Building an HPC Ecosystem in Europe

Linz

July 5th, 2007

Kimmo Koski

CSC – The Finnish IT Center for Science

Topics

1. **Terminology and definitions**
2. **Towards European petaflop computing**
   - Scientific case for petaflop computing
   - HPC in Europe Taskforce (HET)
   - Partnership for Advanced Computing in Europe (PACE)
3. **Middle layer of performance pyramid**
   - Deisa
   - EGI/EGEE
4. **Policy making**
   - e-IRG
   - ESFRI
5. **New HPC Ecosystem in Europe**
6. **Case example: CSC – The Finnish IT Center for Science**

Terminology and pointers

- **HPC**
  - High Performance Computing
- **HET, http://www.hpcineuropetaskforce.eu/**
  - High Performance Computing in Europe Taskforce, established in June 2006 with a mandate to draft a strategy for European HPC ecosystem
- **Petaflop/s**
  - Performance figure 10^{15} floating point operations (calculations) in second
- **e-IRG, http://www.eirg.eu**
  - e-Infrastructure Reflection Group. e-IRG is supporting the creation of a framework (political, technological and administrative) for the easy and cost-effective shared use of distributed electronic resources across Europe - particularly for grid computing, storage and networking.
- **ESFRI, http://cordis.europa.eu/esfri/**
  - European Strategy Forum on Research Infrastructures. The role of ESFRI is to support a coherent approach to policy-making on research infrastructures in Europe, and to act as an incubator for international negotiations about concrete initiatives. In particular, ESFRI is preparing a European Roadmap for new research infrastructures of pan-European interest.

CSC

Terminology and pointers (cont.)

- **PACE**
  - Partnership for Advanced Computing in Europe
  - EU FP7 project for preparatory phase in building the European petaflop computing centers, based on HET work
- **DEISA,** https://www.deisa.org/
  - Distributed European Infrastructure for Supercomputing Applications. DEISA is a consortium of leading national supercomputing centers that currently deploys and operates a persistent, production quality, distributed supercomputing environment with continental scope.
- **EGEE-II,** http://www.eu-egee.org/
  - Enabling Grid for E-sciencE. The project provides researchers in academia and industry with access to a production level Grid infrastructure, independent of their geographic location.
- **EGI,** http://www.eu-egi.org/
  - An effort to establish a sustainable grid infrastructure in Europe
- **GÉANT2,** http://www.geant2.net/
  - Seventh generation of pan-European research and education network

CSC

Performance Pyramid

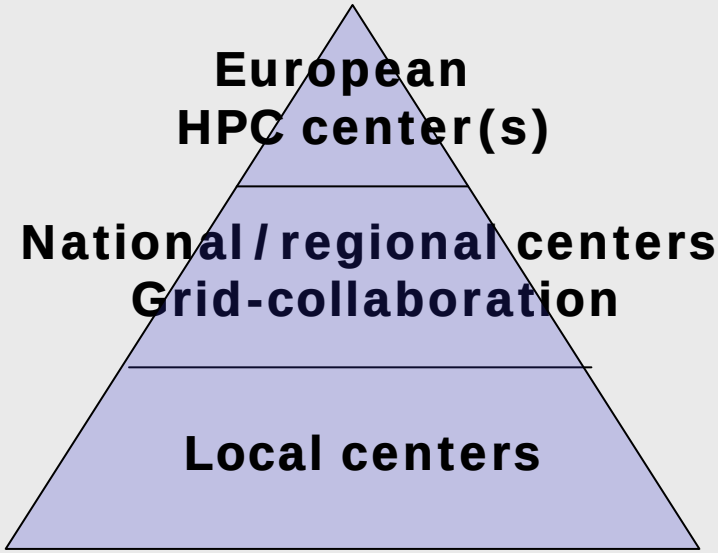

Need to remember about petaflop/s…

- **What do you mean with petaflop/s?**
  1. Theoretical petaflop/s?
  2. LINPACK petaflop/s?
  3. Sustained petaflop/s for a single extremely parallel application?
  4. Sustained petaflop/s for multiple parallel applications?
- **Note that between 1 and 4 there might be several years**
- **Petaflop/s hardware needs petaflop/s applications, which are not easy to program, or not even possible in many cases**
  - Do we even know how to scale over 100000 processors …

Computational science infrastructure

Petaflop computing

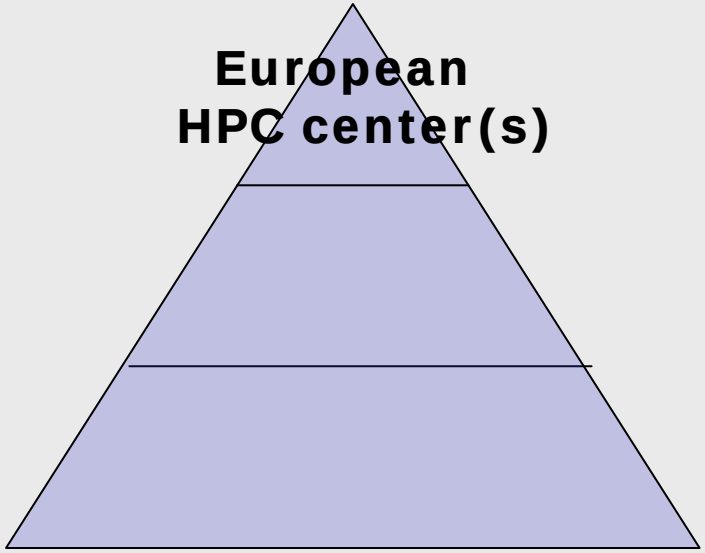

Scientific Case prepared cooperatively by panels of European scientists

November 2005: Barcelona Workshop

Identify key challenge applications

February 2006: Cadarache Workshop

Refine challenges which need a European Tier 0 infrastructure

Panel statements on the necessity & scope of Tier 0 Petascale Resources

- **Multiple projects require capability computing in the order of 1-10 Peta Flops in 2010-2012**
- **Must take into account specific requirements of fields (unlikely a single architecture will fit best all needs)**
- **Infrastructure needs to comprehend user support, grid connectivity, data handling**
- **Infrastructure should be managed similarly to other large scientific infrastructure: requirements, resource allocation, evaluation**

The report through scientific areas

- **Weather, Climatology, Earth Sciences**
- **Astrophysics, HEP and Plasma Physics**
- **Material Science, Chemistry and Nanoscience**
- **Life Sciences**
- **Engineering**

CSC

Weather, Climatology, Earth Sciences Main Scientific objectives (1)

- Quantify uncertainties on the degree of climate warming
- Understand and predict ocean properties and variations at all scales
- Predict weather and flood events with high socio-economic and environmental impact
- In Earth Sciences: mitigation of seismic hazards, treaty verification for nuclear weapons, increased discovery of economically recoverable petroleum and monitoring of waste disposal.

Astrophysics, HEP and Plasma Physics (2)

- **Astrophysics:**
  - Understand the formation of stars and planets, the origin and the evolution of the Universe. Evaluate the data from space experiments such as the European Planck Surveyor satellite.
- **Plasma Physics**
  - The science and technology challenge raised by the construction of the magnetic confinement thermonuclear fusion reactor ITER calls for a major theory and modeling activity.
  - Confining a high temperature plasma poses some of the computationally most challenging problems of nonlinear physics.

Material Science, Chemistry and Nanoscience (3)

- **Simulations of nucleation, growth, self-assembly and polymerization central to the design and performance of many diverse materials**
- **Catalysis is a major challenge in the chemistry of complex materials, with many applications in industrial chemistry.**
- **In biochemistry, a vast number of reactions taking place in the human body are not understood.**
- **Nanoscience**

CSC

Life Sciences (4)

- **Systems Biology: use of sophisticated models to represent the behaviour of cells, tissues, and organs,**
- **The understanding of nucleosome dynamics, will be crucial to understand the mechanism of gene regulation.**
- **Large Scale Protein Dynamics: The study of large conformational changes in proteins.**
- **Representing "in silico" the formation of the different protein complexes associated with a signalling pathway.**

Engineering (5)

Transportation

- Complete rotorcraft simulation
- Green aircraft
- Combustion in engines (automobile, aviation)

Energy

- Power plant design
- Materials

Earth resources (exploration, optimized extraction)

Design of future nano-electronic devices

Natural risks mitigation

- Forest fires
- Seismic, tsunami risks

The Need for European commitment in Capability Computing

In summary, key panel recommendations:

- **« Providing scientists and engineers with access to capability computers of leadership-class must be recognized as an essential strategic priority in Europe. »**
- **« The panel stresses that these resources should be reserved for the most exigent computational tasks of high potential value. »**
- **« This would require putting in place an appropriate process to screen proposals, and to run the resources as permanent research infrastructure. »**

HPC in Europe Taskforce (HET)

- Temporary taskforce shaping the European strategy for petaflop computing
- Founded in June 2006
- Strategy work delivered in January 2007
- Members from 11 European countries
- Chaired by CSC, Finland

- Complete documentation available: http://www.hpcineuropetaskforce.eu/

HET Scope: HPC Ecosystem

- **The upper layers of the pyramid**
  - HPC centers / services
  - European projects (HPC/Grid, networking, …)
- **Activities which enable efficient usage of upper layers**
  - Inclusion of national HPC infrastructures
  - Software development and scalability issues
  - Competence development
- **Interoperability between the layers**

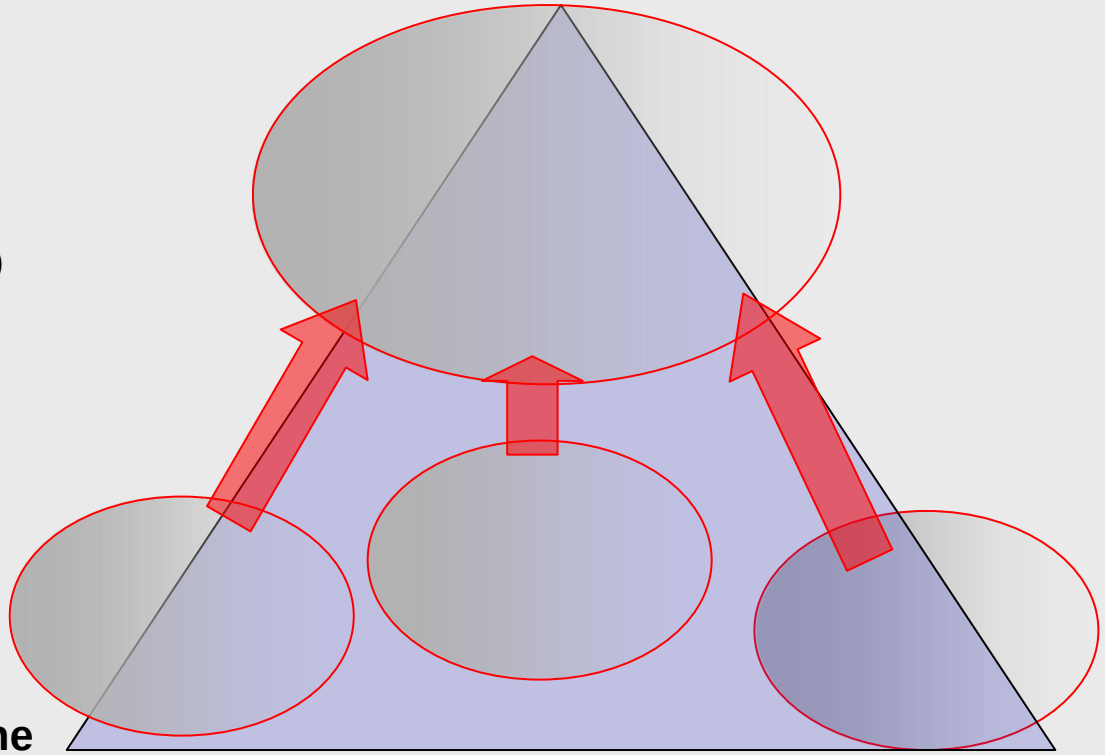

CSC

HET outcome

- **Entry in ESFRI list for petaflop computing**
- **Papers:**
  - Scientific case for European HPC (most work done by previous HPCEUR project)
  - Proposal for funding models
  - Proposal for peer review process
  - Views for HPC Ecosystem
  - Summary paper with recommendations
- **Good team spirit with a common approach**
- **Basis for practical implementation**
  - Consortium for ESFRI Preparatory phase
  - Memorandum of Understanding for European Tier 0 HPC service

HET Recommendations

- **Recommendation for the development and operation of a “top end” infrastructure**
  - HET recommends establishment of a small number of European HPC facilities to provide extreme computing power – exceeding petaflop capability – for the most demanding computational tasks.
- **Recommendation for developing the full European ecosystem**
  - HET recommends increased emphasis on the development of the full HPC ecosystem, including the local infrastructure, national and regional facilities, top-level European computing capabilities and the interoperability of their services.
- **Recommendation to enable petascale computing by supporting the development of novel software architectures**
  - HET recommends starting a range of activities aimed at addressing the key issues in building software that allows exploiting the performance potential of petascale machines in a coherent, efficient, scalable and sustainable manner.

HET Recommendations (cont.)

- **Recommendation to implement an efficient and highly reliable infrastructure for storing large amounts of data**
  - HET recommends increasing emphasis on permanent and persistent data repositories as a part of the HPC ecosystem.
- **Recommendation to support competence development in computational science through extensive training and education activities**
  - HET recommends supporting extensive training and education activities which focus on enabling more efficient and higher quality use of the top-end facilities in the long run
- **Recommendation to raise the visibility of HPC**
  - HET recommends strong activities to increase the visibility and improve the publicity of computational science in order to highlight the strategic impact and need of numerical simulation for most areas of science and engineering

HET Recommendations (cont.)

- **Recommendation to boost collaboration**
  - HET recommends support for collaborative actions with a target to link the major players in HPC Ecosystem – existing grid and HPC projects, national and regional centers, main computational research groups, funding organizations and potential new planned FP7 efforts – in order to maximize synergy for actions.
- **Recommendation to support European industry**
  - HET recommends support for collaborative efforts with European industrial HPC users and European HPC industry at large – from hardware and software R&D to product design and manufacturing.

PACE Partnership for Advanced Computing in Europe

Signing of the MoU in Berlin 17.4.2007

CSC

PACE

- **Target to build the basis for European petaflop/s centers, responding to the ESFRI Roadmap item 'European HPC Service'**
- **FP7 Project proposal for preparatory phase, call deadline 2.5.2007**
- **Memorandum of Understanding: 14 countries signed and more to come**
- **PACE consortium partners (14 countries)**
  - Austria, Finland, France, Germany, Greece, Italy, Norway, Poland, Portugal, Spain, Sweden, Switzerland, The Netherlands, United Kingdom

What is going to happen with PACE?

- **Project proposal deadline was met and proposal is in**
- **If successful, project start assumed November-December 2007**
- **Two years, 10+10 MEUR volume**
- **Prototypes for petaflop computing during 2008-2009**
- **Construction of centers start 2009 aiming at first center operational in 2009-2010**
- **Open issues to be solved during the preparatory phase:**
  - Which companies to prototype and where to place them?
  - Who will host the petaflop centers?
  - Who will pay for construction?
  - Who can use the resources and under which conditions?
  - How to link with other projects, for example DEISA?

Middle layer

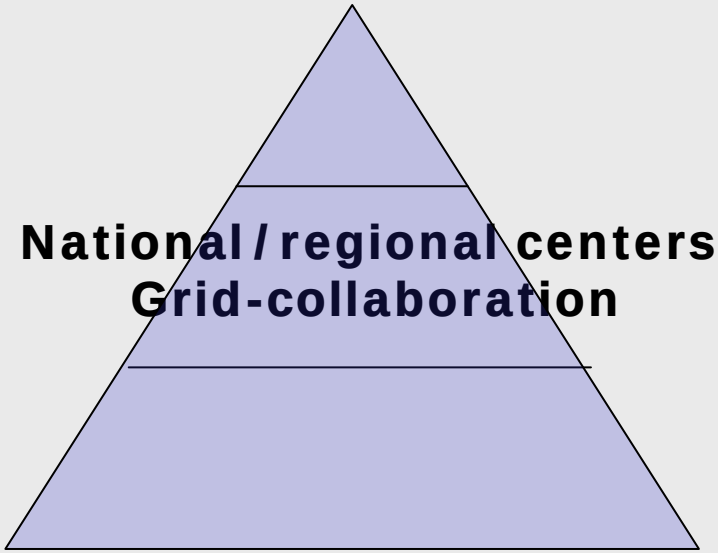

Distributed European Infrastructure for Supercomputing Applications

DEISA –Distributed European Infrastructure for Supercomputing Applications

- A consortium of leading national supercomputing centres deploying and operating a persistent, production quality, distributed supercomputing environment with continental scope
- Grid-enabled FP6 funded Research Infrastructure
- A 4-year-project started on May 2004
- Total budget is *37,1 M€* (incl. DEISA and eDEISA contracts), EU funding - *20.9 M€*

DEISA Supercomputing Grid

Distributed European Infrastructure for Supercomputing Applications

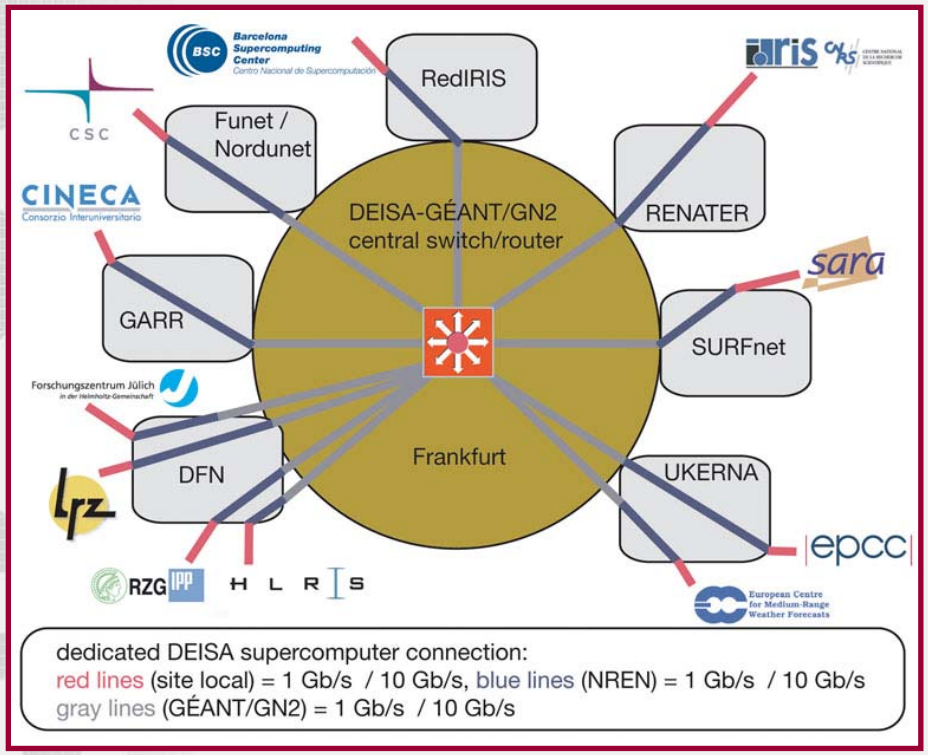

DEISA Network Infrastructure

- DEISA is an European Supercomputing Service built on the top of the existing national services.
- This service is based on the deployment and operation of a persistent, production quality, distributed supercomputing environment with continental scope.
- The integration of national facilities and services, together with innovative operational models, is expected to add substantial value to existing infrastructures.
- Main focus is High Performance Computing (HPC).

Participants: 11 partners, technology and industrial partners

DEISA: May 2004 - April 2008, incl.
eDEISA from May 2006 to April 2008

DEISA Supercomputing Grid

Distributed European Infrastructure for Supercomputing Applications

- Over 200 Tf in 2007
- Systems interconnected with dedicated 1Gb/s network – currently upgraded to 10 Gb/s – provided by GEANT and NRENs

IBM AIX super-cluster	IDRIS, FZJ, RZG, CINECA,ECMW, HPCx	68.6 Tf
IBM PowerPC Linux system	BSC	40 Tf
CRAY XT4	CSC	10.5Tf
Linux cluster 2.7 Tf moving to SGI ALTIX system	LRZ	33 Tf, 70 Tf in 2007
SGI ALTIX Linux cluster	SARA	2.2 Tf
NEC SX8 vector system	HLRS	12,7 Tf

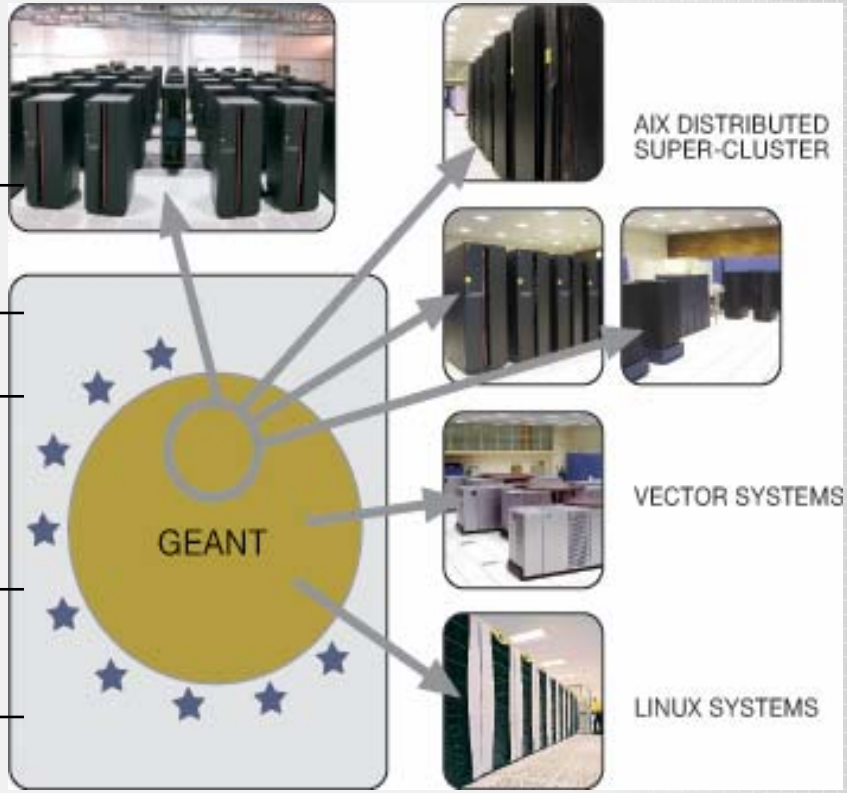

DEISA Extreme Computing Initiative

Distributed European Infrastructure for Supercomputing Applications

- Enabling leading computational science
- Identification, deployment and operation of a number of «flagship» applications requiring the infrastructure services, in selected areas of science and technology.
- European Call for proposals in May-June every year. Applications are selected on the basis of scientific excellence, innovation potential and relevance criteria, with the collaboration of the HPC national evaluation committees.
- There are 23 projects in operation in 2007, and 29 projects operated in 2006
- Supported by the Applications Task Force, whose objective is to enable and deploy the Extreme Computing applications:
  - Hyperscaling of huge parallel applications, data oriented applications
  - Workflows and coupled applications
  - Production of an European Benchmark Suite for HPC systems (in collaboration with the HPC-EUR initiative, to be used in future procurements of European supercomputers).

EGEE-II Applications Overview

Enabling Grids for E-sciencE

- **>200 VOs from several scientific domains**
  - Astronomy & Astrophysics
  - Civil Protection
  - Computational Chemistry
  - Comp. Fluid Dynamics
  - Computer Science/Tools
  - Condensed Matter Physics
  - Earth Sciences
  - Fusion
  - High Energy Physics
  - Life Sciences
- **Further applications under evaluation**

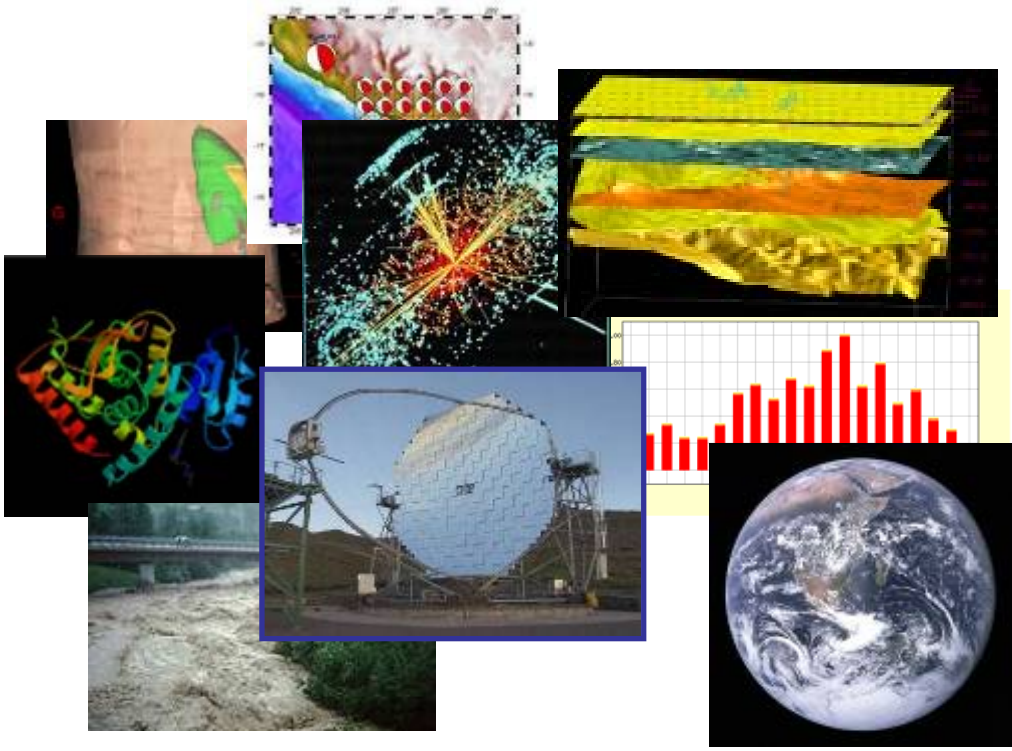

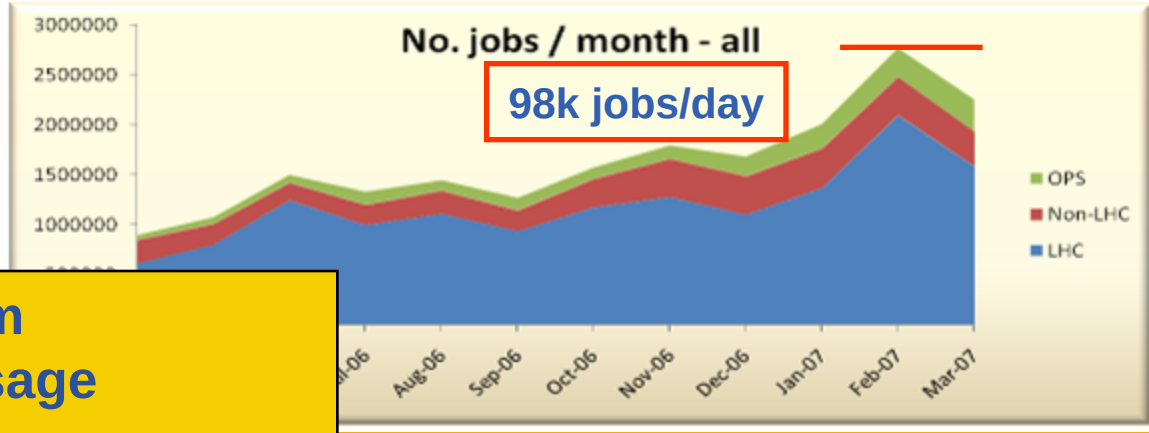

Applications have moved from testing to routine and daily usage

~80-90% efficiency

European Grid Initiative

Goals:

- Ensure the long-term sustainability of the European e-infrastructure
- Coordinate the integration and interaction between National Grid Infrastructures
- Operate the European level of the production Grid infrastructure for a wide range of scientific disciplines to link National Grid Infrastructures

Evolution

National

Sustainable European Grid

Global

Testbeds → Routine Usage → Utility Service

EGI Design Study (EGI_DS)

- Project Proposal, submitted to the European Commission for funding within **FP7-INFRASTRUCTURES-2007-1**, 1.2.1 Design Studies (May 2, 2007)

Participant no.	Participant organisation name	Short name	Country
1 (Coordinator)	**Institut für Graphische und Parallele Datenverarbeitung der Johannes Kepler Universität Linz**	**GUP**	**A**
2	**Greek Research and Technology Network – GRNET S.A.**	**GRNET**	**GR**
3	**Istituto Nazionale di Fisica Nucleare**	**INFN**	**I**
4	**CSC – Scientific Computing Ltd.**	**CSC**	**FI**
5	**CESNET, z.s.p.o.**	**CESNET**	**CZ**
6	**European Organization for Nuclear Research**	**CERN**	**CH**
7	**Verein zur Förderung eines Deutschen Foschungsnet-zen – DFN-Verein**	**DFN**	**D**
8	**Science & Technology Facilities Council**	**STFC**	**UK**
9	**Centre National de la Recherche Scientifique**	**CNRS**	**F**

Policy work

Supporting the creation of the ER(I)A

- **“Capacities” implementing the seventh Framework Programme (2007-2013)**
  - …supporting analysis of emerging needs, at supporting the work of ESFRI and e-IRG, at the effective implementation…

- **e-IRG: e-Infrastructure Reflection Group http://www.e-irg.org/**

- **ESFRI: European Strategy Forum on Research Infrastructures http://www.cordis.lu/esfri/**

e-IRG Mission & Objectives

Mission **(drafted in Rome on December 10, 2003):**

- **to support on the** political**,** advisory **and** monitoring **level, the creation of a policy and administrative framework for the easy and cost-effective shared use of electronic resources in Europe (focusing on Grid-computing, data storage, and networking resources) across technological, administrative and national domains.**

Objectives**:**

- **to support the creation of a framework for easy and cost-effective shared use of distributed electronic resources across Europe**
- **to define and recommend** best practices **for each of the (pan-) European grid efforts.**

➔ **e-IRG White Paper & Roadmap: http://www.e-irg.eu/publ/**

ESFRI

- **Strategy Forum with a consulting role to EU**
- **Wide representation of scientists in various disciplines**
- **Roadmap process for major new European research infrastructures (range of 10-1000 MEUR for an infrastructure)**
- **Roadmap published in 2006**
  - 35 projects labeled mature
  - One of the projects European HPC Service
- **Preparatory projects for each project**
  - 1-4 years
  - Deadline for project call May 2nd, 2007

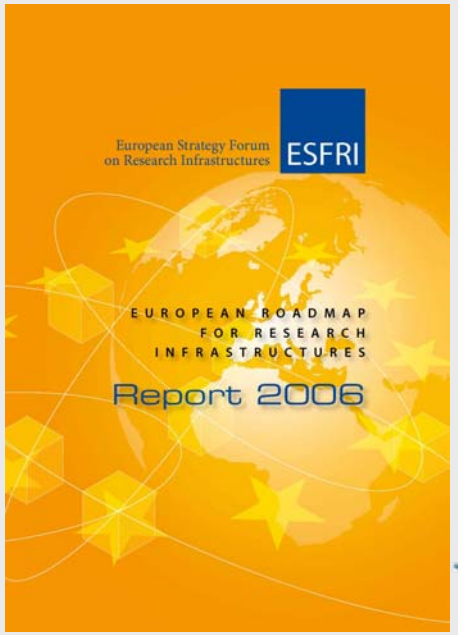

CSC

New HPC Ecosystem is being built…

European HPC after FP6

- **Multiple Grid projects with varying results – learning for collaboration**
- **Early experiences about interoperability between national HPC centers**
- **Communities start to form, in various levels**
- **Research community more active in computational science domain**
- **European Union targets in creating sustainable infrastructures**

- **Petaflop computing raised to European agenda, scientific case for high-end computing available**

Targets for European HPC collaboration 2007 onwards

- **Continuation of existing grid projects (DEISA, EGEE …) and development in GEANT2 network infrastructure**
- **Building European petaflop computing services integrated in the full HPC ecosystem according to the performance pyramid model (PACE)**
- **Maximal synergy between PACE and DEISA (integration after some time?)**
- **Interoperability between PACE and EGI/EGEE**
- **Building up research infrastructure services for ESFRI roadmap**
- **Target to establish an active European community for HPC: infrastructure, resource sharing, communication and collaboration over country borders**

CSC

New market for European HPC

- **35 ESFRI list new research infrastructure projects, most of them starting a preparatory phase project late 2007**
  - 1-4 years
  - 1-7 MEUR * 2 (petaflop computing 10 MEUR * 2)
- **Successful new research infrastructures start construction 2009-2011**
  - 10-1000 MEUR per infrastructure
- **Existing infrastructures are also growing**
- **Results:**
  - Growing RI market, considerably rising funding volume
  - Need for horizontal activities (computing, data, networks, computational methods and scalability, application development,…)
  - Real danger to build disciplinary silos instead of searching IT synergy

Advice for HPC vendors: Europe wants to develop HPC Ecosystem in Europe

- What can you do in Europe?
  - **Manufacturing**
  - **Research**
  - **Software development**
  - **Integration work**
  - **Training**
  - **Other**
- Pre-commercial procurement used increasingly
- Prototyping is part of the petaflop/s project, although the prototypes for tomorrow's petaflop systems are usually today's production systems
  - **Innovation is more in scalable software than hardware**

How to ensure sustainability in the European HPC ecosystem?

- **Integration and compatibility with national infrastructures**
- **Building on the current HPC-related work, such as a successful European research infrastructures projects**
- **An appropriate level of communication, experience and competence sharing between the national centres**
- **Co-ordination by an European group with a strong mandate**
- **Financial support from the national funding bodies and European Union**
- **Increased visibility for computational science**
- **Using the experienced HPC user groups to help the emerging new HPC communities**
- **Training and competence development in HPC/Grid arena**
- **Promoting application development work and optimisation skills**

CSC

Case example: CSC Finland

Work in Finland

- **CSC as a national IT center for science (centralized model, 140 persons, turnover over 16 MEUR plus investments in 2006, fully owned by Finnish state, non-profit limited company)**
- **CSC is the national contact point for the National Grid Initiative (NGI)**
- **Examples of high level national activities**
  - Grid strategy group 2005-2006
  - Infrastructure strategy 2006, working group for implementation 2007
  - Proposal for Competence Clusters 2006 (Forest and ICT clusters about to start)
  - Strategy for computational science 2007
  - Information Society Program until March 2007

CSC Fact Sheet

Operated on a non-profit principle

All shares to the Ministry of Education of Finland in 1997

Reorganized as a limited company, CSC-Scientific Computing Ltd. in 1993

Founded in 1971 as a technical support unit for Univac 1108

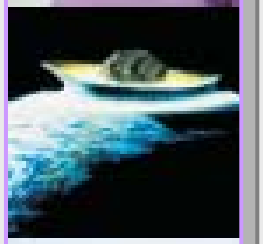

CSC Turnover in 2006 15.6 M€, 150 employees

Since March 2005, facilities in Keilaniemi, Espoo

First supercomputer Cray X-MP/EA 416 in 1989

Funet started in 1984

CSC

CSC

MISSION:

CSC, as a part of the Finnish national research structure, develops and offers high quality information technology services

VISION 2012:

CSC – a leading center of excellence in information technology for science in the European research area

CSC Fields of Services

UNIVERSITIES

POLYTECHNICS

RESEARCH INSTITUTES

COMPANIES

FUNET SERVICES

COMPUTATIONAL SERVICES

DATA SERVICES

APPLICATION SERVICES

INFORMATION MANAGEMENT

TRAINING AND PUBLICATIONS

CSC's supercomputers

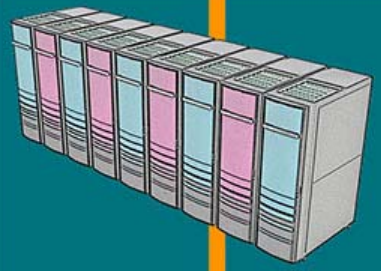

Cray Hood
6736 processor cores
6,7 TB memory
70 Teraflops peak computing power
70 TB disk space

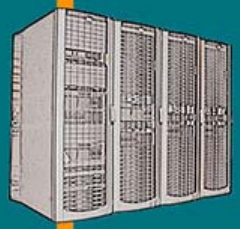

HP ProLiant cluster
2048 processor cores
4 TB memory
10 Teraflops peak computing power
98 TB disk space

Sun Fire 25K
192 processor cores
384 GB memory

Funet

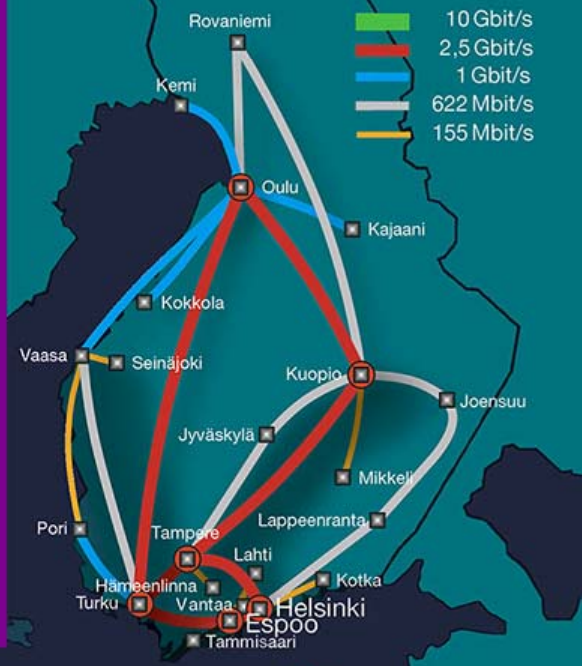

Example of a successful project for national Grid infrastructure: Material Sciences Grid (M-grid)

- **Goal: Computing capacity mainly for the needs of physics and chemistry researchers**
- **Joint project between seven Finnish universities, Helsinki Institute of Physics and CSC**
  - Jointly funded by the Academy of Finland and the participating universities, in production since 2004
  - Coordinated by CSC
- **First large initiative to put grid middleware into production use in Finland**
- **Platform: Linux based PC clusters**
  - Total theoretical computing power 5 TFlops
  - Uses NorduGrid ARC middleware

CSC

For more information:

http://www.csc.fi/

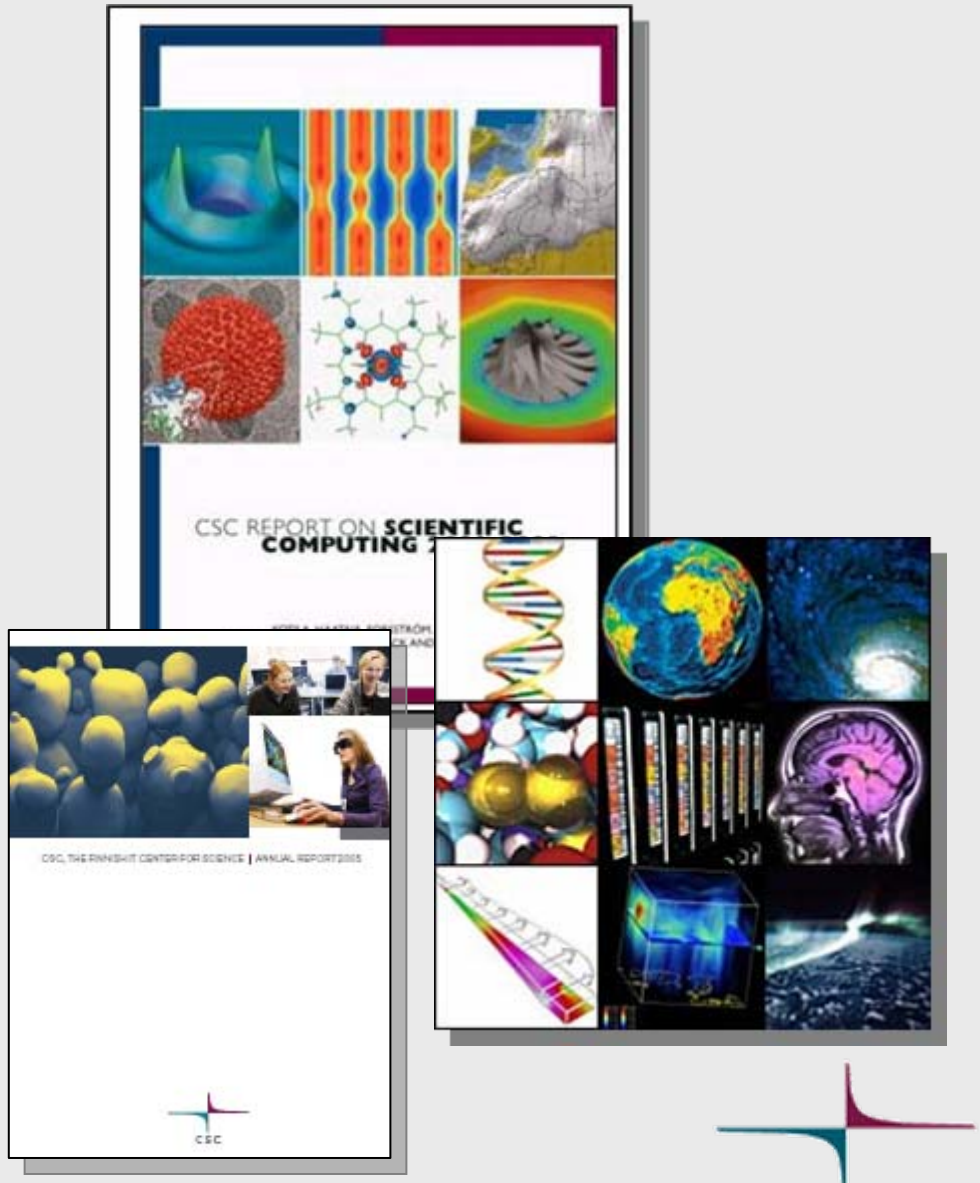

CSC

Conclusions

- **Petaflop computing will be implemented in Europe and integrated as one part of HPC ecosystem**
- **Borders between the layers in performance pyramid will decrease**
- **Competence development in computational science and scalable software development become increasingly important**
- **ESFRI Roadmap will boost the HPC development in Europe**
- **Horizontal activities (EGI, PACE, …) are desperately needed in order to promote collaboration and avoid the creation of disciplinary IT silos**