

DEISA

DEISA2 engaging for a European HPC Ecosystem

www.deisa.eu

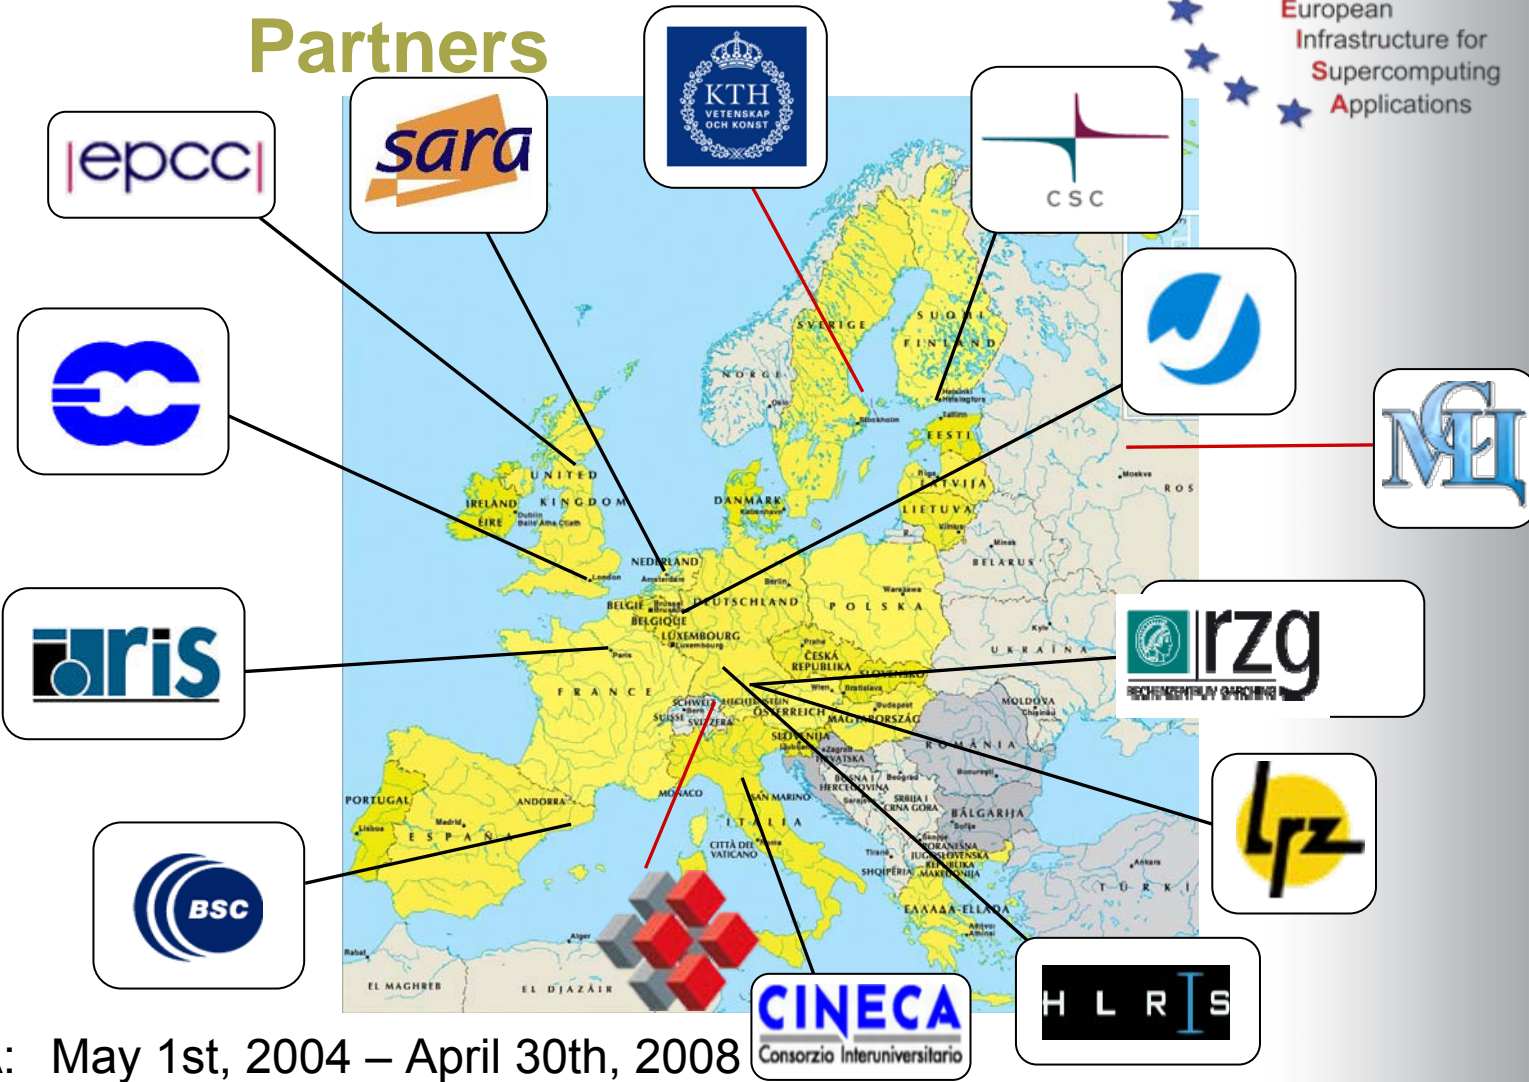
Hermann Lederer
DEISA and RZG, lederer@rzg.mpg.de
SciDAC 2008
July 13-17, 2008



RI-222919



Partners



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

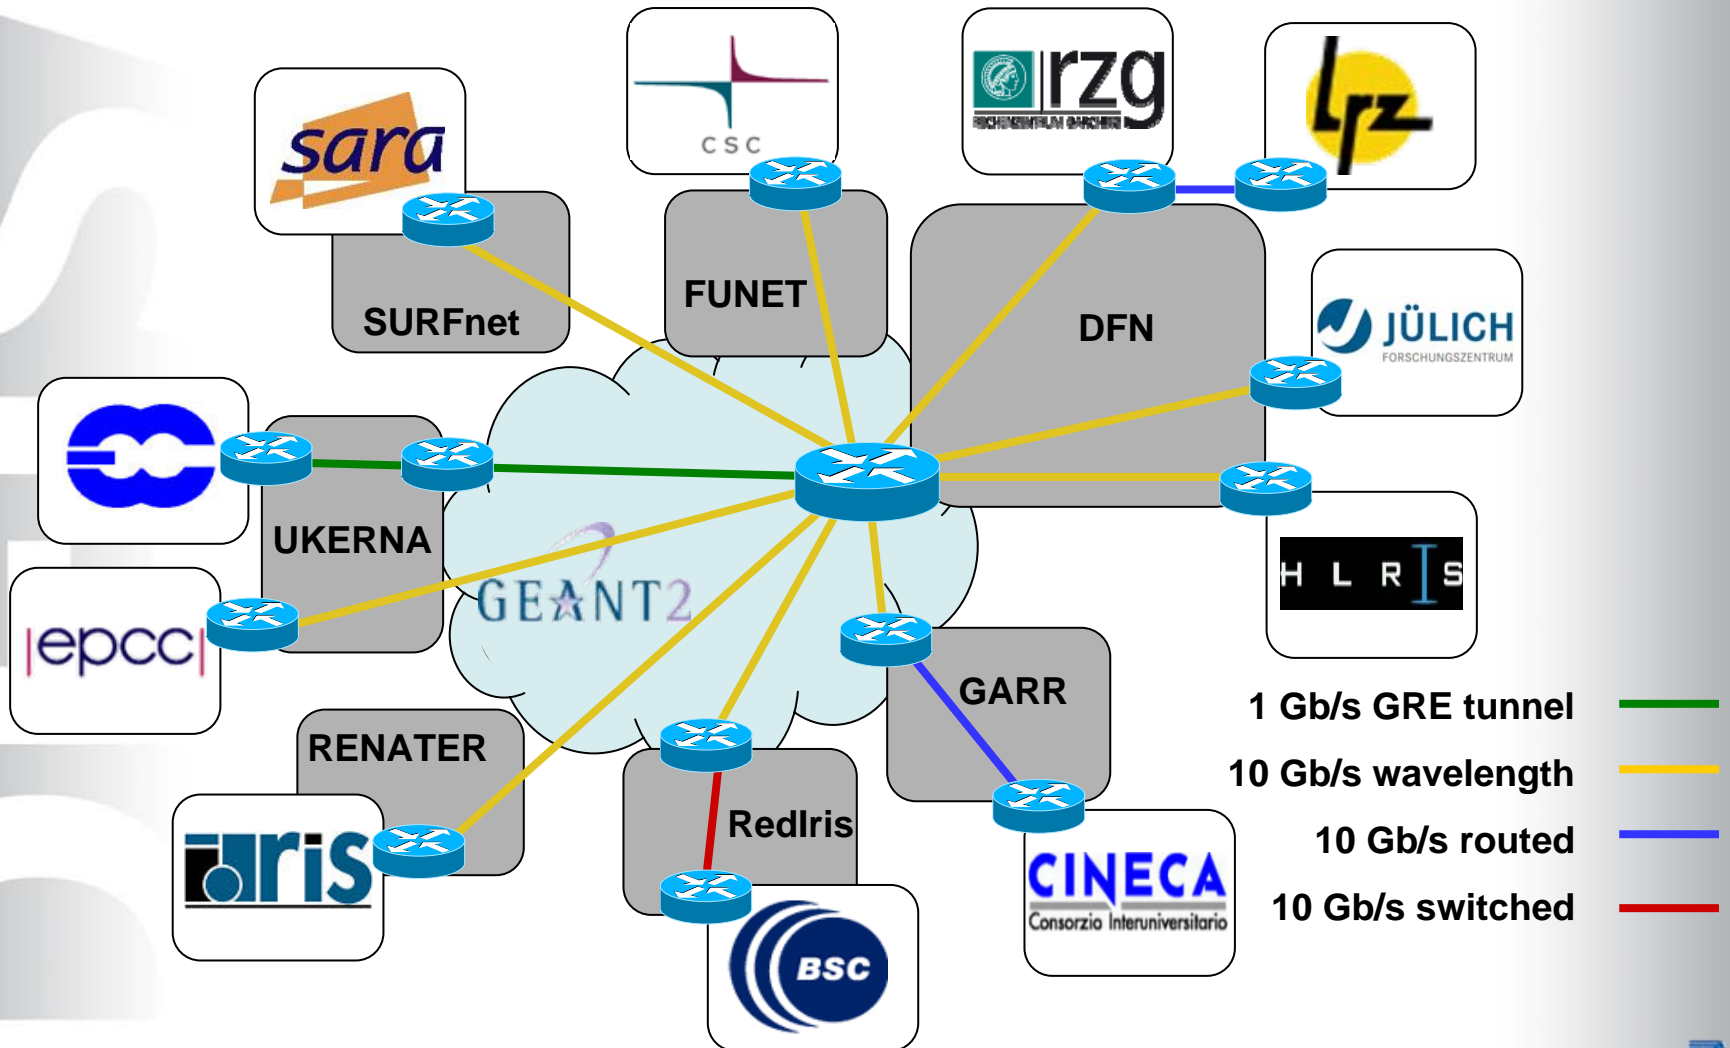
DEISA2: May 1st, 2008 – April 30th, 2011

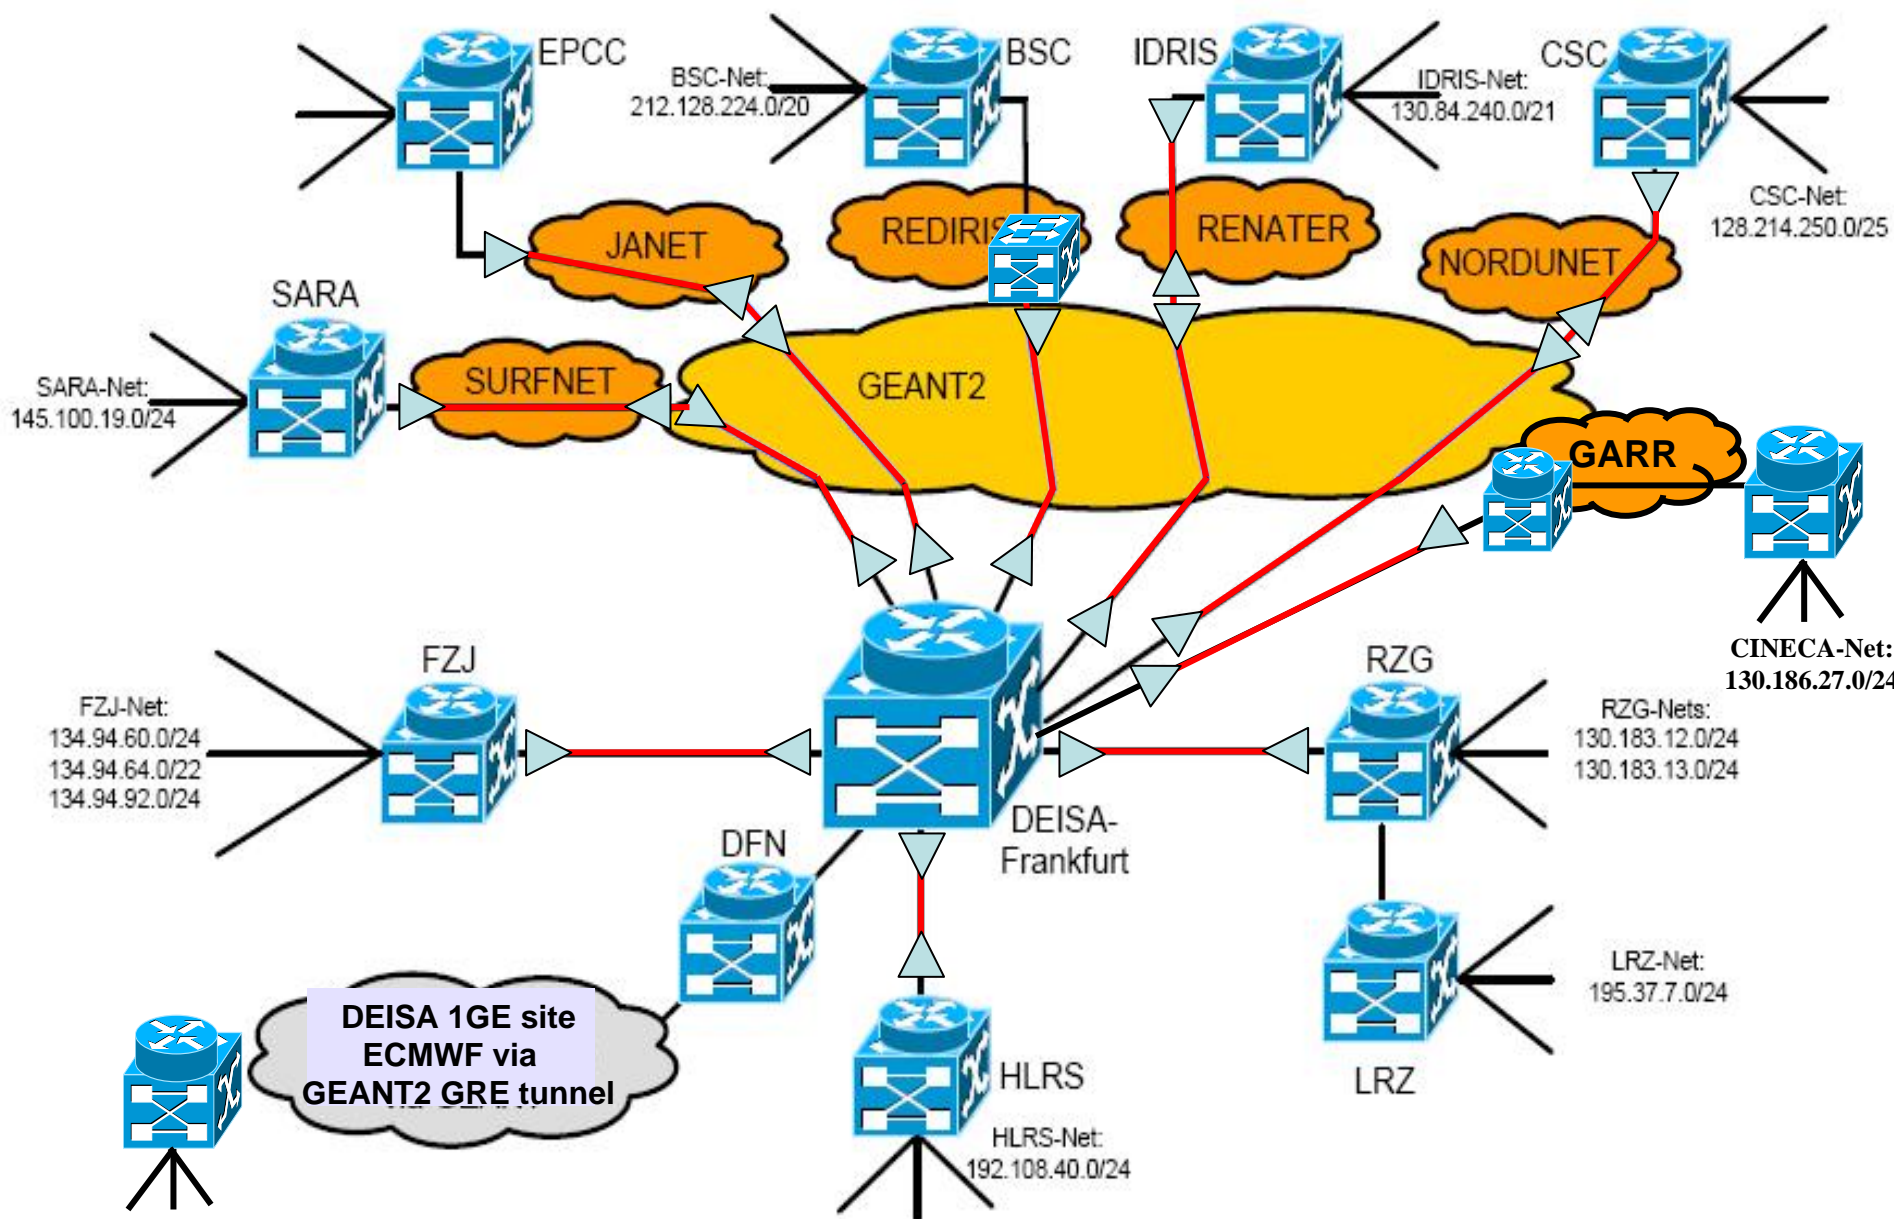


DEISA Partners

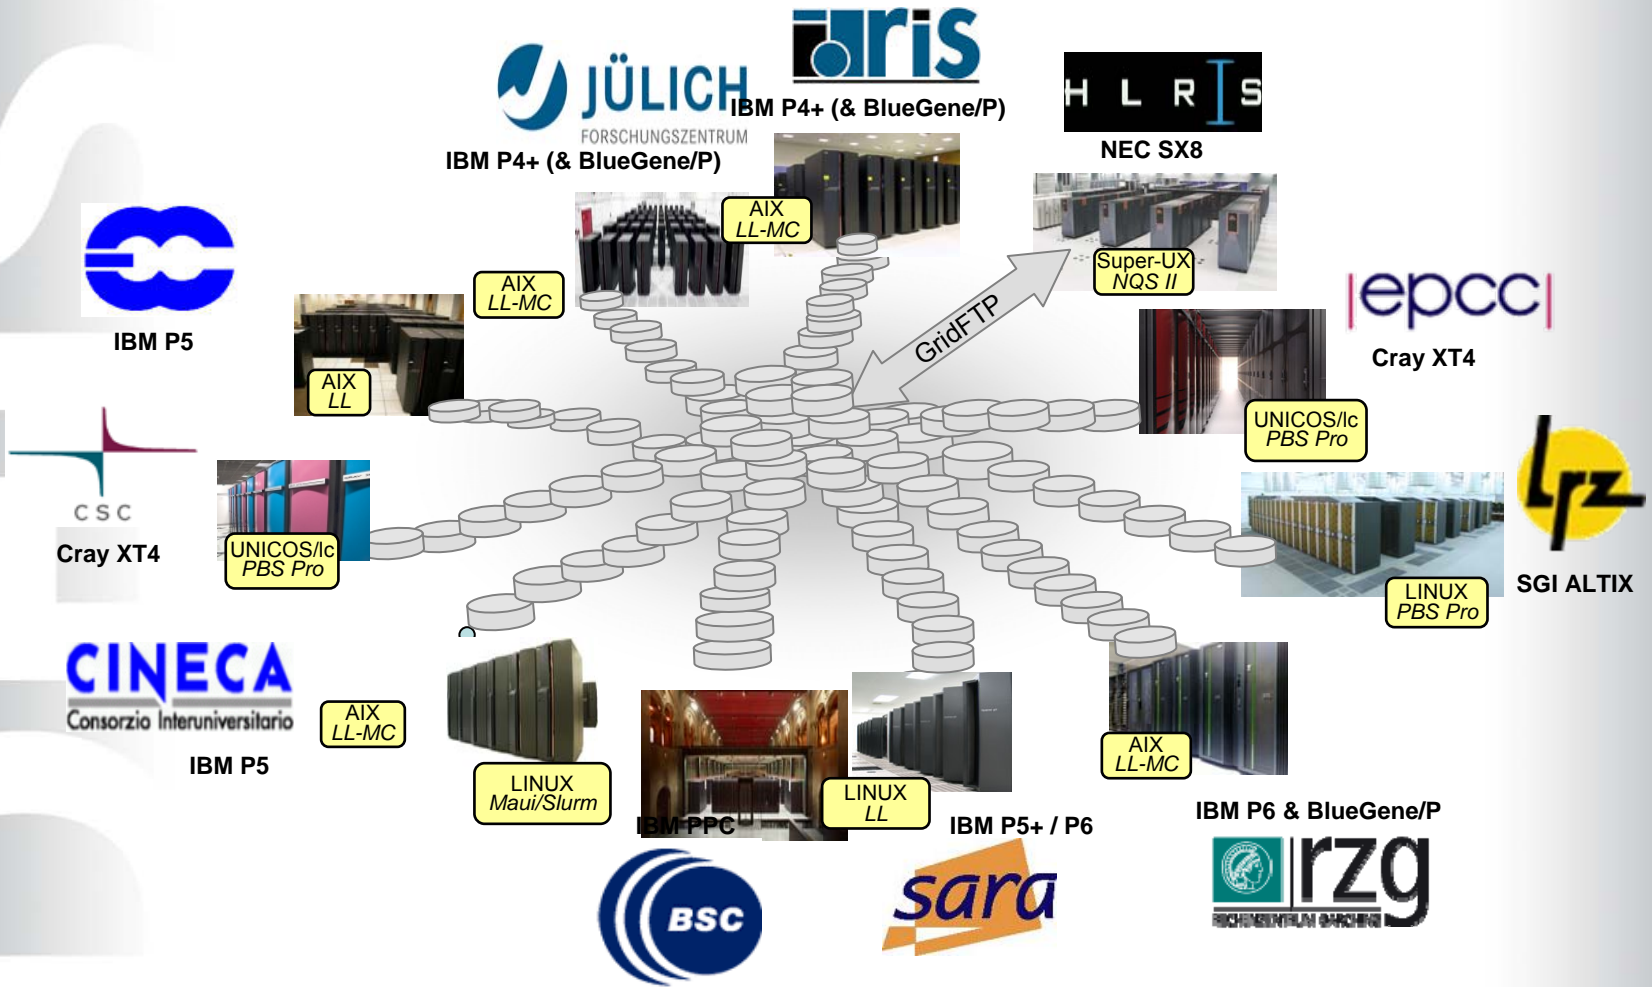
BSC	Barcelona Supercomputing Centre	Spain
CINECA	Consortio Interuniversitario per il Calcolo Automatico	Italy
CSC	Finnish Information Technology Centre for Science	Finland
EPCC	University of Edinburgh and CCLRC	UK
ECMWF	European Centre for Medium-Range Weather Forecast	UK (int)
FZJ	Research Centre Juelich	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 network infrastructure





DEISA Global File System (based on MC-GPFS)



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:
 - Close to 1 PF aggregated Peak performance on state-of-the art supercomputers

- **Cray XT4, Linux**
- **IBM Power4+, 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 fractions of resources dedicated to DEISA usage

DEISA Compute Resources

- Budgets -

- DEISA is operated on top of national supercomputers
- Only a fraction of the HPC resources is dedicated to Extreme Computing Projects at European level
- Major parts of resources are reserved for national projects under national control
- DEISA receives no EU hardware funding and thus only benefits from the regular compute resource upgrades on national budgets from the member states.

DEISA Core Infrastructure and Services

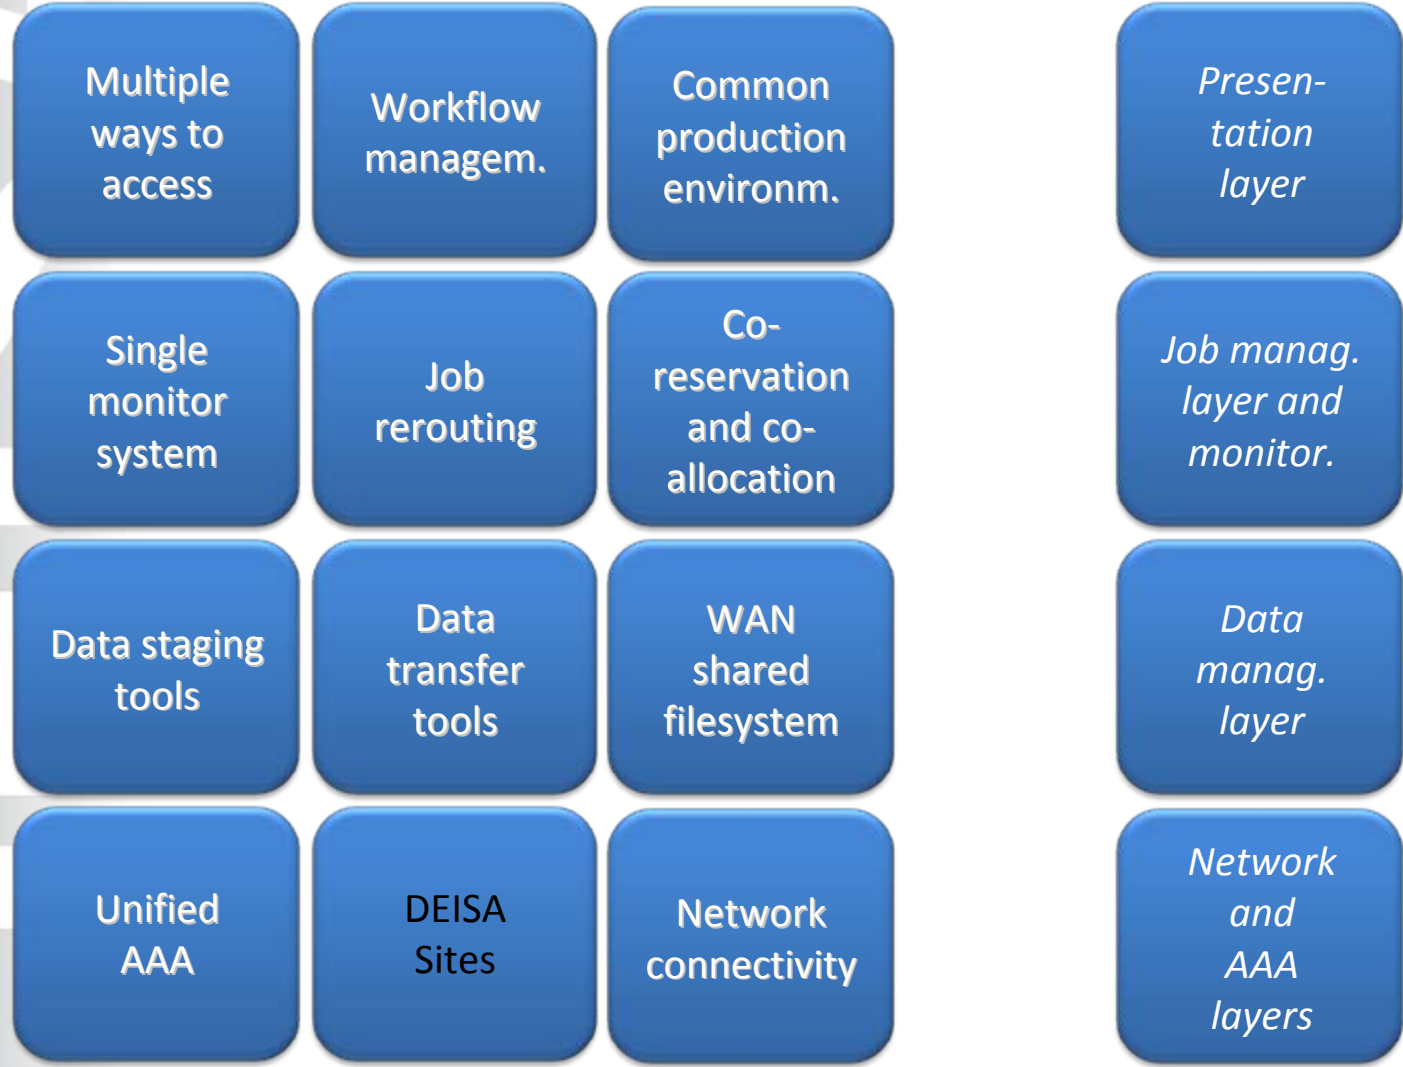


- Dedicated high speed network
- Common AAA
 - Single sign on
 - Accounting/budgeting
- Global data management
 - High performance remote I/O and data sharing with global file systems
 - High performance transfers of large data sets (e.g. GridFTP)
- User Operational infrastructure
 - Distributed Common 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

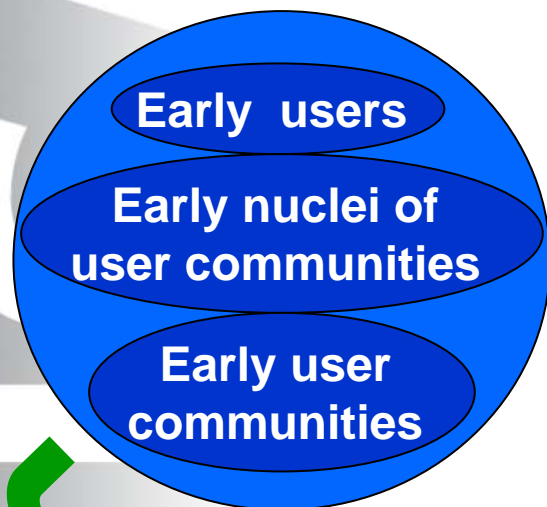
DEISA Additional Services

- Supported Services
 - Workflow management and alternative access method
 - Grid middleware: UNICORE, DESHL
 - Science gateways to supercomputing resource, Portals
 - Job rerouting between identical architectures
- Additional Services
 - Co-Allocation
 - Other middleware components, e.g. from the Globus Toolkit
- Integration of Associate Partners

DEISA Software Layers




DEISA - Users & Application Support



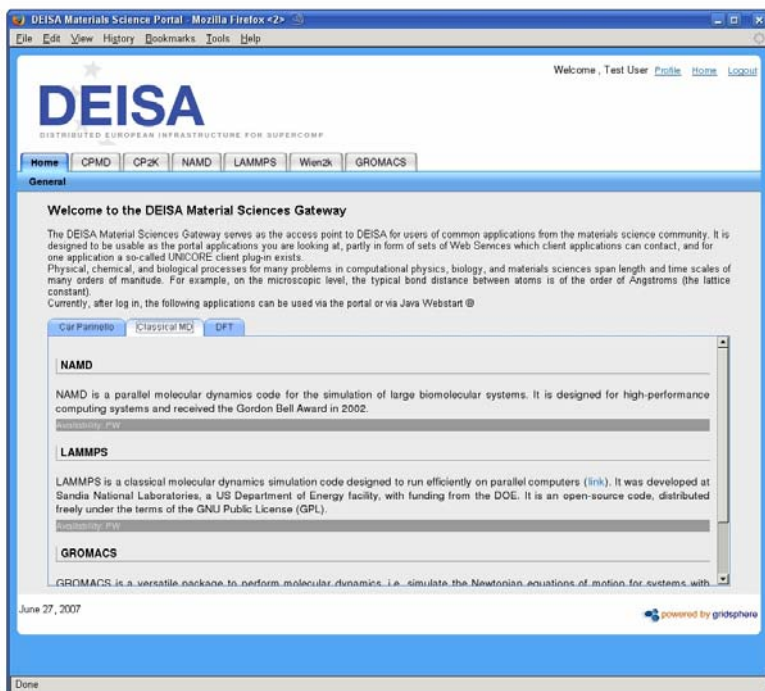
supported
by

Joint Research Activities
JRA1: Materials Science
JRA2: Cosmology
JRA3: Plasma Physics
JRA4: Life Sciences
JRA5: Industry / CFD
JRA6: Coupled Applications
+
Service Activity SA4
User & Application Support



supported
by

Applications Task Force
ATASKF
+ Service Activity SA4
User & Application Support
+ Operation
+ Training Sessions



Dedicated comfortable access to DEISA resources

Special support for important materials science codes




**Supported applications:
CPMD, CP2K, GROMACS, LAMPPS, NAMD, WIEN2K**

Threefold access strategy:

- Rich internet application based access independent of the location of the user
- Web Services, enabling integration into the user's own applications
- UNICORE Client with application-specific plug-ins

JRA in Cosmology

Application support for the VIRGO consortium
and their application portfolio (including pre- and post- processing)
in the heterogeneous DEISA environment



The Virgo Consortium

formed 1995, is a large international collaboration of astrophysicists dedicated to carrying out the largest and most precise simulations of the formation of cosmic structure.

Codes supported &
DEISA optimized:

Gadget

Flash

+ Pre-processing

+ Post-processing

Unicore and DESHL enabled

JRA in Plasma Physics

DEISA Fusion Simulation Suite:

application enabled and optimized for usage on the massively parallel state-of-the-art DEISA hardware

Simulation Codes


TORB / ORB5 GENE,
EUTERPE GEM GYGLES

High relevance for the European fusion community


Significant improvements of both the single processor performance and the scalability of the codes

After successful enabling, all codes run efficiently in the range of 512 up to 32,000 processor-cores.


All enabled codes (but the last just finished one) selected for successful projects in the DEISA Extreme Computing Initiative



ORB5



GENE



JRA in Life Sciences

Promotion of parallel software in the Life Sciences community

The screenshot shows a web browser window displaying the DEISA Life Science Portal. The page title is "DEISA Life Science Portal - Beam@niccy". The URL in the address bar is "https://bioportal.deisa.eu/enginframe/deisa/deisa.xml?_service=BLAST". The page features the DEISA logo, the NICE enginframe logo, and the European Union flag. A navigation bar includes "Home", "Your Jobs", "Spoolers", and "Logout". The account information is "Account: ldrtst02 - Project: prtest01". The main content area is titled "BLAST" and contains a form for submitting BLAST jobs. The form includes fields for "Blast job name", "Blast program", "Blast database", "Blast sequence (FASTA, NCBI Accession numbers or GIs)", "Expect Value", "Filter query sequence (DUST with blastn, SEG with others)", "Query strands to search against database", "Produce HTML output", "Restrict search of database to list of GIs", "Use lower case filtering of FASTA sequence", "Number of concatenated queries", and "Type of job". A "Submit job" button is located at the bottom of the form.

Running big
simulations
on the DEISA
infrastructure
that couldn't
be done locally

Providing ease
of access to
resources

Application
support for life
science portal

DEISA Life Science Portal


JRA: Industrial CFD and CAA applications (via CINECA)

Exploit DEISA infrastructure capabilities for the use of industry relevant commercial codes, in order to:

- **Demonstrate the use of commercial codes on the DEISA infrastructure**
- **Raise the limit of industrial simulations capabilities a step forward**
- **Give hints on how to set up commercial codes ASP service into the DEISA infrastructure**
- **Give a comprehensive overview on use of distributed grid architecture in the field of CFD and CAA**



**Focus on CFD
and CAA
applications**



JRA Coupled Applications

Example:

3D Combustion / Radiation (FOCUS)

- Study the impact of radiative heat transfer (RHT) on the combustion process (3D – more realistic) in DECI context.
- The 2 coupled codes:
 - AVBP solves Navier-Stokes equations and computes the chemical species evolution.
 - DOMASIUM code computes the radiative field coming from the main species.

Achievements

– Built and optimized 2 large coupled configurations (load balancing)

• Gas turbine: emphasize to take account of wall properties (sustain: 350 procs. — 288 + 60)

• 3D diedra: bring out deep changes in the flame behaviour (sustain: 312 procs. — 192 + 120)

– Changed the coupling technology with PALM without difficulties

- PhD thesis, 2 publications in international reviews.

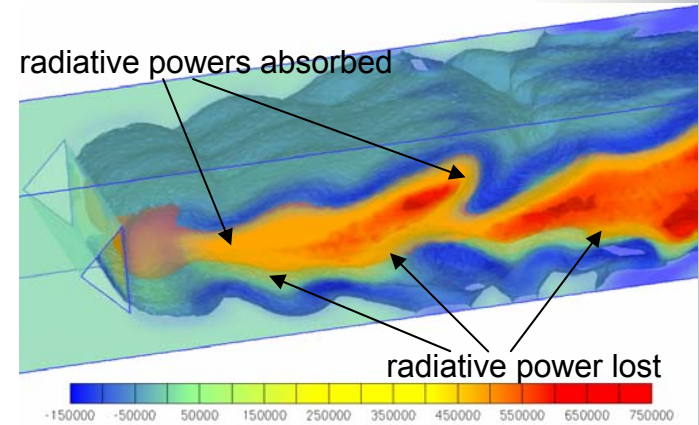


Fig 1: Radiative power in KW.m⁻³




Fig 2: Temperature field in the central plane of the burner ($z = 0$) at $t = 0.55s$

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



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 (ending June 30, 2008)

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

DEISA - Users & Application Support



supported
by

Joint Research Activities

JRA1: Materials Science

JRA2: Cosmology

JRA3: Plasma Physics

JRA4: Life Sciences


JRA5: Industry / CFD

JRA6: Coupled Applications

+

Service Activity SA4

User & Application Support



supported
by

Applications Task Force

ATASKF

+ Service Activity SA4

User & Application Support

+ Operation

+ Training Sessions

DEISA Extreme Computing Initiative DECI



Application Enabling and Optimizations

- Adaptation to the DEISA infrastructure
- Optimizations, also architecture dependent optimizations
- Determination of best suited architecture
- Scaling of parallel programs
- Support for data intensive applications
- Workflow applications to chain several compute tasks (simulation, pre- and post-processing steps)
- Coupled applications e.g. for climate system modelling with separate components for ocean, atmosphere, sea ice, soil, atmospheric chemistry and aerosols

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



North American Institutes involved in projects from DECI calls 2005, 2006, 2007:

Canadian Institute for Theoretical Astrophysics, Toronto, Canada
University of Saskatchewan, Canada
The University of Western Ontario, Canada

University of California, Physics Department, USA

University of California, San Diego, USA

University of Florida, USA

Arizona State University, USA

National Center for Atmospheric Research, USA

Center for Statistical Genetics, Department of Biostatistics II, USA


New Mexico State University, USA

Yale University School of Medicine, USA

Achievements and Scientific Impact 2008


DEISA

ADVANCING SCIENCE IN EUROPE



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






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

DECI Project MolSwitch

First-principles statistical mechanics for molecular switches at surfaces


(K. Reuter, FHI Berlin, Germany)

Azobenzene on copper, silver and gold surfaces




Controlled reversible switching should be possible on Ag surfaces

Building the Infrastructure ... where are we?

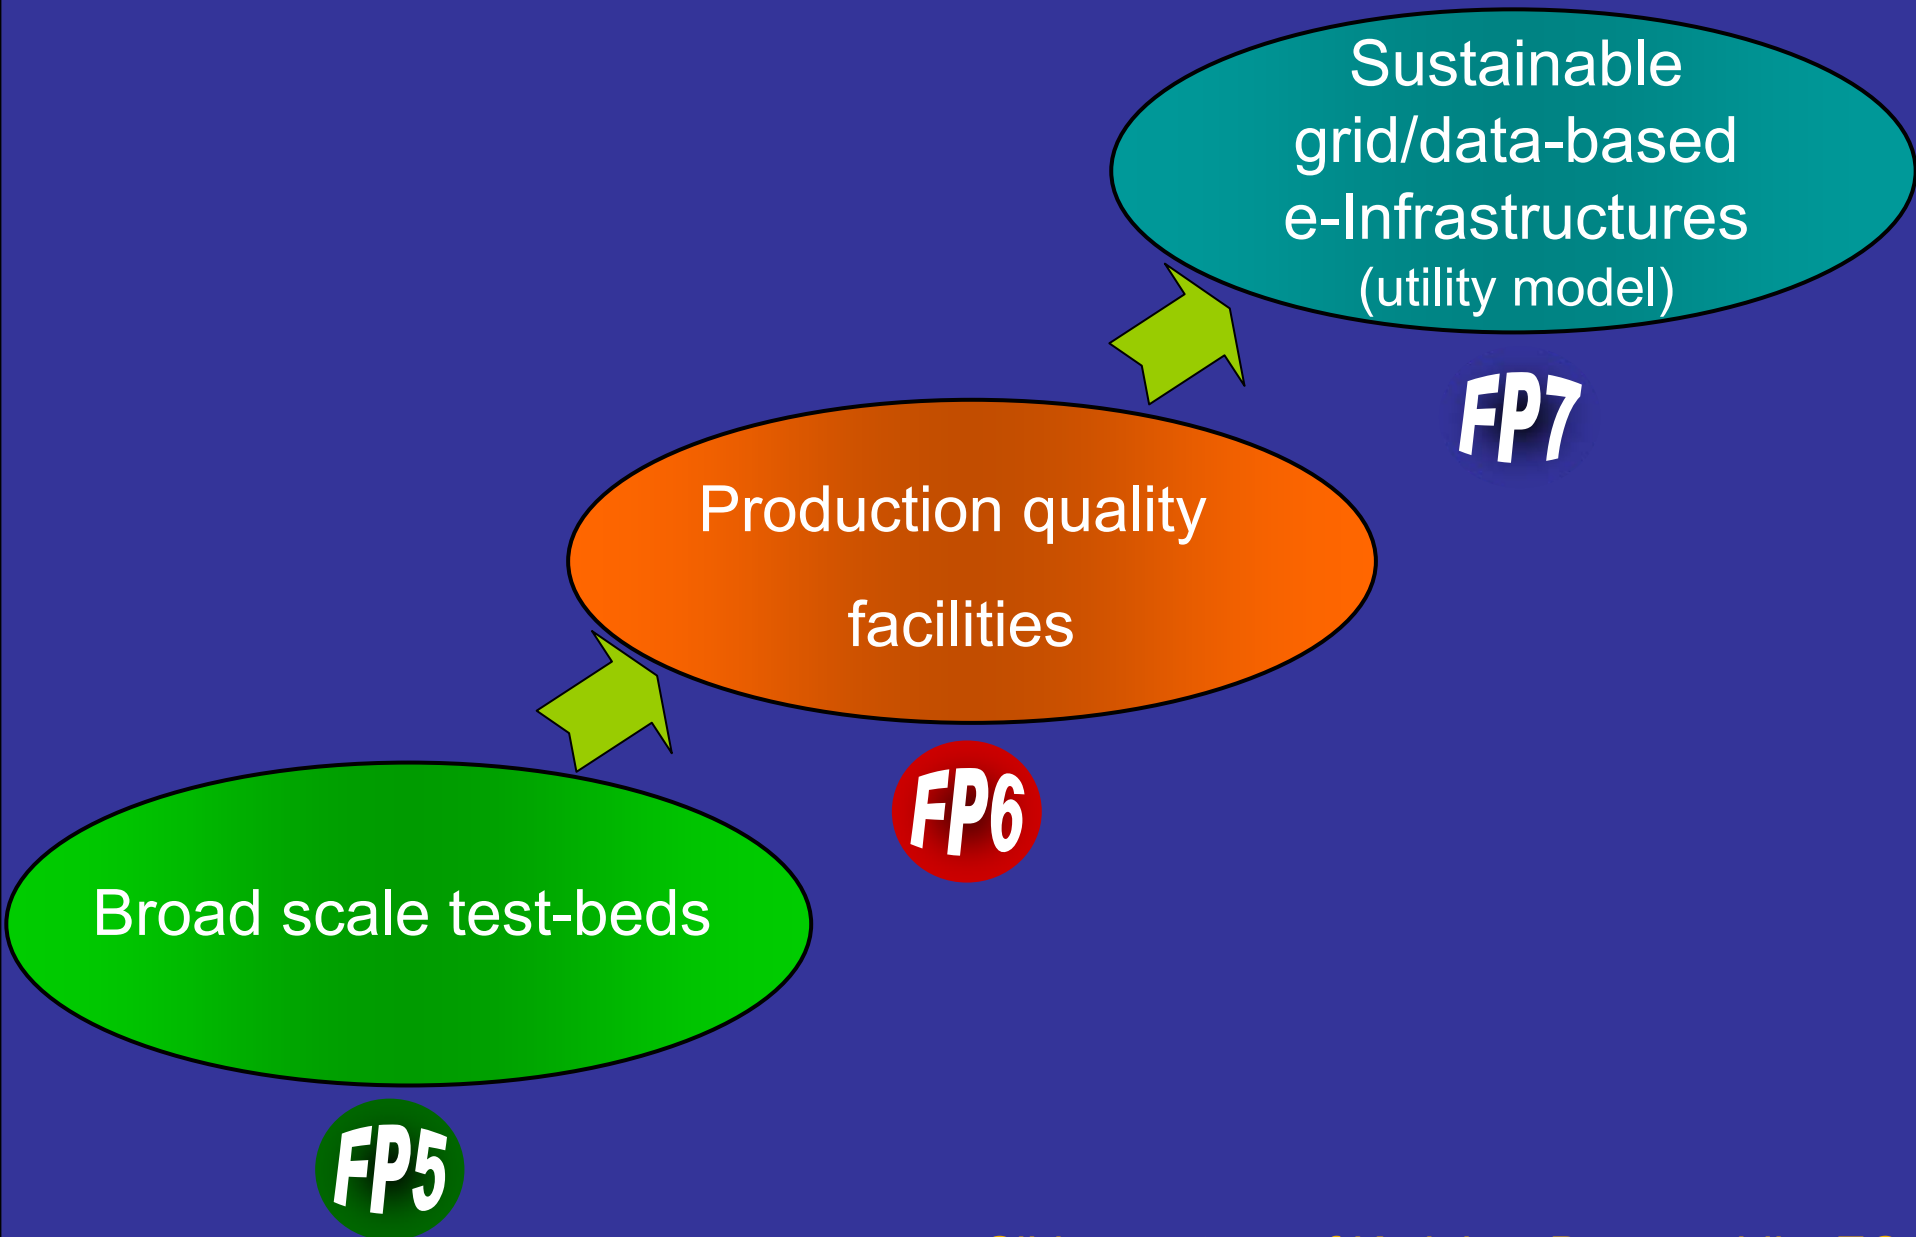


But what about the Ecosystem?

What about the Ecosystem?



Towards sustainable grid-empowered e-Infrastructures



European Strategy Forum on Research Infrastructures

ESFRI Report 2006, p. 65

European High-Performance
Computing Service

The facility

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.



PRACE

Partnership for Advanced Computing in Europe




MoU signed by 14 countries
in April 2007

Extended to 16 countries
In May 2008

EU FP7 PRACE project (Jan 2008 – Dec 2009):
To prepare for the installation of leadership class
supercomputers (Tier-0) in Europe

new "petaflop" supercomputers



15



Mario Campolargo, OGF23, June 2008

In **EU FP7** The DEISA Consortium continues to support and further develop the distributed high performance computing infrastructure and its services through the **DEISA2 project** (May 2008 – April 2011)

- Continuation and further enhancement of activities and services relevant for

Operation Technologies Applications Enabling

- Extension of the service provisioning model from single project support (DECI) to supporting Virtual European Communities.

- Cooperation with PRACE

Of strategic importance is the cooperation with the PRACE project which is preparing for the installation of a limited number of leadership-class Tier-0 supercomputers in Europe.

DEISA Partners & PRACE July 14, 2007



DEISA Centre	Country	PRACE *
BSC	Spain	x
CINECA	Italy	x
CSC	Finland	x
EPCC	UK	x
ECMWF	UK	
FZJ	Germany	x
HLRS	Germany	x
IDRIS-CNRS	France	x
LRZ	Germany	x
RZG	Germany	
SARA	The Netherlands	x
CEA-CCRT	France	x
CSCS	Switzerland	x
KTH	Sweden	x
JSCC	Russia	

*
Either member or member of organisation which is a PRACE member

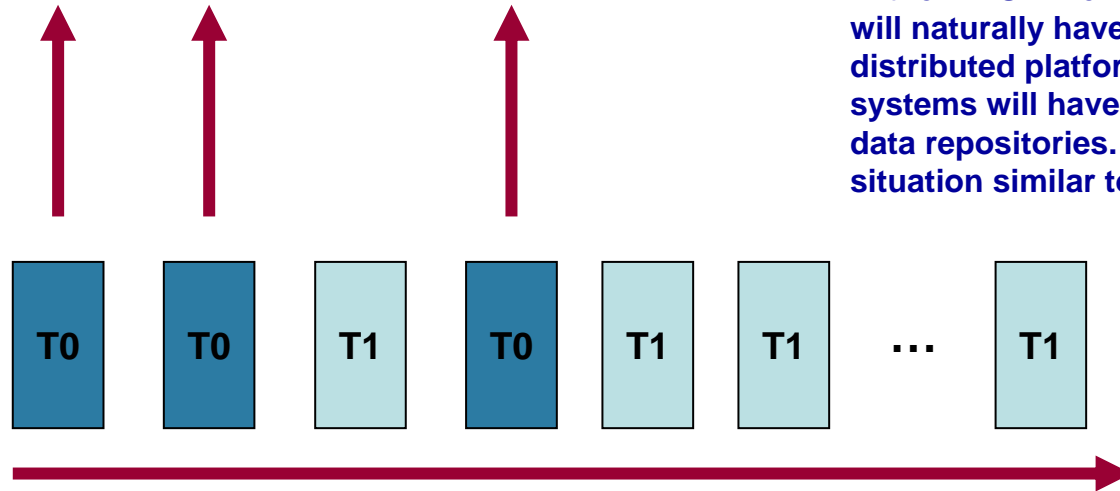
Tier0 / Tier1 Top Layer of the HPC Ecosystem

T0 : future shared petascale European systems
T1 : leading national supercomputing systems

PRACE

Designing an infrastructure that will enable the operation of shared petascale European systems

Enhancing performance in selected sites and providing wide access to shared systems



In the DEISA-2 environment, scientists will naturally have access to several distributed platforms, and shared systems will have to access remote data repositories. We will be in a situation similar to TeraGrid.

Member states sites

DEISA-2 : strong integration of T0 and T1 systems (automatically provides wide, seamless and efficient access to shared systems and data repositories)

The DEISA-1 services have been tailored for this mode of operation. There is a positive feedback between the two orthogonal lines of action:

- *DEISA is paving the way to the efficient operation of T0 systems.*
- *T0 systems will drive the massive adoption of the DEISA services.*

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
- T1 Centers
 - Leading national Centers, cyclically renewed, optionally surpassing the performance of older T0 machines
 - National Governance structure
- The 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)

Real needs of HPC users of T0/T1

HPC users - T0/T1 - are conservative, standard access methods are preferred, no interest in complicated middleware stacks.

Comfortable Data Access among T0/T1/T2 is the key for satisfied users and for the future European HPC Infrastructures!

- Global Login
 - *“HPC users prefer a personal Login in each system”*
- Comfortable Data Access
 - *“HPC users need a global, fast and comfortable access to their data”*
- Common Production Environment
 - *“HPC users do not need an identical but an equivalent HPC software stack”*
- Global Help Desk
 - *“HPC users wish one central point of contact and as local as possible”*
- Application Support
 - *“HPC users need help in scalability and adaptation to different architectures”*

Engagement for Standardization



DEISA has a vital interest and is engaged in standardization matters related to High Performance Computing

DEISA major standardization needs are

Job submission, job and workflow management, data management, data access and archiving, networking and security (including AAA)

DEISA standardization involvement

Participation in OGF standardization groups like JSDL-WG and OGSA-BES for job submission, UR-WG and RUS-WG for accounting purposes, DAIS for Web Services

DEISA collaboration in standardization with other projects

Close collaboration with the GIN community group within the OGF
The Infrastructure Policy Group was established

Members: DEISA, EGEE, TeraGrid, OSG, Naregi

Goal: Achievement of seamless interoperation of leading Grid Infrastructures worldwide

Current Focus: Authentication, Authorization, Accounting and Auditing (AAAA)


Essentials of DEISA2



- Consolidation of the existing DEISA infrastructure
- Evolvement 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 ecosystem
- Bridging worldwide HPC projects: To facilitate the support of international science communities with computational needs traversing existing political boundaries

DEISA 2008

Operating the European HPC Infrastructure




~1 PetaFlop/s
Aggregated peak performance

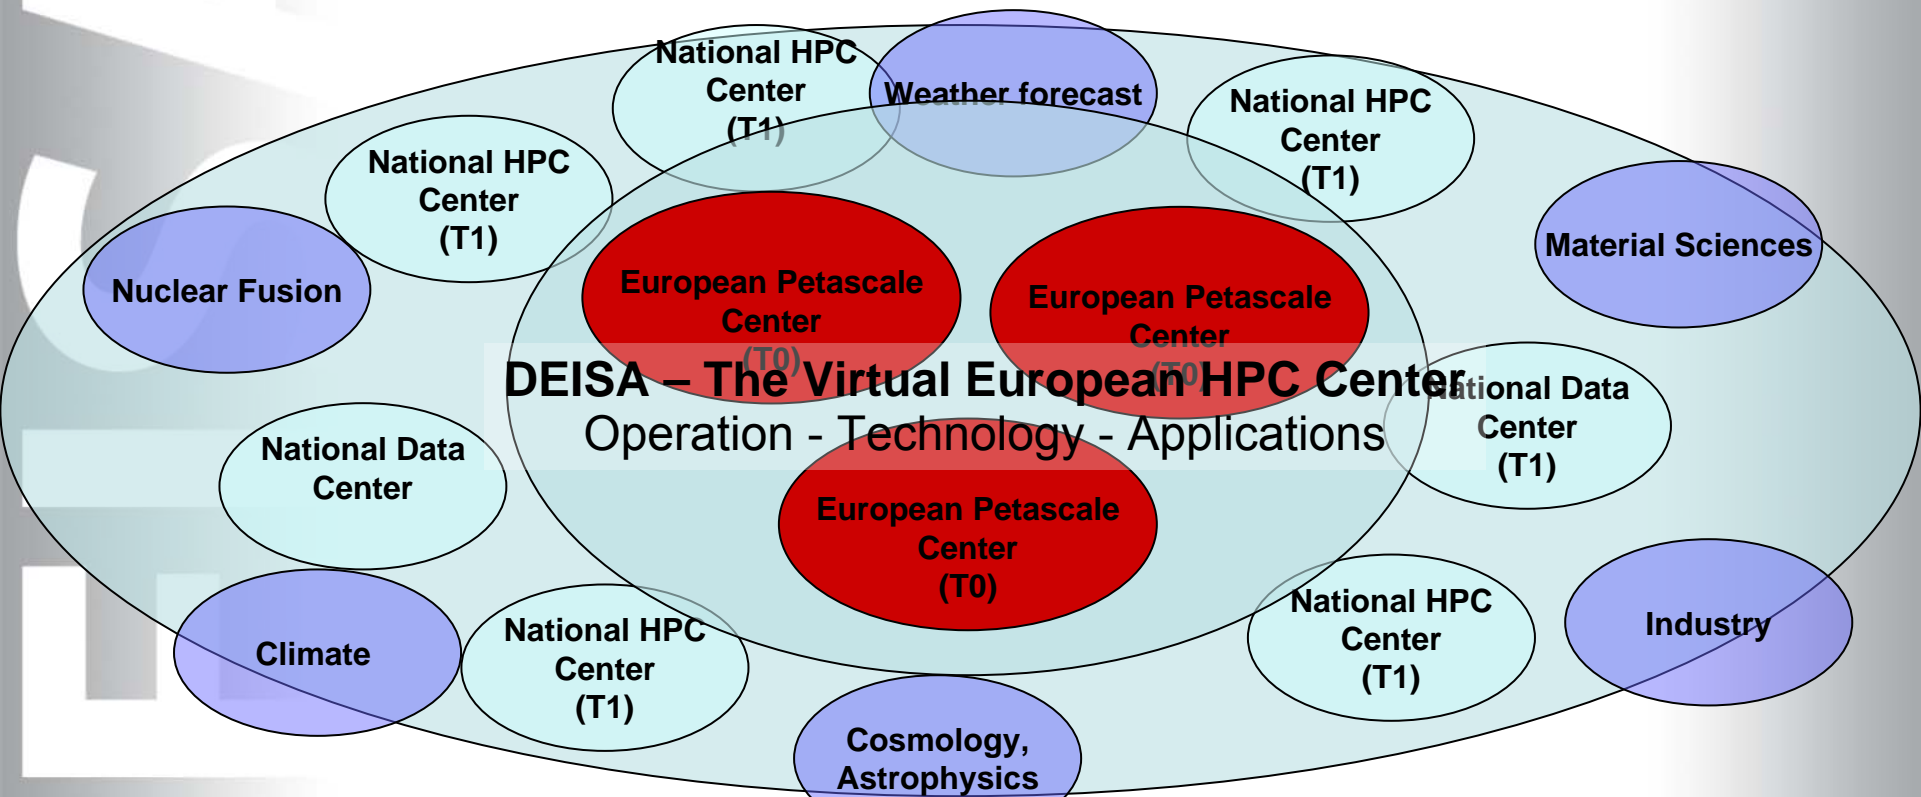
Most powerful European supercomputers for most challenging projects

Top-level Europe-wide application enabling

Grand Challenge projects performed on a regular basis

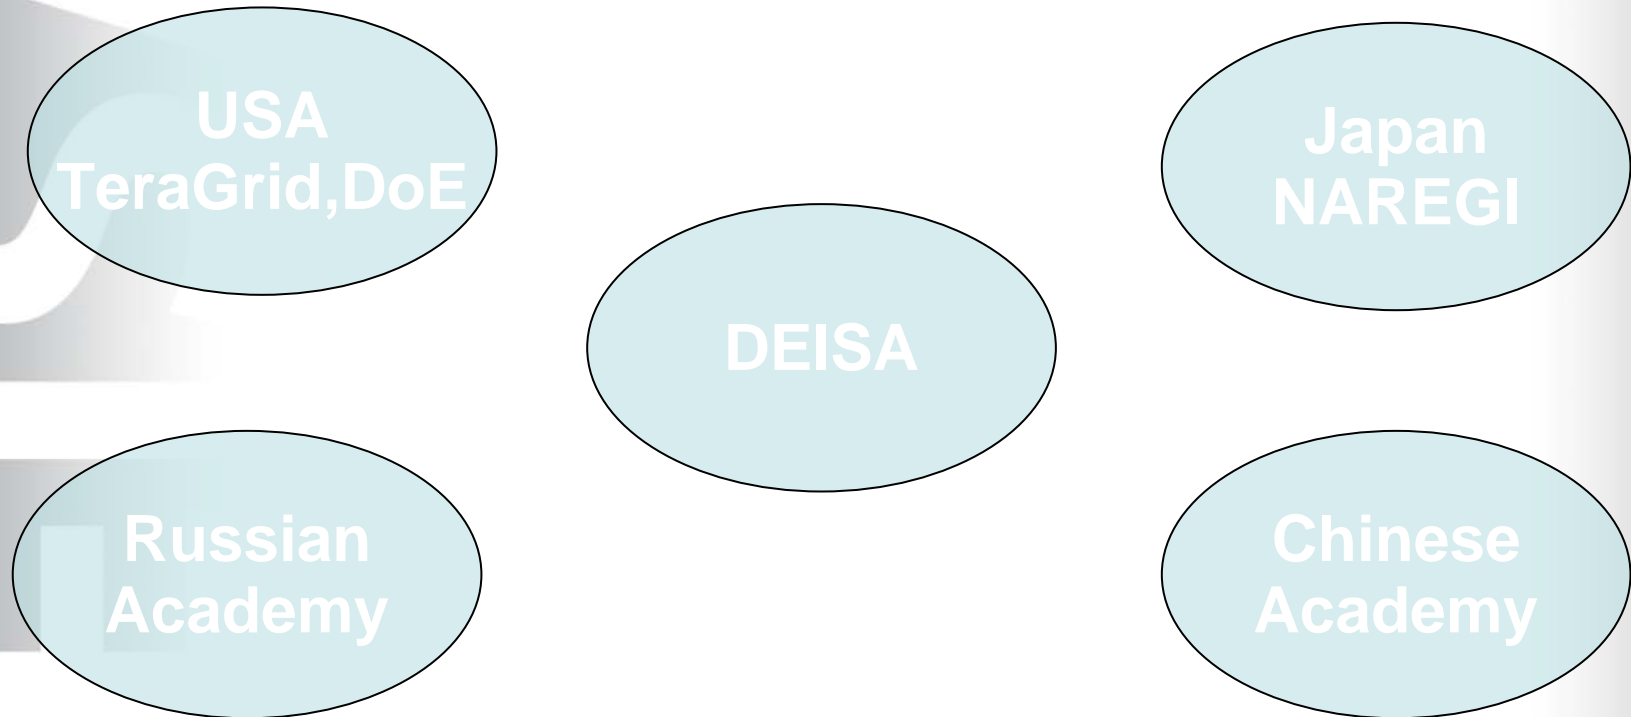


Vision and Strategy



- Enhancing the existing distributed European HPC environment (DEISA) to a turnkey operational infrastructure
- Advancing the computational sciences in Europe by supporting user communities and extreme computing projects
- Enhancing the service provision by offering a complete variety of options of interaction with computational resources
- Integration of T-1 and T-0 centres
- The Petascale Systems need a transparent access from and into the national data repositories
- Bridging worldwide HPC and Grid projects

Vision and Strategy



- Enhancing the existing distributed European HPC environment (DEISA) to a turnkey operational infrastructure
- Advancing the computational sciences in Europe by supporting user communities and extreme computing projects
- Enhancing the service provision by offering a complete variety of options of interaction with computational resources
- Integration of T-1 and T-0 centres
- The Petascale Systems need a transparent access from and into the national data repositories
- Bridging worldwide HPC and Grid projects

Welcome to cooperate!



DEISA team in March 2008, after six years of pioneering work towards a European HPC Ecosystem