EUFORIA
EU for ITER Applications

Enol Fernández

Euforia Consortium

IFCA - CSIC

GridKa School 2009

ITER

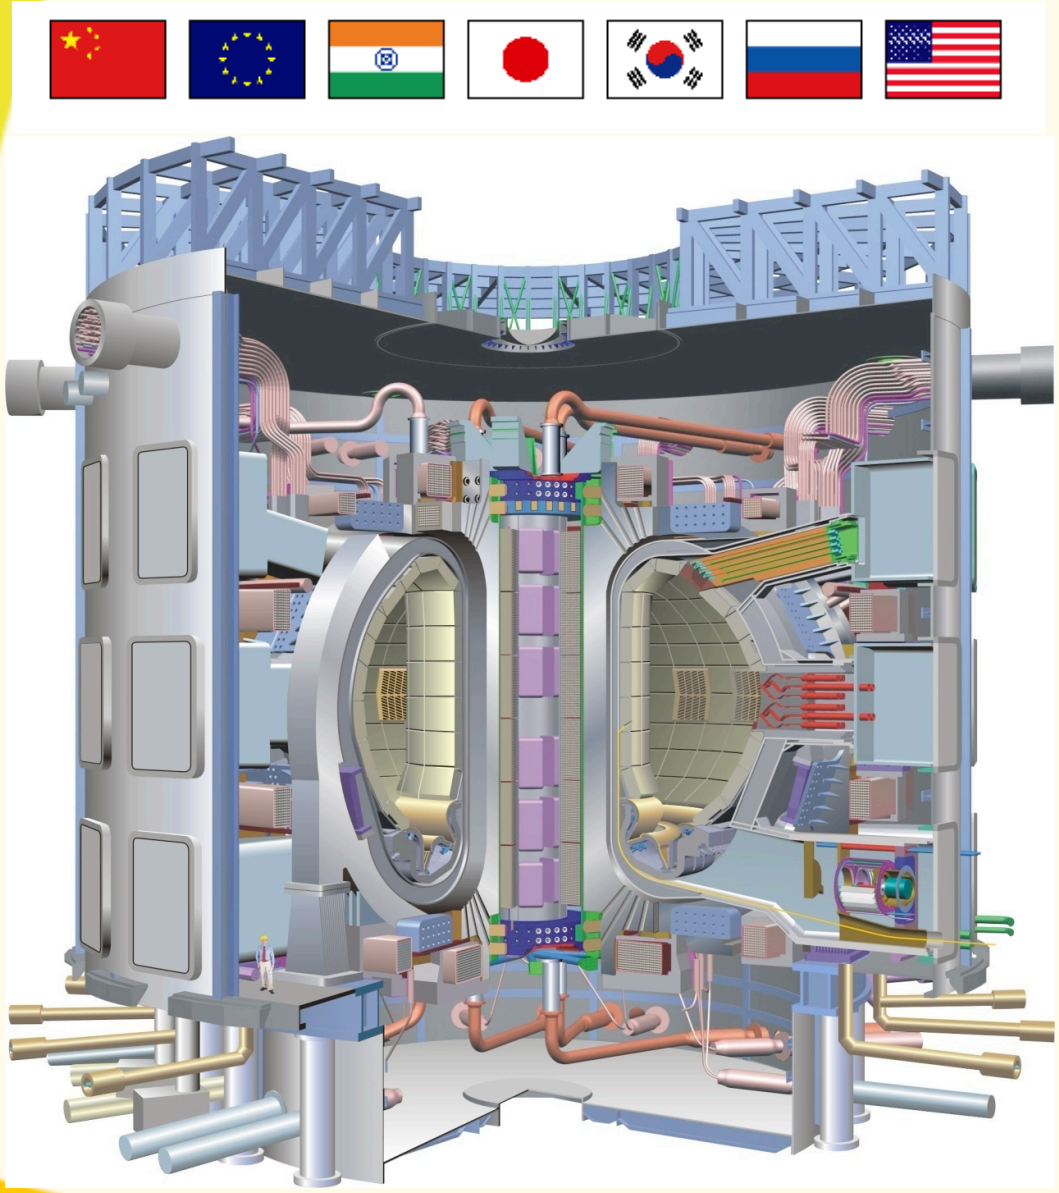

Involves 7 partners representing more than 50% world population
Costs > 10 G$
Under construction in Cadarache, France
Key element on the path to fusion energy production

ITER

- ITER is an experiment that aims to demonstrate that it is possible to produce commercial energy from fusion.
- Requires a high degree of physics modelling and simulation (even in current construction phase)
- Plasmas are complex systems
  - Very demanding from a computational point of view
  - Simulate 1ms takes several days in a 50-100TFlop machine
    - ITER expects to reach 5 minutes

EUFORIA

EUFORIA Objective:

> Develop a comprenhensive framework and infrastructure for the european fusion modelling community.

- Deploy Grid and HPC infrastructures
- Adaptation and optimization of fusion codes
  - Oriented to grid and/or HPC
- Development of advanced tools for:
  - Workflows
  - Visualization
  - Data Mining

EUFORIA

14 member Institutes
36 months

- Management
- Training
- Dissemination
- Grid and HPC infra-structure & support
- Code adaptation & optimization
-Workflows
-Visualization

Consortium Members

Country	Institute	Capabilities
SWEDEN:	*CHALMERS University of Technology (coordinating)*	*Fusion, Grid, (CS)*
FINLAND:	*CSC - Tieteellinen laskenta Oy*	*HPC, (Grid),*
	Åbo Akademi University	*Code Optimization, CS*
FRANCE:	*CEA - Commissariat à l'énergie atomique – Cadarache*	*Workflow, Fusion, CS*
	Université Louis Pasteur	*Visualization, Applied Math*
GERMANY:	*Forschungszentrum Karlsruhe GmbH -FZK*	*Grid, Code parallelisation*
	Max-Planck-Institut für Plasmaphysik - IPP	*Fusion, (HPC, Grid)*
ITALY:	ENEA	Fusion, Grid, HPC, GATEWAY
SLOVENIA:	*University of Ljubljana -LECAD*	*Visualization, CS*
POLAND:	Poznan Supercomputing and Networking Centre	Grid, Migrating Desktop, CS
SPAIN:	*Barcelona Supercomputing Center – Centro Nacional de Supercomputación -BSC*	*HPC, Code optimization*
	Centro de Investigaciones Energéticas Medio Ambientales y Tecnológicas -CIEMAT	*Grid, Code parallelization, Fusion, Grid, NA*
	Consejo Superior de Investigaciones Cientificas - CSIC	*Grid, CS, (NA activities)*
UNITED KINGDOM:	*The University of Edinburgh - EPCC*	*HPC, Code Optimization, NA, User support, (GRID)*

Developing a new paradigm for fusion computing

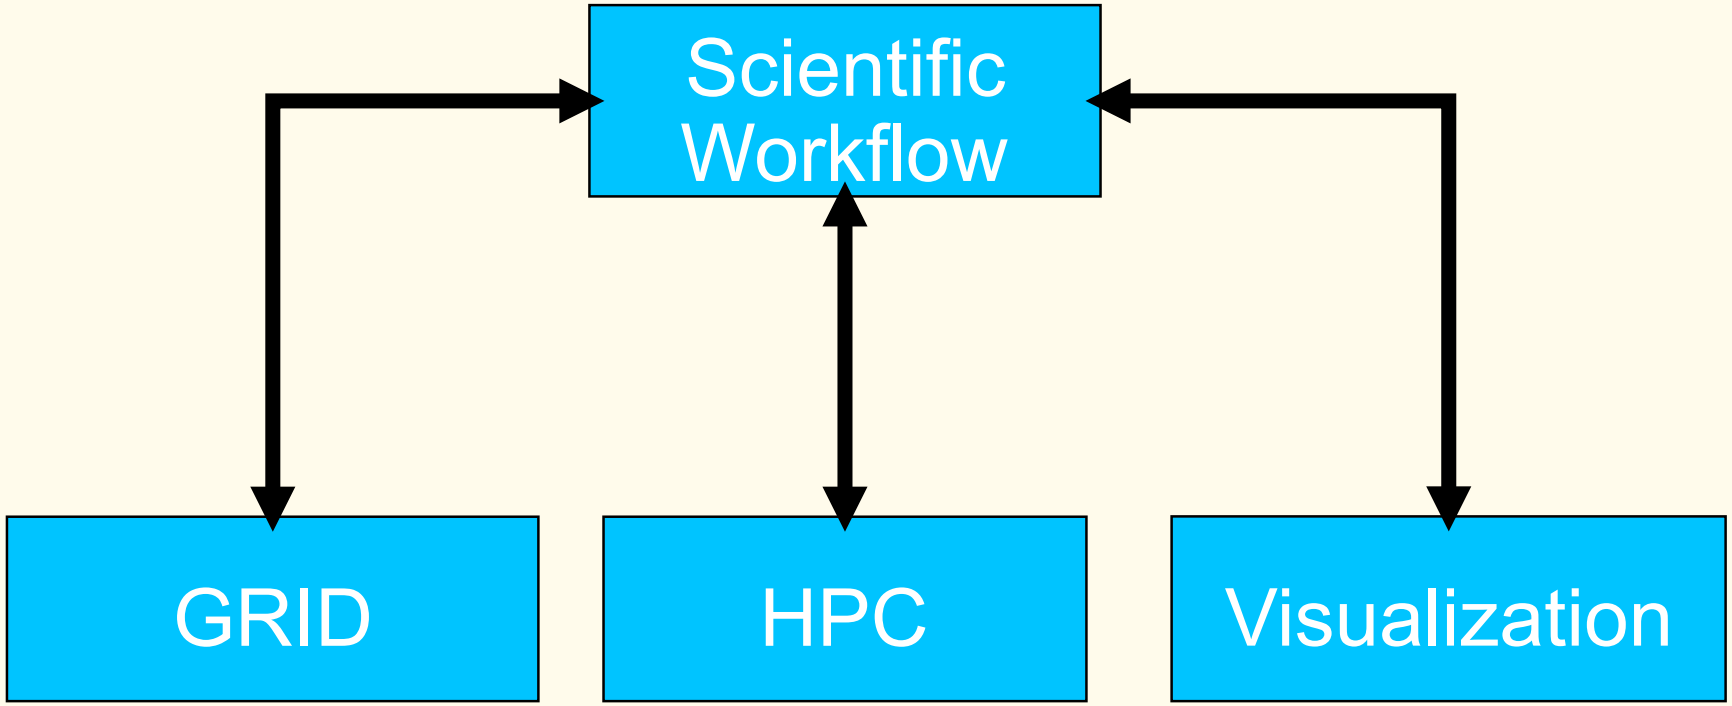

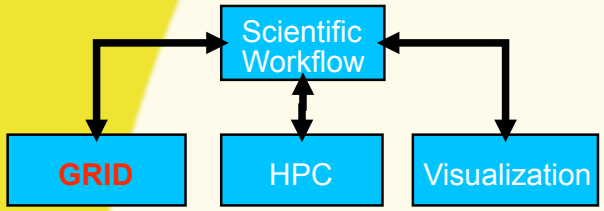

Grid Computing

- Grid computing needs
  - To be transparently coupled into the scientific workflows
  - Needs to be reliable (every launched job should run)
  - Needs to improve performance (if resources are available, a launched job should start rapidly)
  - To inter-operate with the other levels at the data access level

Grid Infrastructure

- Largely based on the developments of the *Interactive European Grid* project (http://www.i2g.eu)
  - Advanced Scheduling Capabilities (CrossBroker)
    - Support to parallel jobs with MPI (mpi-start)
    - Support to interactive jobs via i2glogin
  - RAS in combination with Migrating Desktop
- Full interoperability with other gLite based infrastructures
  - Like EGEE
- Access through the EUFORIA VO

CONSEJO SUPERIOR DE INVESTIGACIONES CIENTÍFICAS

Grid core Services
Crossbroker: MPI , interactivity

- CrossBroker: Int.EU.Grid meta-scheduler
  - Offers the same functionalities as the EGEE Resource Broker, plus:
    - Support for Interactive Applications
      - Interactive agent injection
      - Scheduling priorities;
      - Time Sharing
    - Full support for Parallel Applications
      - OpenMPI, PACX-MPI and MPICH
    - Flexible MPI job startup based on MPI-START

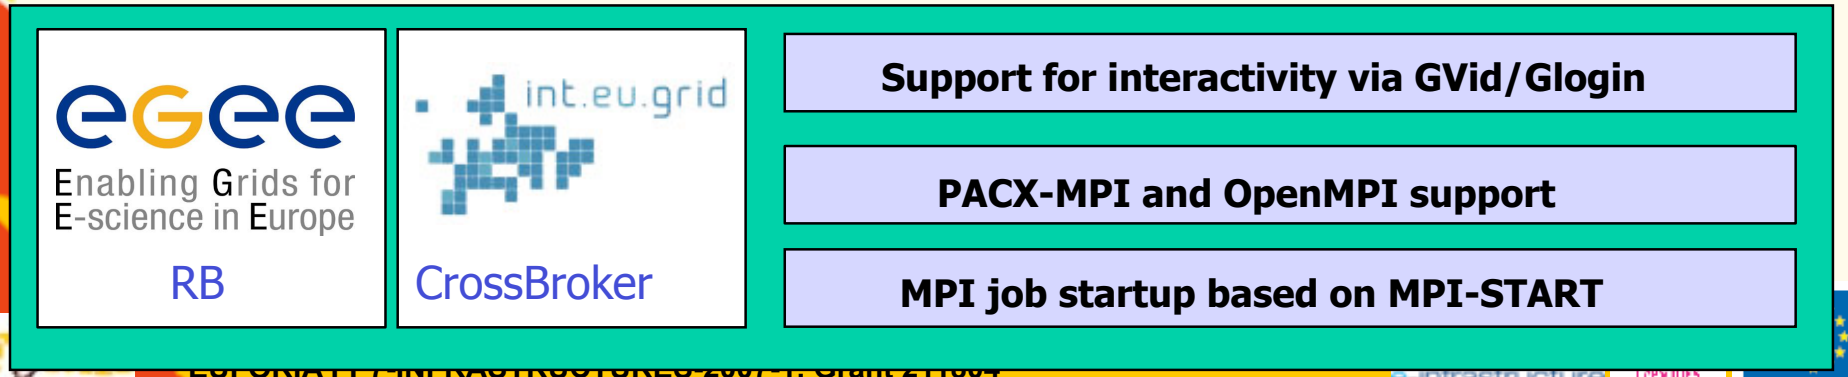

Grid Core Services
RAS / Migrating Desktop

- Roaming Access Server (RAS): web service for Grid Access
  - Used by the Migrating Desktop
  - Performs actions on the grid on behalf of the MD
  - There are two RAS installed
    - Three at PSNC for development
    - One at CSIC for production + one for testing

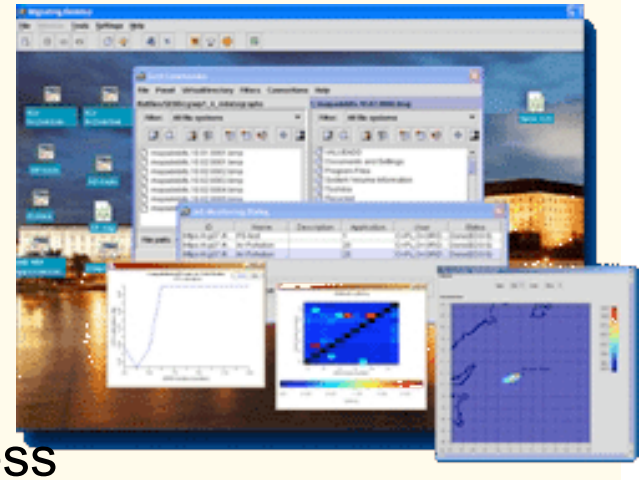

- Migrating Desktop (MD): User Friendly Grid Access
  - Java based GUI that hides the details of the grid infrastructure
  - Provides interactivity and visualization features
    - GVid enables interactivity for OpenGL and X applications
    - Visualization of graphical output

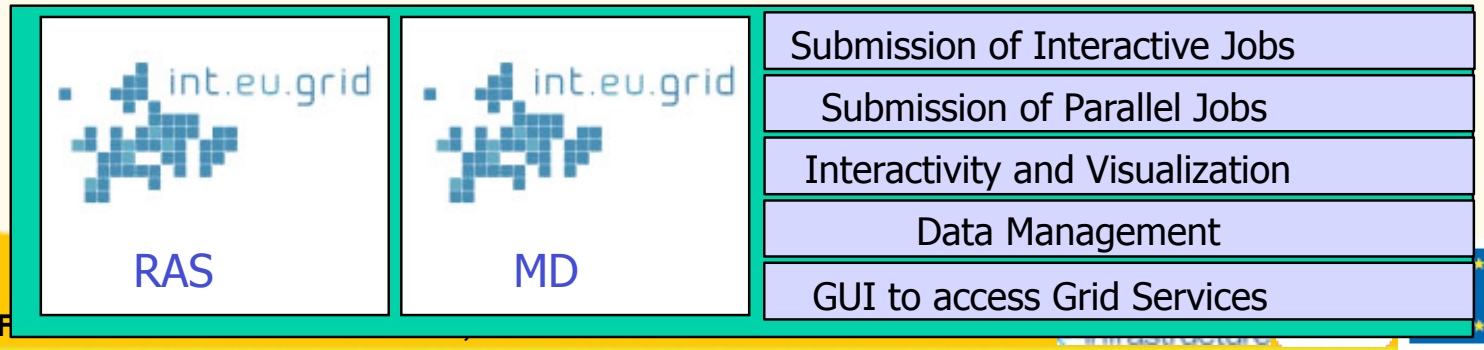

Grid Infrastructure

Core Services

i2g-lfc01.lip.pt
i2glfc01.ifca.es

LFC

VOMS

i2g-voms.lip.pt
i2gvoms01.ifca.es

Top-BDII

i2g-ii01.lip.pt

IBERGRID

ES-NGI + PT-NGI

Local Services

I2G CrossBroker

i2grb01.ifca.es
i2g-rb01.lip.pt

Migrating Desktop

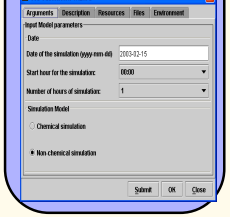

I2G UI

Site-BDII

MonBox

LCG-CE

I2G CE software

I2G WN software
gLite WN

I2G WN software
gLite WN

I2G WN software
gLite WN

Batch Server

SE

Grid Infrastructure

- Centers with infrastructure in Euforia
  - IFCA in Santander, Spain
  - FZK in Karlsruhe, Germany
  - Chalmers University, Sweden
  - Ciemat in Trujillo, Spain
- From the Grid Information System (~2700 CPUs, 40 TB online)

```
#CPU|Free|Total|Jobs|Run|Waiting|ComputingElement
----------------------------------------------------------------------
1488  702    0     0     0    0 i2gce01.ifca.es:2119/jobmanager-lcgpbs-euforia
452   324    0     0     0    0 iwrce2.fzk.de:2119/jobmanager-lcgpbs-i2gpar
640   113    0     0     0    0 svea-gl2.c3se.chalmers.se:2119/jobmanager-lcgpbs-euforia
56    56     0     0     0    0 ce-euforia.ceta-ciemat.es:2119/jobmanager-lcgpbs-euforia

Avail Space(Kb) Used Space(Kb)  Type    SEs
----------------------------------------------------------
38  TB      6TB               n.a     storm.ifca.es
1.8 TB      129931            n.a     iwrse2.fzk.de
142 GB      --                n.a.    svea-gl3.c3se.chalmers.se
```

Grid Infrastructure: towards EGI

- Central services are guaranteed by IBERGRID (LIP and CSIC)
- Main middleware pieces have a continuity path in EGI
  - Migrating Desktop and RAS (PSNC)
  - CrossBroker and MPI support (CSIC)
- Infrastructure ifself is already a part of the NGIs resources
  - CSIC and Ciemat in Spain
  - FZK in Germany
  - Chalmers in Sweden

User Oriented Services

- Euforia VO
  - VOMS server is integrated in the i2g infrastructure
  - ~30 people are inscribed
- Wiki pages and

  User Documentation

http://wiki.eu-euforia.eu

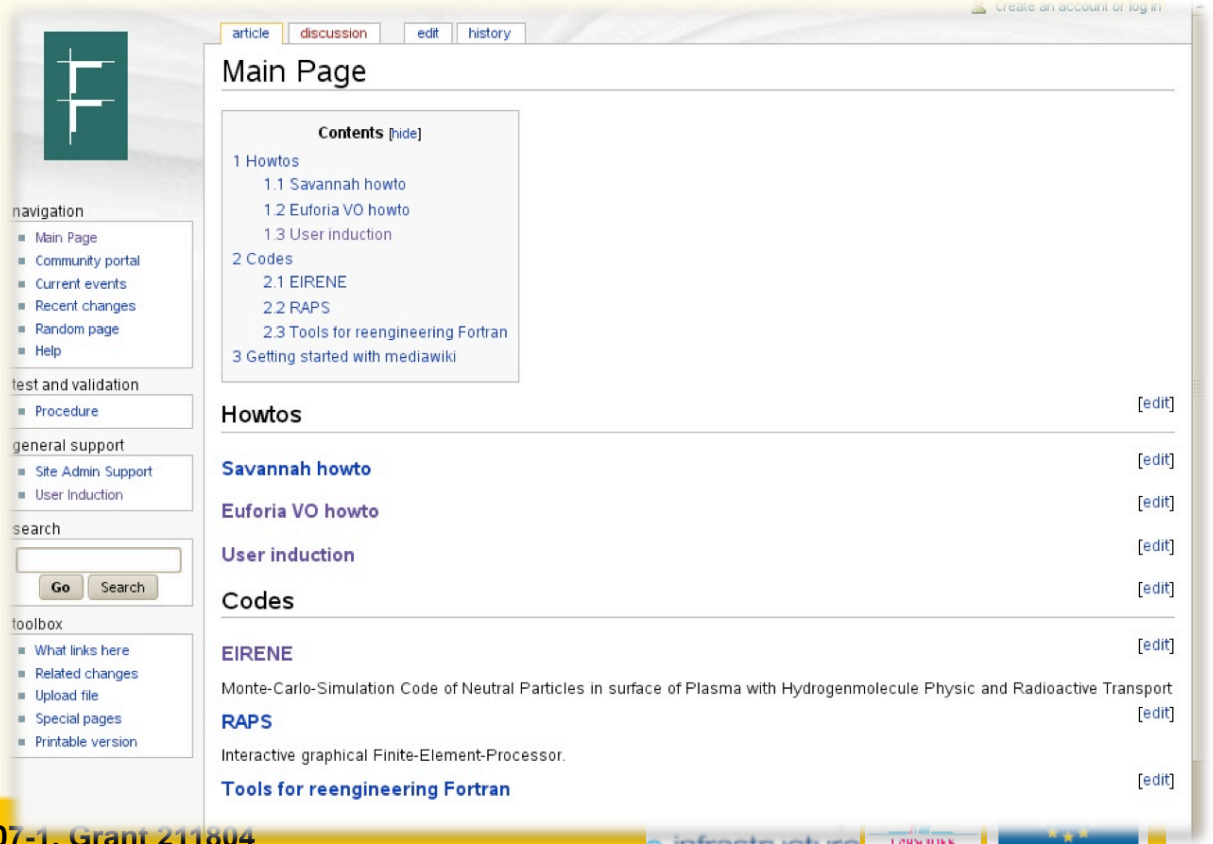

User Oriented Services

- FZK runs a savannah server
  - http://savannah.fzk.de/projects/euforia
  - Grid middleware repository
- The server also includes software repositories for the fusion codes
  - Savannah Projects associated with codes
  - E.g. http://savannah.fzk.de/projects/eirene
  - Access via SVN to code repository
- Compilation service provided by Autobuild
  - Middleware and appliations
  - http://savannah.fzk.de/autobuild/euforia

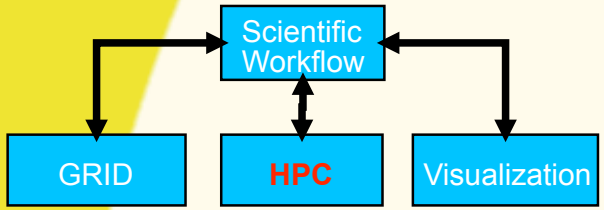

HPC

- Need to be transparently coupled into the workflows
- HPC facilities will need to deal with jobs coming from workflows
  - Negotiations about resource availability and expected response times
  - Deal with communications between the different parts of a workflow
  - Better integrate inter-operability of data between HPC and the other levels

HPC Infrastructure

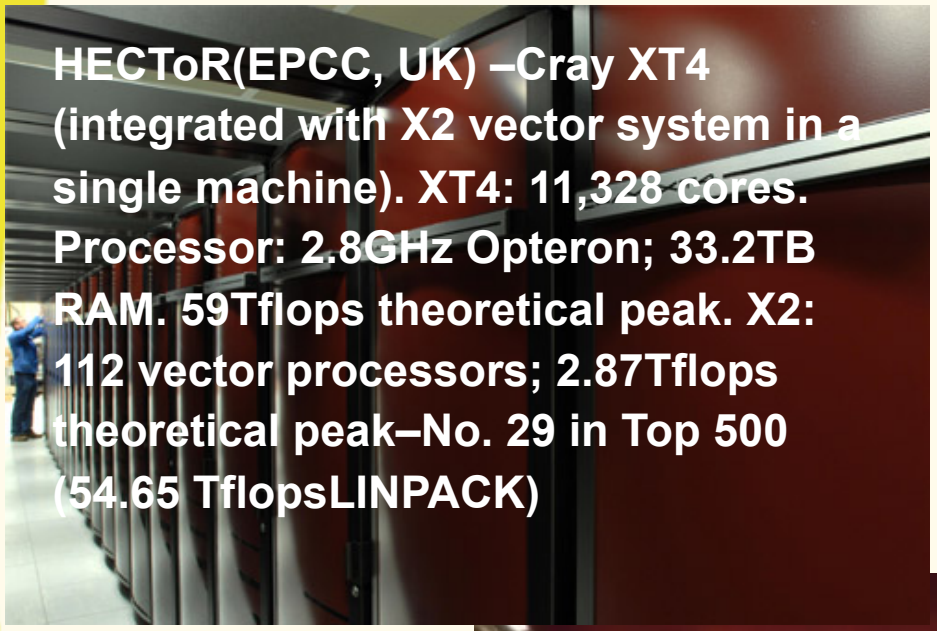

HECToR(EPCC, UK) –Cray XT4 (integrated with X2 vector system in a single machine). XT4: 11,328 cores. Processor: 2.8GHz Opteron; 33.2TB RAM. 59Tflops theoretical peak. X2: 112 vector processors; 2.87Tflops theoretical peak–No. 29 in Top 500 (54.65 TflopsLINPACK)

Louhi(CSC, Finland)–Cray XT4. 4,048 cores. 4.5TB RAM. 37.68Tflops peak.–No. 70 in Top 500 (26.80 TflopsLINPACK)

MareNostrum(BSC, Spain) –IBM Cluster. 10240 cores. Processor: 2.3GHz PPC 970; 20TB. 94.21Tflops peak. –No.26 in Top 500 (63.83 TflopsLINPACK)

HPC Infrastructure

- HPC ecology more diverse than Grid
- Codes optimized for running in the facilities
  - Installation/compilation and profiling done until June 2008
  - Now users can execute their applications (2M CPU Hours available)
- Access through DEISA

Codes	Grid	Ported to (centre)	Libraries needed
CENTORI	No	Cray XT4 (EPCC)	---
GEM	Yes	Cray XT4 (EPCC-CSC) PPC970 (BSC)	FFTW 3.1.1
ERO	No	Cray XT4 (EPCC)	IBM ESSL and PMAPI
EIRENE	Yes	PPC970 (BSC)	---
ELMFIRE	Yes	PPC970 (BSC) Cray XT4 (BSC)	GSL, BLAS, PETSC, PESSL
GENE	Yes	PPC970 (BSC)	ESSL, MKL or FFTW BLAS/LAPACK/ PETSC/SLEPC OpenMP
BIT1	Yes	PPC970 (BSC) Cray XT4 (CSC)	X11
Esel	No	CSC	HDF4(version 4.2r2)
Isdep	Yes	CSC	---
SOLPS	No	CSC	NCARG (version 5.0.0) NetCDF (version 3.6.2) ACML (version 4-0-1) (or some other BLAS/ LAPACK library) X11

HPC Infrastructure

- Unicore interface for Altamira (HPC facility at IFCA)
  - Access with User certificates
  - Job execution
  - File staging

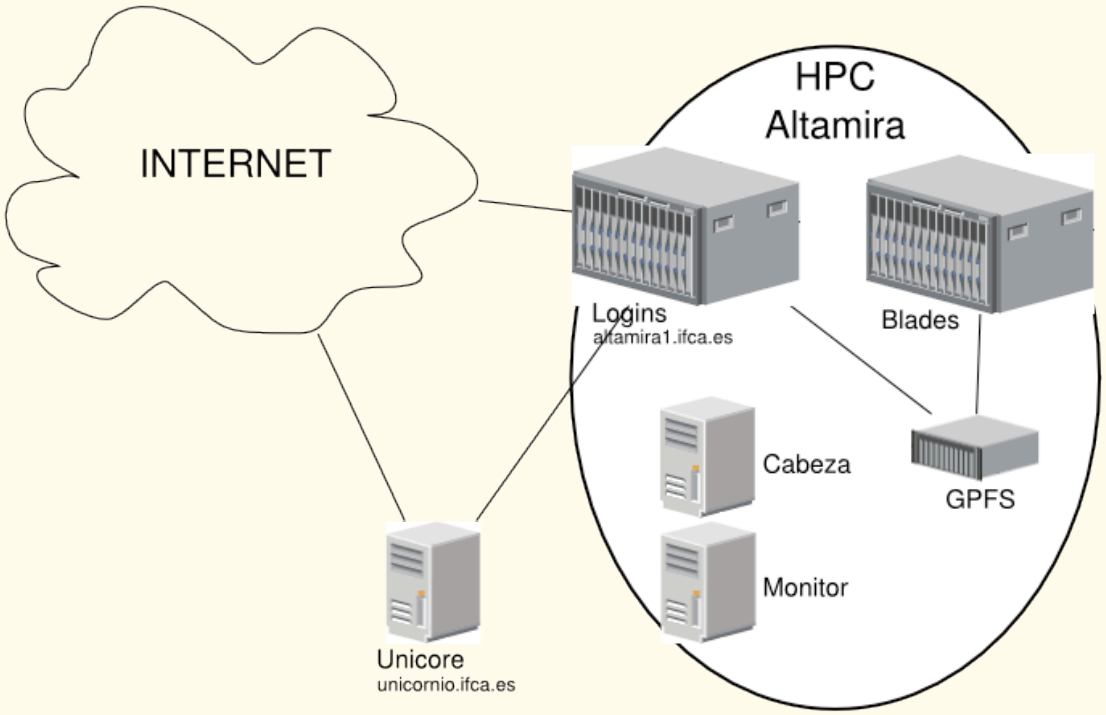

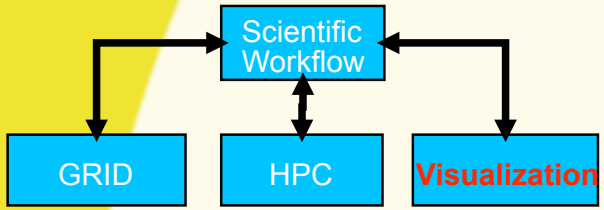

Visualization

- Challenges in visualization include
  - Visualizing the data flowing around a workflow so that the scientist can monitor/diagnose a running job
  - Deal with very large amounts of data produced in a distributed environment
  - Help provide the scientist with a better understanding of his/her results
  - Help the scientist by producing visualizations with that “Wow!” factor

Visualization

- Development of post-processing data analysis and visualization tools based on Python, NumPy, and Matplotlib packages.
- Development of post-processing data analysis and visualization tools based on VisIt.
- Integration of Python, NumPy, Matplotlib into the Kepler workflow.
- Integration of VisIt or VTK into the Kepler workflow.
- Development of a lossy wavelet based compressed data format.
- Develop visualization techniques for 4D and 5D particle distributions.

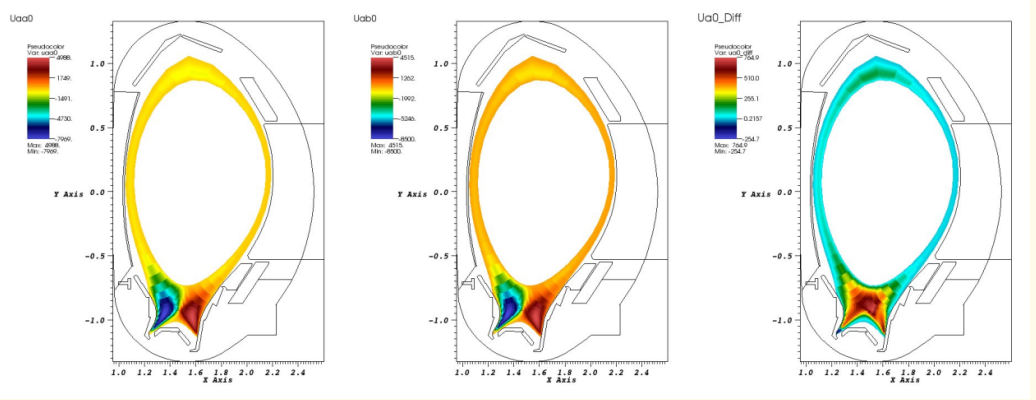

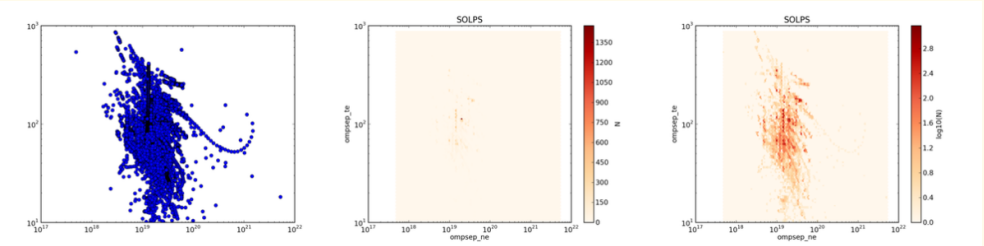

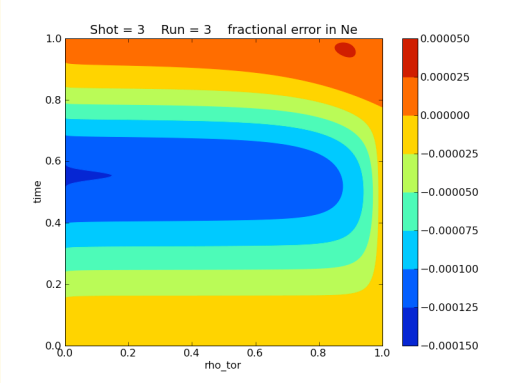

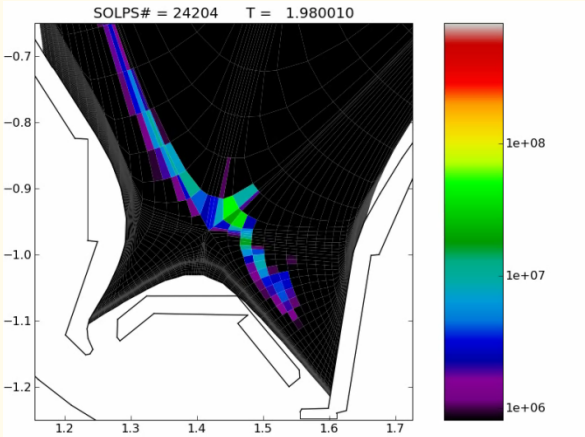

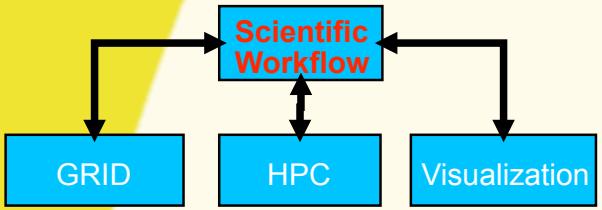

Scientific Workflows

- End-to-end simulation of a fusion device requires integration and coupling of several complex codes
- Scientific Workflows:
  - Enhance productivity
  - Allow for new approaches problems to be solved
  - Allow for traceability, reproducibility, …
  - Allow for a better use of resources

Scientific Workflows

The objective is to schedule jobs on the GRID and HPC infrastructures together with jobs running on other computing facilities.

- launch and control jobs in a transparent manner for the users.
- data communication: data transfers and visualizations are required at run time and to access the experimental and simulated data.
- middleware connectivity for Kepler Workflow (java Library for gLite, HPC)
- scheduler notification tools
- Integration with ITM tools (UAL – Universal Access Layer

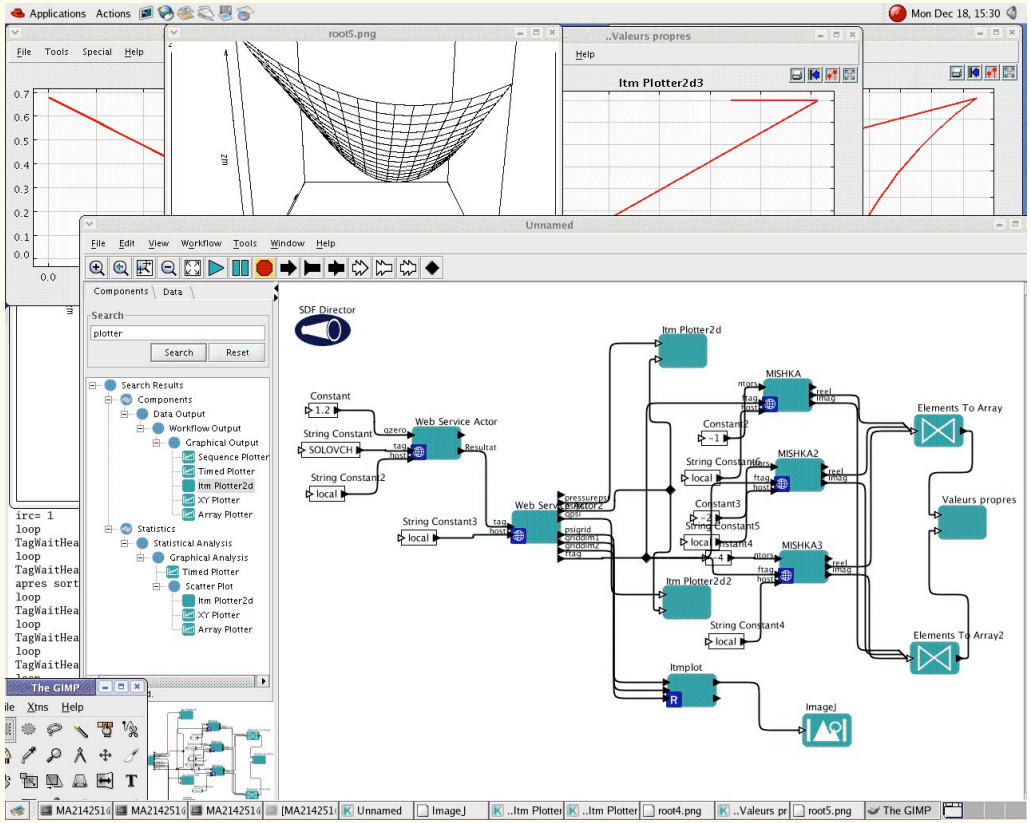

Scientific Workflows

- EUFORIA integrates those codes and existing tools as Keplerworklflows:
  - Kepler already used by the fusion community
  - Workflows seen as a set of actors with I/O ports for communication
  - Actors can be nested (hierarchical workflows)
- New actors developed in EUFORIA
  - Submission to i2g and gLite
  - Submission to HPC (via Unicore)

Scientific Workflows

- Workflows started on a central facility
- Computation done elsewhere (mainly)
  - On the GRID
    - Serial and small scale parallel
    - For scenario scans
  - On an HPC
    - For larger runs
    - Chosen by the workflow on the basis of
      - Suitability
      - Minimum cost
      - Minimum time to completion
- Results (or a “reference”) are brought back to the central facility for permanent storage and further analysis

Scientific Workflows

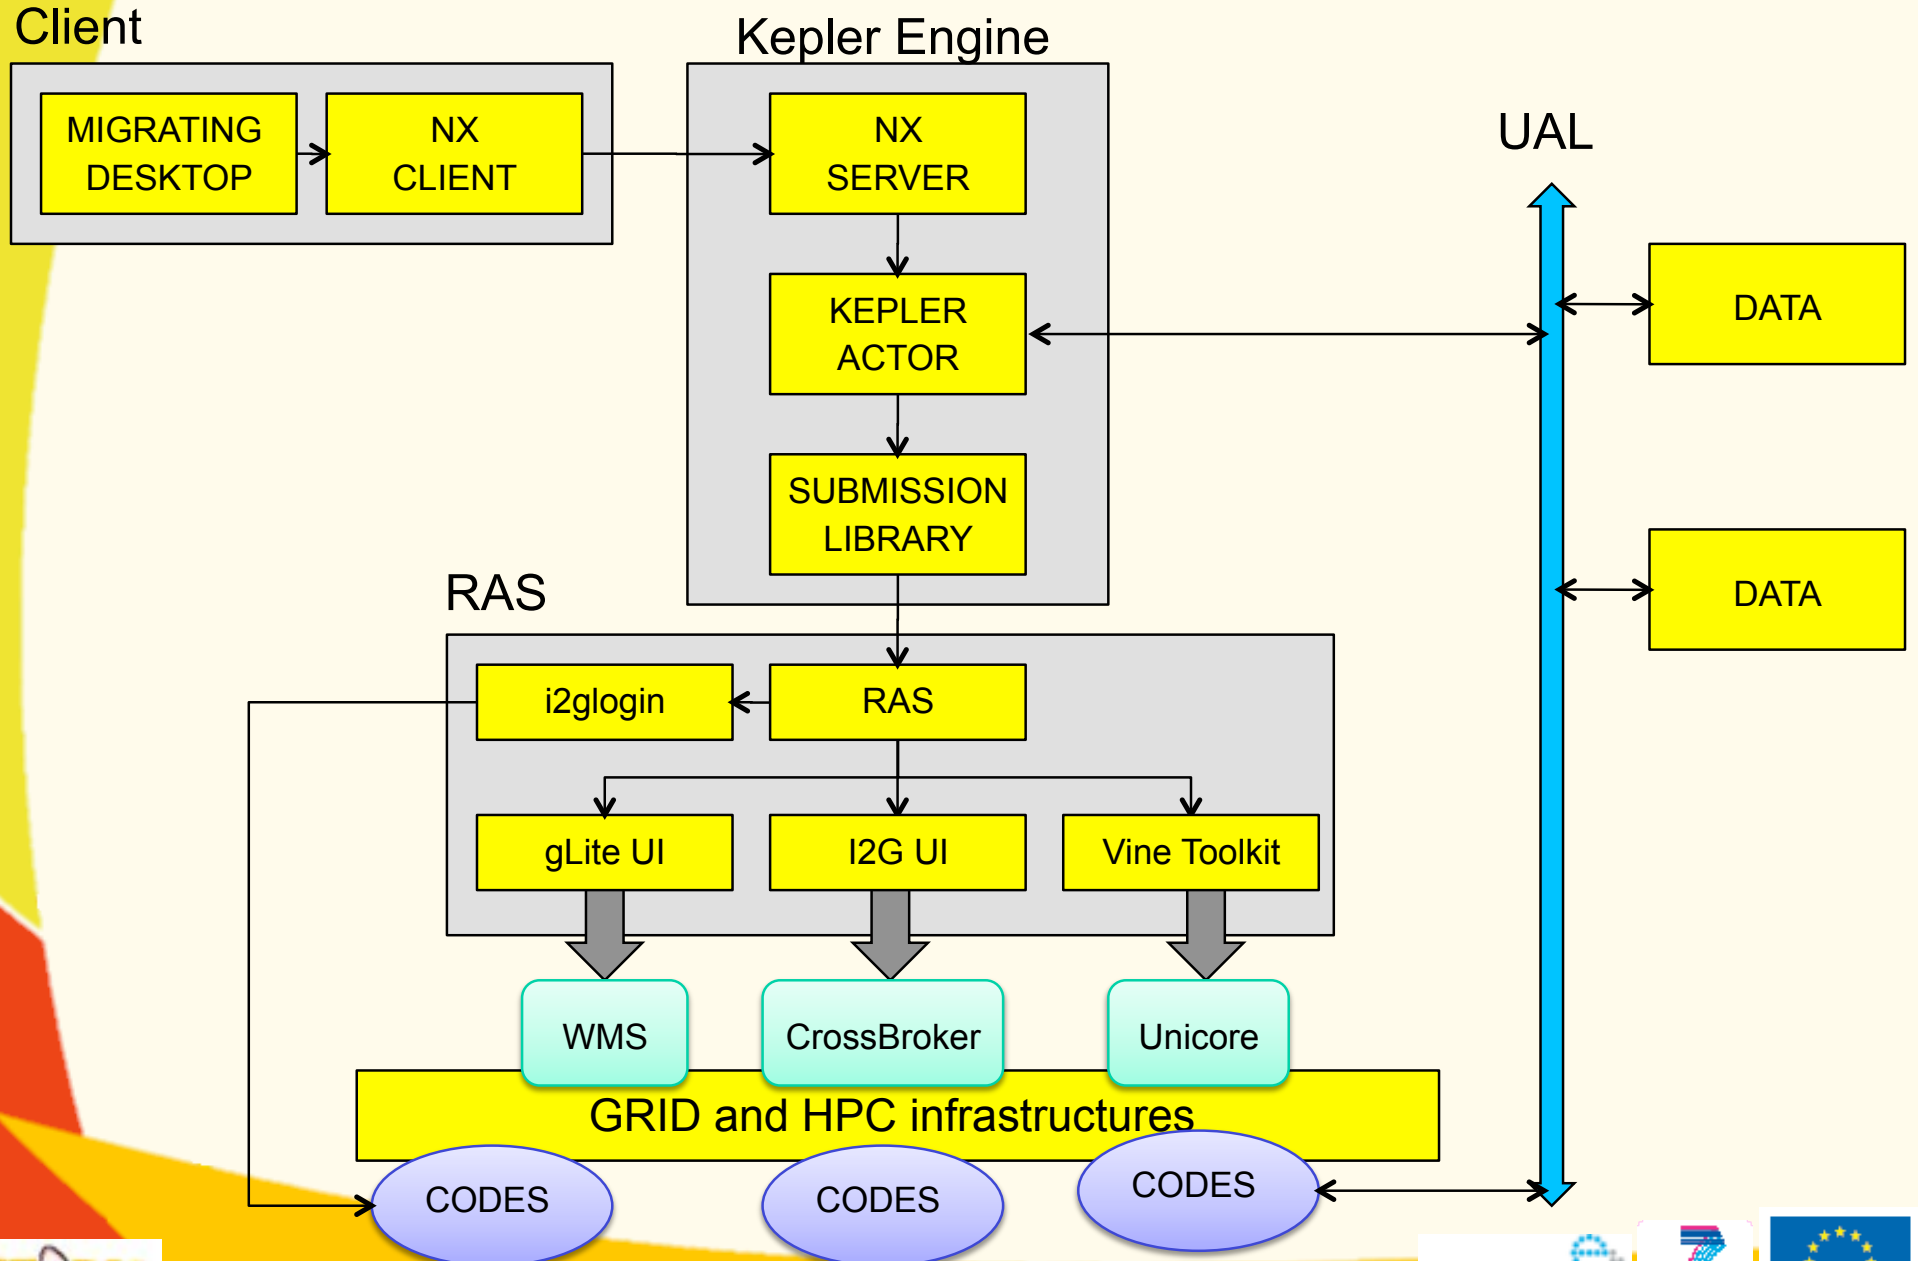

Scientific Workflows

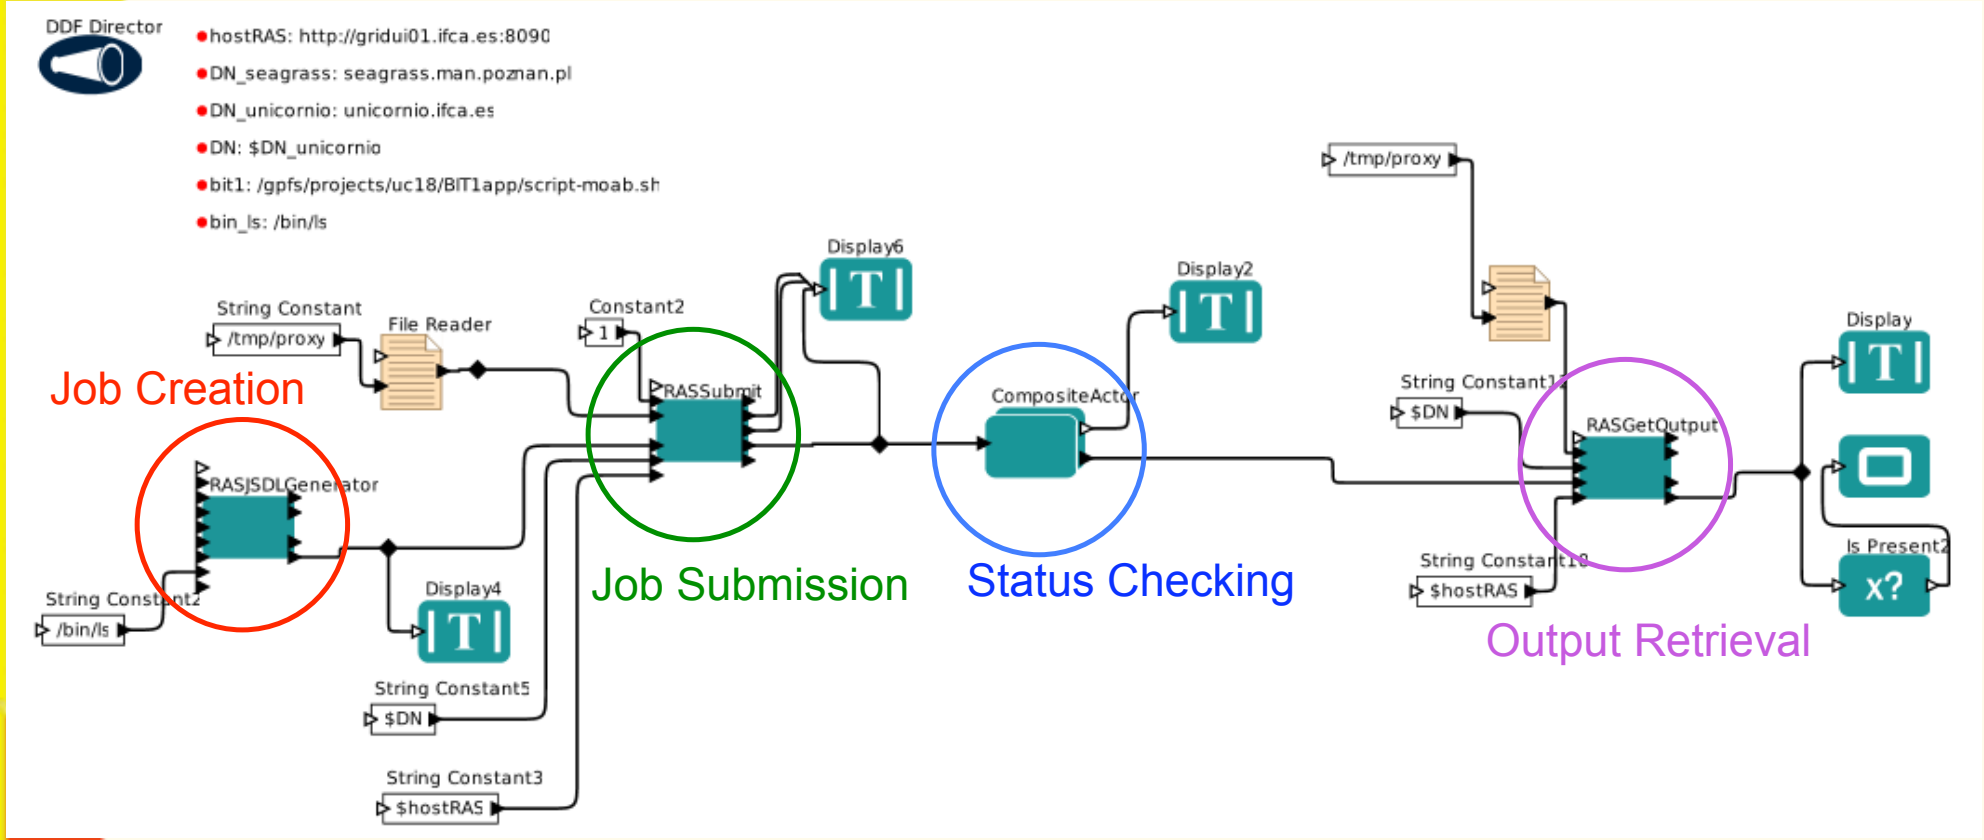

Conclusions

- EUFORIA develops an integrated environment for fusion modelling
  - Grid and HPC infrastructure
  - Use of existing tools from i2g and fusion community
    - RAS/Migrating Desktop, CrossBroker, Vine, Unicore
    - Kepler, UAL
- Scientists able to perform end-to-end simulations of fusion devices with workflows
  - Executing codes in grid and HPC