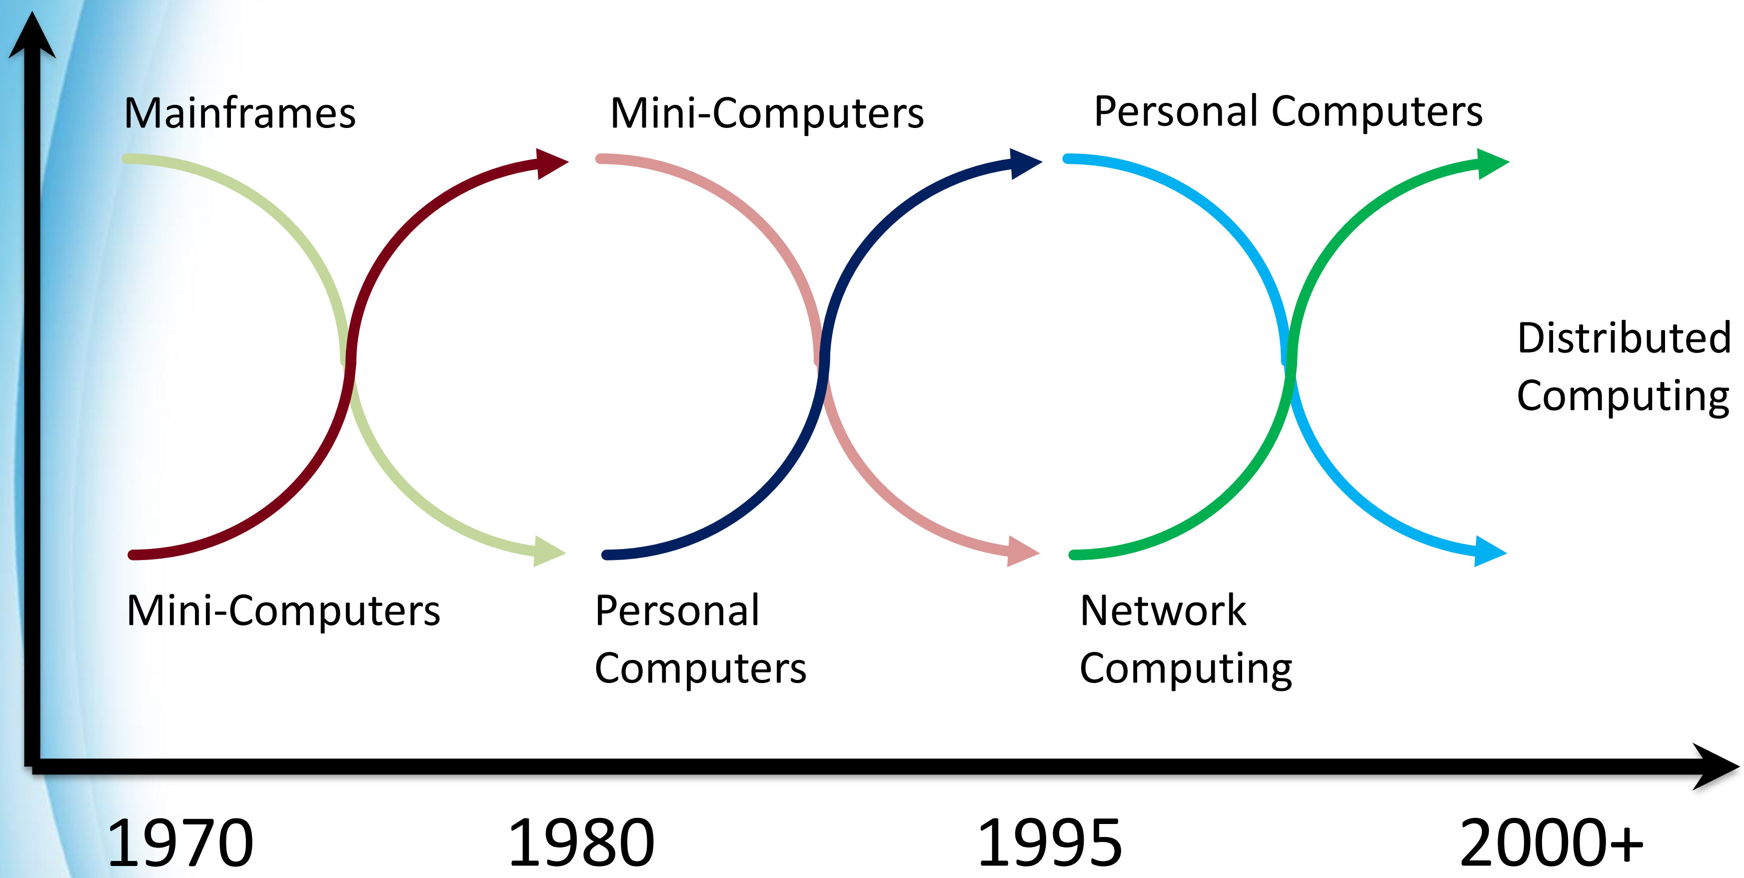

Moore's law

- **Moore's law** is the observation that the number of transistors in a dense integrated circuit doubles approximately every two years. The observation is named after Gordon Moore, the co-founder of Fairchild Semiconductor and Intel, whose 1965 paper described a doubling every year in the number of components per integrated circuit, and projected this rate of growth would continue for at least another decade. In 1975, looking forward to the next decade, he revised the forecast to doubling every two years. The period is often quoted as 18 months because of Intel executive David House, who predicted that chip performance would double every 18 months (being a combination of the effect of more transistors and the transistors being faster).

Cluster Computing

Grid Computing

Cloud Computing

DISTRIBUTED COMPUTING

Cluster Computing

- A cluster is a type of parallel and distributed system, which consists of a collection of *inter-connected stand-alone* computers working together as a *single* integrated computing resource

The Cluster

- A Node
  - A single or multiprocessor system with memory, I/O facilities, & OS

- A Cluster:
  - Generally two or more computers (nodes) connected together
  - In a single cabinet, or physically separated & connected via a LAN
  - Appear as a single system to users and applications
  - Provide a cost-effective way to gain features and benefits

Parallel Computing

- Parallel Computing
  - A form of computation in which many *calculations are carried out simultaneously*
  - Large problems can often be divided into smaller ones, which are then solved concurrently ("*in parallel*")

- Advantages
  - Cost efficient
  - High performance
  - Improved utilization

Parallel Programming Models

- Parallel programming model
  - A set of software technologies to express parallel algorithms and match applications with the underlying parallel systems
- Flynn’s taxonomy
  - Flynn's taxonomy is a classification of computer architectures, **proposed** by Michael J. Flynn in 1966
- Shared-memory or Distributed-memory model

Flynn's Taxonomy

- Programs and computers are classified by
  - whether they were operating using a single set or multiple sets of instructions
  - whether or not those instructions were using a single or multiple sets of data

	Single Instruction	Multiple Instructions
Single Data	SISD	MISD
Multiple Data	SIMD	MIMD

Two Basic Models

- Shared memory model
  - Memory can be simultaneously accessed by multiple process with an intent to provide communication among them or avoid redundant copies

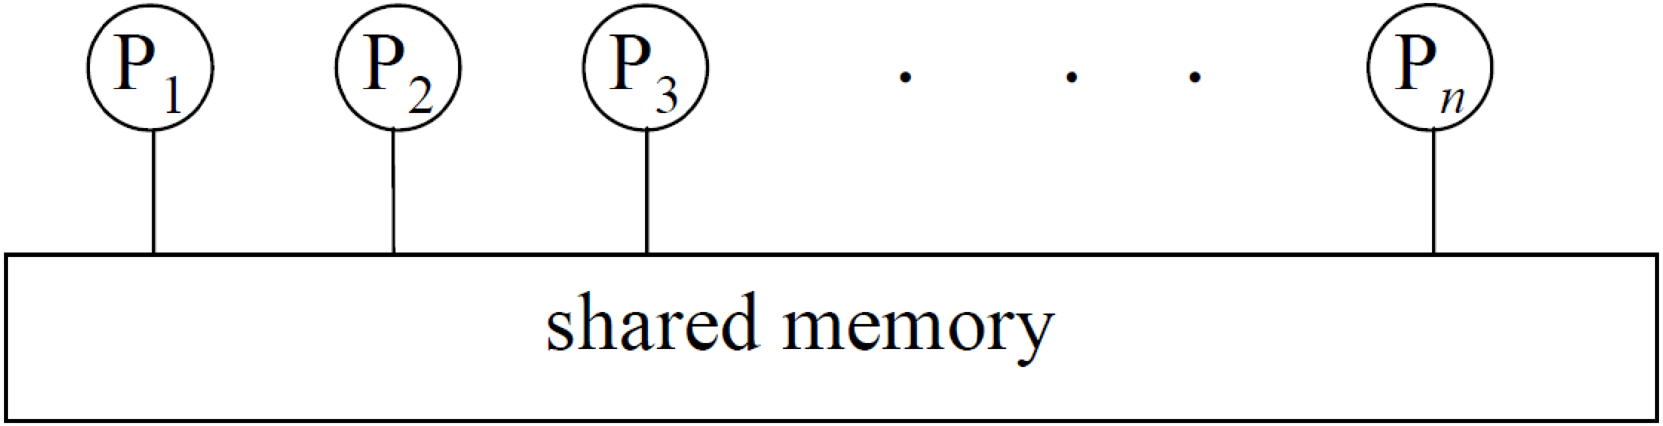

Two Basic Models

- Distributed memory model
  - A multiple-processor computer system in which each processor has its own private memory
  - Computational tasks can only operate on local data, and if remote data is required, the computational task must communicate with one or more remote processors

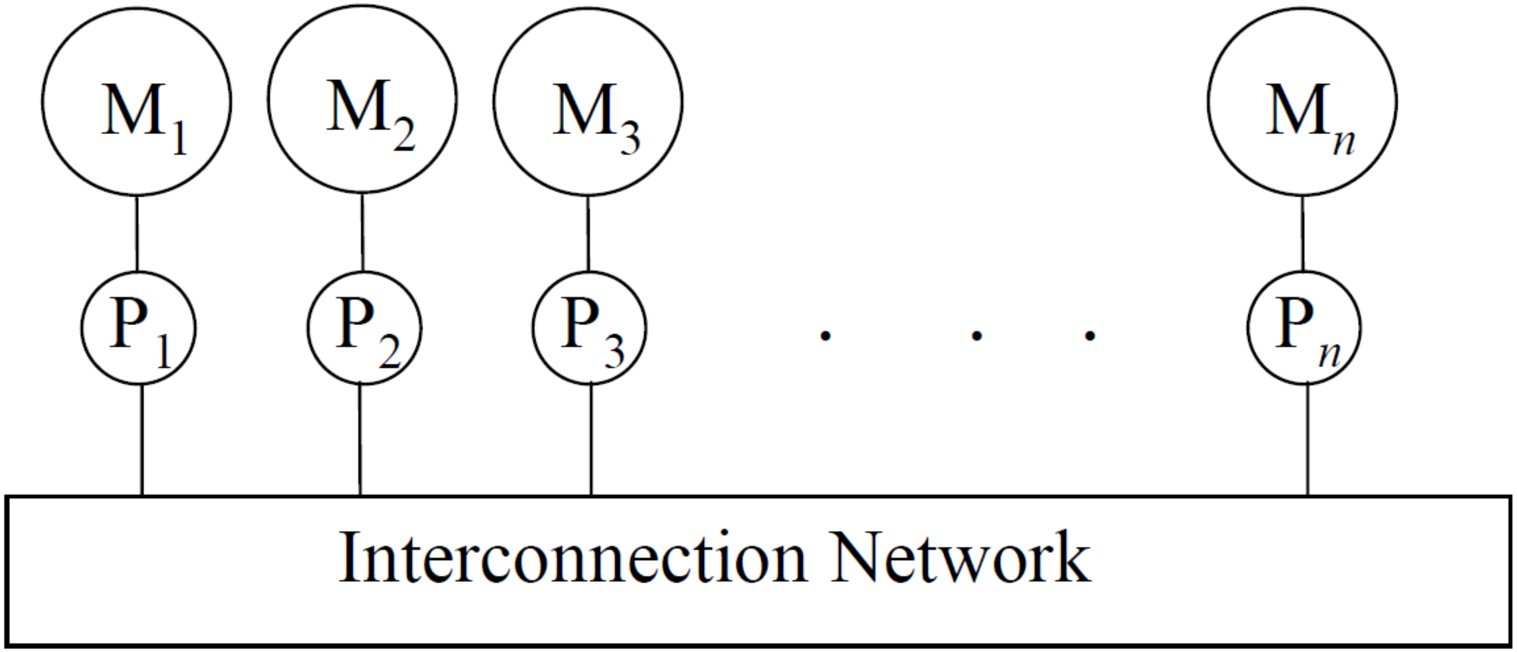

Computer Collaboration in Cluster

- How machines co-work in cluster?
  - Using MPI to inter-connect those stand-alone computers
- Message Passing Interface (MPI)
  - A specification that allows computers to communicate with each other
  - Use Message Passing to do synchronization

Message Passing Interface (MPI)

- Distributed memory programming model
- Complete computers connected through an interconnection network

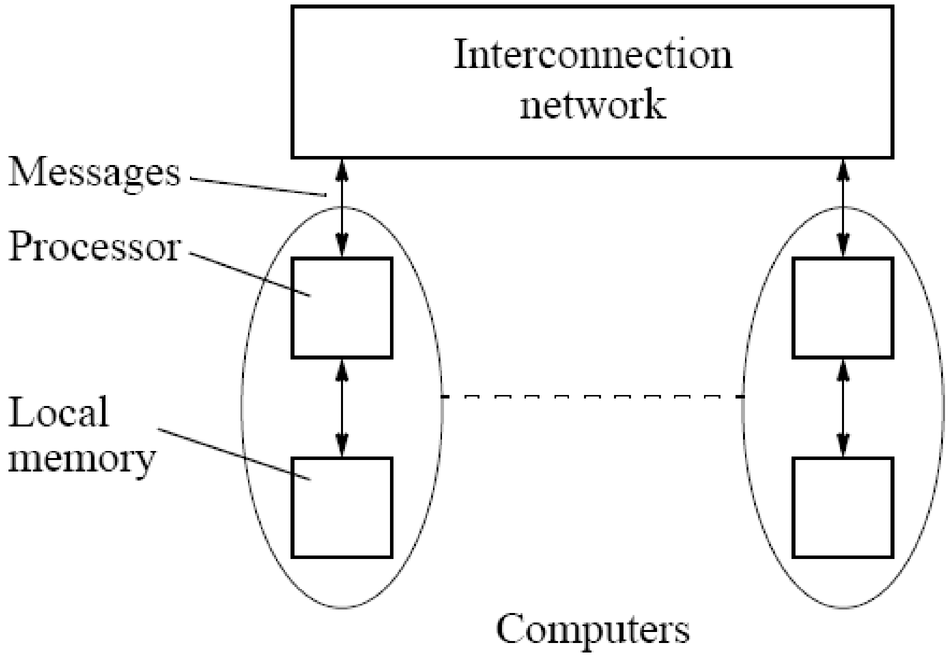

Message Passing Interface (MPI)

- Single program with multiple data (SPMD)
  - Similar to SIMD (single instruction multiple data)

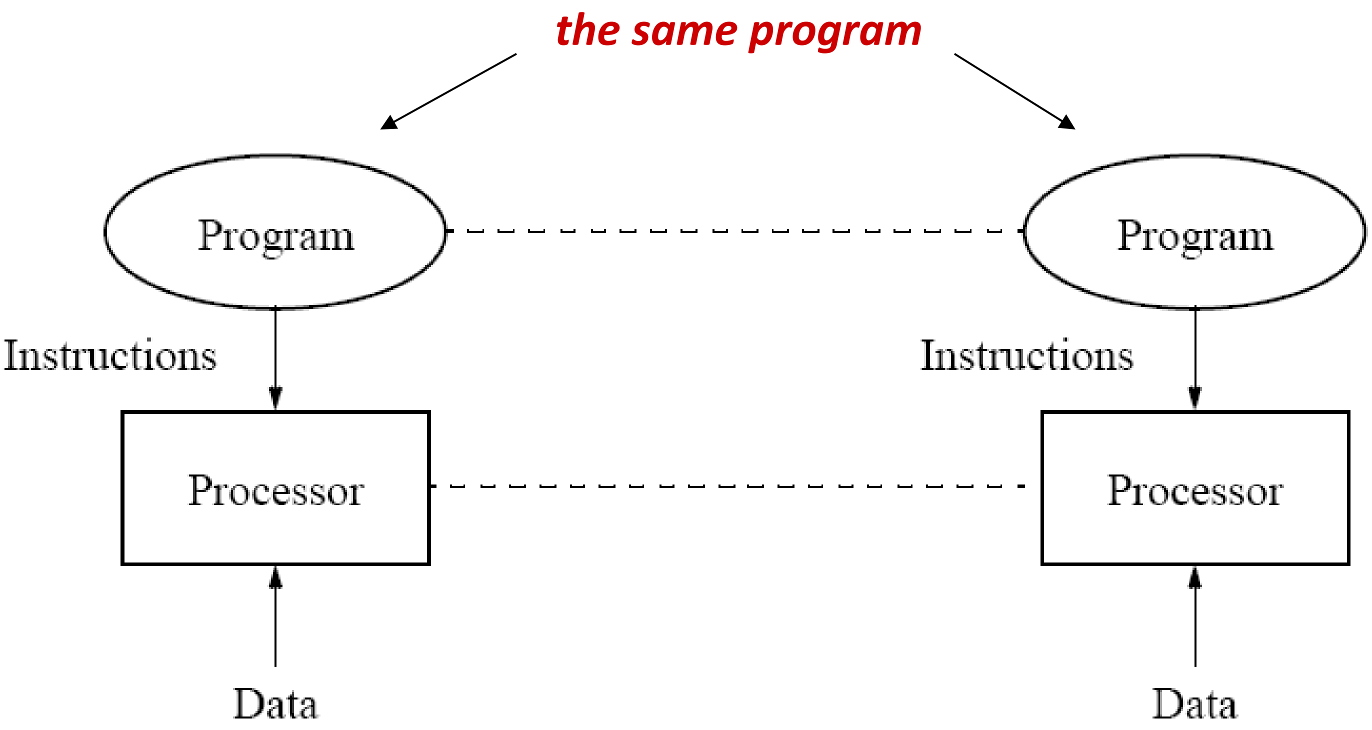

Message Passing Interface (MPI)

- Typical parallel execution flow :

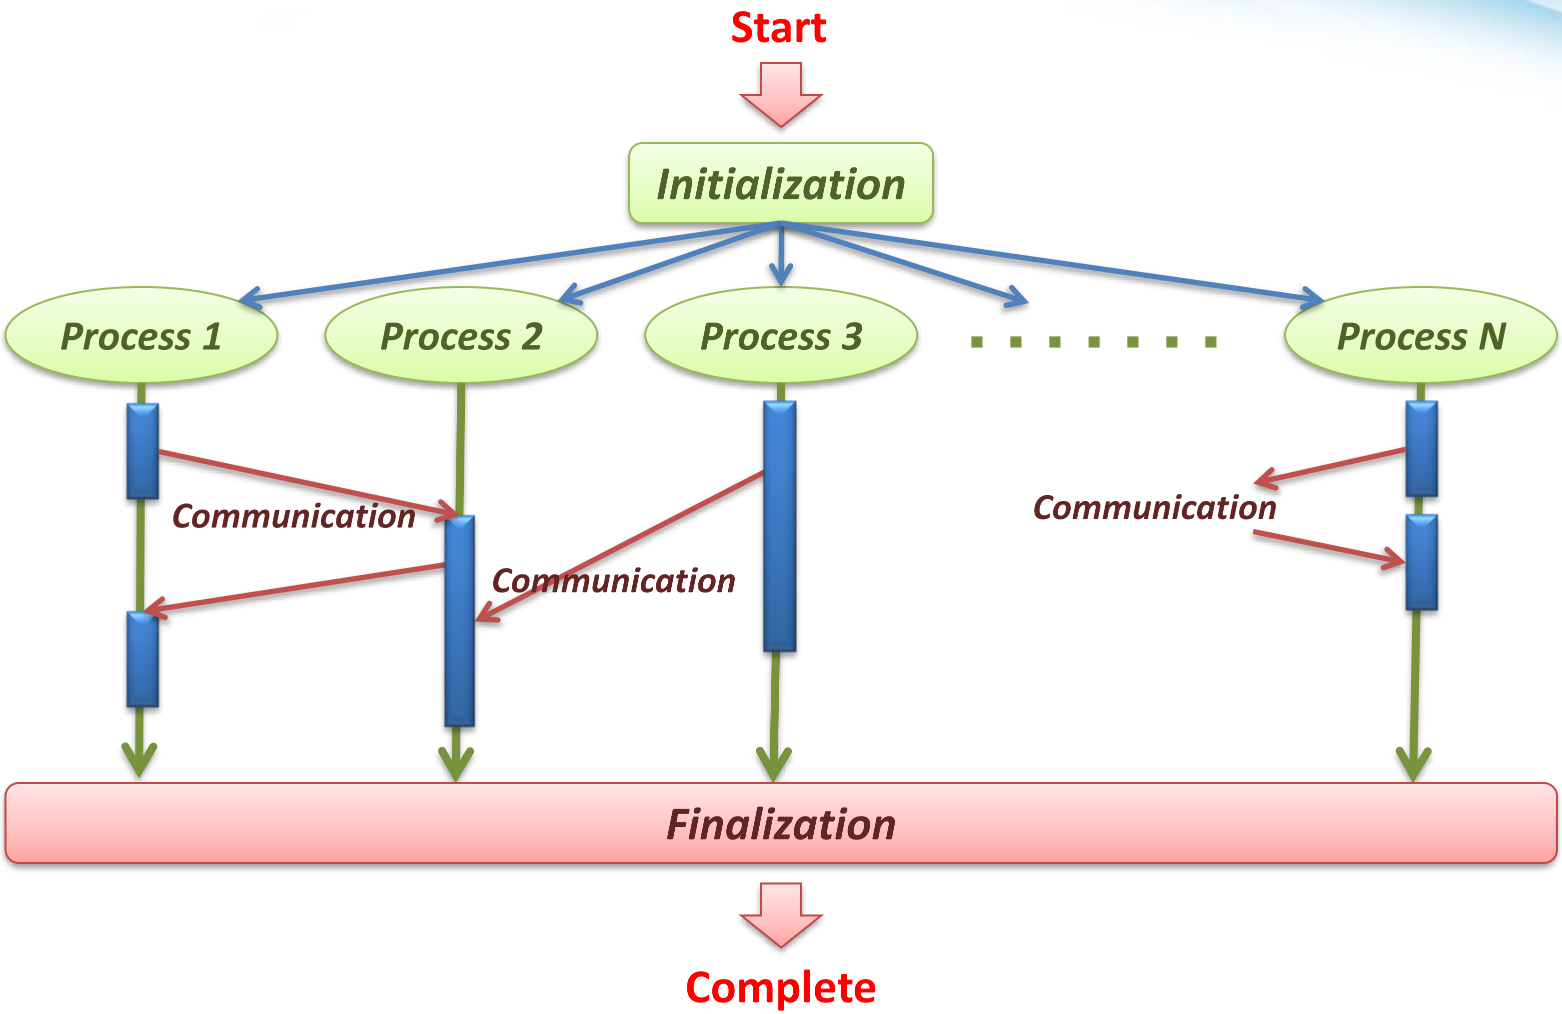

Cluster Applications

- Cluster had become the mainstream computing system

- Apply parallel program to
  - Computing intensive
  - Data intensive
  - Timing critical system
  - … any time you are pleased

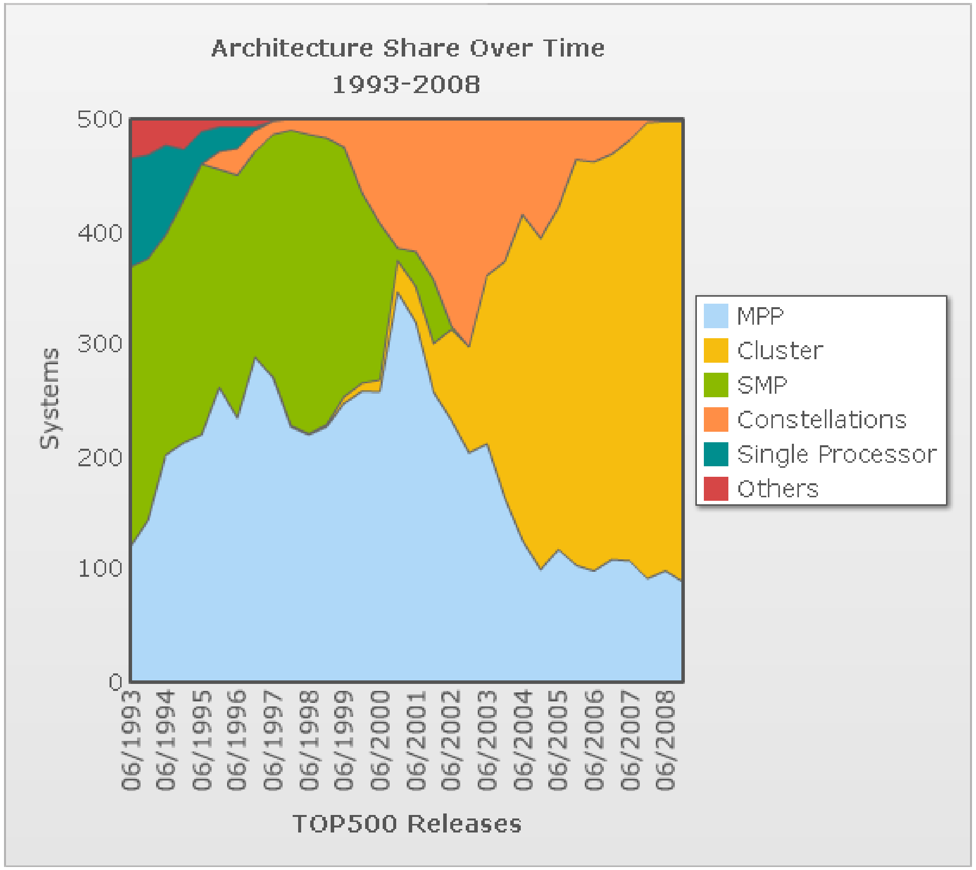

MPP: massively parallel processing
Symmetric multiprocessing (SMP) involves a multiprocessor computer hardware and software
Constellation: an extension of the cluster architecture to include large systems as the computing nodes

Cluster Computing

Grid Computing

Cloud Computing

DISTRIBUTED COMPUTING

Demand for More Computing Power

- The large-scale, computational-/data-intensive scientific applications require more resources

- Applications
  - Scientific application
  - Computer animation
  - Computer games
  - Image processing
  - Data mining
  - … etc

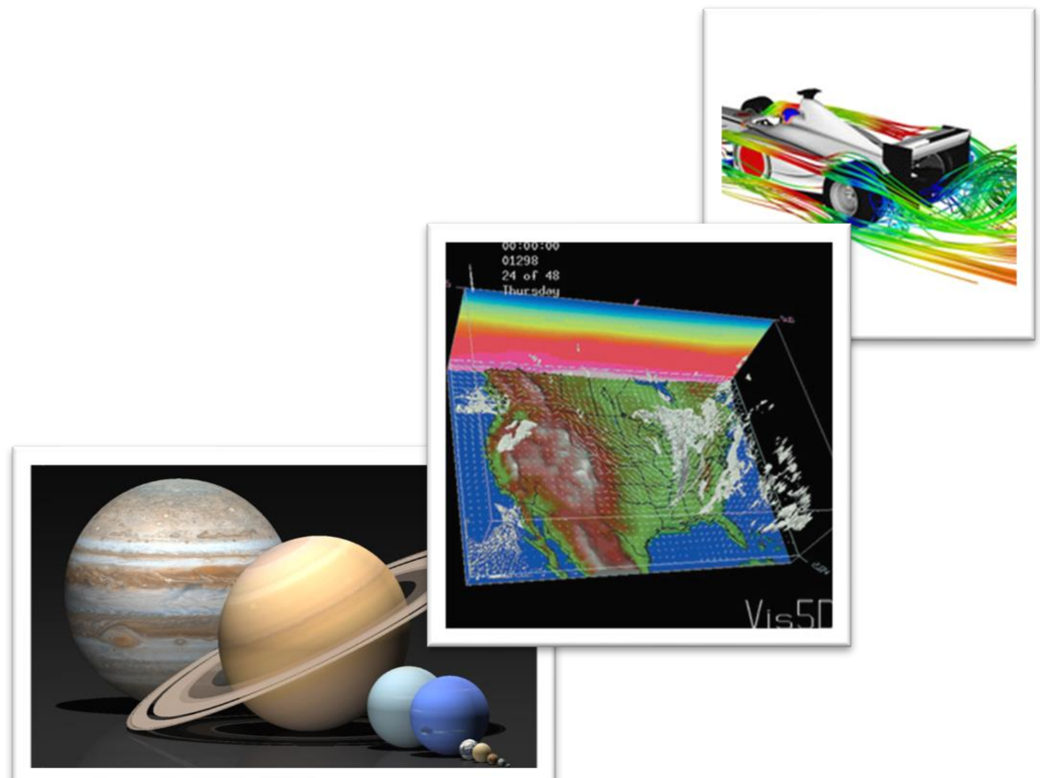

Grid Computing

- Grid computing
  - To coordinate resource sharing and problem solving in dynamic, multi-institutional virtual organization (VO)

- Provide the generic approach for a general resource-sharing framework that address the VO requirement
- Support the creation and use of computation- and data-enriched environments
- Uses open standards and interfaces

The Grid

- Grid
  - An analogy with the electric power grid around 1910
  - In the mid-1990s, this term is coined to denote a proposed *distributed computing infrastructure* for advanced science and engineering

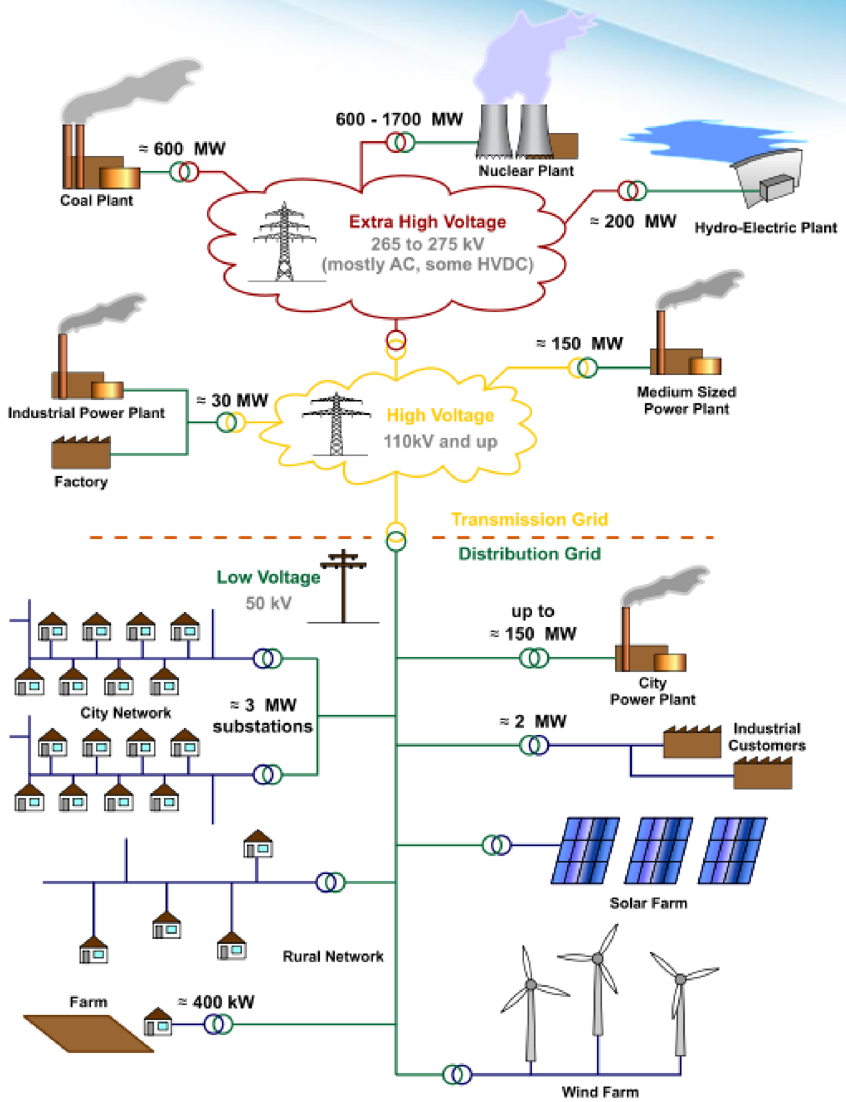

Resource Collaboration in Grids

- Virtual Organization (VO)
  - A set of individuals and/or institutions defined by such sharing rules form
    - what is shared
    - who is allowed to share
    - the conditions under which sharing occurs

- An actual organization can participate in one or more VOs by sharing some or all of its resources

Example of VOs

- Three actual organizations (the ovals) & Two VOs (P, Q)
  - P: links participants in an aerospace design consortium
  - Q: links colleagues who have agreed to share spare computing to run ray tracing computations

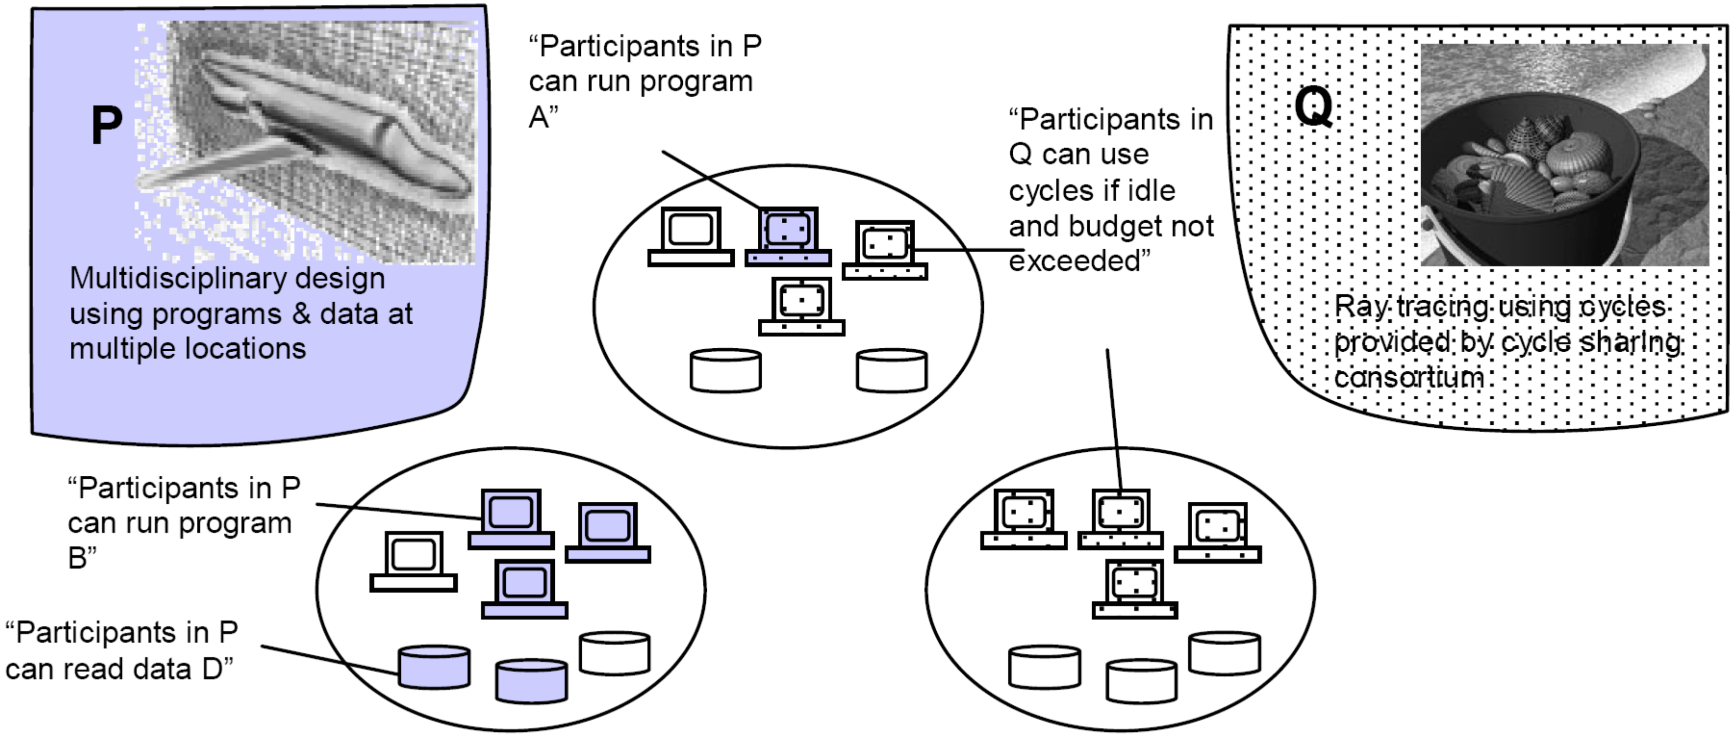

Grid Environments

- Resources are heterogeneous
  - supercomputers, storage systems, data sources, and specialized devices owned by different administrative domains
- Enables the sharing, selection, and aggregation of a wide variety of geographically distributed resources
- To solve large scale resource-intensive problems in science, engineering, and commerce
- Hailed as the next revolution after the Internet and the World Wide Web

UVA Computational Resources

- The **TOP500** project ranks and details the 500 most powerful non-distributed computer systems in the world https://www.top500.org/lists/2017/06/
- UVA cluster computational resources
  - **Rivanna: https://arcs.virginia.edu/rivanna**
  - **Ivy: https://www.rc.virginia.edu/tags/ivy/**

Basic System Architecture

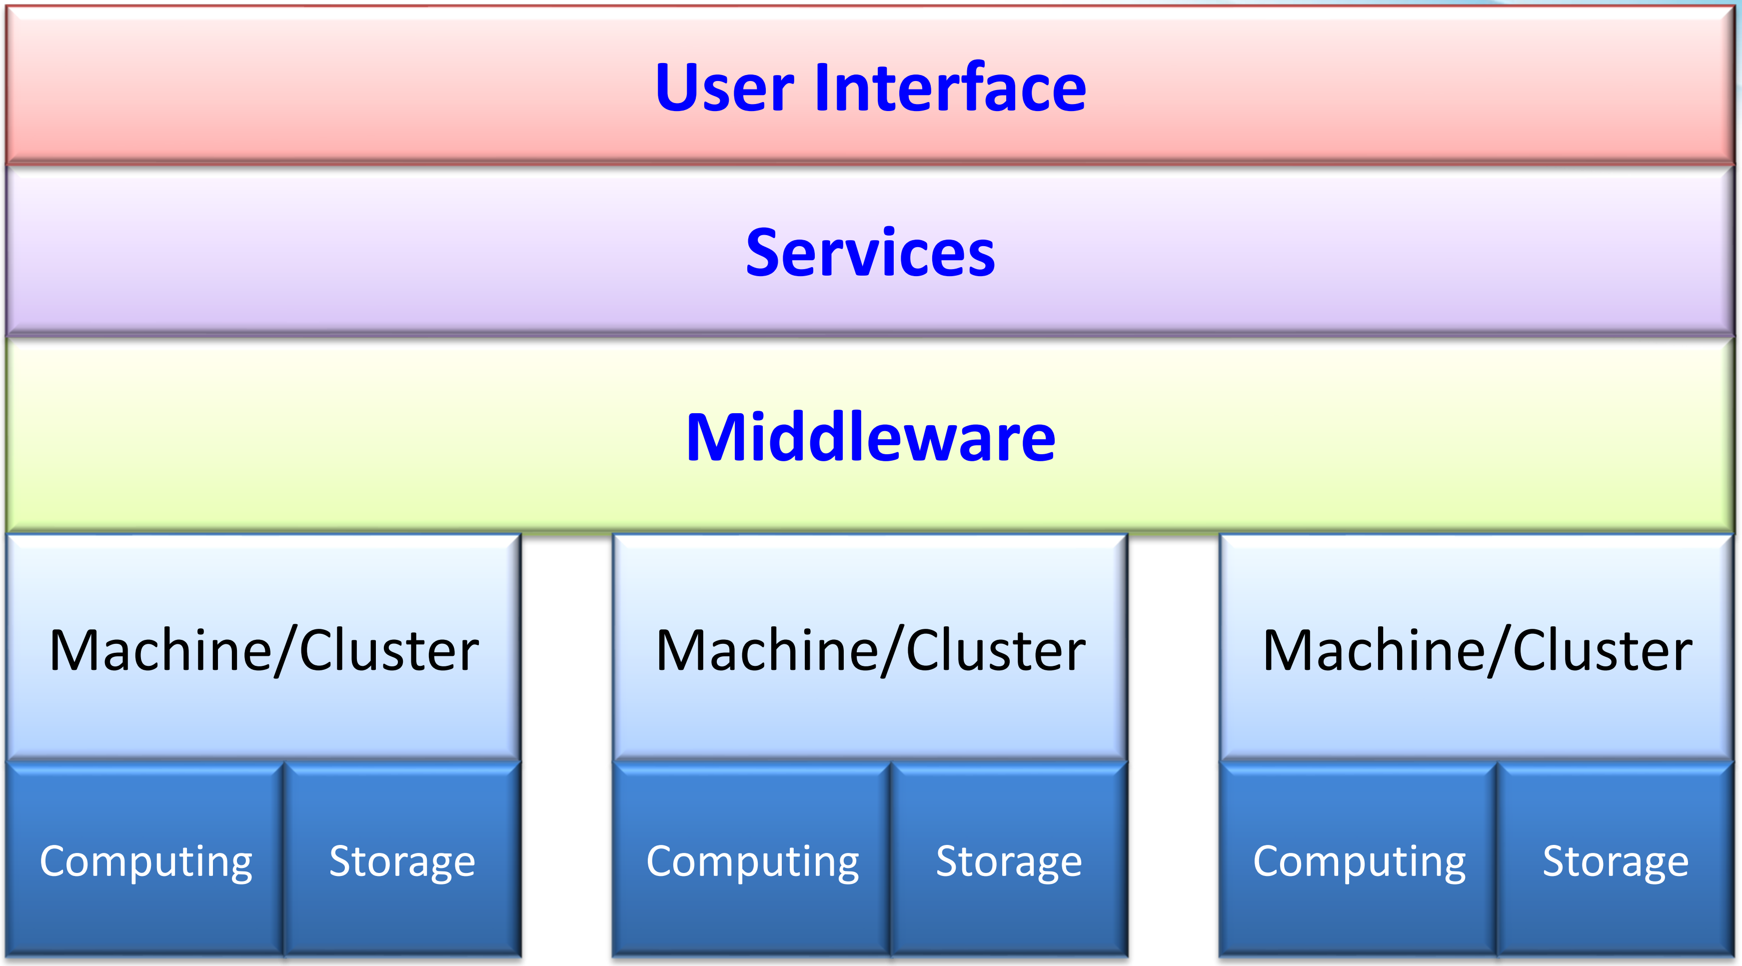

Major Services

- Computing service
  - Deal with the job execution process including the automatically scheduling and management policy
  - Job classification:
    - Higher priority jobs: more expensive
    - Lower priority jobs: less expensive
- Storage service
  - Provide a storage space for users

Computing Paradigm Evolution

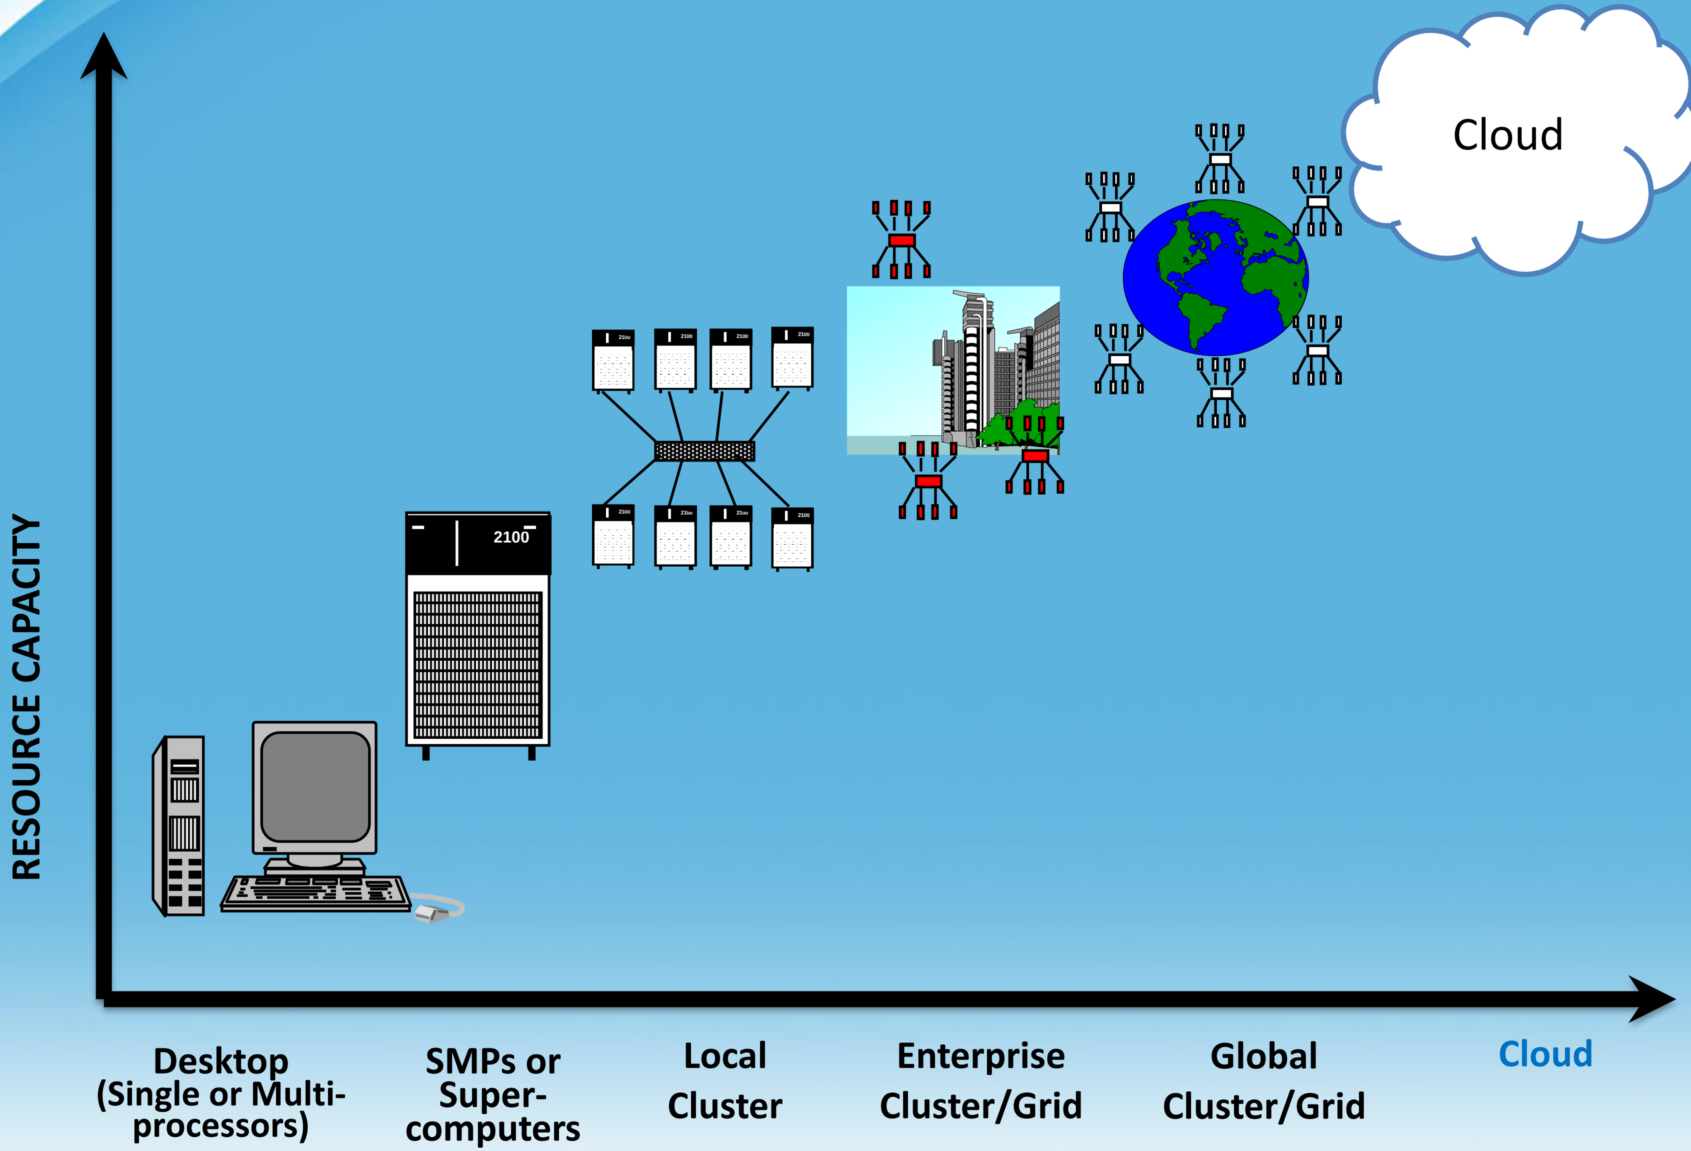

Cluster Computing

Grid Computing

Cloud Computing

DISTRIBUTED COMPUTING

Long united, must divide
Long divided, must unite

Water pond

Carry bucket

Small bottle

Drinking fountain

Past → **Now**

Mainframe

PC

Laptop & Mobile

Cloud Computing

The "Cloud"

- The term "cloud" is often used as a metaphor for the Internet
  - A simplified way to represent the complicated operations in the network
- Currently, the term "cloud" is further used as an abstraction of complexities
  - E.g., servers, applications, data, and heterogeneous platforms

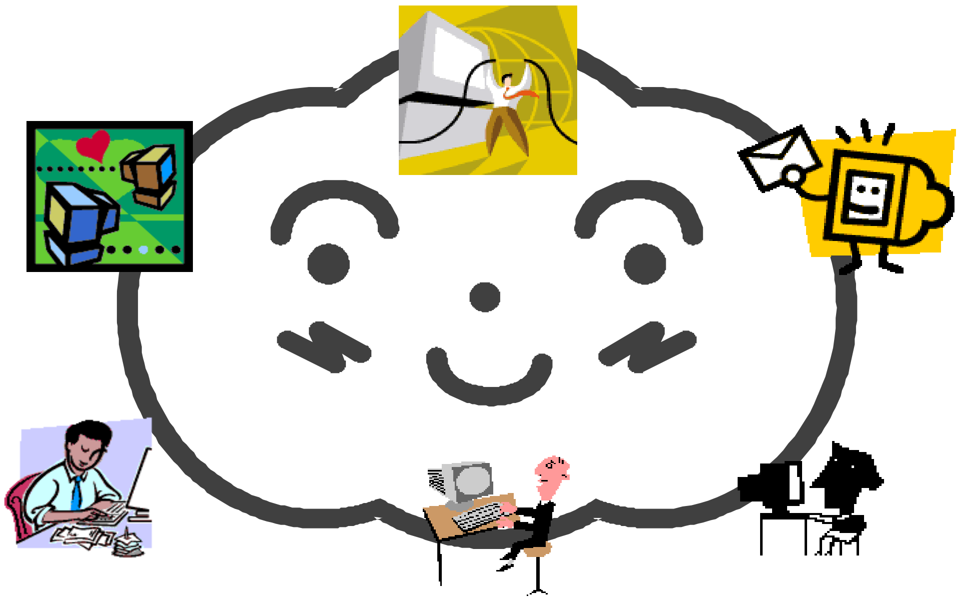

Cloud Computing – A New Paradigm

- An IT service delivered to users that provides:
  - A simple user interface that automatically provisions IT resources
  - Capacity on demand with massive scalability
  - Innovative service delivery models for applications

1990 → 2009

Grid Computing

- Solving large problems with parallel computing

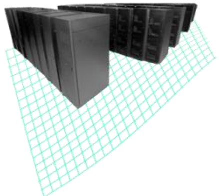

Utility Computing

- Offering computing resources as a metered service

Software as a Service

- Network-based subscriptions to applications

Cloud Computing

- Anytime, anywhere access to resources delivered dynamically as a service

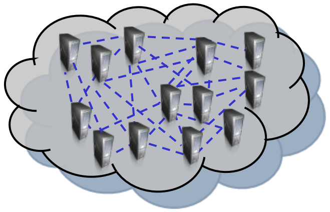

Cloud Computing in Mathematics

- 1
  - One single integrated environment
  - A collection pool of resources and services
- 0
  - Zero management
  - Automatic management and resilience of resource or service up/down/fail
- ∞
  - Endless possibility
  - Scalability, Availability, Accessibility, Manageability, Performance

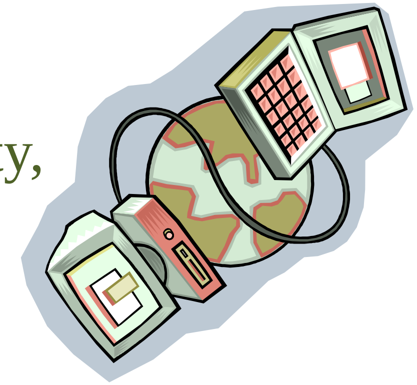

Cloud Computing in IT

- An acquisition and delivery model of IT resources
  - Help improve business *performance* and control the *costs* of delivering IT resources to the organization
- From a user perspective
  - Provides a means of acquiring computing services via the Internet while making the technology beyond the user device almost *invisible*
- From an organization perspective
  - Delivers services for consumer and business needs in a simplified way, *providing unbounded scale* and *differentiated quality of service* to foster rapid innovation and decision making

Emerging Technologies Hype Cycle

Gartner 2009

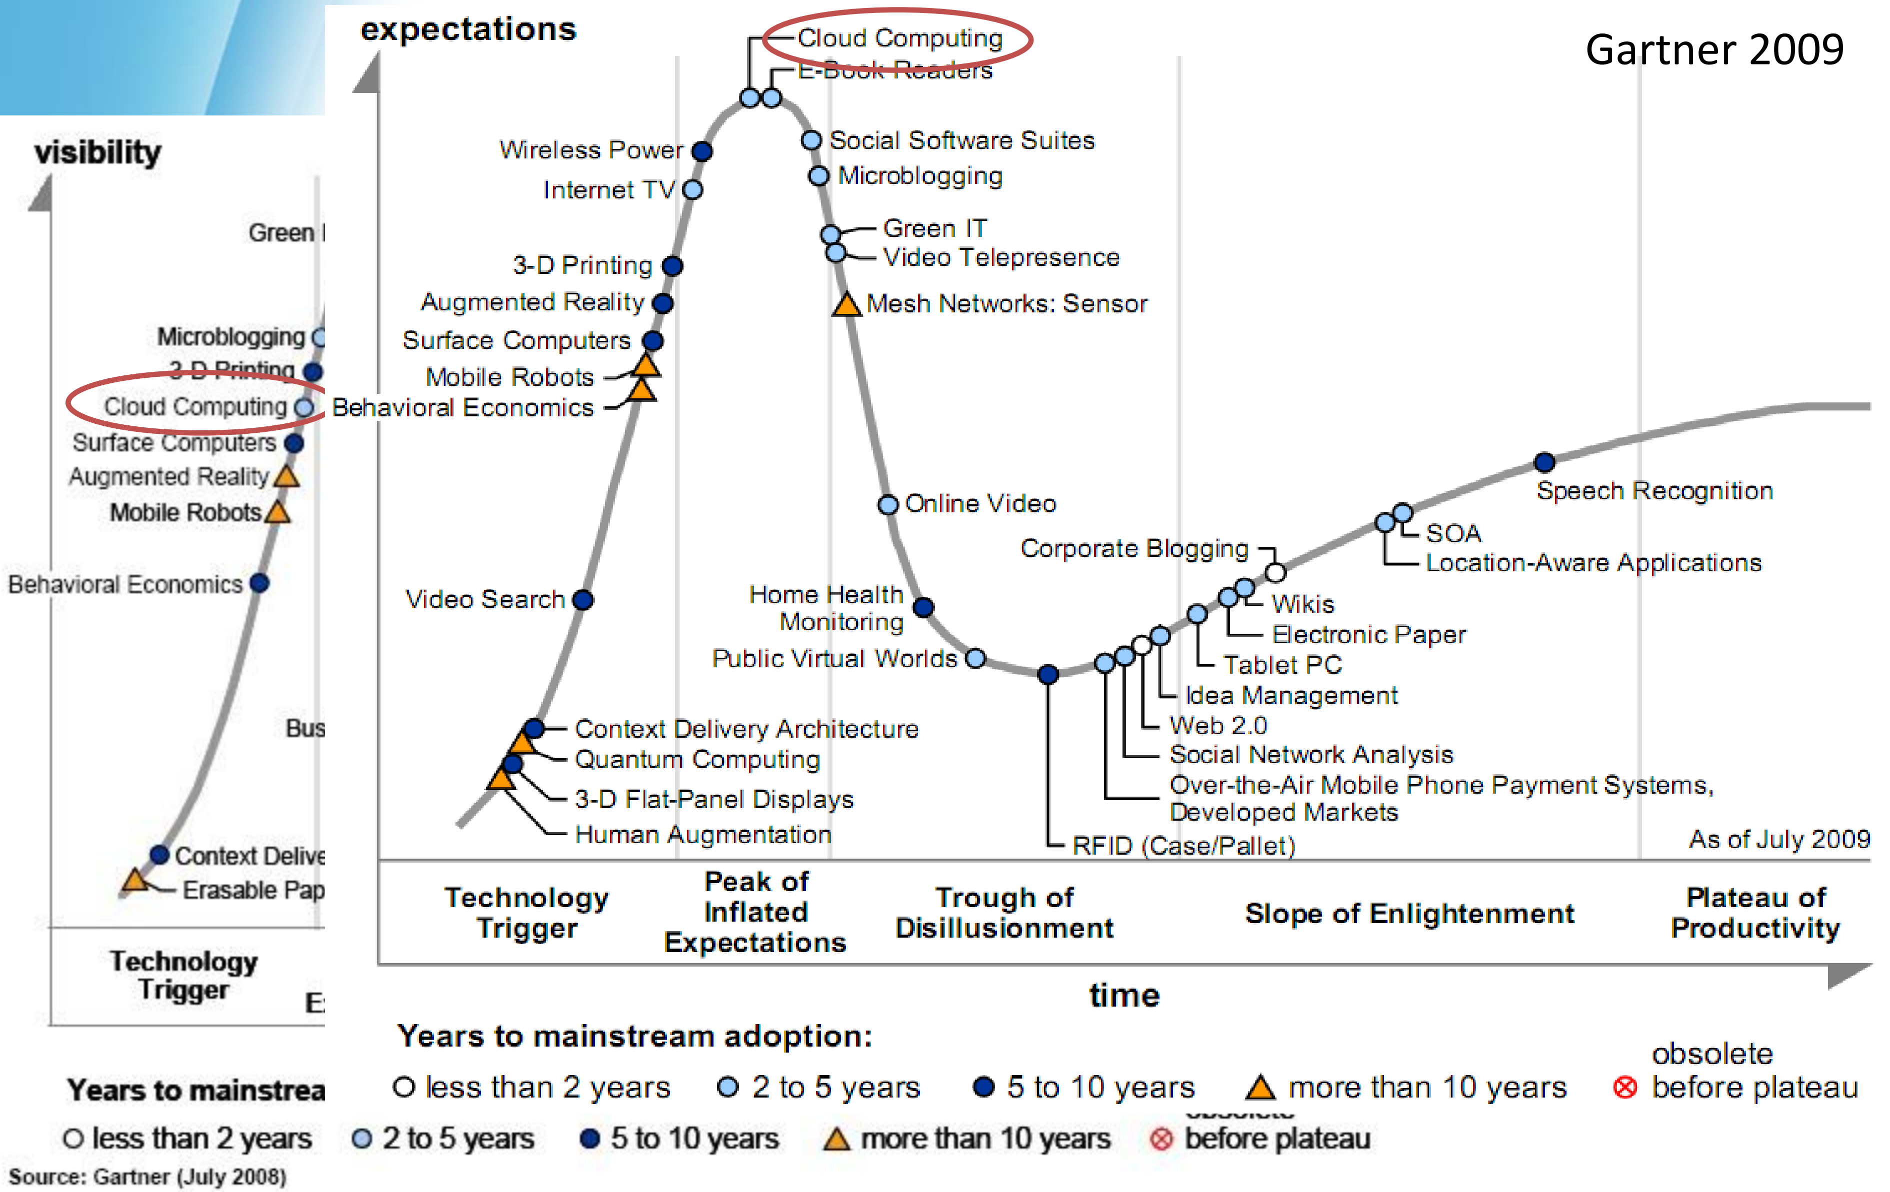

2011 Gartner "IT Hype Cycle" for Emerging Technologies

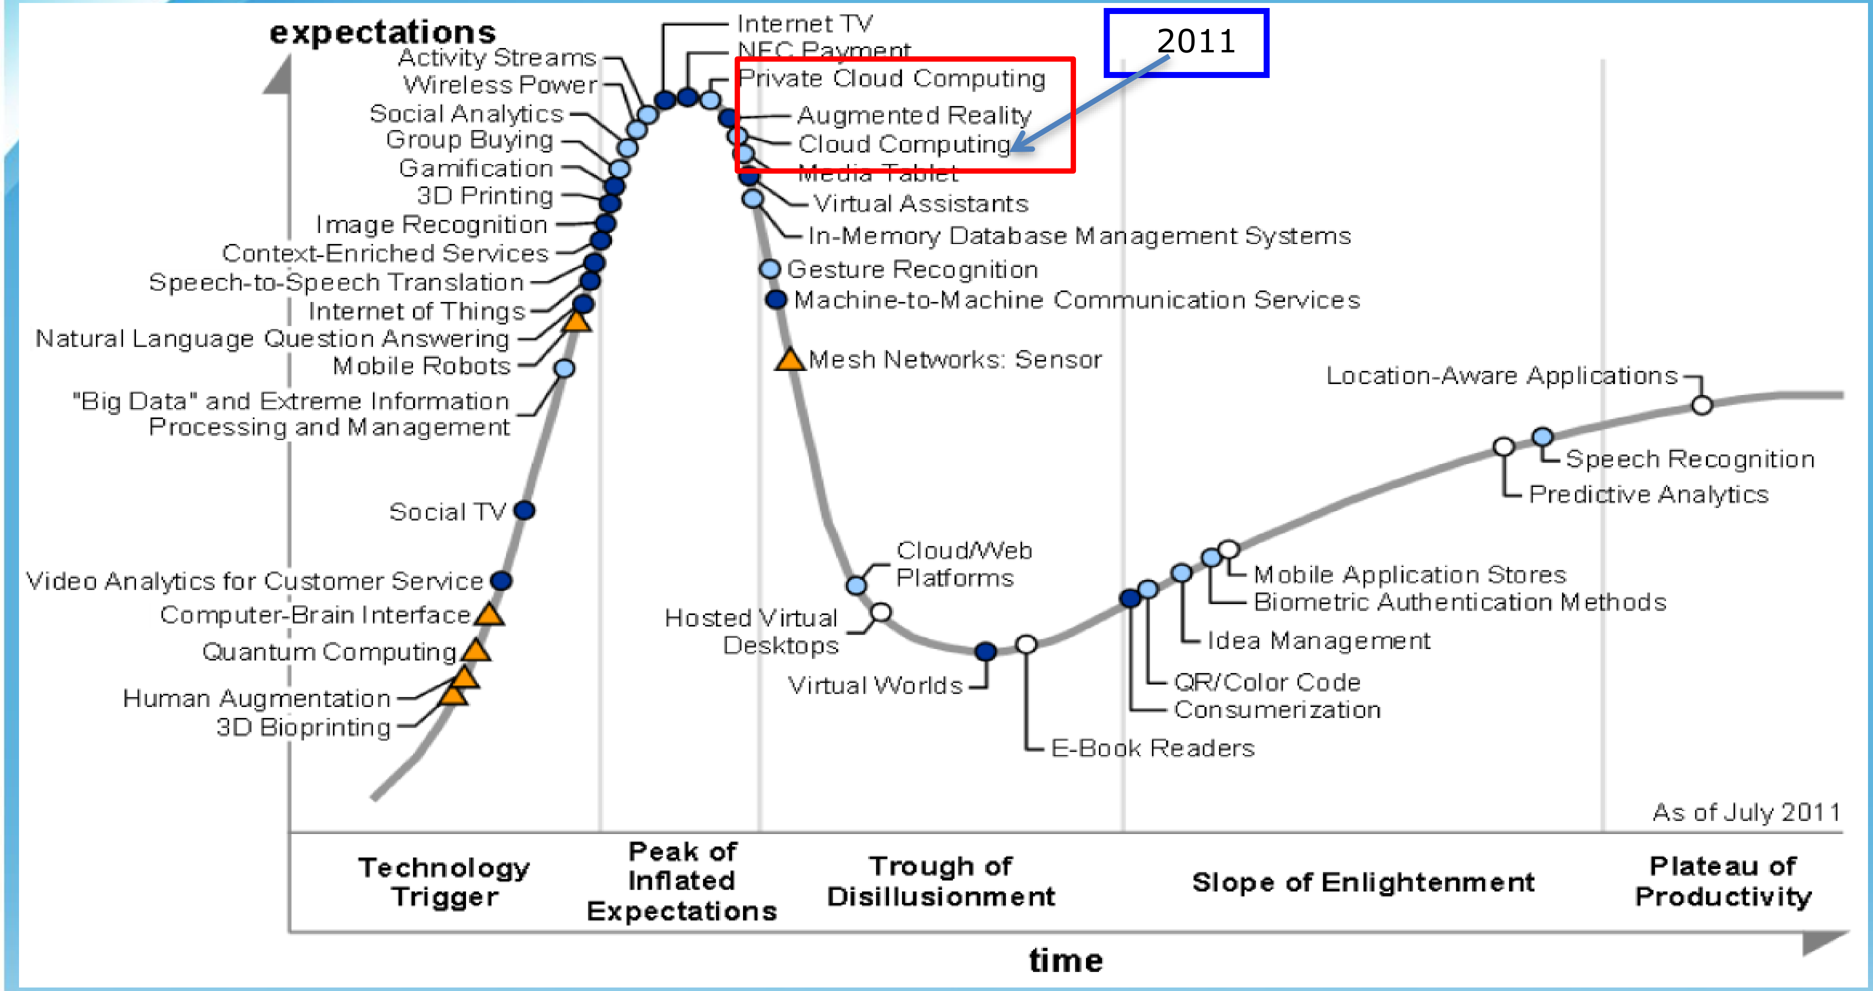

Years to mainstream adoption:

○ Less than 2 years ◎ 2 to 5 years ● 5 to 10 years ▲ More than 10 years ⊗ Obsolete before plateau

2016 Gartner "IT Hype Cycle" for Emerging Technologies

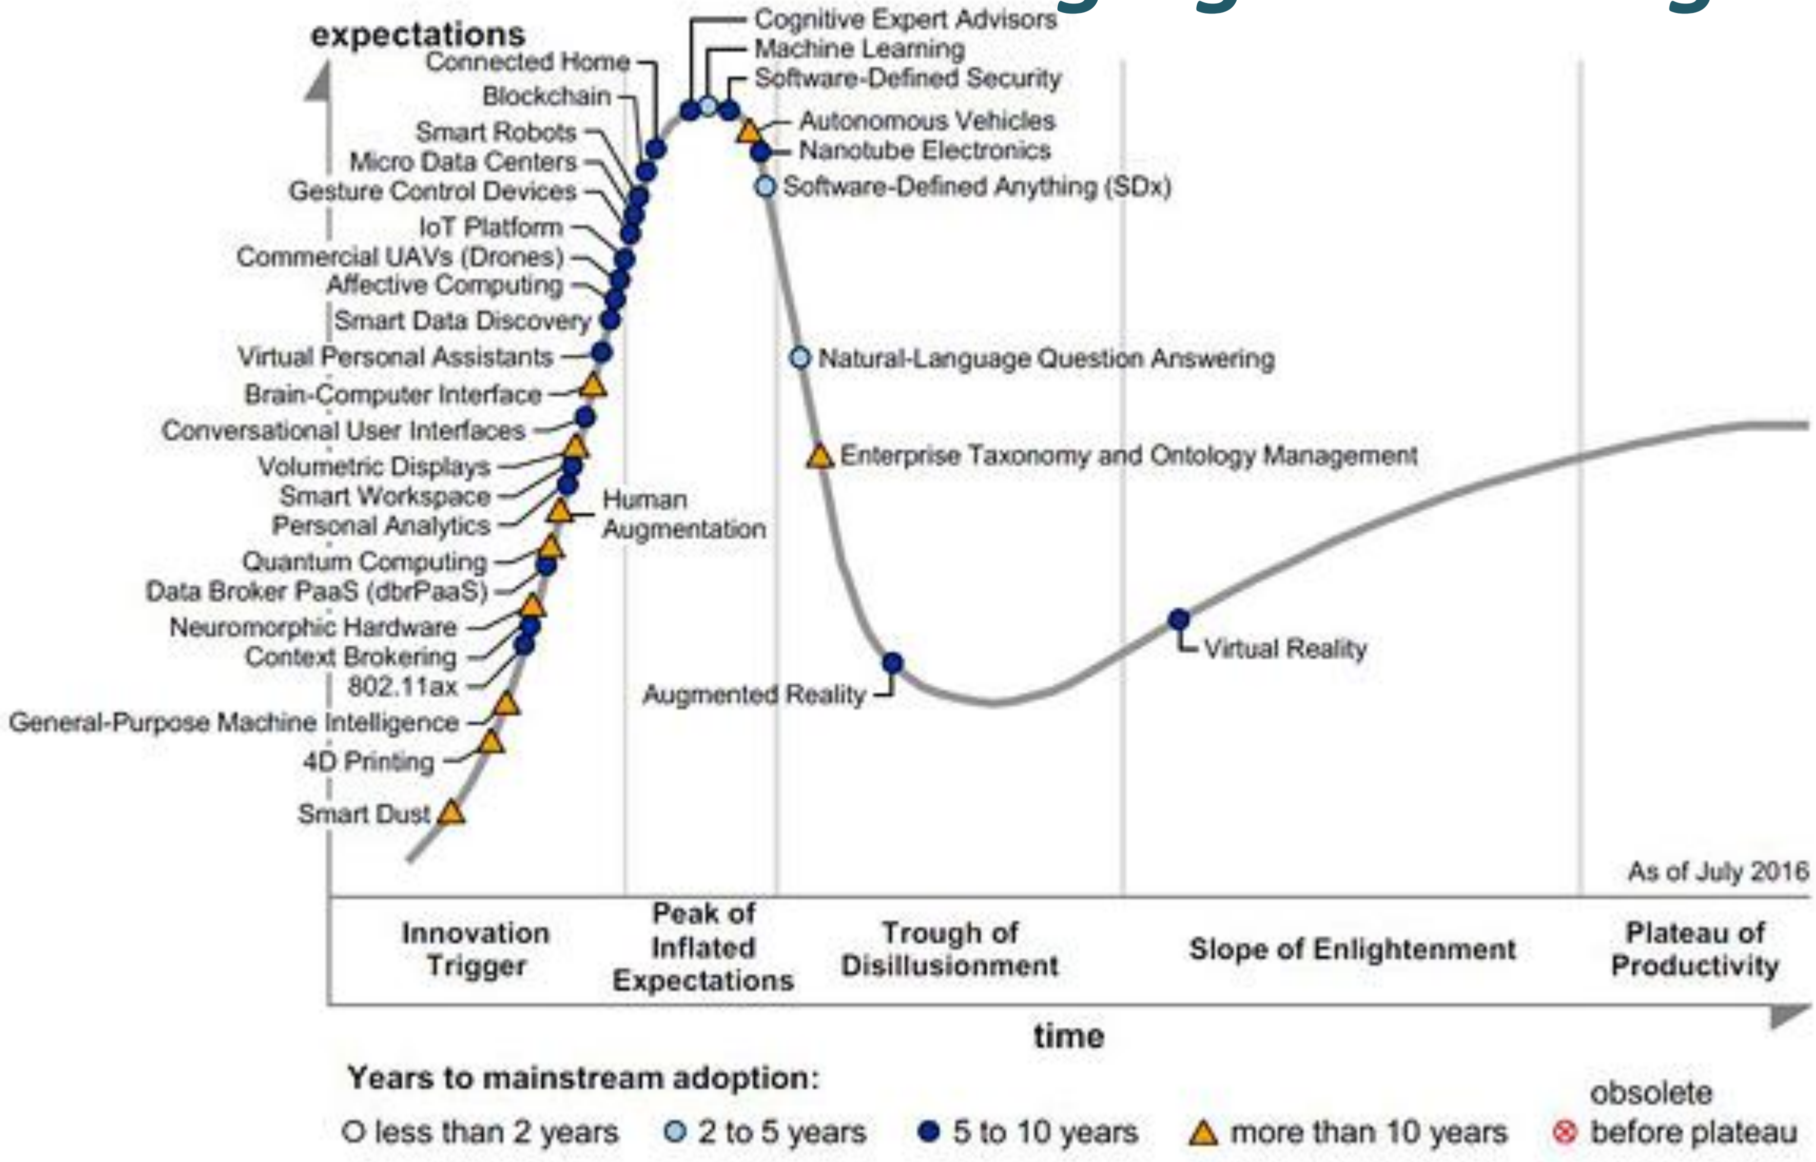

Source: Gartner (July 2016)

2019 and 2022 Gartner “IT Hype Cycle”

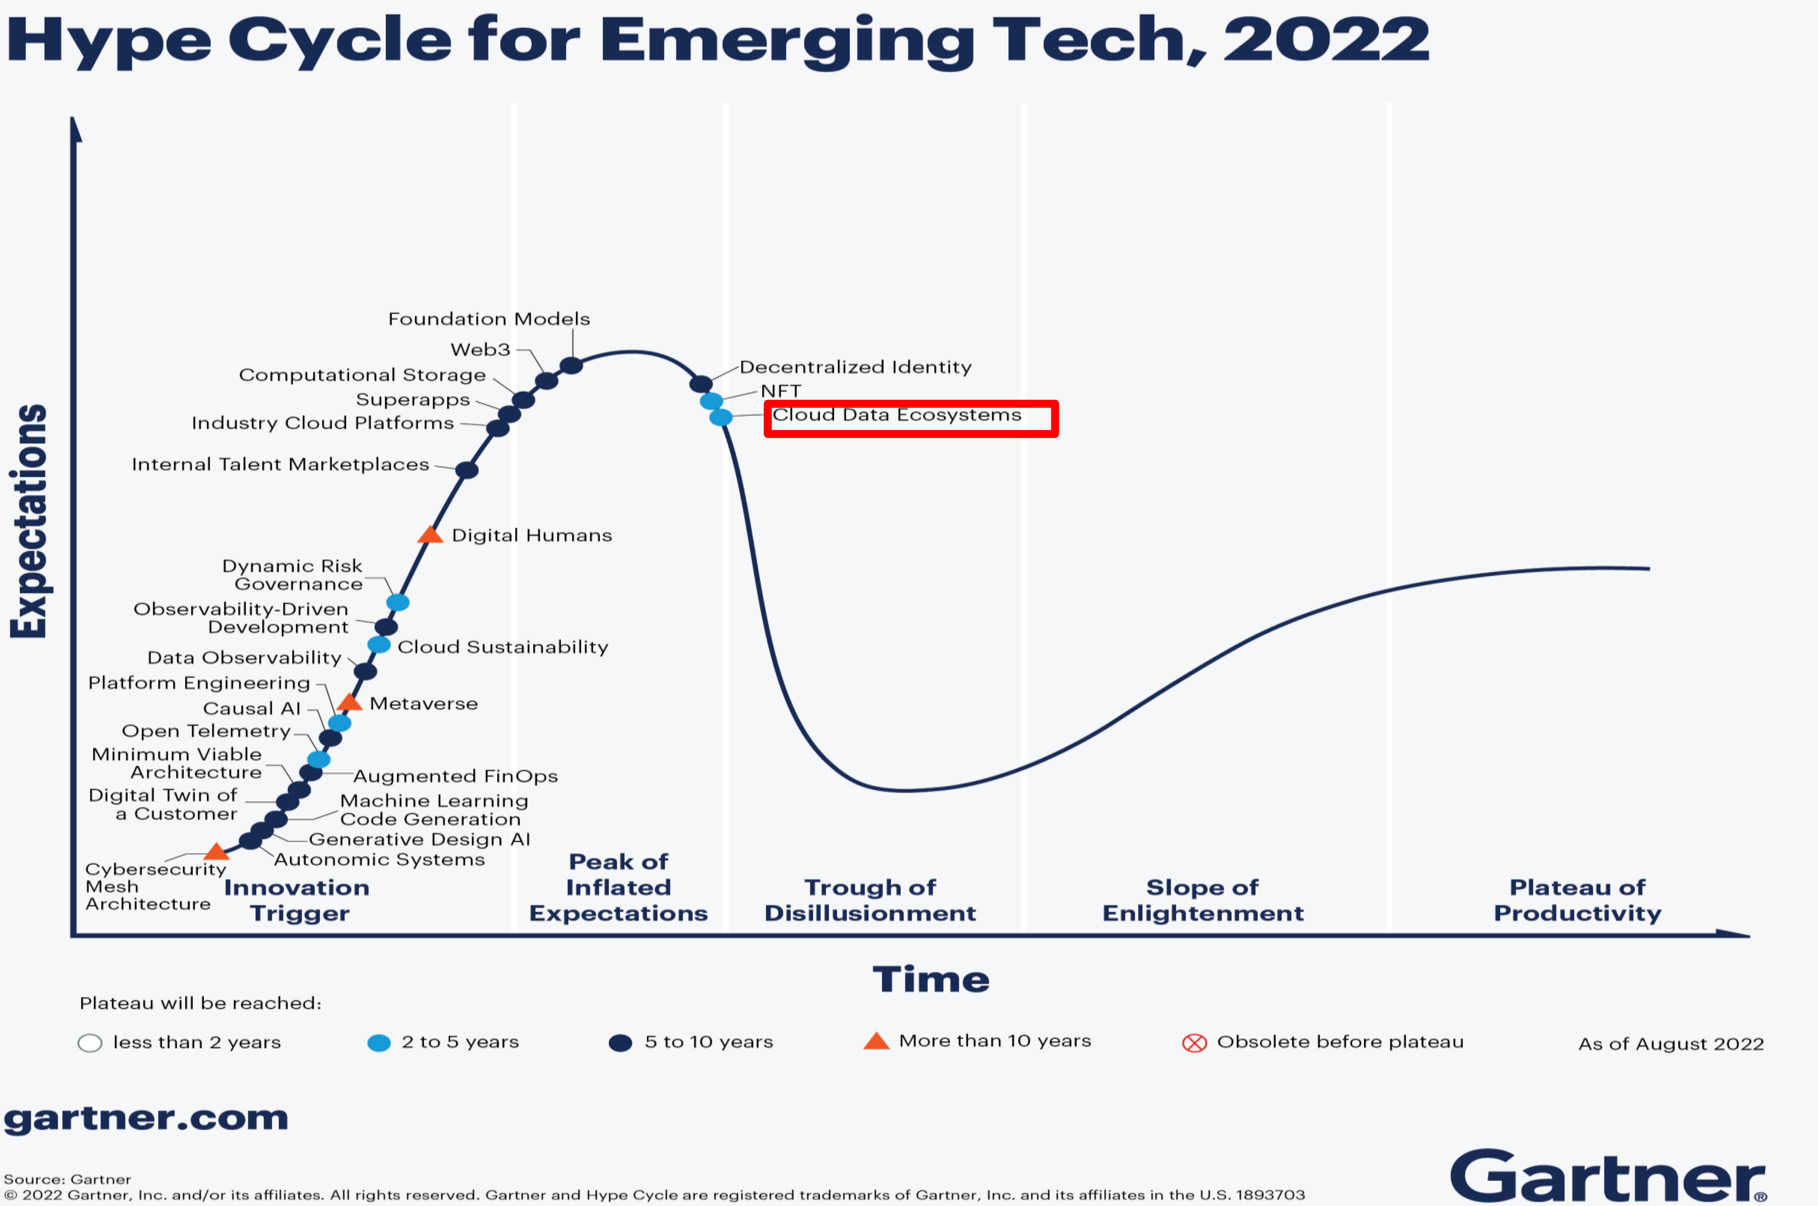

Source: Gartner

Emerging Technologies Hype Cycle

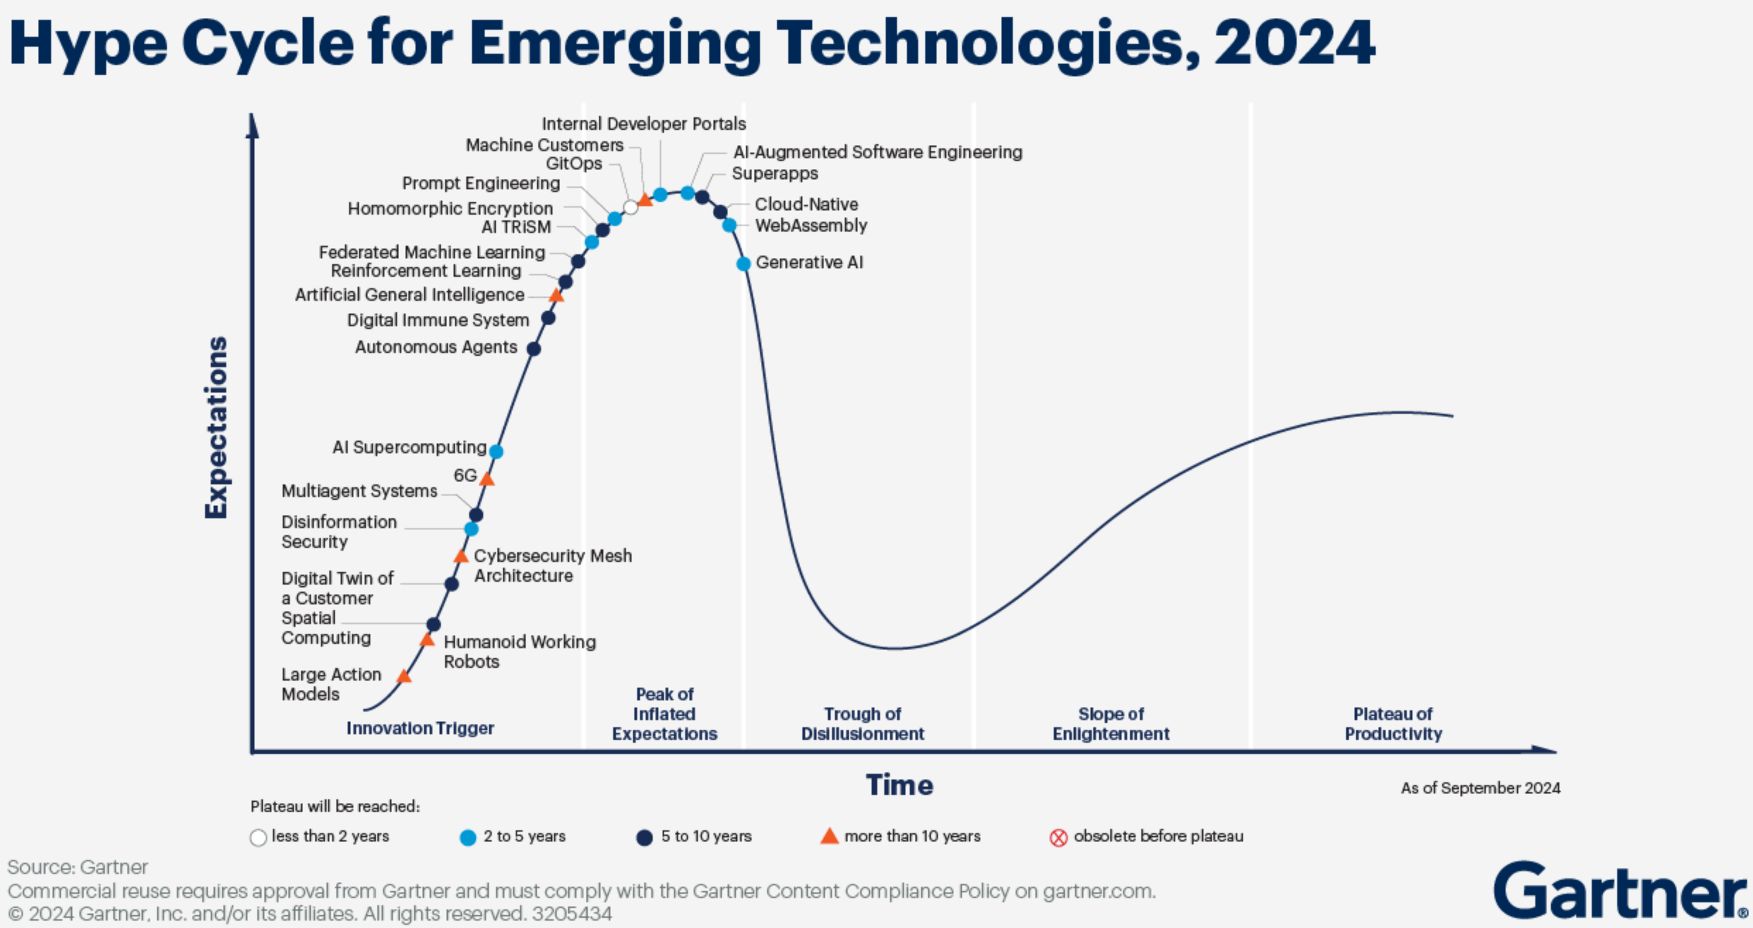

Virtual Reality (VR) Goggles

VR Goggles

VR Goggles

Make Your Dreams Come True

- https://www.oculus.com/?gclid=CJuQgMDrwtMCFRlLDQod2MYGKg

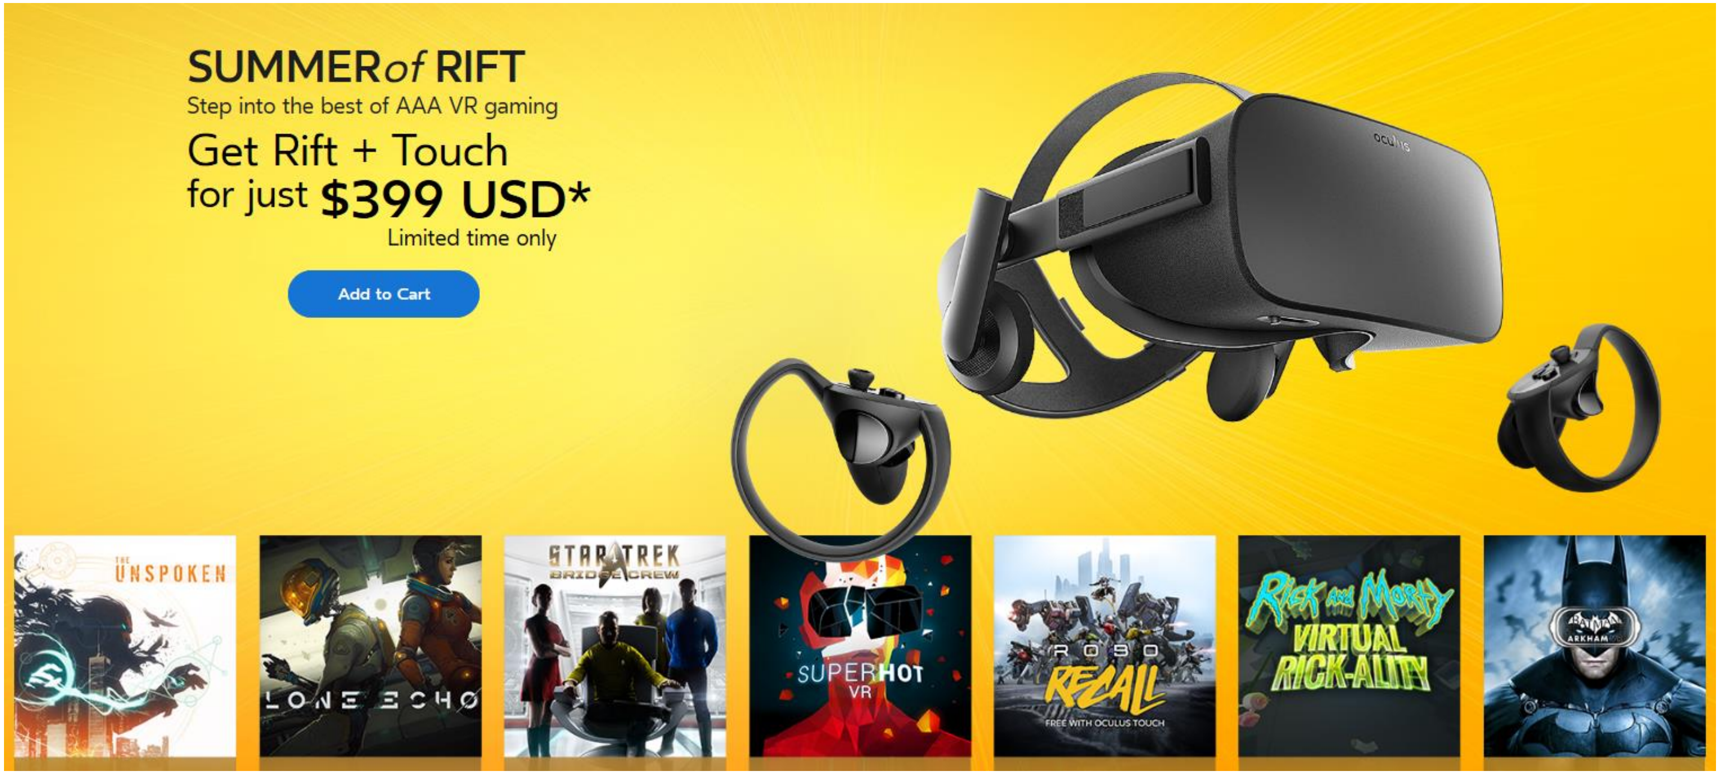

Make Your Dreams Come True

- **https://www.umdrightnow.umd.edu/news/brendan-iribe-oculus-ceo-gives-31m-new-umd-computer-science-building**
- **Brendan Iribe, Oculus CEO, Gives $31M for New UMD Computer Science Building**
- September 12, 2014
- Share Email Print
- Contacts:
- Katie Lawson 301-405-4622
- *Largest Gift in University History Set to Transform Computer Science Education*
- COLLEGE PARK, Md. – The University of Maryland announced today a gift of $31 million from Oculus co-founder and CEO and UMD alumnus Brendan Iribe – the largest gift in the university's history. The majority of the gift, $30 million, will help fund construction of the Brendan Iribe Center for Computer Science and Innovation, a new computer science building designed for cutting-edge work in virtual reality, augmented reality, computer vision, robotics and future computing platforms. The remaining $1 million of the gift will establish the Brendan Iribe Scholarship in Computer Science.

Emerging Technologies Priority Matrix

benefit	years to mainstream adoption			
	less than 2 years	2 to 5 years	5 to 10 years	more than 10 years
transformational	Web 2.0	Cloud Computing Internet TV Public Virtual Worlds SOA	3-D Printing Context Delivery Architecture RFID (Case/Pallet)	Human Augmentation Mobile Robots Quantum Computing
high		E-Book Readers Electronic Paper Green IT Location-Aware Applications Online Video Social Network Analysis Social Software Suites	Augmented Reality Home Health Monitoring Wireless Power	Behavioral Economics Mesh Networks: Sensor
moderate	Corporate Blogging	Idea Management Microblogging Over-the-Air Mobile Phone Payment Systems, Developed Markets Tablet PC Video Telepresence Wikis	3-D Flat-Panel Displays Speech Recognition Surface Computers Video Search	
low				

Gartner 2009

https://www.youtube.com/watch?v=s5I4wcQ6n0g: Microsoft Global Datacenters and Network Infrastructure

https://www.youtube.com/watch?v=4e97g7_qSxA: The world's largest data center

Why Microsoft Has Underwater Data Centers

Benefits from Cloud Computing

Q: Rate the benefits commonly ascribed to the 'cloud'/on-demand model

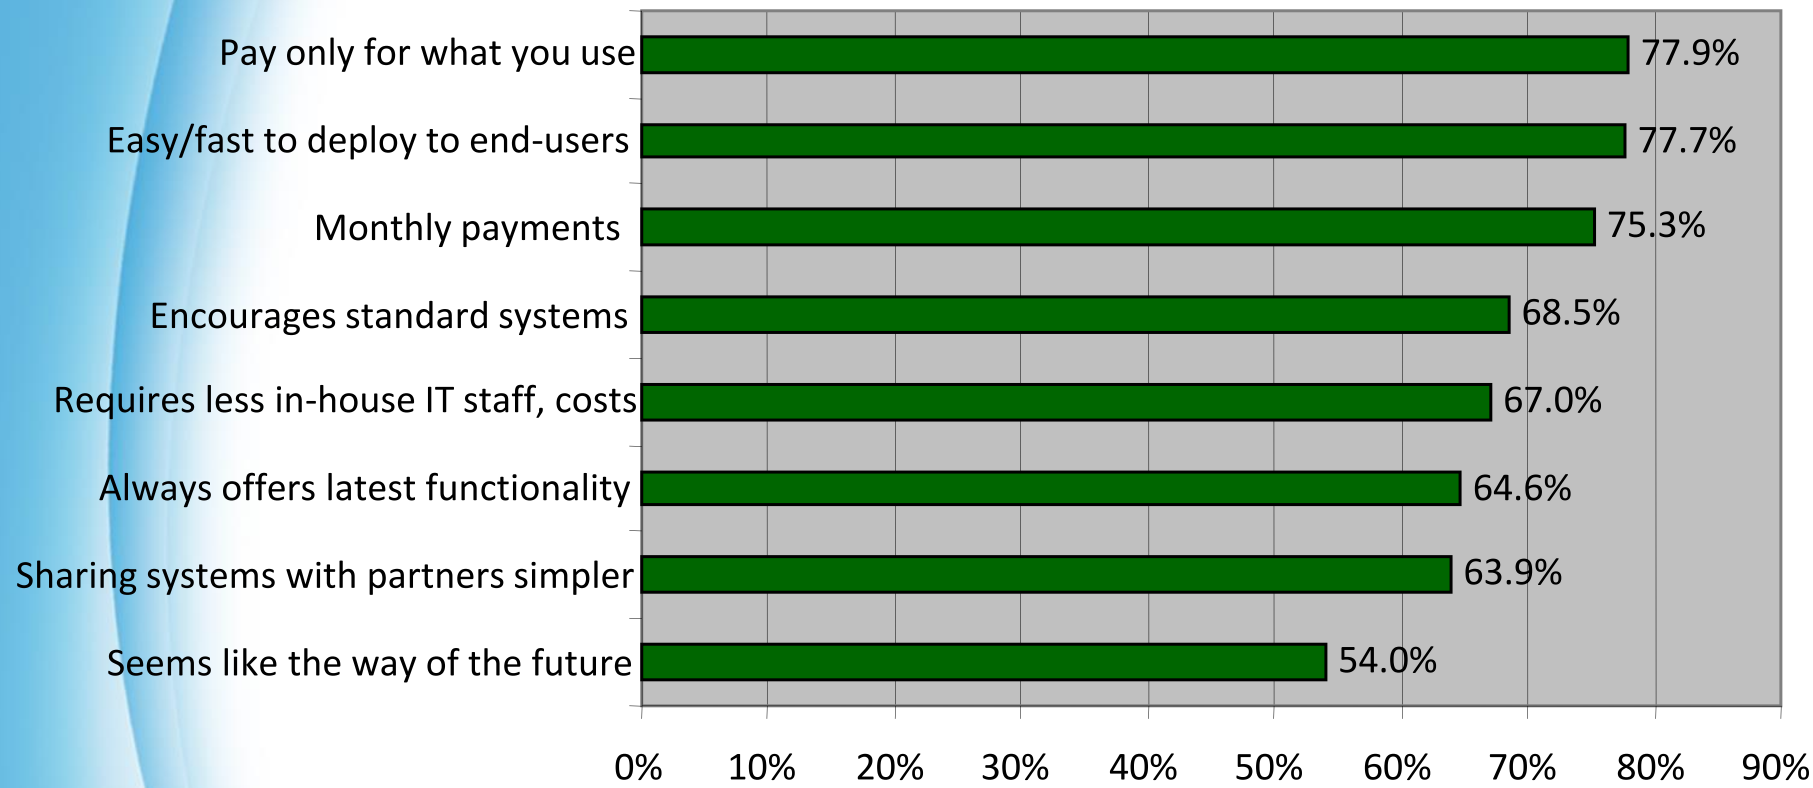

Adoption of Cloud Computing

Q: Rate your likelihood to pursue the cloud model for the following

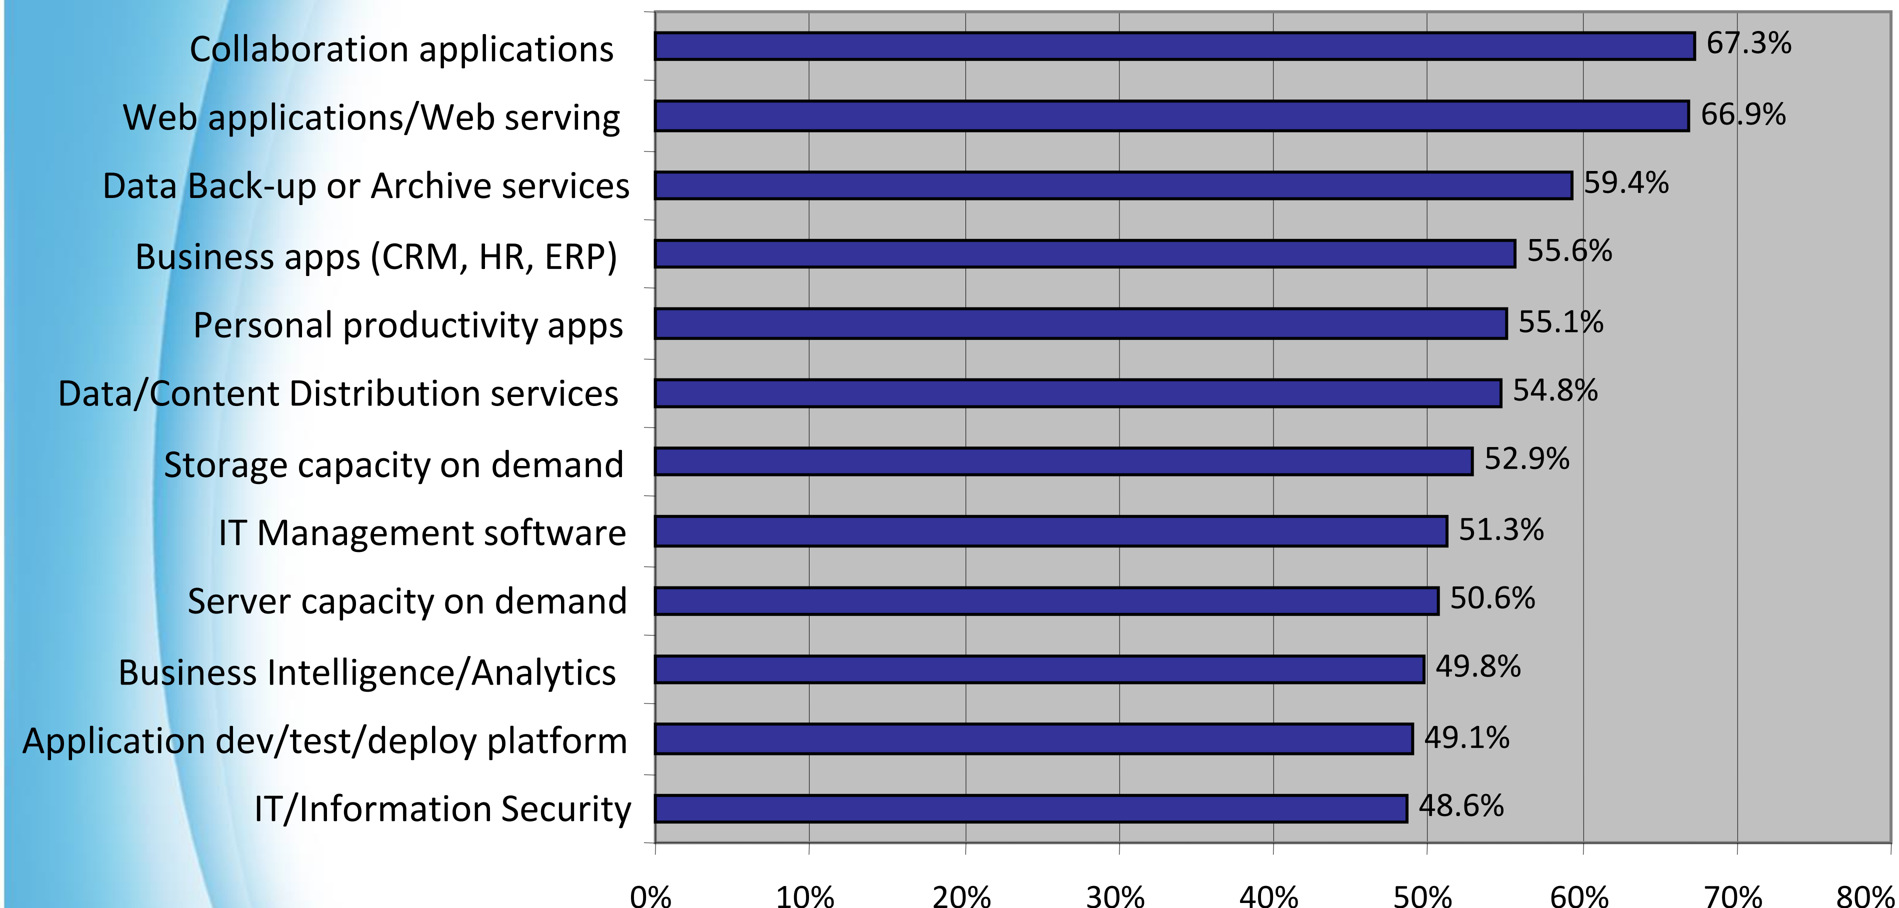

Customer relationship management (**CRM**)

Enterprise resource planning (**ERP**)

Cloud Value Drivers & Effects

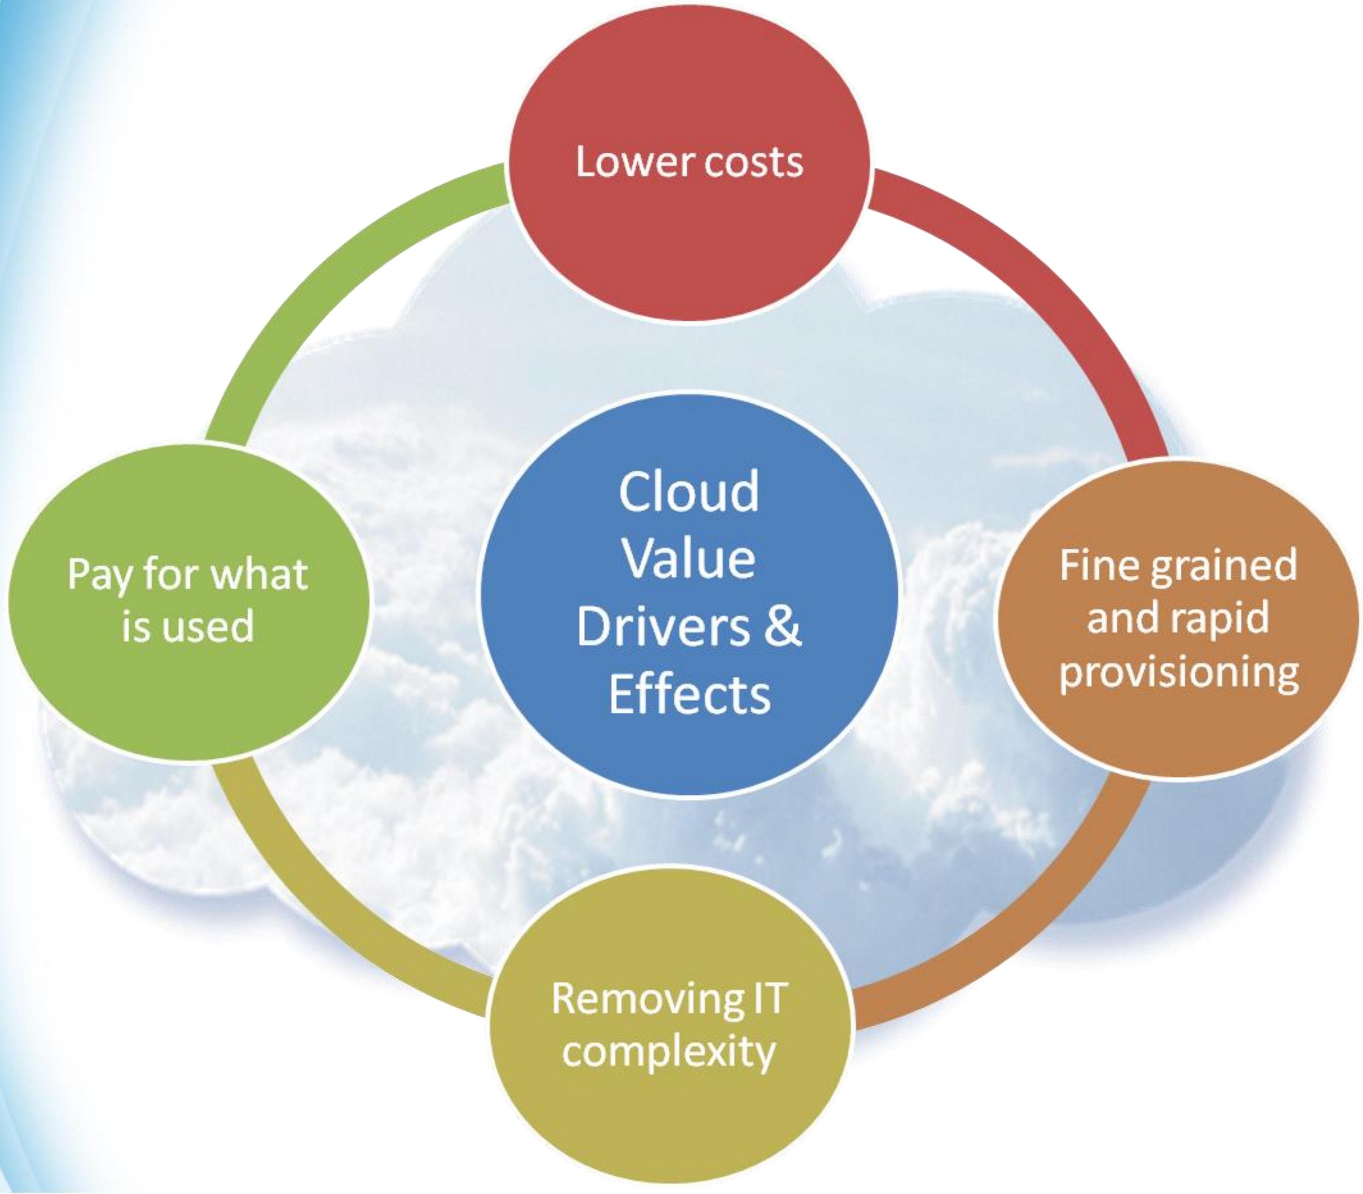

Lower IT operating and capital costs

Lower costs

- The leading value driver is lower IT operating and capital costs
- Lower IT cost
  - Optimize, consolidate and reduce servers
  - Improve capital utilization & quality
  - Reduce energy costs
- Enables implementation of ideas or applications
  - Cheaper pilots encourages experimentation and innovation
  - Reduction in the costs for large, compute and storage intensive applications
  - A Pay for Use model and the lower cloud costs of large computing and storage resources

Fine grained and rapid provisioning

Fine grained and rapid provisioning

- Fine grained IT services with very rapid provisioning change the way IT can acquire capacity

	Traditional IT	Cloud
Servers	Today buy large capacities using multi-year leases/ capital	With cloud capacity on demand, pay-as-you-go
Software	Today multi-year Software Licenses	With cloud SaaS model, pay by the month
Infrastructure capacity in very small increments	Traditional IT capacities come in large increments with up-front capital costs	Fine grained cloud services allow capacity to be obtained on just what is needed then and on a pay-as-you-go basis
Rapid provisioning and scaling up or down easily	Today routine provisioning 2 to 3 weeks	With cloud provisioning in minutes to hours

Removing IT complexity

Removing IT complexity

- Removing IT complexity from end users
- End users can easily access services without worrying about technical details
  - Acquire computing services via the Internet
    - using web-based user interface
  - Cloud enhances user experience through faster and richer cloud services

Pay for what is used

Pay for what is used

- Cloud pricing models based on paying for what is used

- Avoid the upfront cost for infrastructure and the financial risk
  - Cloud pricing models allow pay for what is used
  - Scare capital needed to invest in infrastructure is replaced with an operating expense

Typical IT budget models :

- Applications: 35%
- Infrastructure: 60%
- Other: 5%

Cloud Computing Players

more and more…

Google
amazon web services
salesforce.com
ebay
Microsoft
conduit
curl
EMC² where information lives
FORCE10
Intuit QuickBase
keynote
ULTRADNS
ONEsite
OpSource
S60
Sun microsystems
acquia
ascentium
awareness
BLUEARC
CORADIANT
dixero
jive
kapow TECHNOLOGIES
lithium
PANTHER
SynthaSite
INDUSTRY LIVE
webAsyst
xbo
remy CORPORATION www.remycorp.com
Alcatel·Lucent
Novell.
OPEN TEXT The Content Experts
IBM
hp
HORN GROUP
Socialtext
ORACLE
telligent
trampoline systems
WorkLight
spigit
nGenera
liquid PLANNER
IGLOO CORPORATE SOCIAL NETWORKING
jive
Alfresco
connectbeam
acquia
INQUIRA
GIGASPACES
enomaly
ELASTIC COMPUTING
parascale Powering Cloud Storage
3tera
elastra
Joyent
ServePath
CLOUDWORKS
nirvanix
NETRICS Data Matching Solved
Skytap GO VIRTUAL
Appistry

Players : Users (Consumers) or Providers

- Players can act as Users or Providers in the Cloud Ecosystem

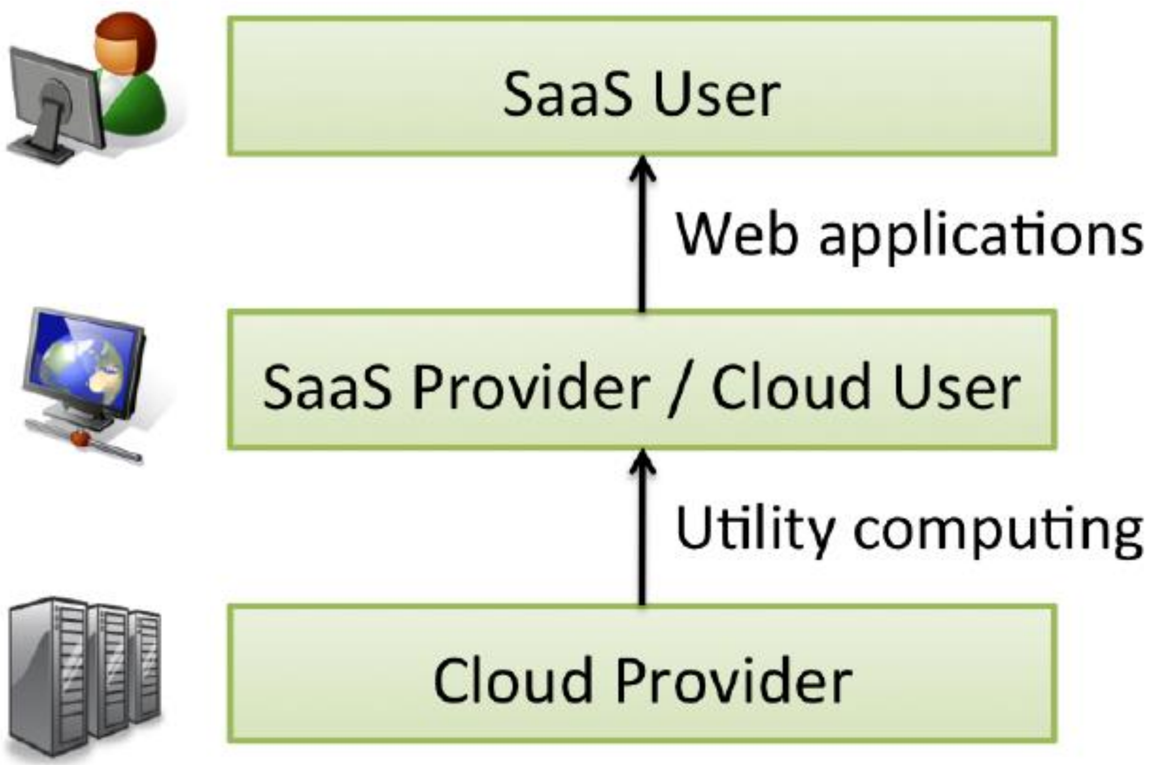

Cloud Computing Economy

- Enable SaaS providers that do not own IT infrastructure

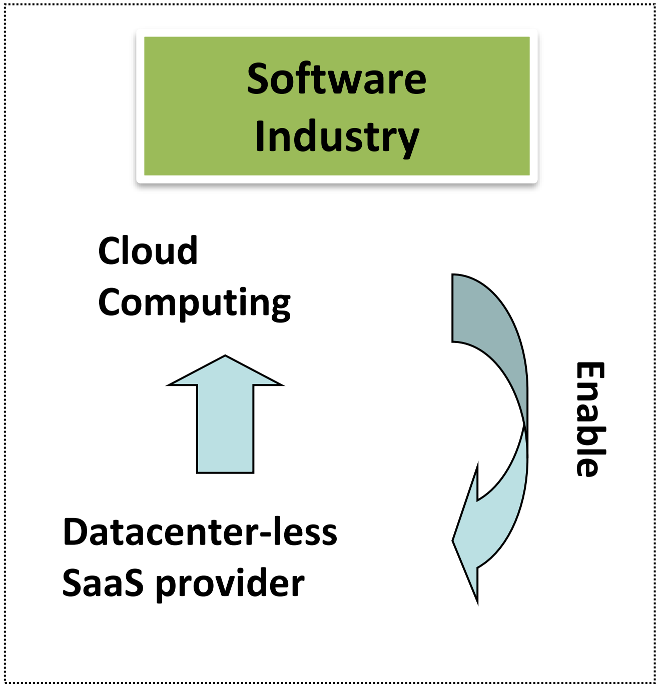

Industry Chain

Cloud Services

- Upstream – IaaS provider
  - Amazon EC2, CHT hicloud, ect.
- Midstream – PaaS provider
  - Google GAE, Windows Azure, etc.
- Downstream – SaaS provider
  - Saleforce.com, Google docs, etc.
- End Users
  - Thin client

Amdahl's law and Gustafson's law

- Amdahl's law:

Speedup (S)=T/[α T+(1-α)T/n]=1/[α+(1- α)/n]

Efficiency (E)=S/n =1/[αn +1−α]

α:percentage of a program that is not parallelized

- Gustafson's law:

W: workload in a given program (i.e., execution time in a cluster)

W': sequential execution time on a single processor

S': scaled-workload speedup

S'=W'/W=[αW+{1-α}nW]/W=α +(1-α)n

Efficiency: E'=S'/n=α/n+(1-α)

System Availability

- System Availability=MTTF/(MTTF+MTTR)
- MTTF: mean time to failure
- MTTR: mean time to repair
- High availability is desired in all cluster, grids, P2P networks and cloud systems.

Summary

- Long united, must divide; long divided, must unite
- Modern IT require
  - To increase capacity or add capabilities to their infrastructure dynamically
    - without investing money in the purchase of new infrastructure
  - All the while
    - without needing to conduct training for new personnel
    - without the need for licensing new software
- Given a solution to the above mentioned demands
  - Cloud computing

References

- IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries, 1990.
- R. Buyya (ed.), High Performance Cluster Computing: Systems and Architectures, Prentice Hall, 1999.
- I. Foster, C. Kesselman (Eds.), The Grid: Blueprint for a Future Computing Infrastructure, Morgan Kaufmann, San Francisco, USA, 1999
- I. Foster et. al., "The Anatomy of the Grid. Enabling Scalable Virtual Organizations," International Journal of High Performance Computing Applications, 2001.
- Computing Curricula 2005
- J. Rittinghouse and J. F. Ransome, Cloud Computing: Implementation, Management, and Security, CRC Press, 2009.
- M. Armbrust et. al., "Above the Clouds: A Berkeley View of Cloud Computing," Technical Report No. UCB/EECS-2009-28, University of California at Berkeley, 2009.
- R. Buyya et. al., "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, 2009.
- Hai Jin and Raj Buyya, High Performance Cluster Computing: Architectures and Systems.
- NIST (National Institute of Standards and Technology). http://csrc.nist.gov/groups/SNS/cloud-computing/
- Gartner, Inside the Hype Cycle: What's Hot and What's Not in 2009.
- Cloud Computing 2010 – An IDC Update.
- IBM Corporation.
- From Wikipedia, the free encyclopedia.
- All resources of the materials and pictures were partially retrieved from the Internet.

Class Exercise

Consider a program for multiplying two large-scale $N \times N$ matrices, where N is the matrix size. The sequential multiply time on a single server is $T_1 = cN^3$ minutes, where c is a constant determined by the server used. An MPI-code parallel program requires $T_n = cN^3/n + dN^2/n^{0.5}$ minutes to complete execution on a n-server cluster system, where d is a constant determined by the MPI version used. Assume the program has a zero sequential bottleneck ($\alpha = 0$). The second term in T_n accounts for the total message-passing overhead experienced by n servers.

Answer the following questions for a given cluster configuration with n=64 servers, c=0.8 and d=0.2. Parts (a, b) have a fixed workload corresponding to the matrix size N=5000, Parts (c, d) have a scaled workload associated with an enlarged matrix size $N' = n^{1/3}N = 64^{1/3} * 5000 = 20000$. Assume the same cluster configuration to process both workloads. Thus, the system parameters n, c, and d stay unchanged. Running the scaled workload, the overhead also increases with the enlarged matrix size N'.

a. Using Amdahl's law, calculate the speedup of the n-server cluster over a single server.
b. What is the efficiency of the cluster system used in Part (a)?
c. Calculate the speedup in executing the scaled workload for an enlarged N'*N' matrix on the same cluster configuration using Gustafson's law.
d. Calculate the efficiency of running the scaled workload in Part (c) on the 64-processor cluster.
e. Observation?

Solutions

- a) Speedup, S = 1/[α + (1 - α) / n]

  Since, a = 0 and n = 64, Speedup, S = 64
- b) Efficiency, E = S / n = 1 = 100%
- c) Speedup, S' = Time taken by a single server / Time taken by the cluster= T_1/T_n= cN'^3 / [(cN'^3 / n) + (dN'^2 / $n^{0.5}$)]=0.8 x 64 x N^3 / [$0.2N^2(4N + 1)$] = 256N / (4N + 1)=64

  since c = 0.8, d = 0.1, N' = $n^{1/3}$ N, N = 20,000, n = 64
- d) Efficiency, E' = α / n + (1 – α), Since α = 0, E' = 1 = 100%
- e) For both cases i.e. fixed workload and scaled workload the speedup and efficiency is the same i.e. S = 64 and E = 100%

Class Exercise

- Consider a multicore processor with three heterogeneous cores labeled A, B, and C. Suppose an application needs to compute the square of each element of an array of 60 elements. Assume that all three cores start executing the application at the same time and no cache misses are encountered in all core operations. Assume 1 unit time for core A to compute the square of an element. Core B takes 3 units time and core C takes 4 units time to compute the square of an element . Given the following division of labor in the three cores:
- Core A 35 elements
- Core B 15 elements
- Core C 10 elements
- (1) Compute the total execution time (in time units) for using the three-core processor to compute the squares of 60 elements in parallel. The three cores have different speeds. Some faster cores finish the job and may become idle, while others are still busy computing until all squares are computed.
- (2) Calculate the processor utilization rate, which is the total amount of time the cores are busy (not idle) divided by the total execution time they are using all cores in the processor to execute the above application.

Solutions

- Core A: 35*1=35, Core B: 15*3=45, Core C: 10*4=40
- Total execution time = 45

- (35+40+45)/(45*3)=0.889

Class Exercise

- Consider a cluster that has little availability support. Upon a node failure, the following sequence of events takes place:
- 1. The entire system is shut down and powered off.
- 2. The faulty node is replaced if the failure is in hardware.
- 3. The system is powered on and rebooted.
- 4. The user application is reloaded and rerun from the start.
- Assume one of the cluster nodes fails every 196 hours. Other parts of the cluster never fail. Steps 1 through 3 take two hours. On average, the mean time for step 4 is two hours.
- What is the availability of the cluster? What is the yearly failure cost if each one-hour downtime costs $60,000?

Solutions

- MTTF = 196 hours, MTTR = 4 hours
- SA = MTTF / (MTTF + MTTR) = 98%
- Cost = (1 - 0.98) * 365 * 24 * 60000=10.7M

Slides for the Speedup problem

Problem

- Consider parallel execution of an MPI-coded C program in SPMD (single program and multiple data streams) mode on a server cluster consisting of n identical Linux servers. SPMD mode means the same MPI program is running simultaneously on all servers but over different data sets of identical workloads. Assume that 25 percent of the program execution is attributed to the execution of MPI commands. For simplicity, assume that all MPI commands take the same amount of execution time. Answer the following questions using Amdahl's law:

(1) Given that the total execution time of the MPI program on a four-server cluster is T minutes, what is the speedup factor of executing the same MPI program on a 256-server cluster, compared with using the four-server cluster? Assume that the program execution is deadlock free and ignore all other runtime execution overheads in the calculation.

Solution approach

- The problem asks to find the speedup factor of executing the same MPI program on a 256-server cluster, compared with using the four-server cluster
- Ultimately, calculate the ratio of speed up of

the 256 server cluster to the four server cluster.

That is;

$$S_{256} / S_4$$

speedup **factor** means how many times the speedup of 256 server is larger than the four server cluster.

Extension to the problem

- Suppose that all MPI commands are now enhanced by a factor of 2 by using active messages executed by message handlers at the user space. The enhancement can reduce the execution time of all MPI commands by half. What is the speedup of the 256-server cluster installed with this MPI enhancement, computed with the old 256-server cluster without MPI enhancement?

Solution Approach

- It mentions all MPI commands are now enhanced by a factor of 2, so the value of alpha mentioned before should be alpha/2.
- However, due to the nature of the MPI commands as mentioned in the hint, 0.25 is the sequential part. No matter how many servers you use, each server has to spend 25% of the time to run MPI commands.

So, the denominator of the Amadahl's law remains same, that is;

$$S = 1/\{0.125 + 0.75/256\}$$

- Finally, calculate the ratio as following

$$S_{256new}/S_{256old}$$