

DEISA : strategies and perspectives

Towards cooperative extreme computing in Europe



Victor Alessandrini
IDRIS - CNRS
www.deisa.org



DEISA Symposium
Munich, May 21-22 2007.

Victor Alessandrini, IDRIS - CNRS



DEISA objectives



- *To enable Europe's terascale science by the integration of Europe's most powerful supercomputing systems.*
- DEISA is an European Supercomputing Service built on top of existing national services. This service is based on the deployment and operation of a persistent, production quality, distributed supercomputing environment with continental scope.
- Main focus is High Performance Computing (HPC).
- After three years of operation, DEISA confirms its original motivation of providing a basic vector of integration of HPC resources at the continental scale. The DEISA services have been tailored to enable seamless access to, and high performance cooperative operation of, a distributed park of leading supercomputing platforms in Europe.
- DEISA services are deployed on top of a dedicated high speed network infrastructure connecting computing platforms, using selected middleware. Their primordial objective is enabling **capability** computing across remote computing platforms and data repositories.

Participating Sites



BSC	<i>Barcelona Supercomputing Centre</i>	Spain
CINECA	<i>Consortio Interuniversitario per il Calcolo Automatico</i>	Italy
CSC	<i>Finnish Information Technology Centre for Science</i>	Finland
EPCC/HPCx	<i>University of Edinburgh and CCLRC</i>	UK
ECMWF	<i>European Centre for Medium-Range Weather Forecast</i>	UK (int)
FZJ	<i>Research Centre Juelich</i>	Germany
HLRS	<i>High Performance Computing Centre Stuttgart</i>	Germany
IDRIS	<i>Institut du Développement et des Ressources en Informatique Scientifique - CNRS</i>	France
LRZ	<i>Leibniz Rechenzentrum Munich</i>	Germany
RZG	<i>Rechenzentrum Garching of the Max Planck Society</i>	Germany
SARA	<i>Dutch National High Performance Computing and Networking centre</i>	The Netherlands

THE DEISA SUPERCOMPUTING ENVIRONMENT



Systems interconnected with dedicated 10 Gb/s Network provided by GEANT.




AIX distributed super-cluster




Vector systems (NEC)



Linux systems (SGI Altix, IBM, ...)



In designing DEISA, we knew that Grid computing is not always HPC ...



... and this is why Grids will never fully replace tera or petascale supercomputers.

The reason is that the speed of light is not big enough

Finite signal propagation speed boosts message passing latencies in a WAN from a **few microseconds** to **tens of miliseconds** (if A is in Paris and B in Helsinki)

If A and B are two halves of a tightly coupled complex system, communications are frequent and the enhanced latencies kill the required high performance.

Grid computing works best for embarrassingly parallel applications, or coupled software modules with limited communications.

Large, tightly coupled parallel applications should be run in a single platform. This is why we still need high end supercomputers.


DEISA implements an operational model based on the seamless access to a distributed set of high end supercomputers accessing **(with high performance)** distributed remote data sets.

How is DEISA enhancing HPC services in Europe?



- **Running larger parallel applications** in individual sites, by a cooperative reorganization of the global computational workload on the whole infrastructure.
- Enabling **workflow applications** with UNICORE (complex applications that are pipelined over several computing platforms)
 - In some cases (whan it makes sense) enabling coupled multiphysics Grid applications.
- Providing a **global data management** service whose primordial objective is the tight integration of distributed data with distributed computing platforms (*and I will argue that this paves the way and is critical to the efficient operation of future shared European supercomputers*):
 - **Enabling efficient, high performance access to remote datasets** (with Global File Systems and stripped GridFTP).
 - Integrating hierarchical storage management and databases in the supercomputing Grid.
- **Deploying portals** as a way to hide complex environments to new users communities, and to interoperate with another existing grid infrastructures.

Workflow simulations using UNICORE



UNICORE supports complex simulations that are pipelined over several heterogeneous platforms (workflows).

UNICORE handles workflows as a unique job and transparently moves the output – input data along the pipeline.

UNICORE clients that monitor the application can run in laptops.

UNICORE has a user friendly graphical interface. DEISA has developed a command line interface for UNICORE.

UNICORE has a self-contained, robust security model.

DEISA Global File System integration in 2006


(based on IBM's GPFS)



AIX IBM domain



Linux SGI



High Performance Common Global File System
various architectures / operating systems
High bandwidth (up to 10 Gbit/s)



LINUX Power-PC

Global File System Interoperability demo during Supercomputing Conference 2005 in Seattle

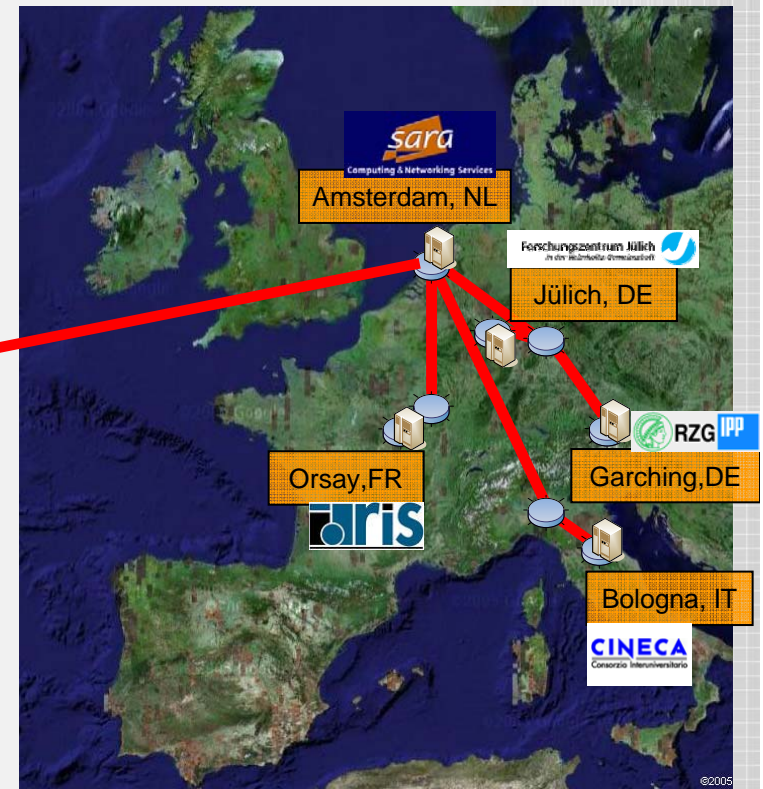


American and European supercomputing infrastructures linked: bridging communities with scalable, wide-area global file systems


TeraGrid Sites



DEISA Sites



Other basic DEISA services



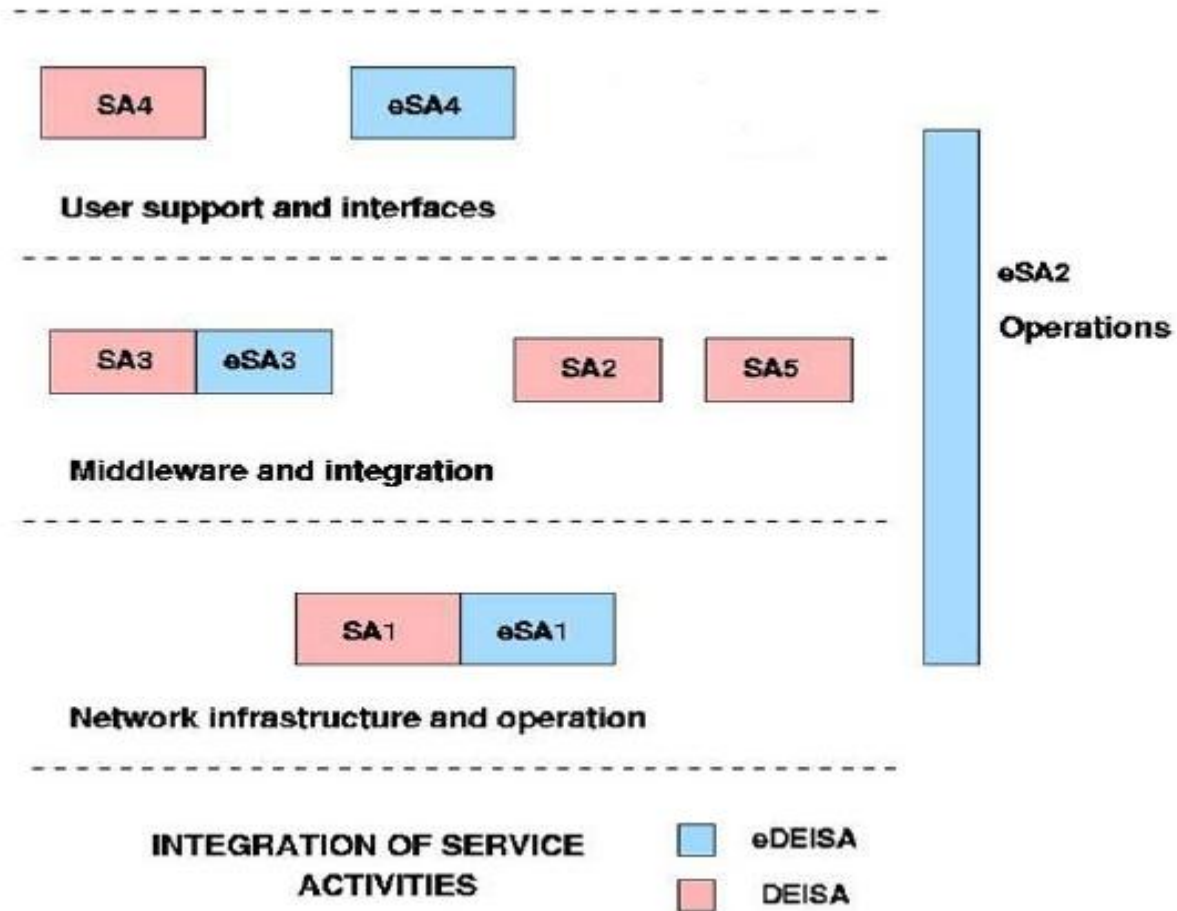
Co-scheduled, parallel data mover tasks

Workflows across different platforms, transparent remote I/O, large file transfers, are starting to operate without inducing performance bottlenecks that would invalidate high performance computing.

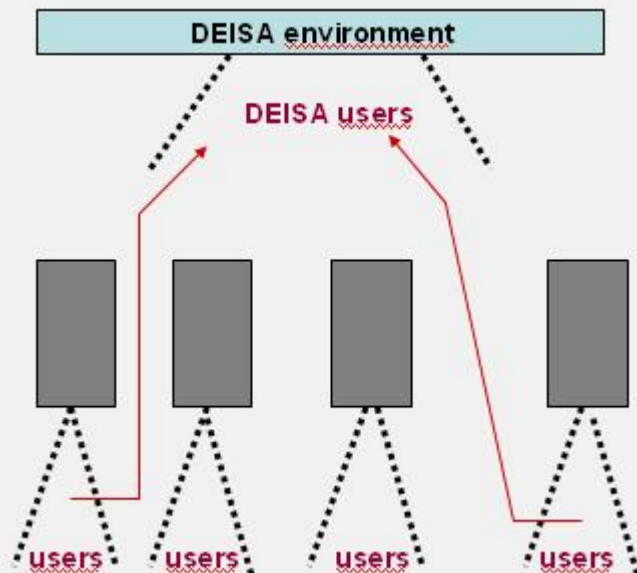
A fundamental and difficult accomplishment is the establishment of a Common Production Environment across all sites that guarantees the easy migration and seamless rerouting of Applications. This is one of the most important added values to the national environments

These services will bloom as the number of scientific users accessing different computing platforms in Europe will grow.

DEISA Service Activities



Enabling science: DEISA users communities



National users communities have accounts on a given site and do not « naturally » see the whole DEISA environment.

Promotion of DEISA users is done via the **Extreme Computing Initiative**.

European call for proposals for grand challenge simulations every year in May since 2005.

About 50 grand challenge projects supported each year since 2005.

Full information about Extreme Computing projects and reports from terminated projects can be found in the DEISA Web server: www.deisa.org

The Extreme Computing Initiative is the current DEISA service provisioning model. This, however, will evolve in FP7.

T0-T1 top layer of the HPC ecosystem



T0 : future shared petascale European systems

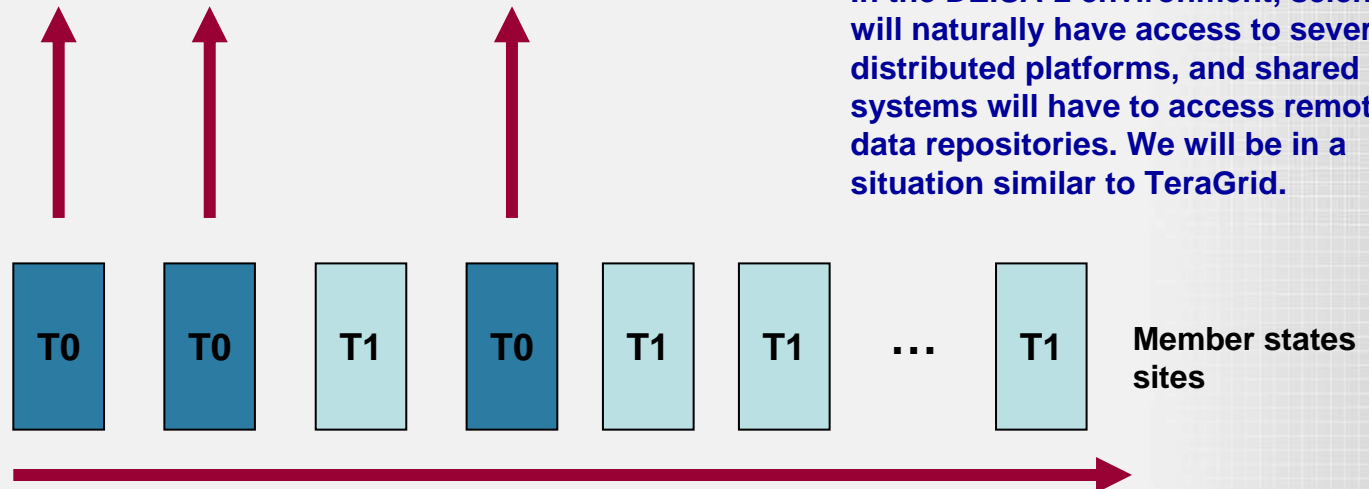
T1 : leading national supercomputing systems

PACE

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.



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.*

DEISA FP7 strategy (1)



- Independently of the implementation and governance details of the future T0 projects, it seems obvious, for the reasons advanced above, that access to T0 and between T0 platforms will be realized through the pre-existing DEISA infrastructure of the T1 sites. Potential T0 sites will be DEISA sites (in the very unlikely case that they will not be, they will have to be integrated immediately). For users the whole T0-T1 environment has to appear as a unique environment, with service provisioning models and best practices as unified as possible. This is not a political statement; this is a common sense statement following from what DEISA considers the everyday best practices for the operation of a global HPC environment in Europe.
- DEISA intends to consolidate its infrastructure by proposing in the very short term – for the 2007 FP7 calls – a more robust organization and improved management and operational and service provisioning models. DEISA is prepared to adapt its new organization to the needs of the T0 layer at the level of global user interfaces, operation and service provisioning models.

DEISA FP7 strategy (2)



- **Other strategic arguments in favour of a strong integration of T0 and T1 layers can be advanced.** T0 systems will necessarily consist of a few systems located in a few leading countries in Europe. Most of the European nations will not host T0 systems. But most of the European nations and organizations will not want to weaken their competence and their leadership in high performance computing. The most reasonable and efficient compromise is to rely on the fact that the actual location of a supercomputer is largely irrelevant. The necessity of sharing investments leads to the necessity of concentrating exceptional computing resources in a few sites, but the HPC specific services, the specific competences and the know how should remain distributed.
- **DEISA is currently the essential infrastructure for the integration of the national supercomputing centres and will be adapted to become the essential infrastructure for a strongly integrated European T0/T1 ecosystem.**

DEISA FP7 organization



- Project Direction - Coordination
- Administrative and Financial Office
- Applications and User Support
- Operations
- Technology

Conclusions



- **The EU as well as the major HPC actors in Europe** are providing visionary leadership for the deployment and operation of an integrated European supercomputing environment.
- There is no doubt that we are moving in the right direction for the establishment of an efficient and persistent European HPC infrastructure
- In FP7, a tight interplay between DEISA and PACE seems to be the right instrument for an aggressive enhancement of the outreach of European HPC.
- **Only the user services need to be persistent, not the instruments adopted to deploy and operate them.**
- Beyond FP7, if an European HPC organization is established, it is possible that the instruments adopted to deploy the DEISA services will have to be reexamined.