


SA1 Grid Infrastructure Report

Marcus Hardt (KIT) | March 30, 2011

ON BEHALF OF SA1

- CSIC, Santander, Spain
- KIT , Karlsruhe, Germany
- CIEMAT, Madrid, Spain
- Chalmers, Gothenburg, Sweden

- Interaction with other parts of Euforia



From the Description of Work

- **Deploy, maintain and operate** the central services necessary to **ensure the integration of the computing resources** into a production Grid infrastructure capable of supporting serial and parallel applications.
- Coordinate the **deployment, maintenance and operation of the Grid resources** provided by the participating sites ensuring a robust, secure and dependable service.
- Provide **support for Virtual Organizations and resource providers** thus helping end-users, VO managers and site managers to achieve their goals and **contribute to a successful usage of the infrastructure.**

Further objectives

- Include **more than 500 processors** in high performance clusters
- Include low **latency clusters** to support applications benefiting from **distributed parallel processing** (in particular **MPI**).
- **Coordinate efforts with the EGEE sites** supporting the **Fusion VO** in order to have an **integrated, compatible infrastructure**, to support the use of the Grid for fusion applications.
- **Grid Appliance**, which will be useful when spreading grid technologies to new communities.

Deliverable	Title	Date	Status
DSA1.1	Deployment of testbed and operational infrastructure	M6	on time
DSA1.2	Grid Infrastructure Status Report	M18	on time
DSA1.3	Deployment of Grid appliances	M24	on time
DSA1.4	Final report	M36	on time
Additional	Cloud Pilot Report	M36	on time

Sites and resources

- All sites used the int.eu.grid extensions to gLite middleware
- Chalmers used NorduGrid with adaptations

#CPUs	Storage[TB]	Site
1488	38	CSIC / IFCA
640	1	Chalmers
540	1.8	KIT
512	28	CIEMAT
56	–	Ceta-CIEMAT (offline after 2009)
3236	68.8	Sum



- All sites supported **EUFORIA** and most **FUSION** Virtual Organisations
- All sites supported MPI parallel jobs
- CSIC/IFCA offered Infiniband interconnect

Utilisation of grid resources

- CPU time only from **NGI.cesga.es**


	NormSumCPU[kh]		
VO name	2008	2009	2010
EUFORIA	7.6	149	238
Fusion	0.15	10	157

- CPU time only from **EGEE.cesga.es**:

	NormSumCPU[kh]		
VO name	2008	2009	2010
EUFORIA	0	38	6
Fusion	1293	1509	1427

- CPU time summing up both from EGEE.cesga.es and NGI.cesga.es

	NormSumCPU[kh]		
VO name	2008	2009	2010
EUFORIA	7.6	197	244
Fusion	1293	1519	1584



CPU time from NGI.cesga.es

Central services and developer tools

Central services (gLite)

Service	CSIC	KIT	PSNC	LIP*	Other
WMS	x	x		x	
CrossBroker	x	x		x	
BDII	x	x			
MyProxy	x			x	
LFC	x	x		x	
User Interface	x	x			ENEA*
RAS	x		x		
MD	x		x		
VOMS	x			x	
Accounting					CESGA*
Monitoring	x				

Developer Tools

Savannah server		x			
Project Wiki		x			
Internal Wiki		x			
Autobuild		x			



* Acknowledgements for hosting services outside the SA1 or EUFORIA effort

EUFORIA FP7-INFRASTRUCTURES-2007-1 Grant 211804



Savannah: For all developers within the fusion community and the project

Web savannah.fzk.de/search/ Search with Google

Savannah

Logged in as **marcus**

- Become Superuser
- My Incoming Items
- My Items
- My Groups
- My Account Conf
- Logout

This Page

- Clean Reload
- Printer Version

Search

in Projects Search

Site Administration


- Main page
- Pending projects
- Site news approval

Hosted Projects

- Register New Project
- Full List
- Contributors Wanted
- Statistics

Site Help

- User Docs: Cookbook
- User Docs: In Depth Guide
- Get Support
- Contact Us

 **Search**

Site Wide

(+) Search Criteria:

Search results for "*" (in Project/Group):

Project	Description	Type
BIT1		EUFORIA
CENTORI		EUFORIA
EMC3-Workflow	The goal ois to simplify extensive parameter studies on the Grid with EMC3-EIRENE.	EUFORIA
EIRENE		EUFORIA
elmfire		EUFORIA
EMC3-eirene		EUFORIA
ERO		EUFORIA
Esel		EUFORIA
European Transport Solver		EUFORIA
Euforia		EUFORIA
Euforia Internal		EUFORIA
Euforia Visualization Integration (JRA4)		EUFORIA
GEM		EUFORIA
GENE		EUFORIA
HELENA		EUFORIA
ISDEP		EUFORIA
model1		EUFORIA
SOLPS		EUFORIA
Tool for Automatic Parameter Studies - TAPaS4grid	This is a Tool for Automatic Parameter Studies on the grid. Adapted on the case of EMC3-EIRENE	EUFORIA
Tyr		EUFORIA

- Tool for automated compilations
 - Goal: **Improve general quality of software**
 - Hourly builds of code from SVN (e.g. in Savannah)
 - Build far away from developers machine
 - Publish output on webpage (<http://savannah.fzk.de/autobuild>)
- ⇒ Newest code automatically available at a central point

Overall status: Good
Date: Tue Mar 8 2011
Start Time: 11:41:01 UTC 12:41:01 CET
Build Timestamp: 1299584461

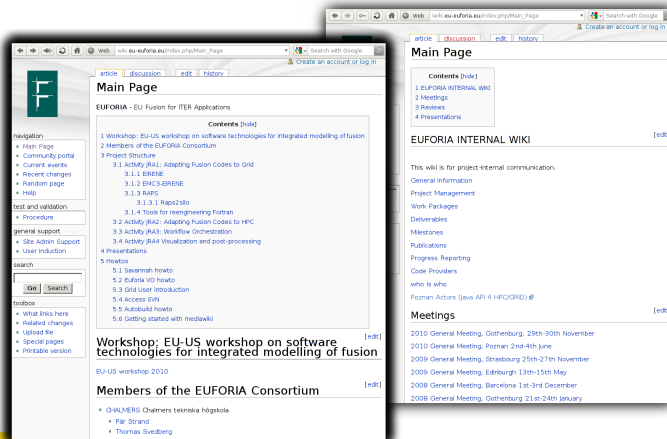
EUFORIA -- Continuous Automatic Builder



Software

EIRENE-Grid-parallel	success
EIRENE-Grid-sequential	success
EIRENE-data	success
EIRENE-trilinhex	success
EMC3-EIRENE	success

- MediaWiki
- Meetings, iteration of deliverables, People
- Documentation, Links




- Success story
 - s.t.a.r.t.: Built because it was required by communities
 - Tapas4Grid: running parameter scans for EMC3-Eirene
 - Well accepted by users
 - Users improved it, and created TAPAS
 - In use for several codes (Bit1, Bee-Algorithms, ...)
 - Parametric jobs were not mentioned in DoW, but required by users

⇒ **Parametric established in the Kepler Workflow Engine**

⇒ **All tools available for use in Fusion VO**

- Interactive Grid Tools (inherited from Int.EU.Grid)
 - MPI-Start: A universal way to start MPI programs on the grid
 - G-Login: Enabling interactivity on the grid
 - CrossBroker: Advanced interactivity and priority on the grid


⇒ **Tools established in the EGI context, available, e.g. in Fusion VO**



- Impact on Fusion researchers:
 - The required tools (as defined in Euforia) are now available on Grid infrastructure (Fusion-VO)
 - Interactive, parallel Jobs are now available to users
 - Cloud prototype is accessible and will be pursued to add features and resources

⇒ **22 Publications** written, using the grid

- In cooperation with and supporting JRA1 and JRA3:
 - Code Porting (supporting JRA1)
 - EIRENE, EMC3–EIRENE, Bit1, Bee Algorithms, Helena
 - Organisation of Special Sessions for "Grid and HPC for Nuclear Fusion Applications" at EuroMirco PDP 2010 & 2011
 - 8 of 18 papers from outside Euforia
 - Integration of Infrastructures into the Workflow Environment
 - Pleasant meetings outside the core project activities
- In cooperation with NA2
 - Participation in two GridKa Summerschools
 - General Talk about EUFORIA
 - Hands on tutorial on grid usage
 - In conjunction with JRA1 for training on Kepler Workflow Integration



Additional activity: Cloud Pilot [1/2]

- Recommendations of the 2nd review:
 - “Try out a commercial cloud system” + “Provide a note on lessons learned”
⇒ Resources shifted from Chalmers and EPCC to PSNC and KIT
- ⇒ setup the SA1 lead **Cloud Pilot project**
- Cloud interest/expertise found at KIT, PSNC, Chalmers and EPCC
- Technical considerations
 - Interface: Amazon vs. open Standards (OCCI)
⇒ EC2-interface but only the subset supported in (EC-funded) OpenNebula-2.0
 - Provider: Commercial (as requested) vs. Scientific (self-interest)
⇒ Both, ONE and AWS
 - Virtual Machine: User provided vs. ready template
⇒ Template VM (with gLite, in case needed)

Additional activity: Cloud Pilot [2/2]

- Architecture
 - Defined together with JRA3
 - Kepler Integration
 - ⇒ Seamless integration of scientific / commercial cloud, grid and HPC
- Benchmarking of the solution
 - In cooperation with JRA2
 - ⇒ Amazon HPC cluster is comparable to scientific HPC Cluster
- Cost Analysis
 - Comparison of Amazon with a typical LHC Tier1 computer centre
 - ⇒ **Amazon is more than twice as expensive as an LHC Tier1 centre**

Additional activity: Cloud Pilot [2/2]

- Architecture
 - Defined together with JRA3
 - Kepler Integration
 - ⇒ Seamless integration of scientific / commercial cloud, grid and HPC
- Benchmarking of the solution
 - In cooperation with JRA2
 - ⇒ Amazon HPC cluster is comparable to scientific HPC Cluster
- Cost Analysis
 - Comparison of Amazon with a typical LHC Tier1 computer centre
 - ⇒ **Amazon is more than twice as expensive as an LHC Tier1 centre**



Search for **euforia cloud**
About 696,000 results (0.33 seconds) Advanced search

Everything

Images

Videos

News

Shopping

More

Cloud Forest Expedition - EUFORIA EXPEDITIONS

The **Cloud** Forest Expedition is an adrenaline filled 24 hour adventure race in the pristine mountaintops of Monteverde, Costa Rica. **Euforia** is known for its ...
www.euforiaexpeditions.com/index.php?option=com...id... - Cached

[PDF] EUFORIA

File Format: PDF/Adobe Acrobat - Quick View
Throughout the project, the expectations of the **EUFORIA** users ...
www.cimat.es/EUFORIA/recursos/doc/.../841040751_103201114312.pdf

- SA1 Objectives
 - Deploy and maintain grid Central Services
 - Deploy and maintain grid sites
 - Provide support for users and VOs
 - Provide support for scientists and developers
 - Accounting, Monitoring
 - Coordinate efforts with EGEE
 - Grid Appliance
- SA1 additional activity
 - Cloud
 - Evaluation
 - Benchmarking
 - Cost analysis

- SA1 Objectives
 - Deploy and maintain grid Central Services
 - Deploy and maintain grid sites
 - Provide support for users and VOs
 - Provide support for scientists and developers
 - Accounting, Monitoring
 - Coordinate efforts with EGEE
 - Grid Appliance
- SA1 additional activity
 - Cloud
 - Evaluation
 - Benchmarking
 - Cost analysis
 - **Publication about the cloud pilot is underway**


Question	Result
Users responding to questionnaire	13 (\approx 95% usage)
User not involved in EUFORIA	30%
Users that used MPI on the Grid	30%
Users that used EUFORIA and Fusion VO	46%
Number of fusion codes that were run on the grid	15
Intended use of the grid: Development / Testing / Production	84% / 92% / 46%
Users that expected scientific results by running on the grid	38%
Publications written by using the grid	22
Access to grid via: gLite-UI / Kepler	84% / 54%
Assistance required:	62%
Happy with the support received	100%
Problems with the grid	23%

Conclusions

- We have provided users and developers with a wealth of services
- We have collaborated with JRA1 and JRA3 to support smooth integration into higher level tools
- The cloud pilot proved the flexibility of the Kepler Workflow Integration
- Use of cloud will be pursued further

- Soft valuers:
 - Work was great fun
 - New collaboration partners found
 - Friends made

- CrossBroker: Crossgrid / Int.EU.Grid meta-scheduler
⇒ Submitted into the EGI project
 - Ported to 64 Bit
 - Offers same functionalities as the EGEE WMS plus:
 - Support for Interactive Applications
 - Interactive agent injection
 - Scheduling priorities
 - Time sharing
 - Full support for **parallel Applications**
 - OpenMPI, PACX-MPI and MPIC
 - Flexible MPI job startup based on MPI-START



- s.t.a.r.t.
 - New commandline interface for gLite
 - Driven by users and admins
 - Usage:

```
submitter --numjobs 10 --openmpi-np 128 --  
--input lfn/grid-vo/data.tar.gz  
--input http://user.org/cfg.zip  
--software http://autobuild.org/32bit/software-1.0.tar.gz  
--libs http://thirdparty.com/libs/mathlib.tar.bz2  
--prepare ftp://user.de/prep.sh  
--      my-application -option=value ...
```

- TAPAS4Grid – Tools for Automated Parameter Scans
 - Uses s.t.a.r.t. to conduct parameter scans
 - In use by several users
 - Support JRA1 with grid access layer for HPC / Unicore on Altamira at CSIC
 - Grid Appliance
 - A set of virtual machines
 - For quickly setting up new grid sites
- ⇒ Worked well, however, large amount of handy work required, due to nature of gLite and **YAIM**