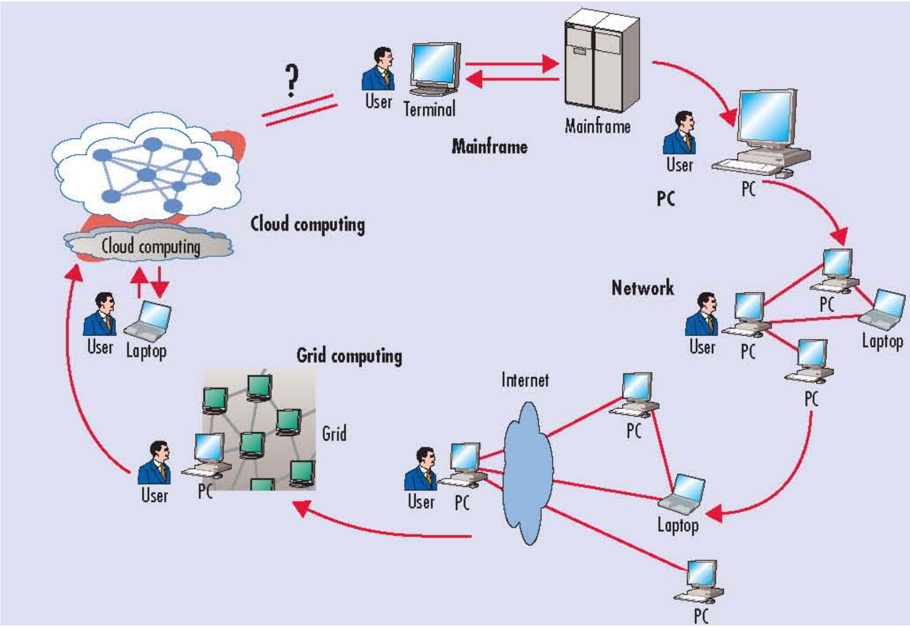

From Desktop to Supercomputer, Cluster, Grid, Cloud, InterCloud …

Where my Research Code Should Run?

Dana Petcu, West University of Timisoara, Romania

Content

- Reasons
- Matching appl and syst characteristics
  - Main characteristics of the systems
  - Main characteristics of the current applications
  - Matching the systems with our applications
- A use case
- Towards programmable e-infrastructures
- Conclusions

Reflection invitation

- Trendy:

“The computer industry is the only industry that is more fashion-driven than women’s fashion” [Oracle]”

1. Should we follow the fashion?
2. What is the profit of being trendy?
3. Are we able to keep the step with these trends?

- This talk – try to answer to “3.”

Well, what is the current trend?

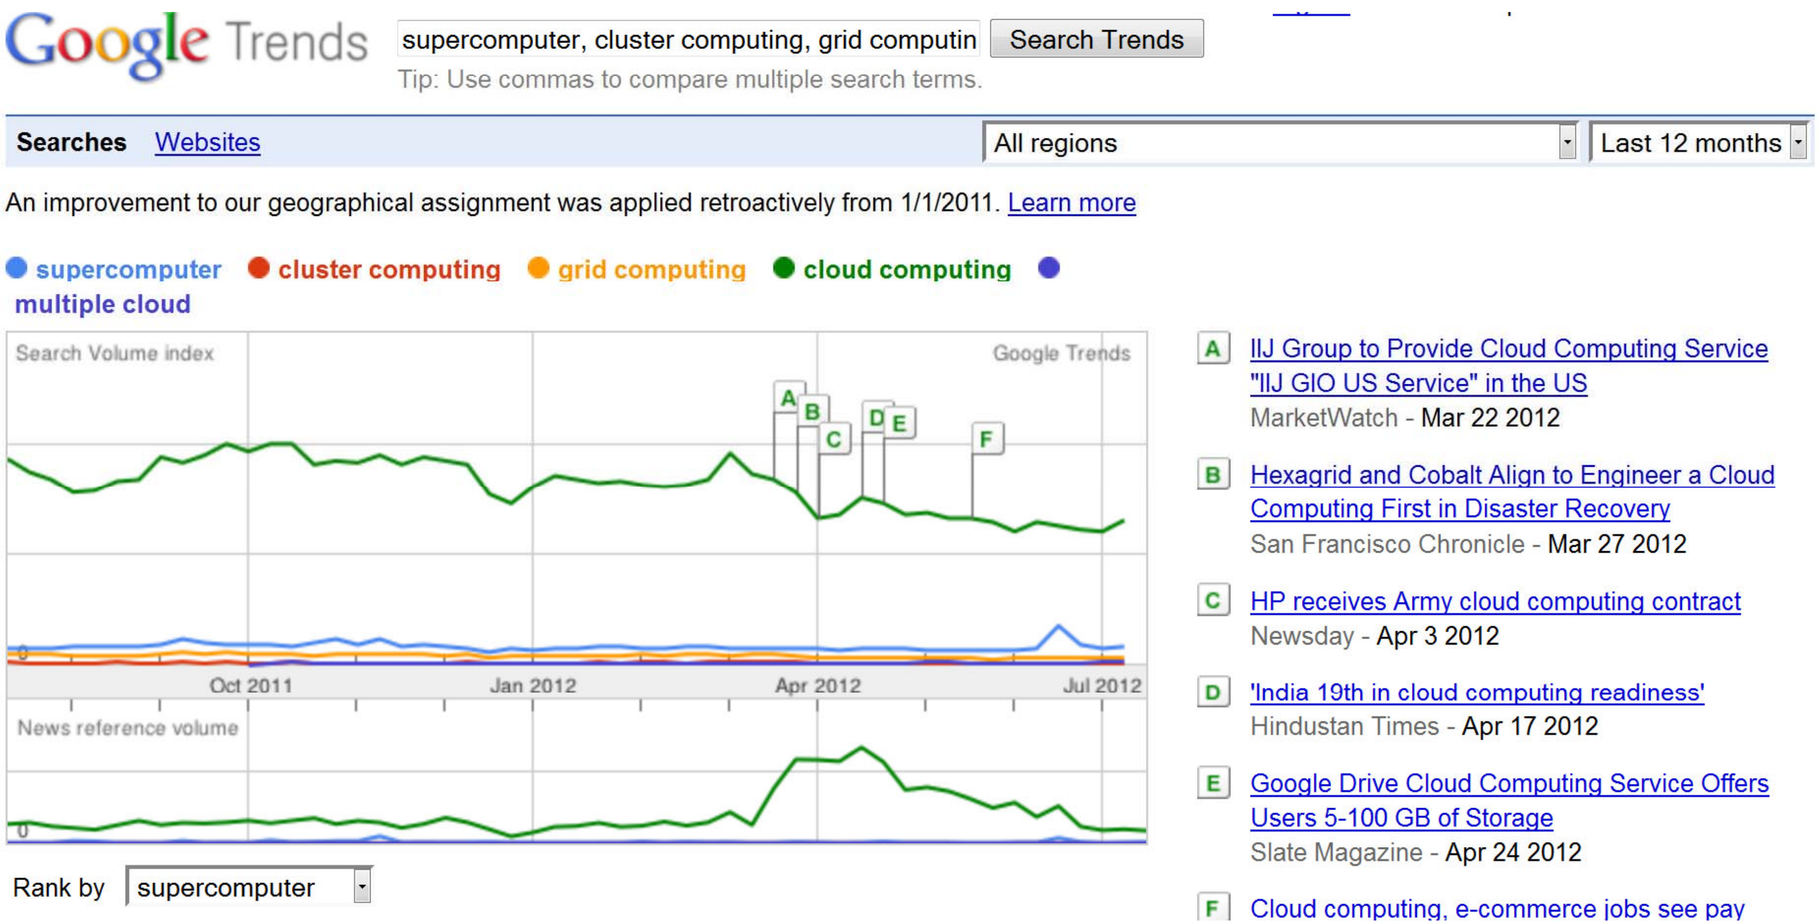

On the back stage

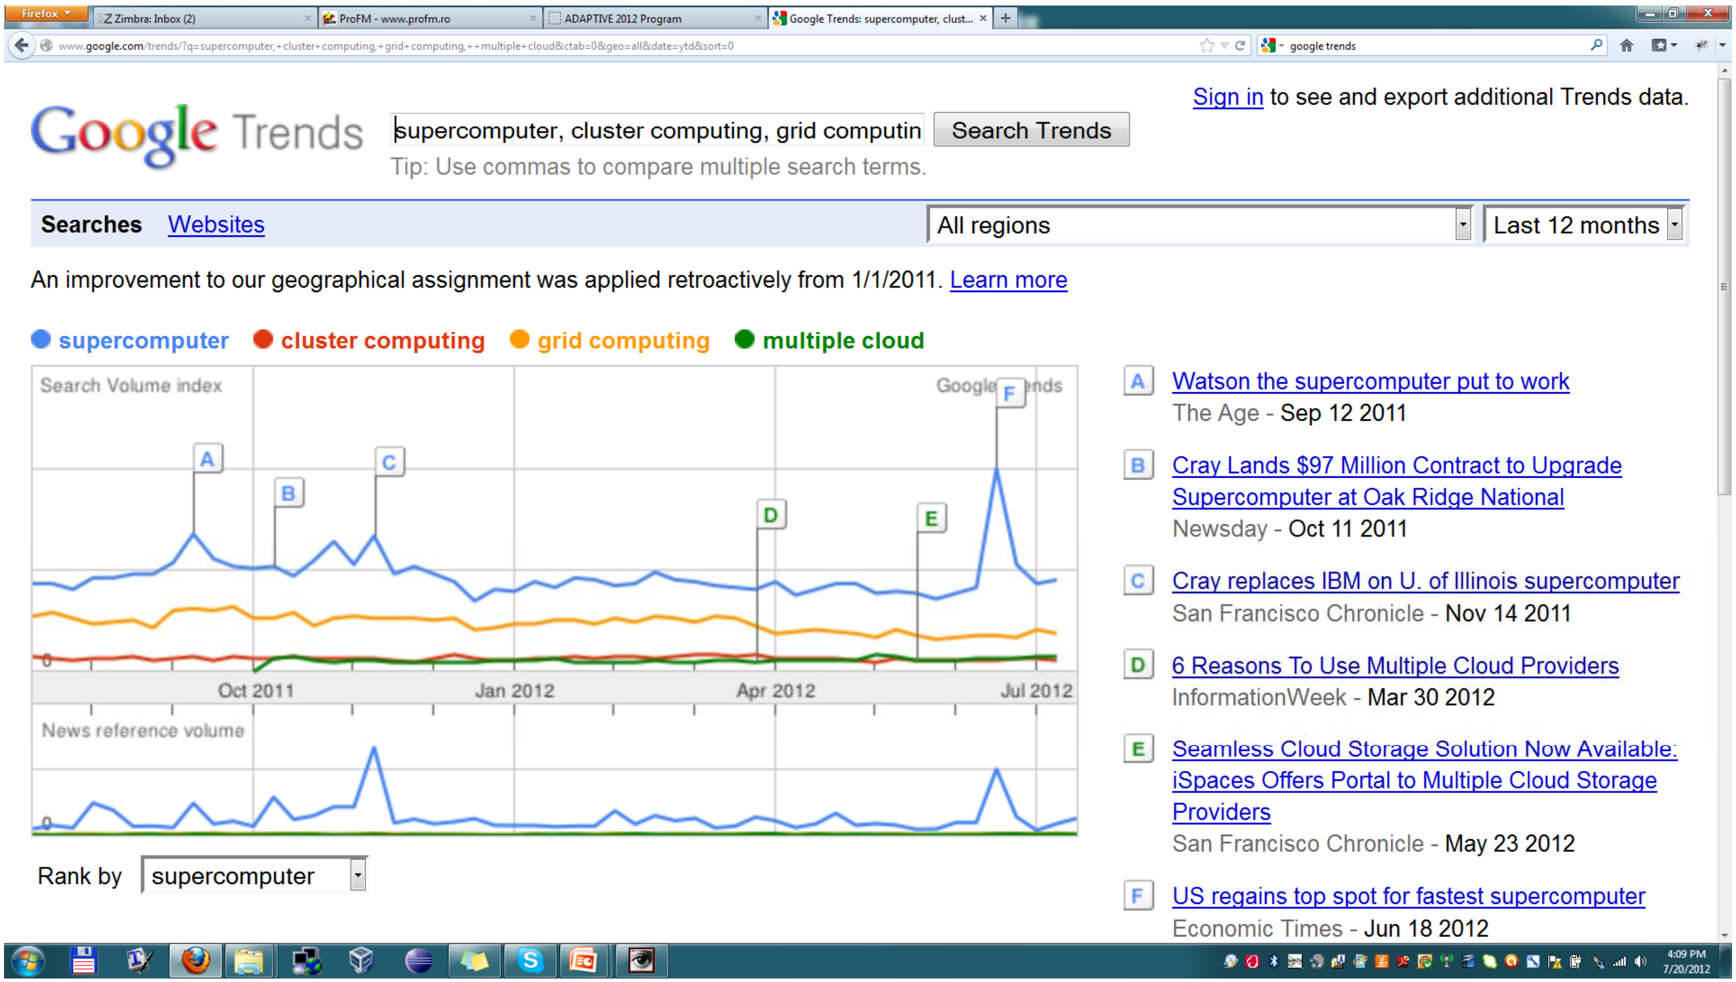

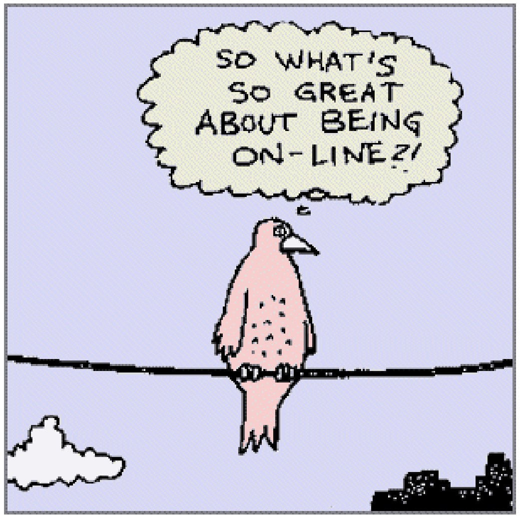

E-infrastructures characteristics

Clusters, Supercomputers, Grids, Clouds, InterCloud …

Characteristics: coupling & administrations

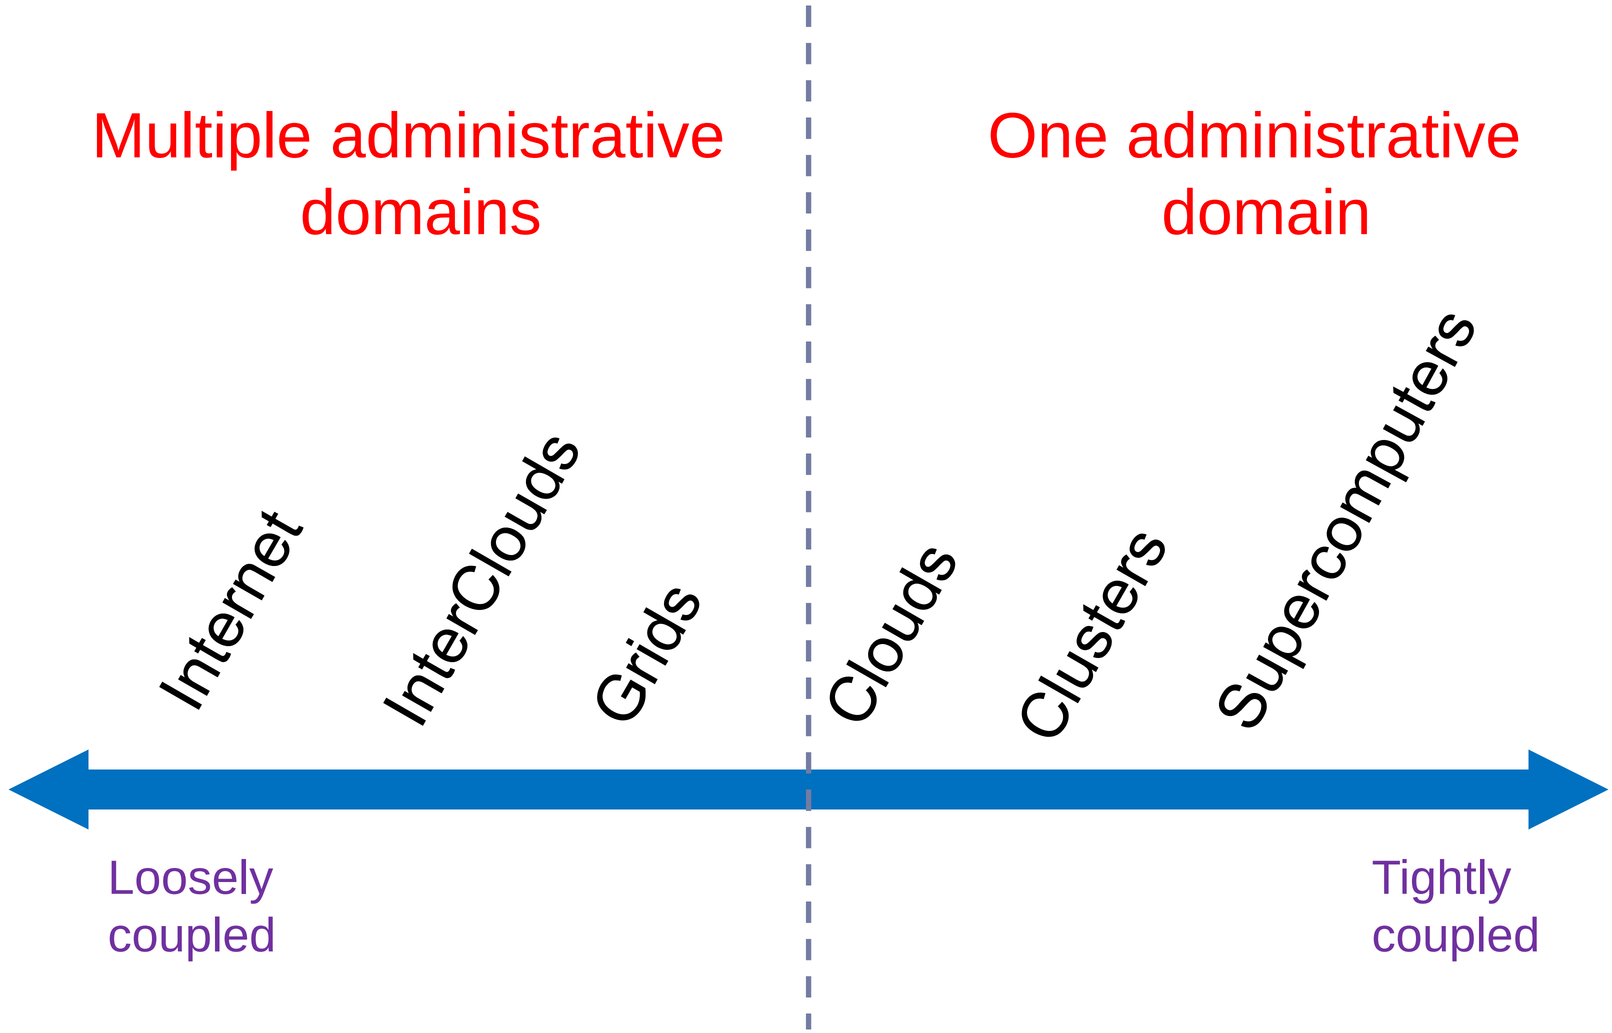

Characteristics: resources number

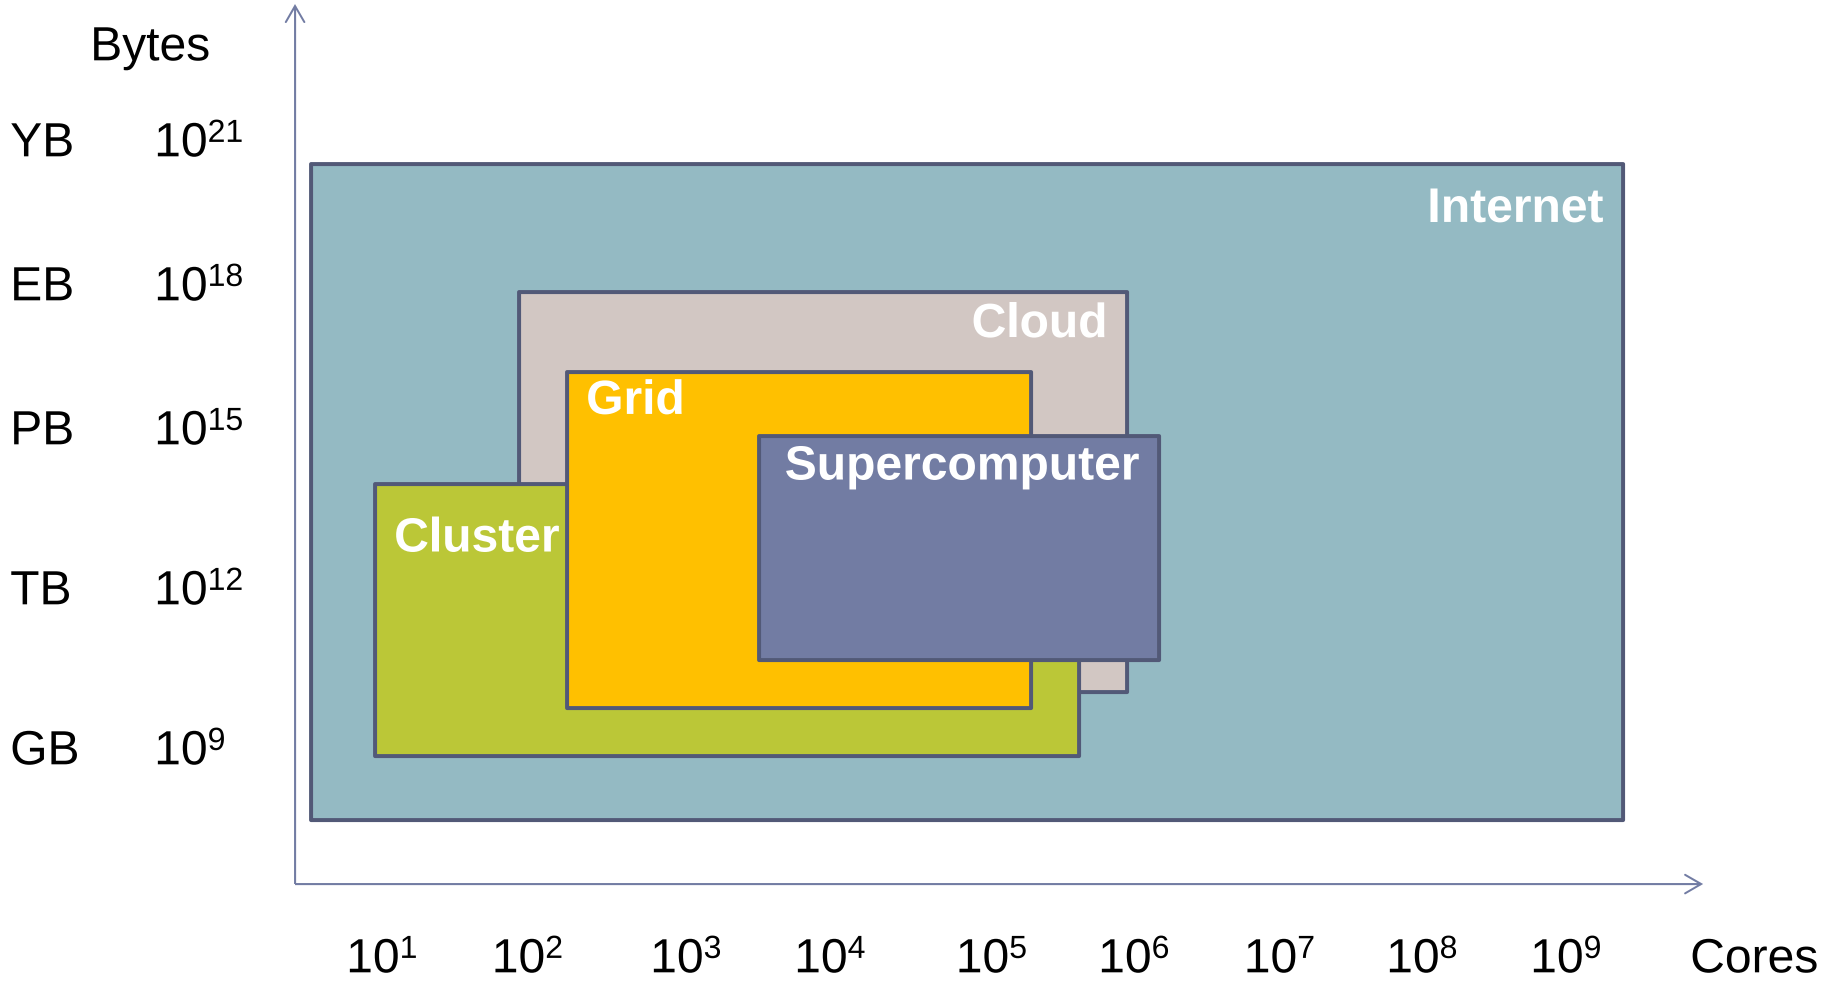

Top 500: the most powerful ones

Rank	Name	Site	Manufacturer	Country	Year	Segmen	Total Cor	Archi	Processor	Processor Tec	Process	OS F	Cores/	System Model	Intercon	Contine
1	Sequoia	DOE/NNSA/LLNL	IBM	United States	2011	Research	1572864	MPP	Power BQC 16C	PowerPC	1600	Linux	16	BlueGene/Q	Custom	Americas
2		RIKEN Advanced Institute for Com	Fujitsu	Japan	2011	Research	705024	Cluster	SPARC64 VIIIfx	Sparc	2000	Linux	8	K computer	Custom	Asia
3	Mira	DOE/SC/Argonne National Labora	IBM	United States	2012	Research	786432	MPP	Power BQC 16C	PowerPC	1600	Linux	16	BlueGene/Q	Custom	Americas
4	SuperMUC	Leibniz Rechenzentrum	IBM	Germany	2012	Academic	147456	Cluster	Xeon E5-2680 8	Intel SandyBridge	2700	Linux	8	iDataPlex DX360M4	Infiniband	Europe
5	Tianhe-1A	National Supercomputing Center i	NUDT	China	2010	Research	186368	MPP	Xeon X5670 6C	Intel Nehalem	2930	Linux	6	NUDT YH MPP	Proprietary	Asia
6	Jaguar	DOE/SC/Oak Ridge National Labo	Cray Inc.	United States	2009	Research	298592	Cluster	Opteron 6274 1	AMD x86_64	2200	Linux	16	Cray XK6	Cray Gemi	Americas
7	Fermi	CINECA	IBM	Italy	2012	Academic	163840	MPP	Power BQC 16C	PowerPC	1600	Linux	16	BlueGene/Q	Custom	Europe
8	JuQUEEN	Forschungszentrum Juelich (FZJ)	IBM	Germany	2012	Research	131072	MPP	Power BQC 16C	PowerPC	1600	Linux	16	BlueGene/Q	Custom	Europe
9	Curie thin	CEA/TGCC-GENCI	Bull SA	France	2012	Research	77184	Cluster	Xeon E5-2680 8	Intel SandyBridge	2700	Linux	8	Bullx B510	Infiniband	Europe
10	Nebulae	National Supercomputing Centre i	Dawning	China	2010	Research	120640	Cluster	Xeon X5650 6C	Intel Nehalem	2660	Linux	6	Dawning TC3600 Blade	Infiniband	Asia
11	Pleiades	NASA/Ames Research Center/NAS	SGI	United States	2011	Research	125980	MPP	Xeon E5450 4C	Intel Core	3000	Linux	4	SGI Altix ICE 8200EX/8	Infiniband	Americas
12	Helios	International Fusion Energy Resea	Bull SA	Japan	2011	Academic	70560	Cluster	Xeon E5-2680 8	Intel SandyBridge	2700	Linux	8	Bullx B510	Infiniband	Asia
13	Blue Joule	Science and Technology Facilities	IBM	United Kingdom	2012	Research	114688	MPP	Power BQC 16C	PowerPC	1600	Linux	16	BlueGene/Q	Custom	Europe
14	TSUBAME	GSIC Center, Tokyo Institute of Te	NEC/HP	Japan	2010	Academic	73278	Cluster	Xeon X5670 6C	Intel Nehalem	2930	Linux	6	Cluster Platform SL390	Infiniband	Asia
15	Cielo	DOE/NNSA/LANL/SNL	Cray Inc.	United States	2011	Research	142272	MPP	Opteron 6136 8	AMD x86_64	2400	Linux	8	Cray XE6	Custom	Americas
16	Hopper	DOE/SC/LBNL/NERSC	Cray Inc.	United States	2010	Research	153408	MPP	Opteron 6172 1	AMD x86_64	2100	Linux	12	Cray XE6	Custom	Americas
17	Tera-100	Commissariat a l'Energie Atomique	Bull SA	France	2010	Research	138368	Cluster	Xeon X7560 8C	Intel Nehalem	2260	Linux	8	bullx super-node S6010	Infiniband	Europe
18	Oakleaf-F	Information Technology Center, T	Fujitsu	Japan	2012	Academic	76800	Cluster	SPARC64 IXfx 1	Sparc	1848	Linux	16	PRIMEHPC FX10	Tofu interc	Asia
19	Roadrunn	DOE/NNSA/LANL	IBM	United States	2009	Research	122400	Cluster	PowerXCell 8i 9	Power	3200	Linux	9	BladeCenter QS22 Clu	Infiniband	Americas
495		IT Services Provider	Hewlett-Packard	United States	2012	Industry	13980	Cluster	Xeon E5620 4C	Intel Nehalem	2400	Linux	4	Cluster Platform 3000 E	Gigabit Et	Americas
496		IT Service Provider	Hewlett-Packard	United States	2010	Industry	10056	Cluster	Xeon X5670 6C	Intel Nehalem	2930	Linux	6	Cluster Platform 3000 E	Gigabit Et	Americas
497	Tsessebe	Centre for High Performance Com	Dell/Oracle	South Africa	2009	Academic	6336	Cluster	Xeon X5570 4C	Intel Nehalem	2930	Linux	4	Blade X6275/ PowerEd	Infiniband	Africa
498		Web Company (F)	Hewlett-Packard	United States	2012	Industry	11040	Cluster	Xeon X5650 6C	Intel Nehalem	2660	Linux	6	Cluster Platform SL160	Gigabit Et	Americas
499		Energy Company (A)	IBM	Italy	2012	Industry	4096	Cluster	Xeon E5-2650 8	Intel SandyBridge	2000	Linux	8	BladeCenter HS23 Clu	Infiniband	Europe
500		IT Service Provider	Hewlett-Packard	United States	2012	Industry	6064	Cluster	Xeon X5672 4C	Intel Nehalem	3200	Linux	4	Cluster Platform 3000 2	Infiniband	Americas

Top 500: the biggest supercomputers & clusters

Sequoia/DOE

RIKEN

Projected performance [J. Dongarra, June'12]

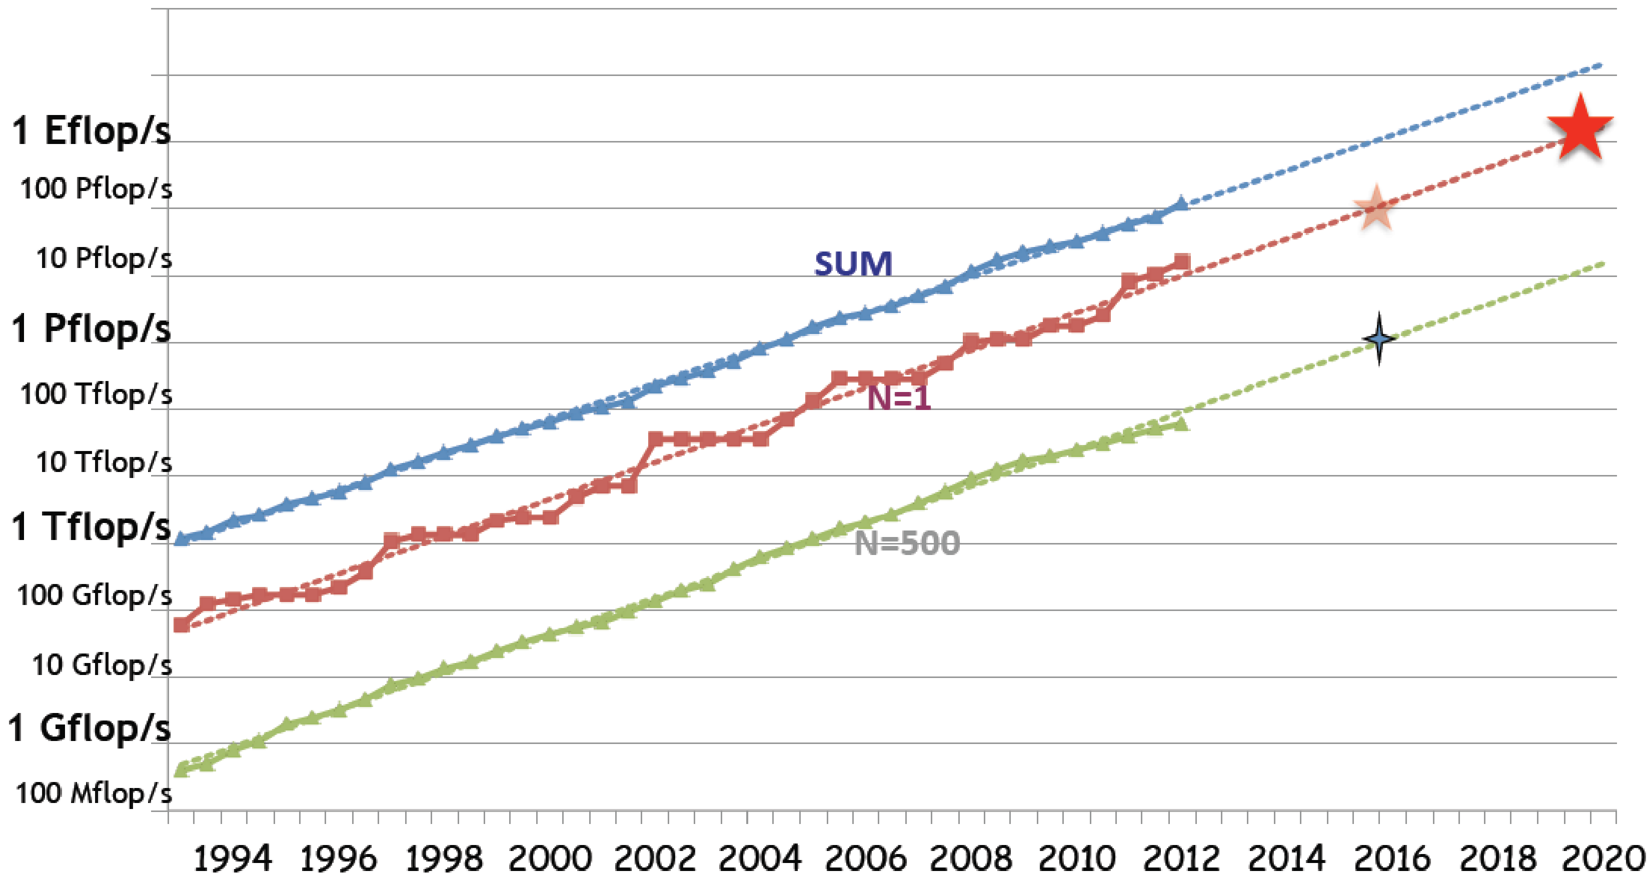

Cloud: the biggest (?)

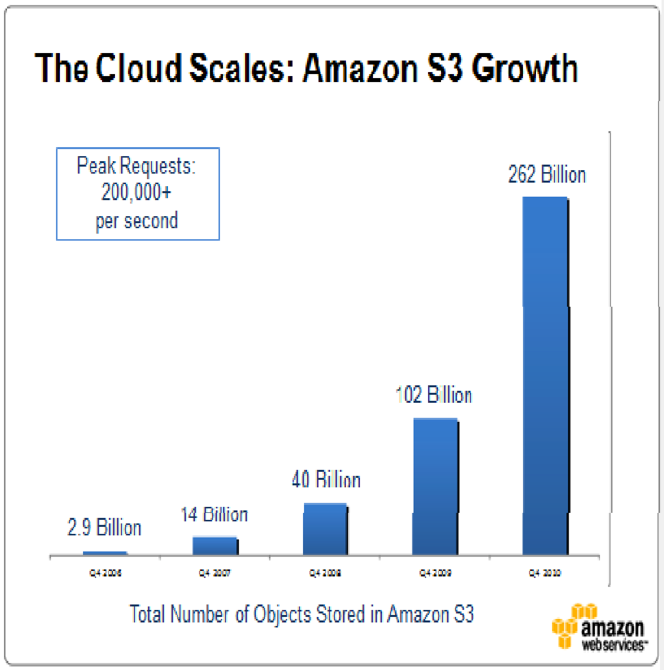

Estimated 900 PB

huanliu.wordpress.com/2012/03/13/amazon-data-center-size/

← Host server CPU utilization in Amazon EC2 cloud — Amazon DynamoDB use cases →

Amazon data center size

MARCH 13, 2012 — 94 COMMENTS

*(Edit 3/16/2012: I am surprised that this post is picked up by a lot of media outlets. Given the strong interest, I want to emphasize what is measured and what is derived. The # of server racks in EC2 is what I am directly observing. By assuming 64 physical servers in a rack, I can derive the rough server count. But remember this is an *assumption*. Check the comments below that some think that AWS uses 1U server, others think that AWS is less dense. Obviously, using a different assumption, the estimated server number would be different. For example, if a credible source tells you that AWS uses 36 1U servers in each rack, the number of servers would be 255,600. An additional note: please visit my disclaimer page. This is a personal blog, only represents my personal opinion, not my employer's.)*

Similar to the EC2 CPU utilization rate, another piece of secret information Amazon will never share with you is the size of their data center. But it is really informative if we can get a glimpse, because Amazon is clearly a leader in this space, and their growth rate would be a great indicator of how well the cloud industry is doing.

Although Amazon would never tell you, I have figured out a way to probe for its size. There have been early guesstimates on how big Amazon cloud is, and there are even tricks to figure out how many virtual machines are started in EC2, but this is the first time anyone can estimate the real size of Amazon EC2.

The methodology is fully documented below for those inquisitive minds. If you are one of them, read it through and feel free to point out if there are any flaws in the methodology. But for those of you who just want to know the numbers: Amazon has a pretty impressive infrastructure. The following table shows the number of server racks and physical servers each of Amazon's data centers has, as of Mar. 12, 2012. The column on server racks is what I directly probed (see the methodology below), and the column on number of servers is derived by assuming there are 64 blade servers in each rack.

data center\size	# of server racks	# of blade servers
US East (Virginia)	5,030	321,920
US West (Oregon)	41	2,624
US West (N. California)	630	40,320
EU West (Ireland)	814	52,096
AP Northeast (Japan)	314	20,096
AP Southeast (Singapore)	246	15,744
SA East (Sao Paulo)	25	1,600
Total	**7,100**	**454,400**

The first key observation is that Amazon now has close to half a million servers, which is quite impressive. The other observation is that the US east data center, being the first data center, is much bigger. What it means is that it is hard to compete with Amazon on scale in the US, but in other regions, the entry barrier is lower. For example, Sao

Cloud:
the biggest (?)

- June 2012:

www.extremetech.com/extreme/131962-google-compute-engine-for-2-millionday-your-company-can-run-the-third-faste

Google Compute Engine: For $2 million/day, your company can run the third fastest supercomputer in the world

By Sebastian Anthony on June 28, 2012 at 3:18 pm | 3 Comments

Share This Article

128 Like | 53 Tweet | 2,613 | 20 +1 | reddit

At the Google I/O conference in San Francisco, Google has announced the immediate availability of Compute Engine, an infrastructure-as-a-service (IAAS) product that directly competes with Amazon EC2 and Microsoft Azure. Citing more than a decade of running and optimizing its own data centers and network infrastructure, Google is claiming that the Compute Engine is more scalable, more stable, and cheaper than the competition.

For this story, we'll focus on scalability and cost (I'm sure that Compute Engine is stable, but Google just hasn't given us any figures to work with). Google says that Compute Engine has access to 770,000 cores — a figure that will surely grow over time. In one demo at Google I/O, a genomics app (it analyzed the human genome) was shown to use 600,000 cores. These cores are made available as Linux virtual machines (VMs), with 1, 2, 4, or 8 cores each. Each core apparently has access to 3.75GB of RAM each — and, of course, each VM is connected together using Google's advanced networking technologies and topologies.

777,000 cores, assuming the entire Compute Engine cluster consists of 8-core CPUs, equates to 96,250 computers. This is a huge number — probably equal to the total number of servers operated by Intel, or data centers such as The Planet or Rackspace, but

Grids: the biggest

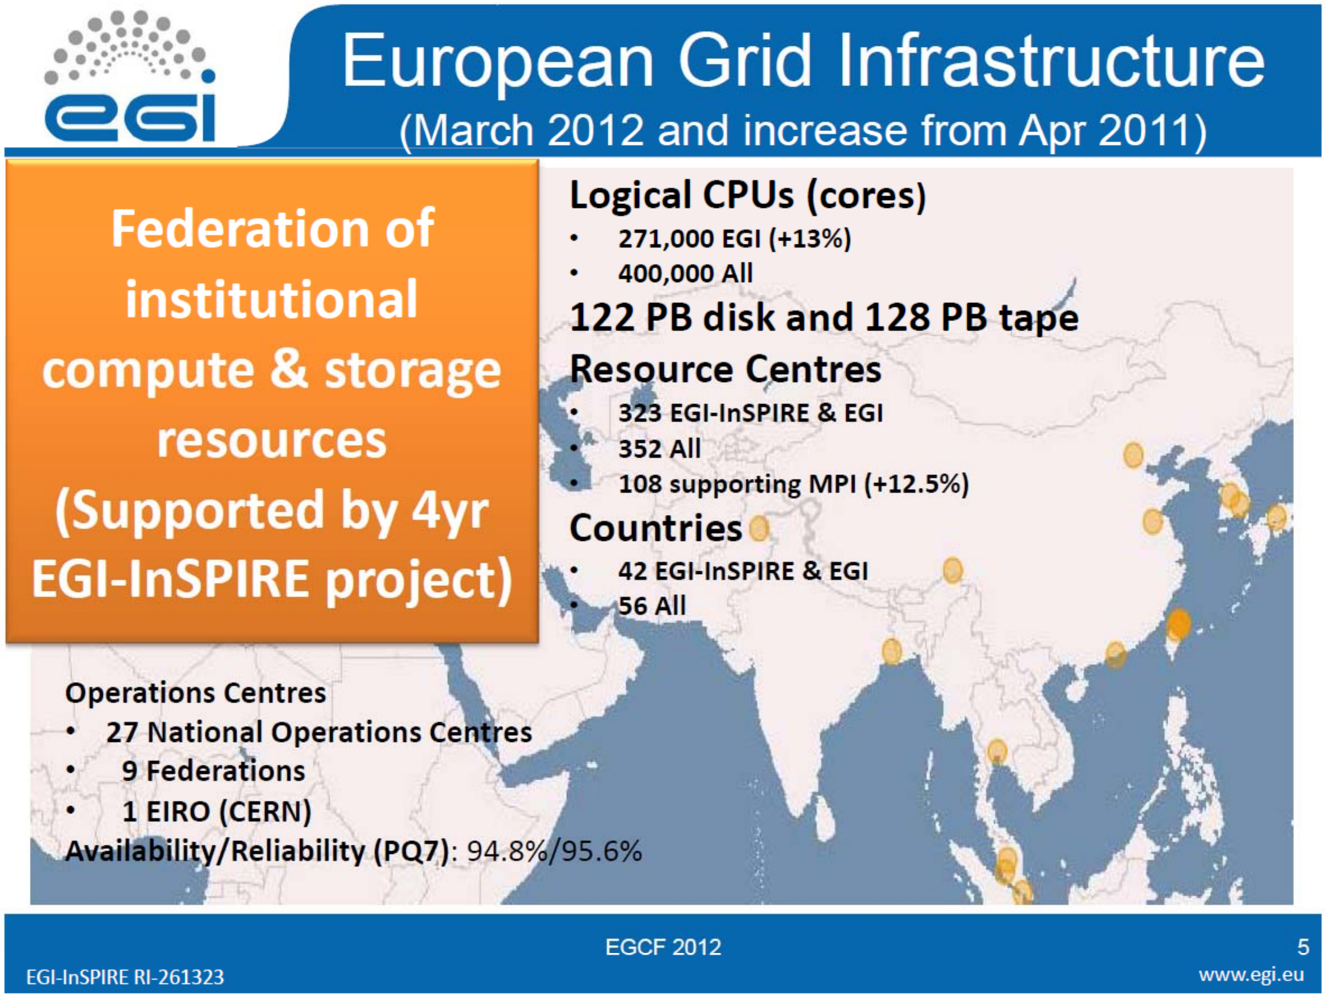

Supercomputer

- A computer at the frontline of current processing capacity in terms of speed of calculation
- How old is the term?
  - 1960s
- What is the aim
  - Solve problems faster
  - Solve problems that cannot be solve
  - on a single server
  - Solve more instances of

    of the same problem

The old: Iliac IV

Cluster

- group of linked computers working together offering an image of a single large resource
- How old?
  - 1990’s, e.g. Beowolf cluster from ‘94
- Aims:
  - “supercomputer at affordable price”
    cost-effectiveness
  - Interactivity
- Nowadays available in medium and large entreprises

Grids

- A form of distributed and parallel computing, building a 'super and virtual computer‘, composed of a cluster of networked, loosely coupled computers acting in concert to perform very large tasks.
- Terms from 1997
- Cluster of cluster
- Virtual organisations
- Join with Web services in 2002

Cloud

- NIST definition:

  Cloud computing
  is a *pay-per-use model*
  for enabling available, convenient,
  on-demand *network access*
  to a shared pool of *configurable computing resources*
  (e.g.networks, servers, storage, appls, services)
  that can be *rapidly provisioned and released*
  with minimal management effort
  or service-provider interaction
- How old?
  - 2007
- Characteristics:
  - On-demand self-service
  - Network access
  - Location-independent resource pooling
  - Rapid elasticity
  - Pay per use
- Aims:
  - Access to e-infras for everyone
  - Serves mainly web applications

InterCloud, multiple Clouds, Sky computing, Cross-Clouds …

- interconnected
  "cloud of clouds”
- extension of the Internet
  "network of networks“
- terms from 2009

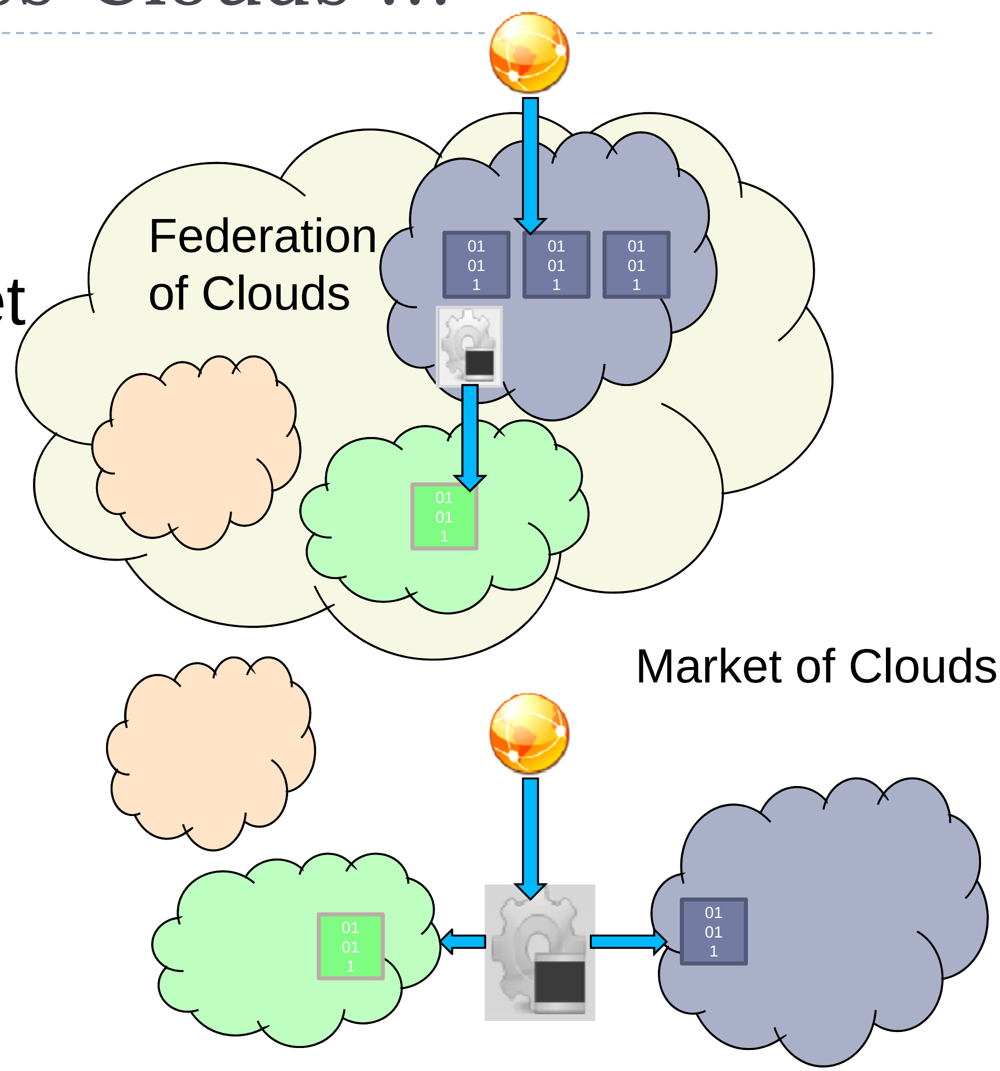

Following the giants: ‘Big’ and Famous Applications

on e-Infrastructures

Appls/supercomputers&big clusters [Top500]

Application Area	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Not Specified	209	41.8	60037590.69	82863853.8	6440642
Research	105	21	40532017.25	53789204.6	4213217
Finance	25	5	1801282.97	3512856.38	335444
Web Services	21	4.2	1755179.7	3249561	295844
Energy	17	3.4	3221250.39	4311936.38	276356
Geophysics	14	2.8	1225886	2918282.4	100624
Services	14	2.8	988820.5	1753013.4	164756
Defense	13	2.6	2588070.4	3138660.08	319536
Weather and Climate Research	13	2.6	3934162	5152868.06	351460
Logistic Services	8	1.6	531532.93	1013975.9	92722
IT Services	8	1.6	566670.5	1098033.52	106572
Entertainment	7	1.4	497856	692428.4	61824
Aerospace	7	1.4	1903523	2528001.47	202508
Environment	6	1.2	754030	1250227.72	59776
Benchmarking	6	1.2	911196	1127694.4	66176
Information Service	5	1	402436.66	722117.44	63452
Information Processing Service	5	1	345266	569035.52	85856
Automotive	2	0.4	177240	200833.92	17136
Telecommunication	2	0.4	150995.72	277047.36	26796
Internet Provider	2	0.4	162555	306390.66	31776
Transportation	2	0.4	126084	237144.32	22288
Semiconductor	2	0.4	180472	253384.72	18360
Electronics	2	0.4	124488	139937.28	13152
Software	1	0.2	172691	209715	16384
Medicine	1	0.2	63830	94208	10240
Cloud Services	1	0.2	89940	155079	4968
Life Science	1	0.2	97071	159948.8	18176
Retail	1	0.2	75649	145705	11904

"Classical" HPC applications

- Type of applications:
  - Weather forecast and climate research
  - Molecular modeling (e.g. crystals, biology, chemistry)
  - Quantum physics and physical simulations (e.g. nuclear fusion)
- Open HPC applications:
  - Bio-informatics:
    - mpiBLAST, MPI-HMMER
  - Molecular Dynamics:
    - GROMCAS, NAMD, Desmond, OpenAtom
  - Environment/Weather
    - POP, WRF, MM5

Applications on Grids [EGI statistics]

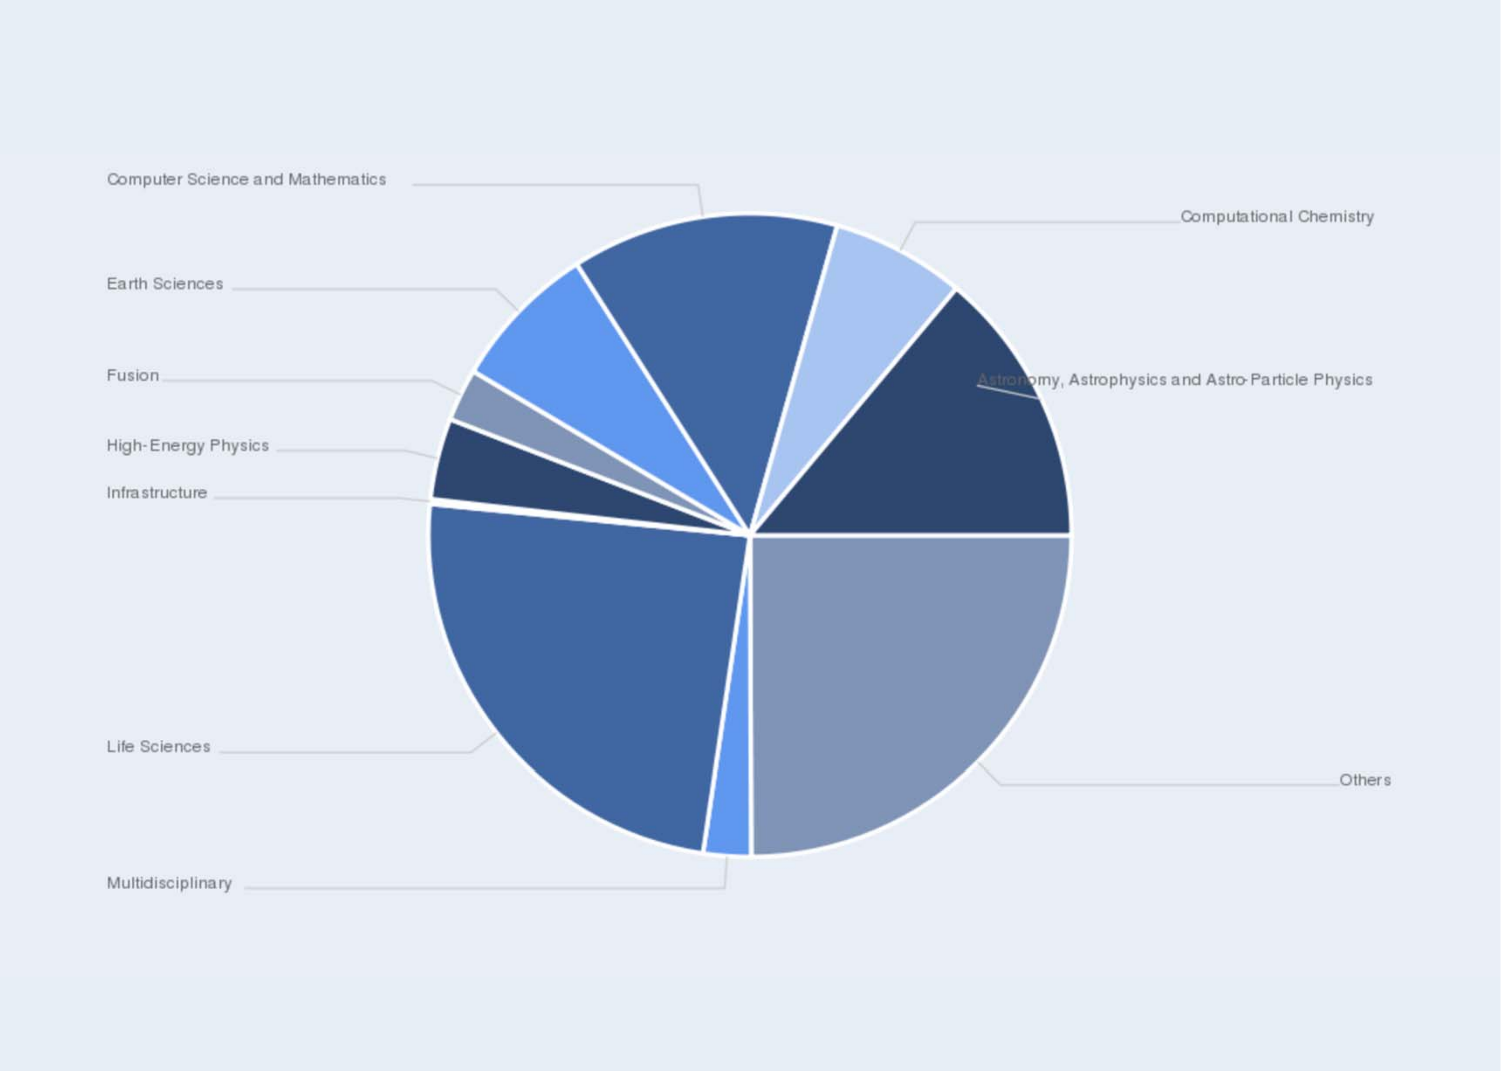

Scientific applications on Clouds

- A typical example:

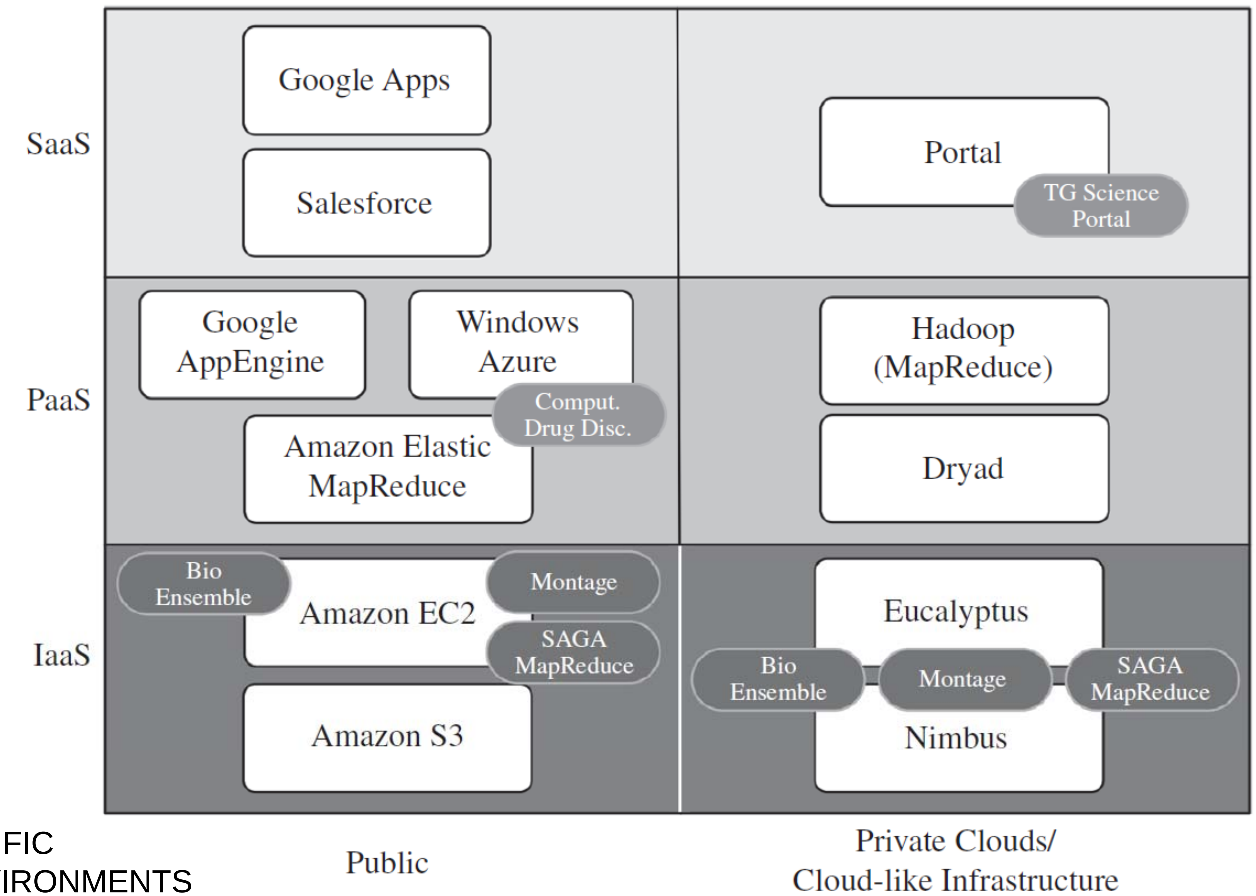

From: UNDERSTANDING SCIENTIFIC APPLICATIONS FOR CLOUD ENVIRONMENTS
S. JHA, D.S. KATZ, A. LUCKOW, A.MERZKY, K. STAMOU,
Cloud Computing: Principles and Paradigms,
Edited by R. Buyya, J. Broberg and A. Goscinski
2011 John Wiley & Sons, Inc.

There are appls which can reach exascale?

- E.g. ExaScience Lab, Leuven
  - Space weather predictions
- DOE – Grand challenge workshop 2011

http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/

  - Blue Brain project

Driving Applications Areas	
Circa 1990	**Circa 2018-2025**
Weather & Ocean Modeling	Climate Modeling
Chemistry & Materials	Chemistry & Materials
Plasma Modeling	Fusion Energy Sciences
Computational Biology & Bioinformatics	Computational Biology & Bioinformatics
High Energy/Quantum Physics	High Energy/Quantum Physics
Combustion Systems	Combustion Systems
Computational Electromagnetics	National Security Applications
Computational Fluid Dynamics (various)	Computational Fluid Dynamics (various)
Semiconductor Modeling & Design	
Superconductor Modeling	
Pharmaceutical Design	
Speech & Natural Language	
Vision & Cognition	
	Nuclear Physics
	Nuclear Energy Systems

FAR TO GO

The Blue Brain Project has steadily increased the scale of its cortical simulations through the use of cutting-edge supercomputers and ever-increasing memory resources. But the full-scale simulation called for in the proposed Human Brain Project (red) would require resources roughly 100,000 times larger still.

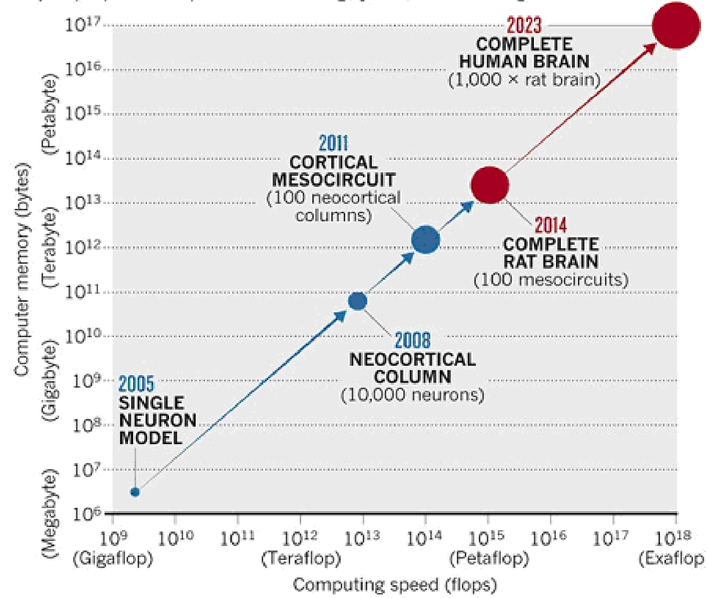

To port or not to port my application

on e-Infrastructures

Researcher vs. Supercomputer admin

- My appl responds too slow on my desktop
- My application is parallel
- My parallel appl scales well for a small no. cores
- My appl needs human interaction
- My appl is written in a specific language or OS
- My appl uses special tools
- My appl uses special libraries
- My application uses real time data streaming
- I need the results today
- I want to test it to learn

- There is a parallel version of it?
- Your appl uses MPI?
- Is your problem scalable to thousands cores?
- Have you a batch version?
- Have you a version on Unix, in C or Fortran?
- Have you licenses for these tools?
- These libraries are available for supercomputers?
- I/O is not supported well
- Are you able to wait in the queue?
- Are you in a project? Who supervise you? Have you apply for an account?

Researcher vs. Cluster admin

- My appl is running on a certain OS
- My appl is using large data
- My appl is using a proprietary code
- My appl is interactive
- My appl should be faster
- My appl is a distributed one
- I have an urgency
- I want to test it to learn

- Is this OS matching ower?
- They need pre- and post-processing? Store?
- Have you multiple/cluster/network licenses?
- Where the interface will sit?
- There is a parallel version?
- How the components are interacting over the network?
- Which priority has your task?
- Are you our employee/student?

Researcher vs. Grid admin

- My appl is running on certain OS
- Is this OS Linux type XX?
- My appl is an executable
- Have you the source codes and us the compilers?
- My appl uses special libs or tools
- Are these libs/tools pre-installed at our Grid sites?
- My appl is interactive
- Have you a batch version?
- I have an urgency
- Can you wait in the queue?
- I want to test it to learn
- Have you a Grid certificate? In which Virtual Organization you are?

Researcher vs. Cloud “admin”

- My appl is running on certain OS
- My appl is parallel
- My appl is an executable or uses special libs
- My appl uses special tools
- My appl is interactive
- I want to test it to learn

- You see this OS on our registry?
- Can we do cluster-on-demand? It is elastic?
- Can you pack it in VM with all dependences?
- Have you licenses for multiple copies?
- Can you comply with the security rules?
- Have you a credit card?

Appl supported on own e-Infras:
- Crystal growing simulations
- Airfoil design
- Datamining in medical databases
- Expert systems for numerics
- Membrane computing simulations
- Earth observation services
- …

Tools for supporting appls:
- EpODE, NESS
- PVMMaple, Maple2Grid
- ParallelJess
- GiSHEO, ESIP
- mOSAIC
- …

A use case

UVT team experience

What we can do with these? [UVT equipments]

Blue Gene/P

4096 cores

Cluster

400 cores

Earth Observation problems

- Both computational and data intensive
- Real time processing confronts several difficulties in one single computer and even impossibility
- Need of a computational environment handling
  - hundreds of distributed databases,
  - heterogeneous computing resources,
  - and simultaneous use

From the small to the big

- Simple algorithms: merge

Band 3 Band 4 Band 5 Pseudo-image

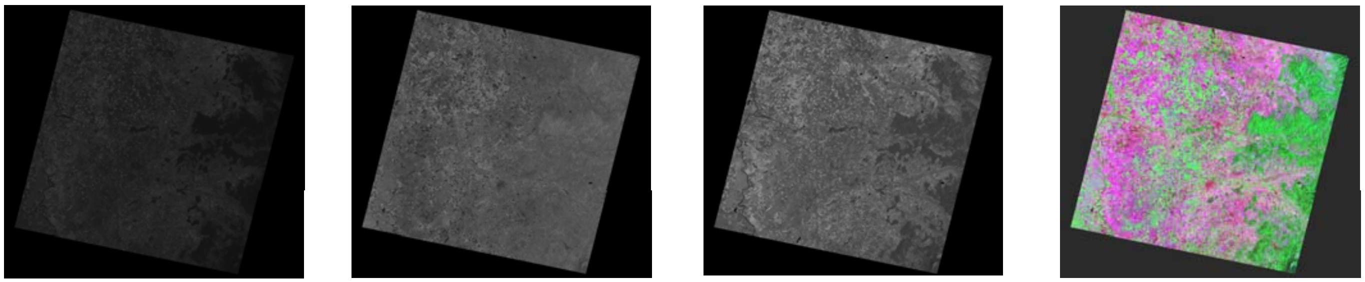

- Computational intensive algorithms

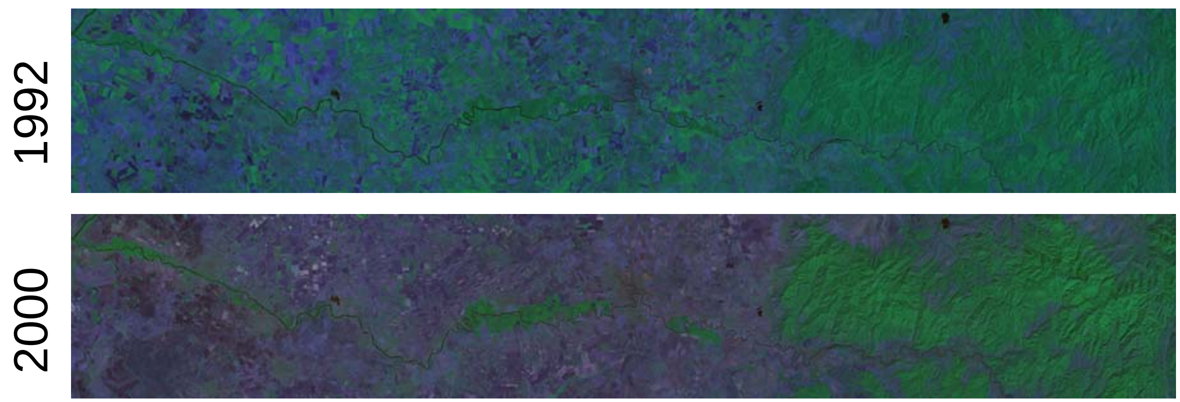

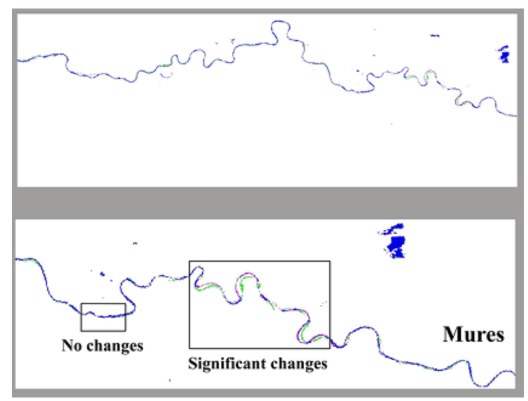

E.g. D.Petcu, V. Iordan, Service based on GIMP for Processing Remote Sensing Images, SYNASC 2006

Why Clusters

- Store the big data
- Process the data where they are

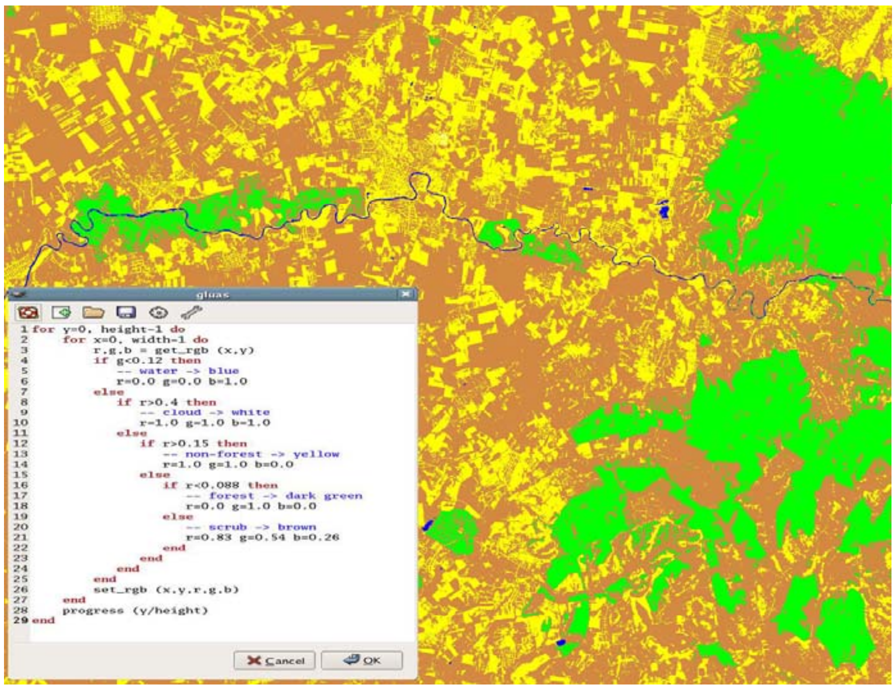

No. of processors	1	2	4
Time (s)	457	256	168
Speedup	-	1.78	2.72
Efficiency	-	89%	68%

Why Supercomputer

Algorithm 4 The general structure of the parallel Fuzzy c-Means

Read image slice $X(p) = X_{i \in S_p}$
Initialize the local membership values $u_{ij}(p)$, $i \in S_p$, $j = \overline{1,c}$
$iter = 0$
repeat
  Compute $C_j(p) = \sum_{i \in S_p} u_{ij}^m(p) X_i(p)$, $j = \overline{1,c}$
  Compute $C'_j(p) = \sum_{i \in S_p} u_{ij}^m(p)$, $j = \overline{1,c}$
  Call `MPI_Allreduce` to compute $C_j = C_j(1) + \ldots + C_j(P)$ for all $j = \overline{1,c}$
  Call `MPI_Allreduce` to compute $C'_j = C'_j(1) + \ldots + C'_j(P)$ for all $j = \overline{1,c}$
  Compute $V_j = C_j / C'_j$, $j = \overline{1,c}$
  Update the local membership values u_{ij}^{new}, $i \in S_p$, $j = \overline{1,c}$
  Compute $Err(p) = \max_{i \in S_p, j=\overline{1,c}} |u_{ij}^{new}(p) - u_{ij}(p)|$
  Call `MPI_Allreduce` to compute $Err = \max\{Err(1), \ldots, Err(P)\}$
  $iter = iter + 1$
  $u_{ij} = u_{ij}^{new}$, $i \in S_p$, $j = \overline{1,c}$
until $Err < \epsilon$ or $iter > iterMax$
Compute the cluster validation measure(s)
if p==1 **then**
  Construct the classified image
end if

- Scalable algorithms

Table 8. Results on BlueGene/P for the parallel version of SFCM (100 iterations, 5 clusters, neighborhood size equal to 5). Test image: AVIRIS image (224 spectral bands, 1087 × 614 pixels)

No. Proc	K_w	K_h	P/16	Time(16)/ Time(P)	Total Time(s)	Time Send(s)	Time Reduce(s)	Time Send(%)	Time Reduce(%)
1024	32	32	64	40.38	4.94	0.10	1.64	2.09	33.27
1024	2	512	64	17.02	11.71	7.30	0.06	62.35	0.48
512	16	32	32	27.55	7.24	0.97	0.05	13.34	0.76
512	1	512	32	10.23	19.50	10.75	0.05	55.14	0.26
256	16	16	16	15.84	12.58	0.10	0.05	0.83	0.43
256	256	1	16	8.59	23.21	7.83	0.05	33.75	0.23
128	8	16	8	7.68	25.96	1.38	0.05	5.31	0.19
128	1	128	8	6.63	30.09	3.97	0.05	13.21	0.16
64	8	8	4	3.90	51.08	1.95	0.05	3.82	0.10
64	1	64	4	3.81	52.34	3.04	0.05	5.81	0.10
32	4	8	2	2.02	98.80	0.08	0.05	0.08	0.05
32	1	32	2	2.01	99.11	0.55	0.04	0.56	0.04
16	4	4	1	1.01	197.65	0.08	0.05	0.04	0.03
16	16	1	1	1.00	199.39	0.69	0.04	0.35	0.02

D. Petcu et al,
Fuzzy Clustering of Large Satellite Images using High Performance Computing,
In Procs of SPIE Volume 8183, no. 818302 (2011), SPIE Remote Sensing Conference: High-Performance Computing in Remote Sensing, 19-22 September 2011, Prague,
Doi:10.1117/12.898281

Why Grids

- Remote services that can be combined
- Process the distributed data where they are

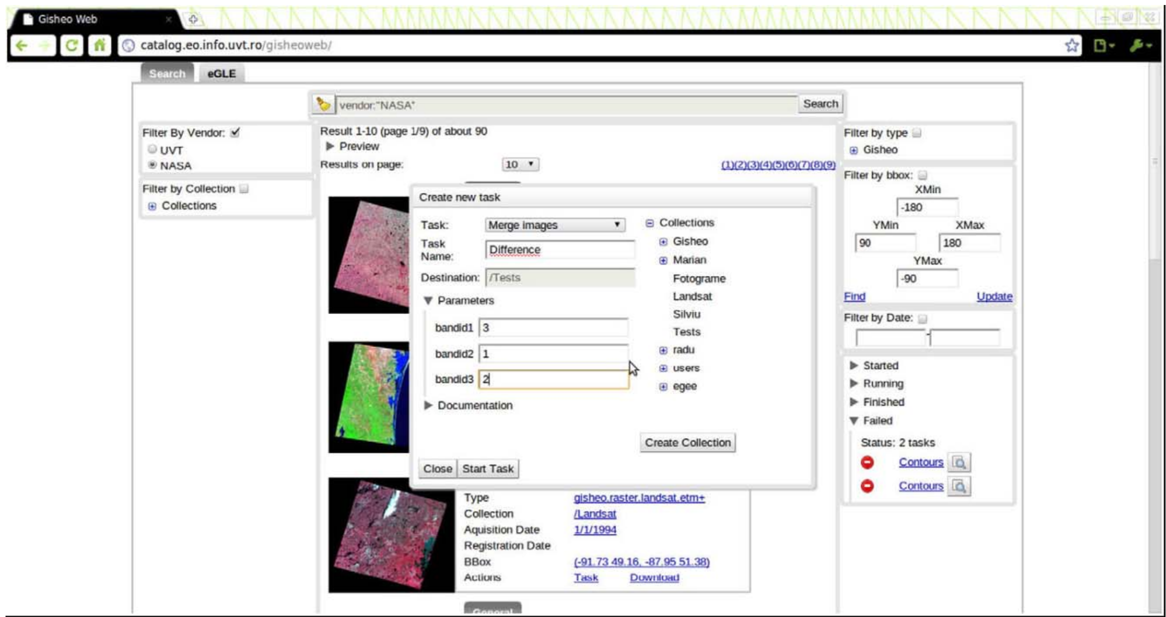

http://gisheo.info.uvt.ro

D. Petcu et al, Experiences in building a Grid-based platform to serve Earth observation training activities, Computer Standards Vol. 34 (6), 2012, 493-508, 10.1016/j.csi.2011.10.010.

Why Clouds

- Store old data
- Share the data
- Reprocess according new algs

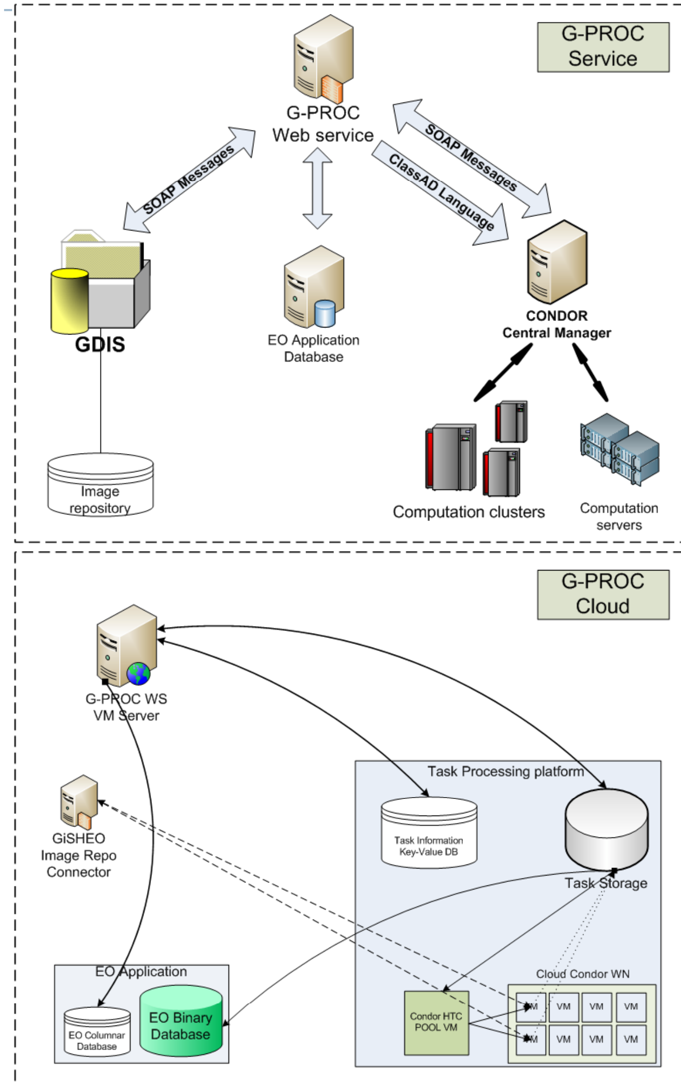

- Roberto Cossu, Claudio Di Giulio, Fabrice Brito, Dana Petcu, Cloud Computing for Earth Observation to appear in the book *Data Intensive Storage Services for Cloud Environments*

Why InterCloud

- Independence from the Cloud

Federation of Clouds:

Horizontal
or
InterClouds

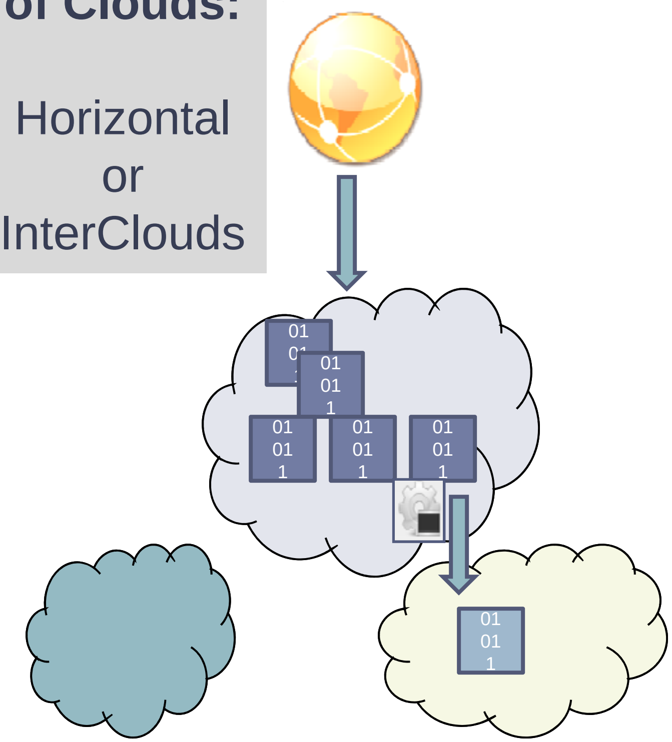

On-the-fly Multiple Clouds:

Cross-Cloud
or
Sky computing

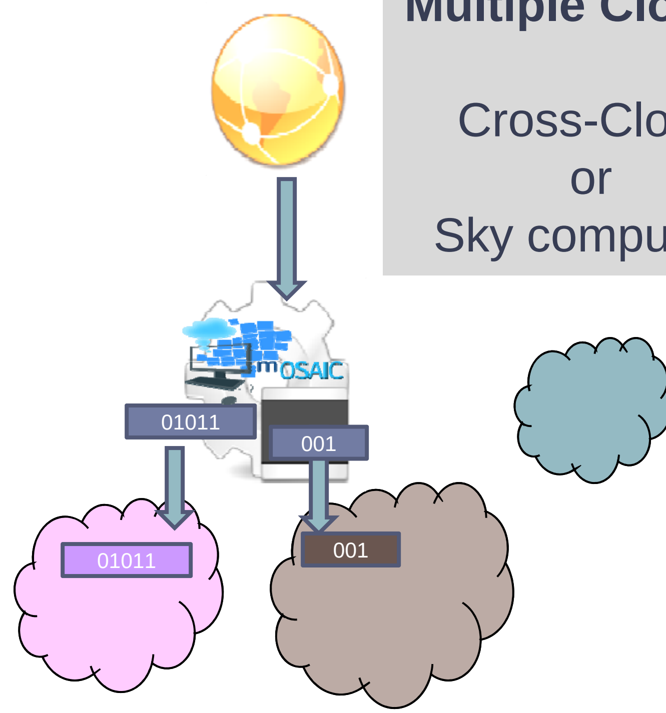

D. Petcu, S. Panica, M. Neagul, From Grid Computing Towards Sky Computing. Case Study for Earth Observation, Proceedings Cracow Grid Workshop 2010, *“If Mohammed will not go to the mountain, the mountain must come to Mohammed.”*

Programmable e-Infrastructures and automated deployment

Why to care about where to go?

Programmable Infrastructures?

- Programmatic access
  to the devices connected to the Internet
- Identified until recently with programmable networks
- Should involve:
  - network switches (Cloud networking),
  - simple gadgets or instruments (Internet-of-Things)
  - data center resources (Cluster, Grid, Cloud computing)
- Problem:
  - Manually intervention is still required in several processes involving e-infras settings

Supercomputer & cluster cases

- Supercomputer:
  - Lowest
  - Queuing system
    - Needs to know in advance the resources
  - Access restricted to registered users
    - Based on complex approval
  - Deployment based on CLIs
    - for file transfers, ssh commands etc
- Cluster
  - Increasing the no. of threads is possible
  - HPC in the Cloud can help?

Grid case

- Globus WSRF
  - Stateful services e.g. in Java
- gEclipse
  - an integrated, Grid enabled workbench tool for Grid appl users, developers and operators based on the Eclipse platform
- SAGA
  - job handling and monitoring, file transfer and management, distributed orchestration mechanisms.
  - Python and C++
  - Uniform Access-layer to DCI (EGI, XSEDE, DATAONE, UK NGS, NAREGI/RENEKI) and Clouds (recent)

Cloud case

- REST interfaces for e-infra resources
  - Designed for application-to-adminAppl interaction
- Programmatic auto-scaling
  - Partially supported by IaaS and PaaS
  - Hidden at SaaS level
- Scheduling mechanism
  - Usually not controlled by the appl developer
- Steps made by Clouds:
  - (programmatic) elasticity in term of resources
  - uniform treatment of infrastructure, software, networking as (programmable) services
  - Cloud networking is essential for the success of Cloud appls

InterCloud case

- **Barrier to pass:**
  - portability and interoperability
- **mOSAIC – Open API and PaaS**
  - Project partially funded by European Commission (2010-2013)
  - Portability of applications build on top of IaaS
  - Appl should by build from components and use message passing
  - References:
    - Official site (including paper list): http://www.mosaic-cloud.eu
    - Codes: https://bitbucket.org/mosaic,
    - Docs: http://developers.mosaic-cloud.eu
    - Demos: on YouTube: search “mOSAIC Cloud computing”
    - *Yesterday training & today talk on auto-scaling*
- **MODAClouds**
  - To start from 1 Oct 2012, www.modaclouds.eu
  - Project partially funded by EC commission
  - Model-driven architecture + Cloud =
    “from the model generated codes in the Cloud”

Conclusions (to take-away)

- Diversity in e-infras still pushes the appl developer to think about the support e-infras
- The processing and storing capacity is huge in any category
- There are some appl that can use exascales in terms of flops and bytes
- We hope in a future when e-infras are programmable