

Grid computing: yesterday, today and tomorrow?

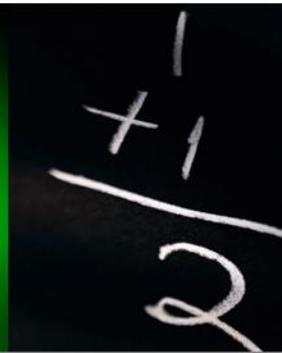
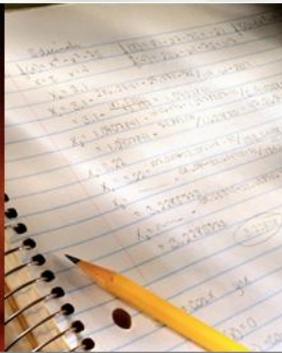
Dr. Fabrizio Gagliardi

EMEA Director

External Research

Microsoft Research

Cracow Grid Workshop 2008
Cracow, October 14th



Outline

- Yesterday and today:
 - Achievements in the area of e-Infrastructures and Grid computing
 - Examples beyond e-Science
 - Issues : Complexity, Cost, Security, Standards
- The future:
 - Cloud Computing, Virtualisation, Data Centers, Software as a Service, Multi-core architectures, Green IT

Conclusions



The European Commission strategy for e-Science

e-Infrastructures: strategy

Connecting the finest minds
Sharing the best scientific resources
Building global virtual communities



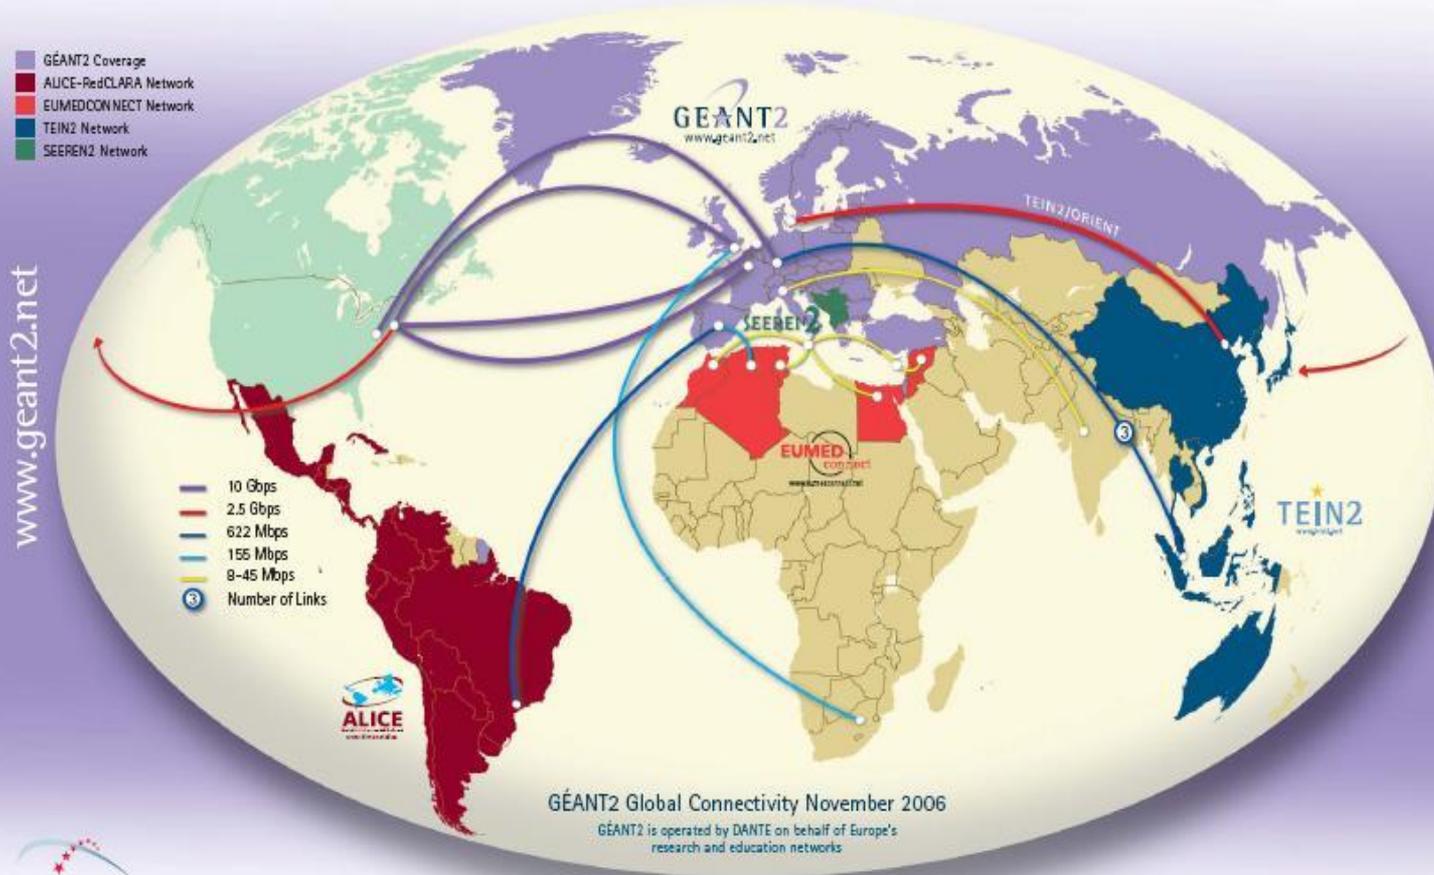
Mario Campolargo, EC, DG INFSOM, Director of Directorate F: Emerging Technologies and Infrastructures



http://cordis.europa.eu/fp7/ict/programme/events-20070524_en.html

e-Infrastructure achievements: Research Networks

GEANT2 At the Heart of Global Research Networking



★ Connect ★ Communicate ★ Collaborate



e-Infrastructure HPC achievements: EGEE and DEISA

egEE
Enabling Grids
for E-science

EGEE

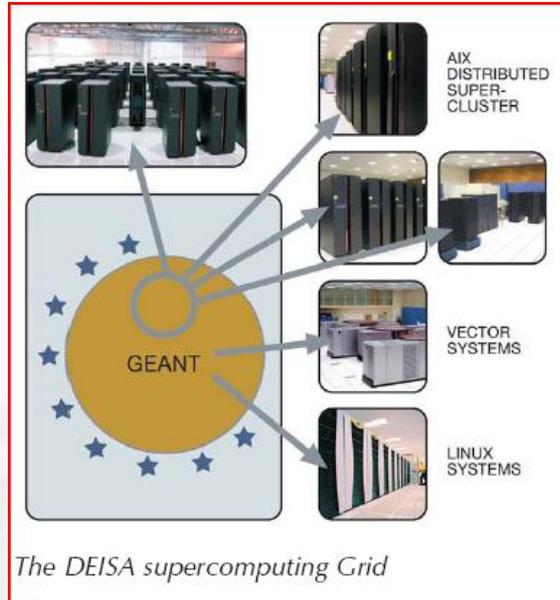
Site: SN-UCAD
SN-UCAD
 Universit? Cheik Anta Diop, Dakar
 Scheduled - 0 / Running - 6
[show CE gridce.cclucad.sn](#)

All SN-UCAD Computational Elements (Last 6 hours)

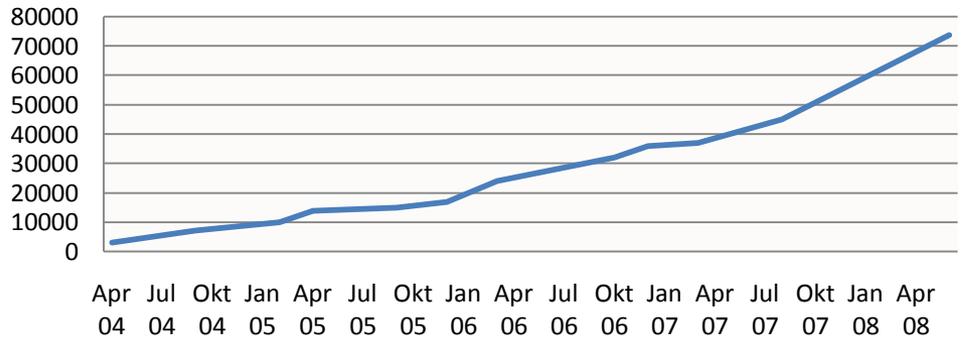
Time	Value
12:00	0.35
13:00	0.65
14:00	0.65
15:00	0.65
16:00	0.65
17:00	0.65

Imperial College

DEISA

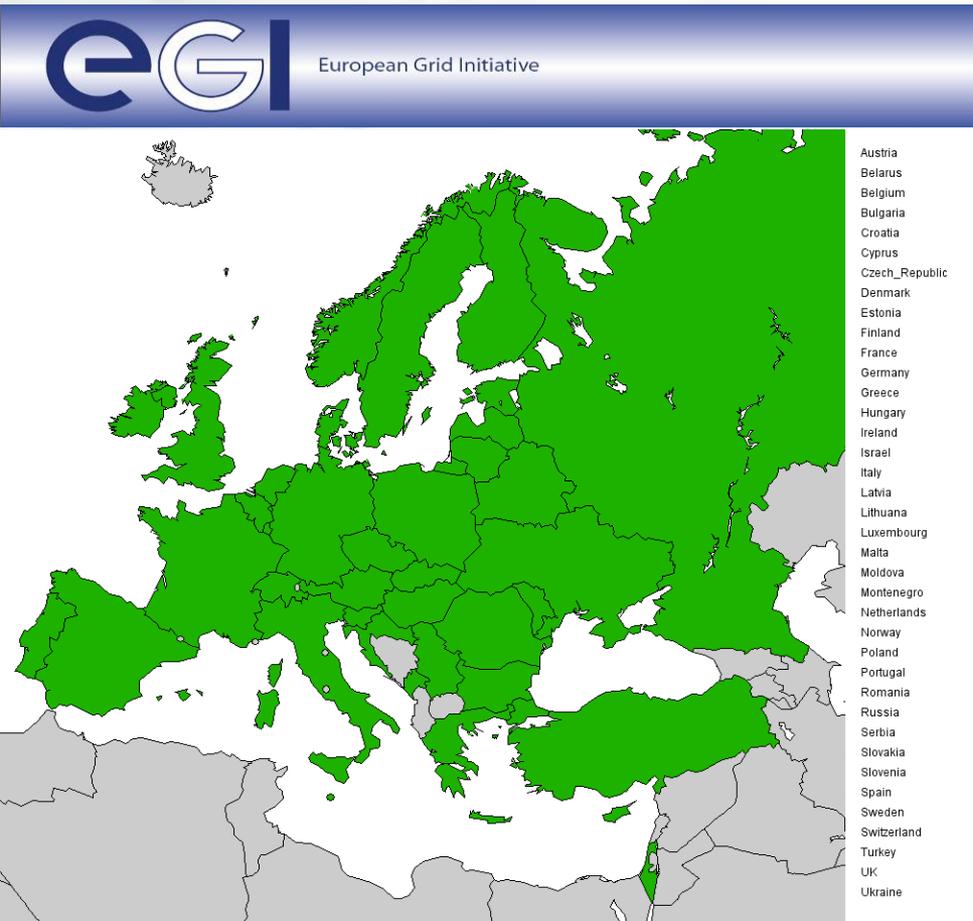


No. Cores



Distributed European Infrastructure for Supercomputing Applications

e-Infrastructure HPC next steps: EGI and PRACE



EGI European Grid Initiative

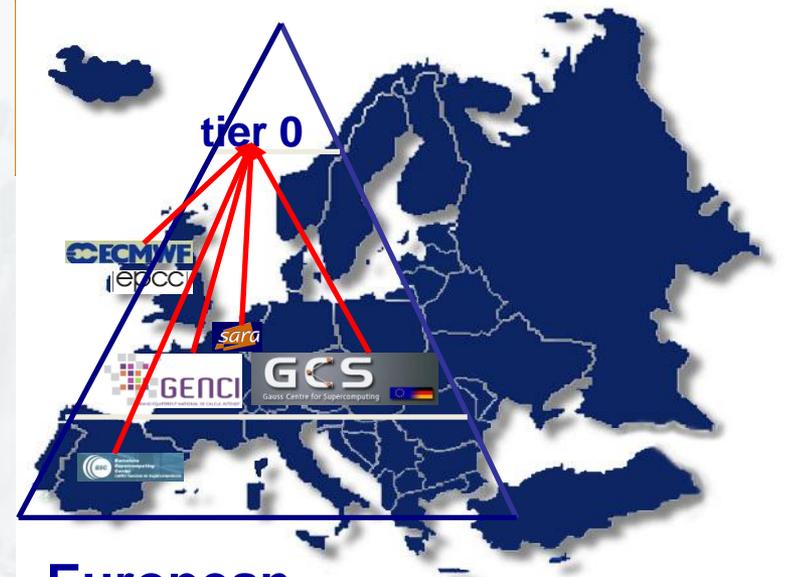
A map of Europe with countries highlighted in green. To the right of the map is a list of countries:

- Austria
- Belarus
- Belgium
- Bulgaria
- Croatia
- Cyprus
- Czech_Republic
- Denmark
- Estonia
- Finland
- France
- Germany
- Greece
- Hungary
- Ireland
- Israel
- Italy
- Latvia
- Lithuania
- Luxembourg
- Malta
- Moldova
- Montenegro
- Netherlands
- Norway
- Poland
- Portugal
- Romania
- Russia
- Serbia
- Slovakia
- Slovenia
- Spain
- Sweden
- Switzerland
- Turkey
- UK
- Ukraine



PRACE

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE



tier 0

Diagram illustrating the European Ecosystem, showing a central node labeled "tier 0" connected to various HPC centers: CECMWE, epcc, sara, GENCI, and GCS (Gauss Centre for Supercomputing).

European Ecosystem

In summary:

Grid achievements for e-Science



- Grid for e-Science: mainly a success story!
 - Several maturing Grid Middleware stacks
 - Many HPC applications using the Grid
 - Some (HEP, Bio) in production use
 - Some still in testing phase: more effort required to make the Grid their day-to-day workhorse
 - e-Health applications also part of the Grid
 - Some industrial applications:
 - Early deployment mainly in different EC projects

omii europe
open middleware infrastructure institute



HEALTHGRID

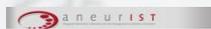


Achieving global e-Science

Session convener



Participant



Grid achievements beyond e-Science

- Grid beyond e-Science?
 - Slower adoption: prefer different environments, tools and have different TCOs
 - Intra grids, internal dedicated clusters , cloud computing
 - e-Business applications
 - Finance, ERP, SMEs and Banking!
 - New economic and business models **NEW**
 - Industrial applications
 - Energy, Automotive, Aerospace, Pharmaceutical industry, Telecom
 - e-Government applications
 - Earth Observation, Civil protection:
 - e.g. The Cyclops project



Examples beyond e-Science

CitiGroup (Citigroup Inc., operating as Citi, is a major American financial services company based in New York City) adopted Grid computing

http://www.americanbanker.com/usb_article.html?id=20080825IXTFW8BS

- Citi chose **Platform Computing's Symphony grid product** to consolidate its computing assets into a single resource pool with increased utilization
- At Citi, since the grid was implemented, individual business units are charged for the processing power they use, creating a **shared services environment**
- Citi is now using near **20,000 CPUs** and there are periods of the day where the **utilization rate is 100 percent**
- Citi is planning of using the cloud in cases their data centers do not suffice (**overflow model** or cooperative data centers)



The screenshot shows a webpage from USBanker. At the top, there is a banner for Adobe LiveCycle ES with the text "Enter Adobe LiveCycle ES" and the Adobe logo. Below the banner is the USBanker logo with the tagline "Beyond Business as Usual". A search bar is visible on the right. The main content area features a navigation menu with links like "Today's Paper", "Tools & Data", "Topics", "AB / TV", "My Account", "Custom Email", "Feedback", and "Cont". The article title is "Taking The Grid Across The Entire Enterprise" by Rebecca Sausner, dated September 2008. The article text discusses Citigroup's Capital Markets division facing a problem with processing power for risk management and pricing applications, and how they implemented a grid computing project to handle peak demand. A sidebar on the left lists various banking services like "Cards/Payments", "Community Banking", "Markets", "Mortgages", "National/Global", "Retail Delivery", "Technology", "Viewpoints", "Washington/Regulatory", "Wealth Management", "Mobile Banker", "News by State", and "Blogs & Scans".

Grid achievements in industry

- IT Industry demonstrated interest in becoming an Grid infrastructure provider and/or user (intra-grids):
 - On-demand infrastructures:
 - Cloud and Elastic computing, pay as you go...
 - Data centers: Data getting more and more attention
 - Service hosting: outsourced integrated services
 - Software as a Service (SaaS)
(e.g. Salesfoce.com services)
 - Virtualisation being exploited in Cloud and Elastic computing
(e.g. Amazon EC2 virtual instances)
- “Pre-commercial procurement”
 - Research-industry collaboration in Europe to achieve new leading-edge products
 - Example: PRACE building a PetaFlop Supercomputing Centre in Europe



The HPC view from ...the clouds!

Not All Clouds are Alike

- Flexible but complex: the Grid

Grids imply dynamic arrival/departure

Grids may include specialized nodes

- Cost-effective but confined: the Cluster

Clusters are typically monocultures: just one type of node

Applications may require tuning to a particular cluster size

- Responsive but repetitive: the Hypervisor

Virtualized servers can be quickly provisioned

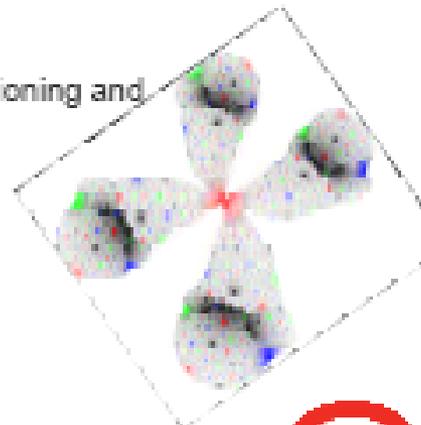
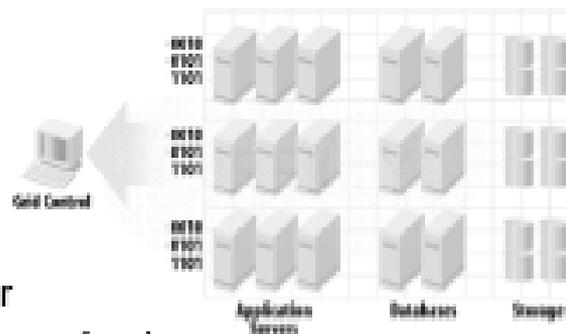
Software stack within the virtual server retains issues of versioning and configuration

Virtual appliances demand monitoring/management

- Enterprise cloud computing implies API leverage

Immediate focus on function; immediate delivery of value

Concerns arise around perception of platform lock-in



salesforce.com



Courtesy Peter Coffee, Salesforce.com

Today and the future: Green IT, pay per CPU/GB virtualisation and/or HPC in every lab?

- Computer and data centers in energy and environmental favorable locations are becoming important
- Elastic computing, Computing on the Cloud, Data Centers and Service Hosting - Software as a Service, are becoming the new emerging solutions for HPC applications
- Many-multi-core and CPU accelerators are promising potential breakthroughs
- Green IT initiatives:
 - The Green Grid (www.thegreengrid.org) consortium (**AMD, APC, Dell, HP, IBM, Intel, Microsoft, Rackable Systems, Sun Microsystems and VMware**)
 - IBM Project Big Green (**a \$1 billion investment to dramatically increase the efficiency of IBM products**) and other IT industry initiatives try to address current HPC limits in energy and environmental impact requirements

Today and the future: Cloud computing and storage on demand

- Cloud Computing: http://en.wikipedia.org/wiki/Cloud_computing
- Amazon, IBM, Google, Microsoft, Sun, Yahoo, major 'Cloud Platform'

poten



ing and

service (S3)

- Sun Network.com www.network.com (1\$/CPU hour, no contract cost)
- IBM Grid solutions (www.ibm.com/grid)
- GoGrid – a division of ServePath company (www.gogrid.com) Beta released (pay-as-you-go and pre-paid plans, server manageability)

EGEE cost estimation (1/2)

Capital Expenditures (CAPEX):

a. Hardware costs: 80.000 CPUs ~ in the order of 120M Euros (80-160M)

Depreciating the infrastructure in 4 years: 30M Euros per year (20M to 40M)

b. Cooling and power installations (supposing existing housing facilities available)

25% of H/W costs: 30M, depreciated over 5 years: 6M Euros

Total: ~ 36M Euros / year (26M-46M)

EGEE cost estimation (2/2)

Operational Expenditures (OPEX):

a. 20 MEuros per year for all EGEE costs (including site administration, operations, middleware etc.)

b. Electricity ~10% of h/w costs: 12M Euros per year (other calculations lead to similar results)

c. Internet connectivity: Supposing no connectivity costs (existing over-provisioned NREN connectivity)

**If other model is used (to construct the service from scratch), then network costs should be taken into account*

Total 32M / year

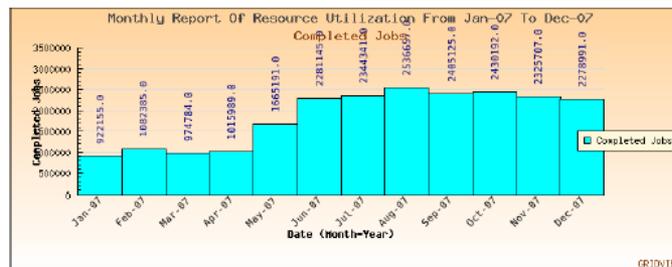
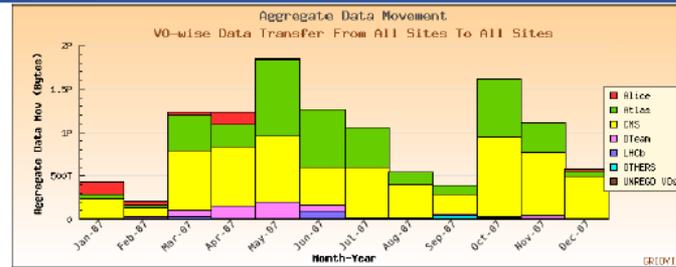
CAPEX+OPEX= 68M per year (58-78M)

EGEE if performed with Amazon EC2 and S3

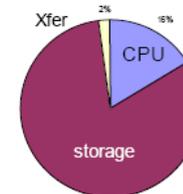
In the order of ~50M Euros, probably more cost effective of EGEE actual cost, depending on the promotion of the EC2/S3 service

EGEE workload in 2007

Data:
25Pb stored
11Pb transferred



CPU:
114 Million hours



Estimated cost if performed with Amazon's EC2 and S3: € 47,486,548

http://gridview.cern.ch/GRIDVIEW/same_index.php

<http://calculator.s3.amazonaws.com/calc5.html?>

EGEE-II INFOS-RI-031688

Bob Jones - EGEE User Forum, 11-14 February 2008 10

Cloud mature enough for big sciences?

Probably not yet, as not designed for them; Does not support complex Scenarios: *“S3 lacks in terms of flexible access control and support for delegation and auditing, and it makes implicit trust assumptions”*



The screenshot shows the header of a magazine titled "symmetrybreaking" with the subtitle "extra dimensions of particle physics" and "A joint Fermilab/SLAC publication". It indicates "VOLUME 05 ISSUE 02" and includes a "VIEW CURRENT ISSUE" link. The main article title is "Are commercial computing clouds ready for high-energy physics?" dated May 23, 2008. The article discusses the use of Amazon's S3 and EC2 services by physicists at Fermilab's Tevatron accelerator. A search bar and RSS feed are also visible on the right side of the page.

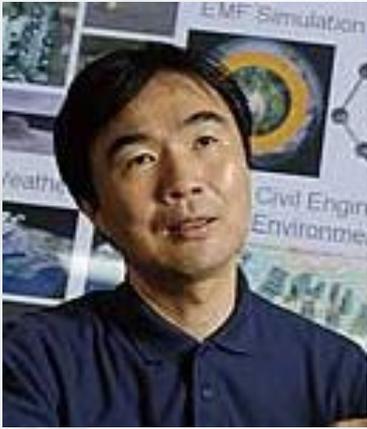
<http://www.symmetrymagazine.org/breaking/2008/05/23/are-commercial-computing-clouds-ready-for-high-energy-physics/>

<http://www.csee.usf.edu/~anda/papers/dadc108-palankar.pdf>

CGW'08, Cracow, Poland

The future:

“To Distribute or Not To Distribute”



- In the late 90s, petaflops were considered very hard and at least 20 years off ...
- while grids were supposed to happen right way
- After 10 years (around now) petaflops are “real close” but there’s still no “global grid”

• What happened:

→ ***It was easier to put together massive clusters than to get people to agree about how to share their resources***

→ ***For tightly coupled HPC applications, tightly coupled machines are still necessary***

→ ***Grids are inherently suited for loosely coupled apps or enabling access to machines and/or data***

- **With Gilder's Law*, bandwidth to the compute resources will promote thin client approach**

* “Bandwidth grows at least three times faster than [computer power](#).” This means that if computer power doubles every eighteen months (per [Moore's Law](#)), then communications power doubles every six months

- **Example: Tsubame machine in Tokyo**

- Prof. Satoshi Matsuoka, TITech
- Keynote at Mardi Gras Conference Baton Rouge, 31 Jan 2008

Multi-core architectures

- Computer CPUs have adopted multi-core architectures with increasing number of cores
 - 2-4 cores in PCs and laptops
 - 8-32 cores in servers, 64-80 cores under development
 - Intel announced a 6 core Xeon



- The trend is driven by many factors:
 - Power consumption, heat dissipation, energy cost, availability of high bandwidth computing at lower cost, ecological impact
- The entire software ecosystem will need to adapt including related applications

Conclusion (1/2)

- We are at a flex point in the evolution of distributed computing
 - **nothing new under the sun...!**
- **Grid** has delivered an affordable HPC infrastructure
 - to scientists all over the world to solve intense computing and storage problems within constrained research budget (and often for social/political reasons) Grid computing
 - leveraging international research networks to deliver an effective and irreplaceable channel for international collaboration
- This has also been effectively used by **industry**
 - to increase the usage of their HPC infrastructure and reduce Total Cost of Ownership (TCO)
- Major issues with wide adoption of Grids have to do with:
 - Cost of operations, complexity, not a solution for all problems (latency, fine grain parallelism are difficult), reliability, security..

Conclusion (2/2)

- **Cloud computing** and **hosted services** are emerging as the next incarnation of distributed computing with some obvious additional advantages but not really designed with scientific applications in mind
- Many changes are happening in the basic underpinning technology (**parallel everywhere!**)
- New boundary constraints and very much **energy** are becoming the limiting factor to the otherwise still valid Moore's Law...
- If we will be able to harness the potential enormous power of parallel computing (not so good story so far) then we might be able to provide better computing solutions for research in energy and eventually better energy solutions for our computing needs!

Thanks

Thanks to the organizers for the kind invitation and to all of you for your attention

Contact me at:
Fabrig @ microsoft.com

