

DEISA - Towards a European HPC Infrastructure for Computational Science

www.deisa.eu

Andreas Schott
DEISA and RZG (Max Planck Society)
Andreas.Schott@rzg.mpg.de

SC'08 Austin

November 19th, 2008



RI-222919



European Strategy Forum on Research Infrastructures


ESFRI Report 2006, p. 65

European High-Performance Computing Service

A European strategic approach to high-performance computing, concentrating the resources in a limited number of world top-tier centres in an overall infrastructure connected with associated national, regional and local centres, forming a scientific computing network to utilise the top-level machines.



new "petaflop" supercomputers



15



European Commission
Information Society and Media

Mario Campolargo, OGF23, June 2008

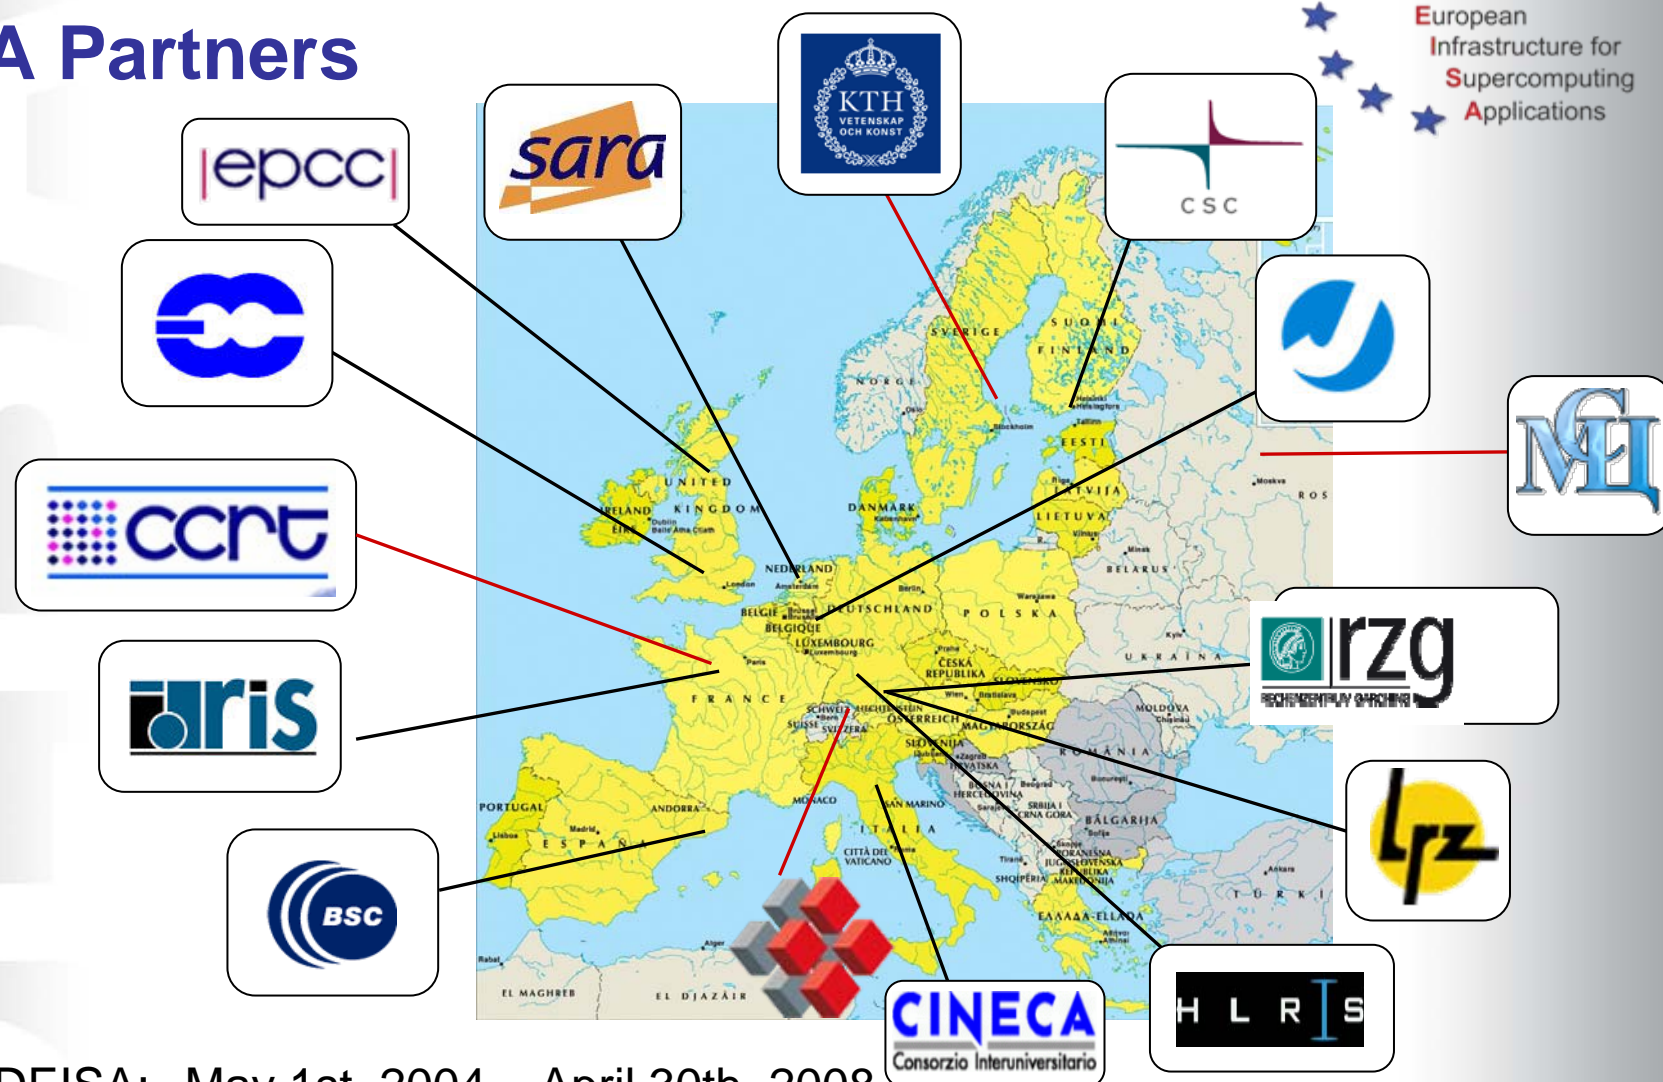
DEISA Partners



BSC	Barcelona Supercomputing Centre	Spain
CINECA	Consortio Interuniversitario per il Calcolo Automatico	Italy
CSC	Finnish Information Technology Centre for Science	Finland
EPCC	University of Edinburgh and CCLRC	UK
ECMWF	European Centre for Medium-Range Weather Forecast	UK (int)
FZJ	Research Centre Juelich	Germany
HLRS	High Performance Computing Centre Stuttgart	Germany
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique - CNRS	France
LRZ	Leibniz Rechenzentrum Munich	Germany
RZG	Rechenzentrum Garching of the Max Planck Society	Germany
SARA	Dutch National High Performance Computing	Netherlands
CEA-CCRT	Centre de Calcul Recherche et Technologie, CEA	France
KTH	Kungliga Tekniska Högskolan	Sweden
CSCS	Swiss National Supercomputing Centre	Switzerland
JSCC	Joint Supercomputer Center of the Russian Academy of Sciences	Russia

DEISA Partners

Distributed
European
Infrastructure for
Supercomputing
Applications




DEISA: May 1st, 2004 – April 30th, 2008

DEISA2: May 1st, 2008 – April 30th, 2011



DEISA 2008

Operating the European HPC Infrastructure



>1 PetaFlop/s
Aggregated peak performance


Most powerful European supercomputers for most challenging projects

Top-level Europe-wide application enabling

Grand Challenge projects performed on a regular basis

epcc

sara



ICH
SZENTRUM

DEISA Core Infrastructure and Services

Dedicated High Speed Network

Common AAA

- Single sign on
- Accounting/budgeting

Global Data Management

- High performance remote I/O and data sharing with global file systems
- High performance transfers of large data sets

User Operational Infrastructure

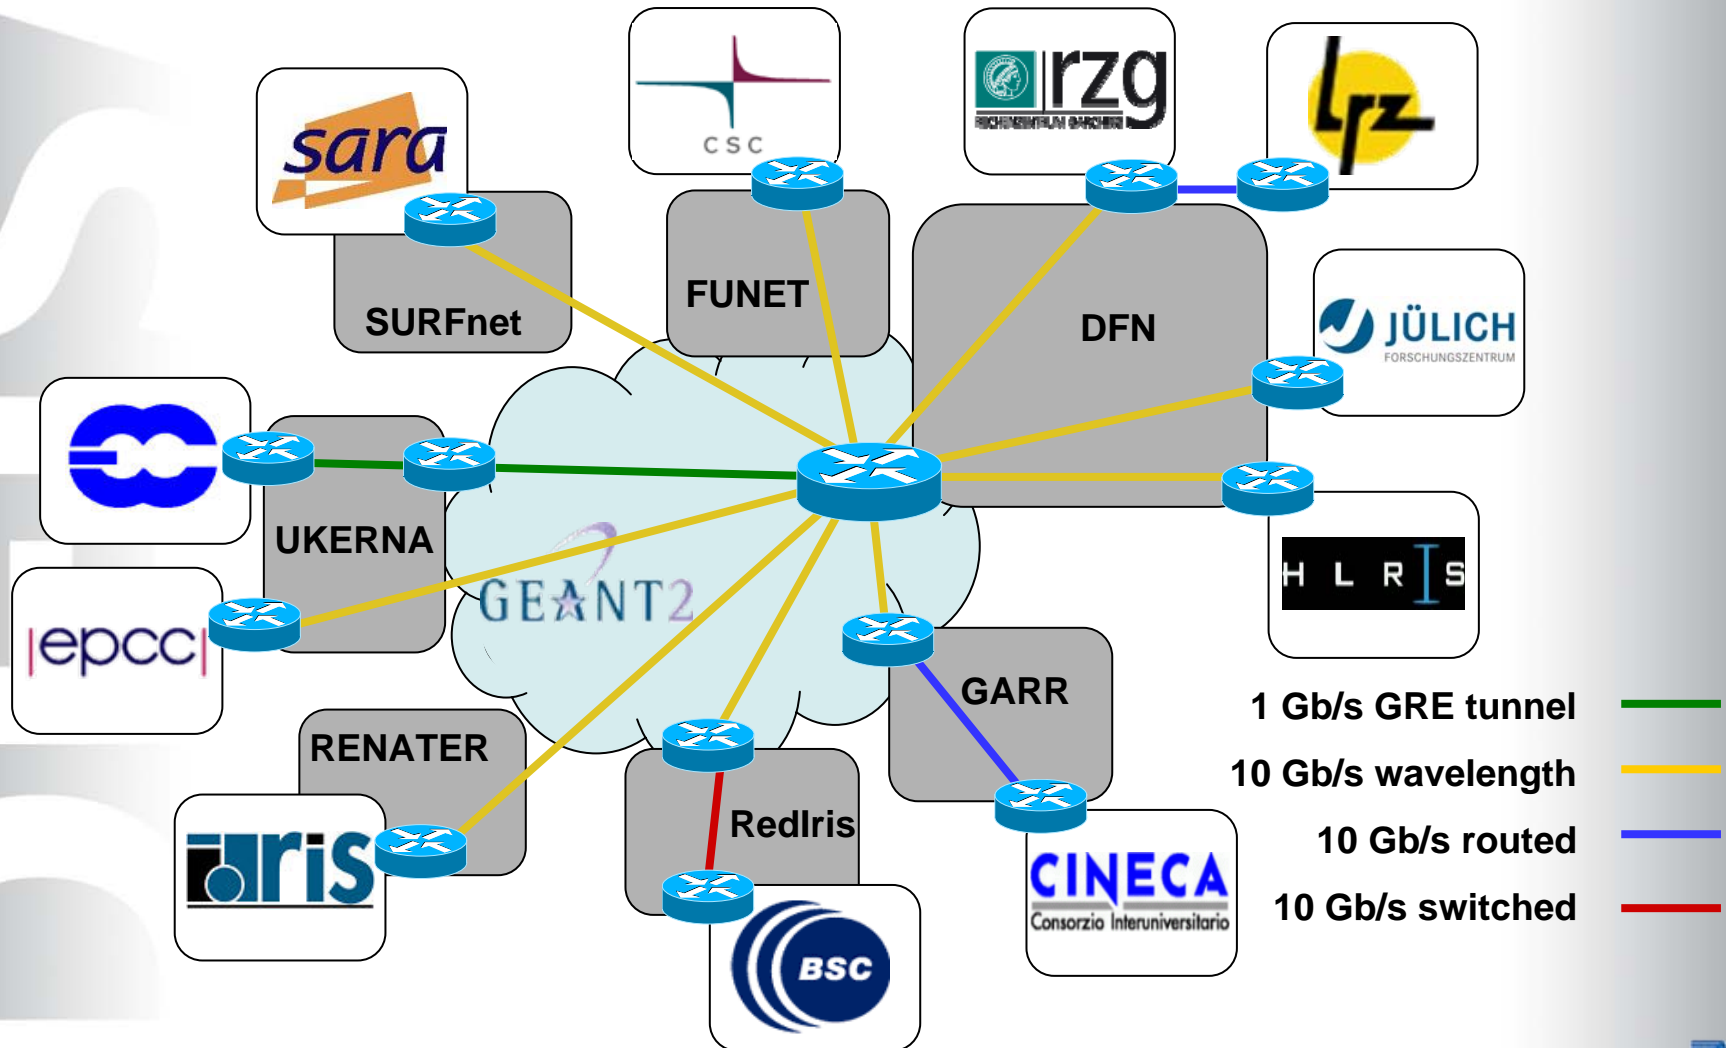
- Distributed Common Production Environment (DCPE)
- Job management service
- Common user support and help desk

System Operational Infrastructure

- Common monitoring and information systems
- Common system operation

Global Application Support

DEISA dedicated high speed network



DEISA Global File System (based on MC-GPFS)



IBM P6



IBM P6 (& BlueGene/P)



IBM P6 (& BlueGene/P)



NEC SX8



AIX LL



AIX LL-MC




AIX LL-MC



Super-UX NQS II



Cray XT4



UNICOS/lc PBS Pro



SGI ALTIX



LINUX PBS Pro



Cray XT4 & XT5



UNICOS/lc PBS Pro



Consorzio Interuniversitario

IBM P5



LINUX Maui/Slurm



IBM PPC




LINUX LL

IBM P6

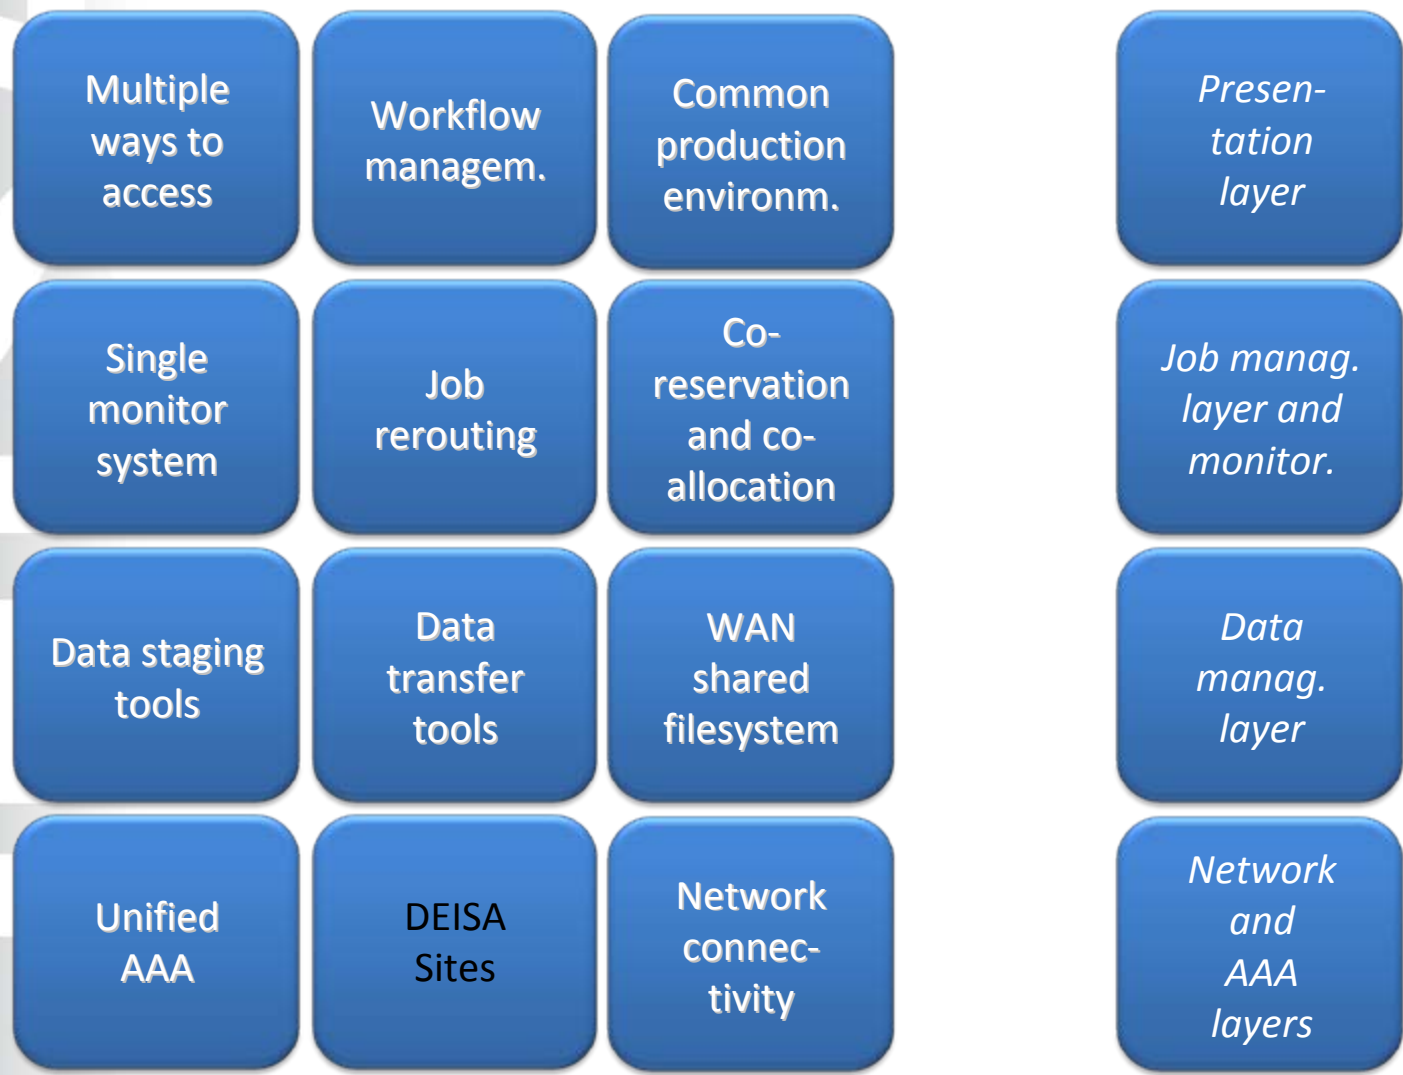


AIX LL-MC

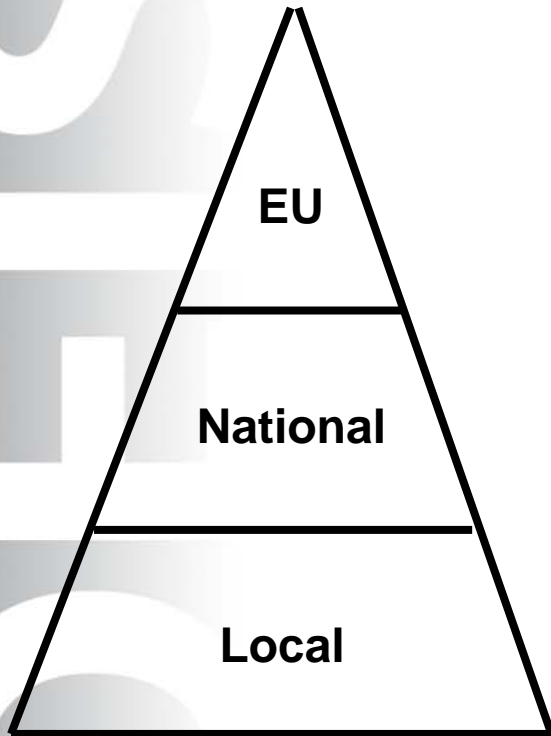
IBM P6 & BlueGene/P




DEISA Software Layers



Supercomputer Hardware Performance Pyramid



Supercomputer Application Enabling Requirements Pyramid



**Capability computing
will always need
expert support for
application enabling
and optimizations**

**The more resource
demanding one single
problem is, the higher
are generally the
requirements for
application enabling
including enhancing
scalability**

DEISA Organizational Structure

WP1 – Management

WP2 – Dissemination, External Relations, Training

WP3 – Operations

WP4 – Technologies

WP5 – Applications Enabling

WP6 – User Environment and Support

WP7 – Extreme Computing (DECI) and Benchmark Suite

WP8 – Integrated DEISA Development Environment

WP9 – Enhancing Scalability

Evolution of Supercomputing Resources

2004 DEISA partners' compute resources at DEISA project start:
~ 30 TF aggregated peak performance

2008 DEISA partners' resources at DEISA2 project start:
Over 1 PF aggregated peak performance on state-of-the art
supercomputers

Cray XT4 and XT5, Linux

IBM Power5, Power6, AIX / Linux

IBM BlueGene/P, Linux (frontend)

IBM PowerPC, Linux (MareNostrum)

SGI ALTIX 4700 (Itanium2 Montecito), Linux

NEC SX8 vector system, Super UX

Systems interconnected with dedicated 10Gb/s network links
provided by GEANT2 and NRENs

Fixed fraction of resources dedicated to DEISA usage

DEISA Extreme Computing Initiative (DECI)



DECI launched in early 2005 to enhance DEISA's impact on science and technology

Identification, enabling, deploying and operation of "flagship" applications in selected areas of science and technology

Complex, demanding, innovative simulations requiring the exceptional capabilities of DEISA

Multi-national proposals especially encourage

Proposals reviewed by national evaluation committees

Projects chosen on the basis of innovation potential, scientific excellence, relevance criteria, and national priorities

Most powerful HPC architectures in Europe for the most challenging projects

Most appropriate supercomputer architecture selected for each project

Mitigation of the rapid performance decay of a single national supercomputer within its short lifetime cycle of typically about 5 years, as implied by Moore's law

DEISA Extreme Computing Initiative



Involvements in projects from DECI calls 2005, 2006, 2007:

157 research institutes and universities

from

15 European countries

Austria
Italy
Russia

Finland
Netherlands
Spain

France
Poland
Sweden

Germany
Portugal
Switzerland

Hungary
Romania
UK

with collaborators from

four other continents

North America, South America, Asia, Australia

DEISA Extreme Computing Initiative

Calls for Proposals for challenging supercomputing projects from all areas of Science

DECI call 2005

51 proposals, 12 European countries involved, co-investigator from US)
30 mio cpu-h requested
29 proposals accepted, 12 mio cpu-h awarded (normalized to IBM P4+)

DECI call 2006


41 proposals, 12 European countries involved
co-investigators from N + S America, Asia (US, CA, AR, ISRAEL)
28 mio cpu-h requested
23 proposals accepted, 12 mio cpu-h awarded (normalized to IBM P4+)

DECI call 2007

63 proposals, 14 European countries involved, co-investigators from
N + S America, Asia, Australia (US, CA, BR, AR, ISRAEL, AUS)
70 mio cpu-h requested
45 proposals accepted, ~30 mio cpu-h awarded (normalized to IBM P4+)

DECI call 2008

66 proposals, 15 European countries involved, co-investigators from
N + S America, Asia, Australia
134 mio cpu-h requested (normalized to IBM P4+)
Evaluation in progress




DECI Project POLYRES

Cover Story of Nature - May 24, 2007

Curvy membranes make proteins attractive


For almost two decades, physicists have been on the track of membrane mediated interactions. Simulations in DEISA have now revealed that curvy membranes make proteins attractive.



- a) proteins (red) adhere on a membrane (blue/yellow) and locally bend it;
- b) this triggers a growing invagination.
- c) cross-section through an almost complete vesicle


B. J. Reynwar et al.: **Aggregation and vesiculation of membrane proteins by curvature mediated interactions**, NATURE Vol 447|24 May 2007| doi:10.1038/nature05840

Achievements and Scientific Impact



Brochures can be downloaded from <http://www.deisa.eu/publications/results>

Evolution of User Categories in DEISA



Tier0 / Tier1 Centers

Are there implications for the services?

Main difference between T0 and T1 centers: policy and usage models !

**T1 centers can evolve to T0 for strategic/political reasons
T0 machines automatically degrade to T1 level by aging**

T0 Centers

Leadership-class European systems in competition to the leading systems worldwide, cyclically renewed

Governance structure to be provided by European organization (PRACE)

T1 Centers

Leading national Centers, cyclically renewed, optionally surpassing the performance of older T0 machines

National Governance structure

Services have to be the same in T0/T1

Because of the change of the status of the systems, over time

For user transparency of the different systems

(Only visible: Some services could have different flavors for T0 and T1)

Summary



Evolution of this European infrastructure towards a robust and persistent European HPC ecosystem

Enhancing the existing services, by deploying new services including support for European Virtual Communities, and by cooperating and collaborating with new European initiatives, especially PRACE

DEISA2 as the vector for the integration of Tier-0 and Tier-1 systems in Europe

To provide a lean and reliable turnkey operational solution for a persistent European HPC infrastructure

Bridging worldwide HPC projects: To facilitate the support of international science communities with computational needs traversing existing political boundaries

DEISA - Towards a European HPC Infrastructure for Computational Science

www.deisa.eu

Andreas Schott
DEISA and RZG (Max Planck Society)
Andreas.Schott@rzg.mpg.de

SC'08 Austin

November 19th, 2008



RI-222919

