EGI-InSPIRE

ANNUAL REPORT ON THE TOOLS AND SERVICES OF THE HEAVY USER COMMUNITIES

EU DELIVERABLE: D6.6

Document identifier:	EGI-D6.6-final
Date:	**12/03/2012**
Activity:	**SA3**
Lead Partner:	**CERN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/742

Abstract

This document is the Annual Report on the HUC Tools and Services: a public report describing the current status of the tools and services provided within SA3, their activity over the past year, and the future activities including measures to assure sustainability of the current work.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Jamie Shiers et al.	CERN/SA3	21/02/2012
Reviewed by	Moderator: Steve Brewer Reviewers: Martin Veicherts, Sergio Maffioletti	EGI/NA2 UCPH/NA2 SWITCH/NA2	25/04/2012
Approved by	AMB		27/04/2012

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	21/02/2012	First draft	Jamie Shiers et al.
2	07/03/2012	Second draft	Jamie Shiers et al.
3	25/04/2012	Final version	Jamie Shiers et al.

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE "Document Management Procedure" will be followed: https://wiki.egi.eu/wiki/Procedures

VI. TERMINOLOGY

A complete project glossary is provided at the following page: http://www.egi.eu/about/glossary/.

VII. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed − both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting 'Grids' of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop Grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the European Strategy Forum on Research Infrastructures (ESFRI) projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

- The continued operation and expansion of today's production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
- The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
- The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
- Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
- Mechanisms to integrate existing infrastructure providers in Europe and around the world into a production infrastructure, so as to provide transparent access to all authorised users.
- Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop Grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.

The production infrastructure supports Virtual Research Communities (VRCs) − structured international user communities − that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

Table of Contents

1 INTRODUCTION

This document is the Annual Report on the HUC Tools and Services for Project Year 2 (PY2[1]) of the SA3 (WP6) work package of the EGI-InSPIRE Project. It is a public report that describes the current status of the tools and services provided within SA3, their activity over the preceding year, and the future activities including measures to assure sustainability of the current work.

Additional information, positioning this work in the context of EGI-InSPIRE as a whole, can be found in the project quarterly reports, MS112 [MS112], MS113 [MS113] and MS114 [MS114] for project quarters 5 – 7 respectively.

[1] Project Year 2 runs from May 2011 until April 2012 inclusive.

2 HEAVY USER COMMUNITIES

Heavy User Communities (HUCs) are Virtual Research Communities (VRCs) that have been using EGEE and now EGI routinely and thus have become more structured and advanced in terms of their Grid usage.

The communities identified as HUCs within this work package are:

- High Energy Physics (HEP)
- Life Sciences (LS)
- Astronomy and Astrophysics (A&A)
- Computational Chemistry and Materials Sciences and Technologies (CCMST)
- Earth Sciences (ES)
- Fusion (F)

Community	Description, capabilities and services offered
All HUCs (TSA3.2)	This task provides support for tools and services that are used by more than one community, including Dashboards, applications such as Ganga, Services such as HYDRA and GRelC, Workflows and Schedulers (SOMA2, Kepler, Taverna) and MPI.
High Energy Physics (TSA3.3) [MS611]	The High Energy Physics (HEP) HUC represents the 4 LHC experiments at CERN, which are fully relying on the use of grid computing for their offline data distribution, processing and analysis. Increasing focus is placed on common tools and solutions across these four large communities together with their re-use by other HEP experiments as well as numerous different disciplines and projects.
Life Sciences (TSA3.4) [MS611]	The Life Science (LS) HUC originates from the use of grid technology in the medical, biomedical and bioinformatics sectors in order to connect worldwide laboratories, share resources and ease the access to data in a secure and confidential way through health-grids.
Astronomy and Astrophysics (TSA3.5)	The A&A HUC is devoted to the evaluation of different solutions for the gridification of a rich variety of applications, as well as the accomplishment of a good level of interoperability among different technologies related to supercomputing, i.e. High Performance Computing and High Throughput Computing, Grid and Cloud.
Earth Sciences (TSA3.6)	Earth Science (ES) applications cover various disciplines like seismology, atmospheric modelling, meteorological forecasting, flood forecasting and many others. Their presence in SA3 is currently centred in the implementation, deployment and maintenance of the EGDR service to provide access from the grid to resources within the Ground European Network for Earth Science Interoperations - Digital Repositories (GENESI-DR). The ES HUC includes also researchers and scientists working in the climate change domain. In particular most of them actively participate in the Climate-G use case. This use case exploits the GRelC service for distributed metadata management and the Climate-G portal as scientific gateway for this collaboration.

Table 1: Summary of the Heavy User Communities

3 SERVICES FOR HIGH ENERGY PHYSICS

3.1 Introduction

The Services for High Energy Physics (HEP) task continues to focus primarily but not exclusively on the 4 LHC VOs (experiments) centred at CERN: ALICE, ATLAS, CMS and LHCb (and hence for and via the WLCG project and collaboration). Services and tools developed or extended by these VOs are also used by other HEP experiments and/or are under consideration for the future.

The work of SA3 goes hand in hand with that of the WLCG project in general: what is described below specifically describes the contribution of SA3 however it should be considered in the broader context of the entire project by considering its role within the WLCG project as well.

From a technical viewpoint, the LHC machine delivered well above expectations during 2011, as did the LHC detectors. The fact that the worldwide grid computing infrastructure not only kept up with load and capacity demands that exceeded the plans but was also accompanied by a measurable improvement in service delivered is testimony to the efforts of the entire community.

Not only were SA3-supported tools and services an important cornerstone of these achievements but the drive for common solutions made further inroads during this period – an impressive result given that stability is the norm during data taking.

During the past year WLCG launched a number of Technical Evolution Groups (TEGs) in which EGI-InSPIRE SA3 members (as well as others from EGI-InSPIRE) played key roles. The final report from these groups is not expected until May 2012 and is expected to drive the further evolution of WLCG and hence have an impact on grid sites and projects involved.

A brief overview of these TEGs, the involvement of EGI-InSPIRE, as well as the main recommendations of each group are given below.

The remainder of this section covers the tools and services that are further described in MS610 [MS610].

3.2 Analysis Tools and Support

3.2.1 HammerCloud

HammerCloud (HC)[2] is a Grid site testing service developed around Ganga. HC uses frequent short jobs to validate a site's availability and functionality, and also delivers on-demand stress tests to aid in site commissioning or general benchmarking. HC was developed with the ATLAS experiment but is used today by CMS and LHCb experiments in addition.

The main achievement of the past year was a merging of different codebases between the three VO users to one generic codebase, which is customizable per VO and will allow extensions to new VOs in the future. CMS and LHCb were already running a common codebase at the start of the year, whereas ATLAS was running on older independent code. During the ATLAS instance migration to the generic codebase, a number of performance issues were discovered; some sub-optimal database queries were found and the VO-agnostic model layer was found to hurt overall performance of the web interface and backend logic.

[2] https://savannah.cern.ch/projects/hammercloud/

After some Django optimizations were applied (including adding related-table hints), performance was restored (in some cases the number of database queries required to generate a page view decreased by 3 orders of magnitude.)

A second large achievement in the past year was an extension of the “job robot” functionalities to validate the VO Grid sites. This HC testing has been seen as very successful by the ATLAS VO; improvements to the auto-exclusion service led to ~25% decrease in Grid failures for ATLAS users. This motivated additional testing of the Monte Carlo Production Grid queues, this work was developed during the past year and put into production in mid-February. With this achievement, HC is now validating all ATLAS Grid sites and job workflows. For CMS, the job robot functionalities have been progressively improved over the past ~6 months the tests are regularly queried to understand the site reliabilities. Feedback from these tests presently depends on human operators, in future it is hoped that some automatic feedback into the experiment's workload management systems will be possible. LHCb has been working in the last quarter to integrate the HC testing with the LHCbDIRAC resource status system, this work is still in progress.

One of the plans for the past year was to find a more stable database service to host the HC databases (previously the HC team has been running its own MySQL servers). It was originally planned to port the HC object layer to Oracle, however a new hosted MySQL service became available in CERN IT during this quarter, so the databases have been migrated to this service. The plan to use Oracle has therefore been scratched.

Moving forward, the plans for 2012 include:

- Integration of LHCb testing with the LHCbDIRAC resource status system;
- Working on a standard benchmark (or benchmarking suite) to evaluate the storage elements (e.g. standard set of ROOT test jobs and datasets);
- Investigation of error and performance correlations across VOs (at multi-VO sites);
- Development of a multi-VO testing interface for users to easily test all experiment workflows;
- HC installation simplification and packaging for other-VO usage.

Regarding organization and sustainability, the HC project is currently led by an EGI-InSPIRE funded FTE, and developments and operations are complemented by at least 1 partial FTE from each user VO (ATLAS has 2 extra support persons (partial)). It is foreseen that the continued life of the project does not depend one any one of these partners.

3.2.2 CMS CRAB Client

The CMS Remote Analysis Builder (CRAB) was the first analysis tool in CMS to aid users in configuring CMS applications for distributed use, by discovering the location of remote datasets and submitting jobs to the Grid infrastructure. CRAB has progressed from a limited initial prototype nearly 5 years ago to a fully validated system that is critical to the production of published physics results. CMS currently observes more than 400 unique users submitting CRAB jobs per week, with close to 1000 individuals per month. The CMS Computing Technical Design Report (CTDR) estimated roughly 100k Grid submissions per day. In the second half of 2010 the job submissions routinely exceeded the estimate by 40-50% and CRAB coped well with the increased load.

After the development of the CRAB submission client a CRAB server was developed, which has increased the scalability of submission and added capabilities of automatic resubmission. The CRAB server also provides a development platform for additional capabilities.

During the past year effort has been spent both to maintain the production version of CRAB and to develop & commissioning the new one. Particularly the maintenance of CRAB version 2 has focused on:

- Fixing problems with JobSplitter as well as deprecating srmv1[3] protocol;
- Guaranteeing the compatibility of the newer version of the CMS Analysis framework (CMSSW) and the underlying condor and gLite middleware changes;
- Supporting the CERN analysis facility migration to the EOS system.

Regarding the CRAB3 development and commissioning activities the main areas of work were:

- Development of the RESTFul based CRABServer interface. APIs for basic functionality have been implemented; a long test and integration phase has been managed. Thanks to the feedback received over the past year further optimization of the REST APIs have been addressed both to improve the functional and technical aspects. In particular, all the security related implementations needed for the final production version are now ready. The validation of the whole set of inputs provided to CRAB by the user is included in the upcoming version;
- Implementation of the User Sandbox system. The User Sandbox is the archive of all user own library, files, configurations which must be shipped to the worker node. A central cache HTTPS based system is now available;
- The new crab client has been made available, supporting all the functionality associated to the resources exposed by the REST layer. The client is now implementing the support for the CRAB2 to CRAB3 configuration file translation. Once moved to CRAB3, the user configuration file will be only python based with the previous text based approach to be removed;
- The first version of the software with the AsyncStageOut functionality, which asynchronously moves the user produced output, has been tested, integrated and released. A basic specific monitoring has been implemented to support the related operations;
- Both the user and code documentation for all the CRAB3 components have been developed using sphinx (http://sphinx.pocoo.org/);
- The first version of the deploy/manage system has been implemented. It is based on the RPM, which indeed has been produced for all the single components of the new CRAB system;
- A lot of effort was spent to support the first multi-user integration test. This was done in the middle of October and the scope of the test is to collect user feedback needed to consolidate the tool and decide the priorities in order to move from development to production quality;
- Several cycles of four days face-to-face meeting have been organized in order to include discussion of both follow-up to the beta tests and development and scheduling of specific items.

The main items of development for the next year are:

- To rewrite the whole monitoring implementation currently developed, moving from a drill-down to a push model approach. The aim is to build a centralized monitoring service, based on CouchDB, which is able to collect monitoring documents from the distributed agents;
- To implement advanced functionality like user data publication, support for the user output

[3] https://www.gridpp.ac.uk/wiki/SRM

merging, and all what will come-up from the users feedback;

- To evolve the current client to a python library. The aim is to enable the possibility to be used by external bot/services like HammerCloud, which will be the first case we will support;
- To evaluate, eventually integrate, and use frontier on worker node to cache UserFileCache;
- To perform a large scale test, including the whole stack which is in the CRAB3 system, namely: Client, REST Interface, UserFileCache, WorkQueue, AsyncStageOut and distributed Agents.

3.3 Data Management Tools and Support

3.3.1 ATLAS Distributed Data Management

ATLAS, one of the LHC experiments, fully relies on the use of Grid computing for offline processing and analysis. This processing is done worldwide using the well-known tier model across heterogeneous interoperable Grids and the ATLAS Distributed Data Management (DDM) project is responsible for the replication, access and bookkeeping of ATLAS data across more than 100 distributed Grid sites.

The work during the last year has focused on the consolidation of the current production system in different fronts:

- The DDM infrastructure monitoring is responsible for the early detection of system failures and degradations of DDM services. The infrastructure has been extended to a client-server model that communicates over message queues and where the server publishes the health reports to the CERN IT Service Level Status monitoring solution. Based on this new implementation, significant effort has been invested in developing and improving the service monitoring for a variety of the DDM subcomponents.
- Optimization of the DDM Site Services, which is the set of agents responsible for the ATLAS data placement using the underlying EGI middleware (mainly FTS, LFC and SRM). Improvements for this component include:
    - Optional submission of FTS jobs to GridFTP (bypassing SRM). This was required for the improvement of the data import and export on the new EOS storage element at CERN.
    - Adaptation of Site Services for Tier3s in order to provide a throttled way of transferring data to these sites. The requirements of these sites are:
        - Submission of FTS jobs to the GridFTP endpoints.
        - No file/dataset registration in the LFC and Central Catalogues: Sites want to have local control of the data without taking care of catalog synchronization.
        - Since there is no LFC registration, file look-up has to be done using *uberftp-ls*.
    - Adaptation to use the recently implemented ATLAS Grid Information System.
    - There have been several releases with general enhancing, fine-tuning and bug fixing.
        - Optimized usage of temporary storage areas for cross-cloud transfers.
        - Pinning of source replicas to ensure their availability.
        - Implemented new policies for tape staging according to latest operational experiences.
    - Migration of the Site Services instance serving US sites from BNL to CERN.

- Evaluation of running the service on virtual machines with lower CPU and memory profiles.

- Support has been provided for a variety of DDM components, with strong focus on the day-to-day operations in data replication.

The current ATLAS DDM software is now in a mature state and the present work is focused on maintenance and support operations. Inside the ATLAS Distributed Computing community there is an on-going discussion about the proposal of the CERN-PH-ADP group of developing a new DDM system (the Rucio project) to solve the current shortcomings and scalability issues in the Central Catalogues. The details about the future of the project and the implications in other groups are unclear at this point.

3.3.2 CMS Data Management

Building on the previous experience acquired by the ATLAS experiment, the CMS Popularity Service has been developed to monitor the experiment's data access patterns (i.e. frequency of data access, access protocols, data tiers, users, sites, CPU usage). The understanding of this framework provides a crucial step ahead towards the automation of data cleaning and data placement.

In addition, a fully automated popularity-based site-cleaning agent has been deployed in order to scan the Tier2 sites that are reaching their space quotas and suggest obsolete, unused data that can be safely deleted without disrupting analysis activity. The implementation of this agent is based on the initial ATLAS code, which has been re-factorized to a plug-in architecture with a common core for the CMS and ATLAS experiments. The system has been significantly improved by adding a database backend that stores the suggested cleaning and a new web frontend.

Future work in this area has mainly two points of focus:

- Firstly, CMS is currently exploring how to extend the Popularity system with the usage patterns of the data accessed through the XRootD protocol.
- Secondly, CMS wants to analyse the possibilities of dynamic data placement strategies and demonstrate that the current Workload Management infrastructure could call out to perform data replication of hot data to facilitate its analysis. This would require interactions:
    - With the Data Management and Workload Management systems for the hook to replicate the data and decide to which sites the replications should be done.
    - With the Popularity Service and dashboard for automatically making additional replica of hot data and to not make additional replicas at sites that currently are oversubscribed.
    - Interactions with the Data Distribution system (PhEDEx) to discover available resources and make the replication requests.

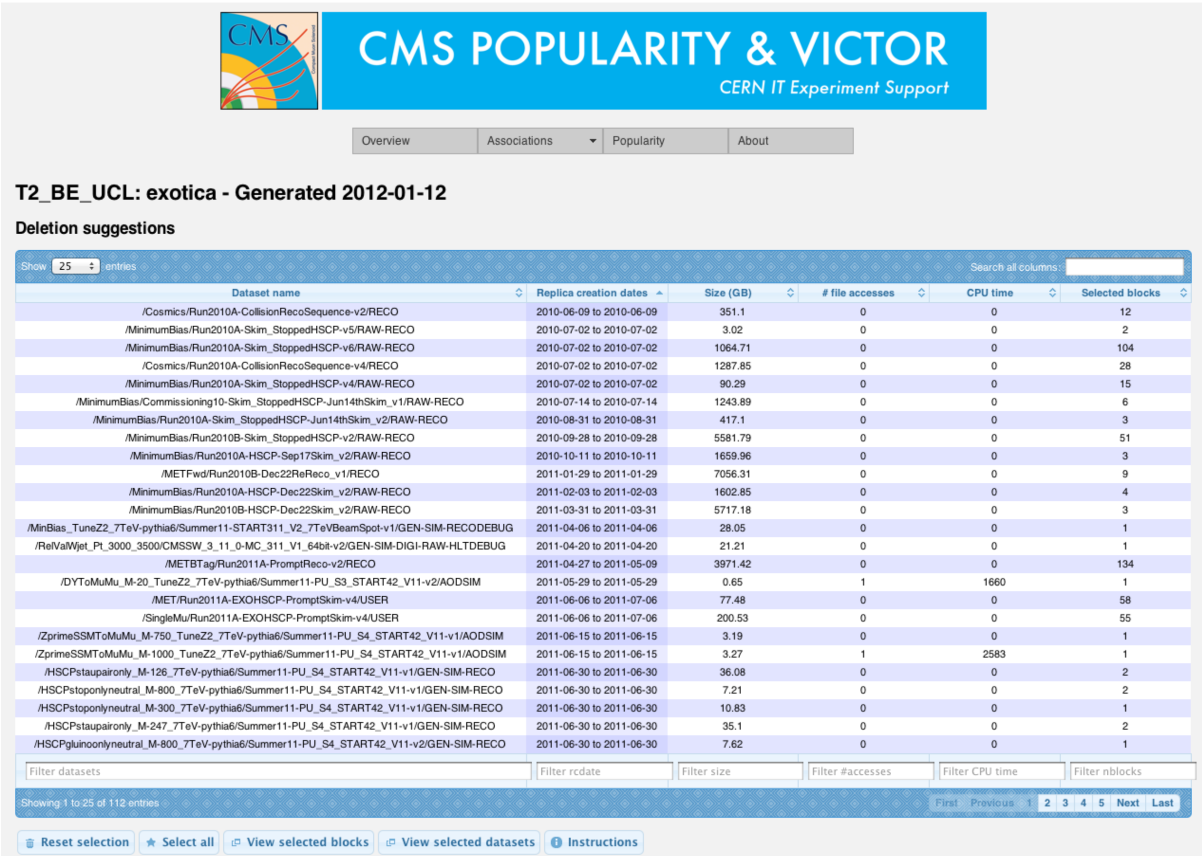

CMS POPULARITY & VICTOR

CERN IT Experiment Support

Overview | Associations | Popularity | About

T2_BE_UCL: exotica - Generated 2012-01-12

Deletion suggestions

Show 25 entries — Search all columns:

Dataset name	Replica creation dates	Size (GB)	# file accesses	CPU time	Selected blocks
/Cosmics/Run2010A-CollisionRecoSequence-v2/RECO	2010-06-09 to 2010-06-09	351.1	0	0	12
/MinimumBias/Run2010A-Skim_StoppedHSCP-v5/RAW-RECO	2010-07-02 to 2010-07-02	3.02	0	0	2
/MinimumBias/Run2010A-Skim_StoppedHSCP-v6/RAW-RECO	2010-07-02 to 2010-07-02	1064.71	0	0	104
/Cosmics/Run2010A-CollisionRecoSequence-v4/RECO	2010-07-02 to 2010-07-02	1287.85	0	0	28
/MinimumBias/Run2010A-Skim_StoppedHSCP-v4/RAW-RECO	2010-07-02 to 2010-07-02	90.29	0	0	15
/MinimumBias/Commissioning10-Skim_StoppedHSCP-Jun14thSkim_v1/RAW-RECO	2010-07-14 to 2010-07-14	1243.89	0	0	6
/MinimumBias/Run2010A-Skim_StoppedHSCP-Jun14thSkim_v2/RAW-RECO	2010-08-31 to 2010-08-31	417.1	0	0	3
/MinimumBias/Run2010B-Skim_StoppedHSCP-v2/RAW-RECO	2010-09-28 to 2010-09-28	5581.79	0	0	51
/MinimumBias/Run2010A-HSCP-Sep17Skim_v2/RAW-RECO	2010-10-11 to 2010-10-11	1659.96	0	0	3
/METFwd/Run2010B-Dec22ReReco_v1/RECO	2011-01-29 to 2011-01-29	7056.31	0	0	9
/MinimumBias/Run2010A-HSCP-Dec22Skim_v2/RAW-RECO	2011-02-03 to 2011-02-03	1602.85	0	0	4
/MinimumBias/Run2010B-HSCP-Dec22Skim_v2/RAW-RECO	2011-03-31 to 2011-03-31	5717.18	0	0	3
/MinBias_TuneZ2_7TeV-pythia6/Summer11-START311_V2_7TeVBeamSpot-v1/GEN-SIM-RECODEBUG	2011-04-06 to 2011-04-06	28.05	0	0	1
/RelValWjet_Pt_3000_3500/CMSSW_3_11_0-MC_311_V1_64bit-v2/GEN-SIM-DIGI-RAW-HLTDEBUG	2011-04-20 to 2011-04-20	21.21	0	0	1
/METBTag/Run2011A-PromptReco-v2/RECO	2011-04-27 to 2011-05-09	3971.42	0	0	134
/DYToMuMu_M-20_TuneZ2_7TeV-pythia6/Summer11-PU_S3_START42_V11-v2/AODSIM	2011-05-29 to 2011-05-29	0.65	1	1660	1
/MET/Run2011A-EXOHSCP-PromptSkim-v4/USER	2011-06-06 to 2011-07-06	77.48	0	0	58
/SingleMu/Run2011A-EXOHSCP-PromptSkim-v4/USER	2011-06-06 to 2011-07-06	200.53	0	0	55
/ZprimeSSMToMuMu_M-750_TuneZ2_7TeV-pythia6/Summer11-PU_S4_START42_V11-v1/AODSIM	2011-06-15 to 2011-06-15	3.19	0	0	1
/ZprimeSSMToMuMu_M-1000_TuneZ2_7TeV-pythia6/Summer11-PU_S4_START42_V11-v1/AODSIM	2011-06-15 to 2011-06-15	3.27	1	2583	1
/HSCPstaupaironly_M-126_7TeV-pythia6/Summer11-PU_S4_START42_V11-v1/GEN-SIM-RECO	2011-06-30 to 2011-06-30	36.08	0	0	2
/HSCPstoponlyneutral_M-800_7TeV-pythia6/Summer11-PU_S4_START42_V11-v1/GEN-SIM-RECO	2011-06-30 to 2011-06-30	7.21	0	0	2
/HSCPstoponlyneutral_M-300_7TeV-pythia6/Summer11-PU_S4_START42_V11-v1/GEN-SIM-RECO	2011-06-30 to 2011-06-30	10.83	0	0	1
/HSCPstaupaironly_M-247_7TeV-pythia6/Summer11-PU_S4_START42_V11-v1/GEN-SIM-RECO	2011-06-30 to 2011-06-30	35.1	0	0	2
/HSCPgluinoonlyneutral_M-800_7TeV-pythia6/Summer11-PU_S4_START42_V11-v2/GEN-SIM-RECO	2011-06-30 to 2011-06-30	7.62	0	0	1
Filter datasets	Filter rcdate	Filter size	Filter #accesses	Filter CPU time	Filter nblocks

Showing 1 to 25 of 112 entries — First Previous 1 2 3 4 5 Next Last

Reset selection | Select all | View selected blocks | View selected datasets | Instructions

Figure 1: Deletion suggestions for the exotica physics-group at T2_BE_UCL

3.3.3 LHCb DIRAC

DIRAC framework provides a complete solution for using the distributed computing resources of the LHCb experiment. DIRAC is a framework for data processing and analysis, including workload management, data management, monitoring and accounting (more details have been given in document [MS610]). LHCbDIRAC framework is the DIRAC extension specific to the LHCb experiment, which has been formally separated from DIRAC in order to streamline the implementation of features requested by LHCb community.

The support of the LHCbDIRAC Data Management system (DMS) has started in EGI-InSPIRE project in October 2010.

The main lines of activity for the next year are the following:

- Consolidation of the service for data consistency checks between file catalogues and SEs, especially in the operational aspects. In particular, procedures should be established for notifying the VO's data management team about the files which were found in an inconsistent state, and decide about their removal. This will streamline the task of consistency checks and will enable the data manager to reduce the time needed to fulfill the task.
- Support and possible improvement of the accounting plots of storage resources usage developed and put in production during the past year. This system, already extensively used by members of the collaboration, could undergo some improvement, taking advantage of

new features of the LHCb DIRAC DMS that are planned to be implemented during next months.

- Adapt when necessary the DIRAC DMS to changes in the data management middleware (in particular SRM and LFC).
- Continue the development of the popularity service, just started during PQ7. Such a service should provide metrics to asset the data-sets popularity and provide a ranking of the most popular data-sets (i.e. data most frequently accessed by users). The final goal is to use the information about data popularity to implement a dynamic data placement model, where the number of replicas of a given data-set is related to its popularity. This would considerably help the VO's data managers to optimize the usage of storage resources on the Grid and would automate many operations that are currently done manually and thus time expansive.
- Provide general support for LHCb computing operations on the Grid, both for production and private user's activity.

3.4 Monitoring and Dashboards

Monitoring of the distributed infrastructure and the activities of the user communities on this infrastructure is a vital for ensuring its quality and performance. Monitoring is of particular importance for Heavy User Communities (including HEP) due to the scale of their activities and the quantity of resources that they are using. There are two main tasks that have to be addressed by the monitoring systems used by HEP VOs: monitoring of the distributed sites and services, and the monitoring of the VO activities, namely job processing and data transfer. The Experiment Dashboard was developed in order to address the monitoring needs of the LHC community, but in contrast to other monitoring systems it provides common solutions that work transparently across various middleware platforms and are not coupled with VO-specific frameworks, offering instead a common way to instrument those frameworks for publishing monitoring data.

3.4.1 Experiment Dashboard

During the referenced period the Experiment Dashboard applications were heavily used by the LHC virtual organizations (VOs), in particular by ATLAS and CMS. The system plays an important role for everyday operations, for site commissioning activity and for the distributed computing shifts. More than 200 CMS physicists daily access CMS Dashboard task monitoring in order to follow processing of their tasks on the distributed infrastructure. ATLAS DDM Dashboard is being actively used for monitoring of ATLAS Data transfers. All LHC experiments use Dashboard SAM portal which provides information for evaluating site usability from the VO perspective.

The performance, scalability and functionality of the system are steadily improving following the growing scale of the LHC computing activities and the requirements of the LHC community.

Most of the Dashboard cluster had migrated to the virtual machines. It facilitated deployment and maintenance tasks.

All Dashboard applications use ORACLE for persistency. In the beginning of 2012, ORACLE cluster at CERN which hosts Dashboard DBs migrated to new hardware and ORACLE 11g. This resulted in substantial improvement of the performance of the Dashboard applications.

One of the important goals of the Experiment Dashboard development is to offer common monitoring solutions for the LHC VOs. Achieving this goal permits the monitoring infrastructure to be simplified and contributes to its sustainability. All Dashboard development in 2011 and beginning of 2012 followed this strategy. During the referenced period, the global transfer monitoring system which provides cross-VO monitoring of all data transfers on the WLCG infrastructure was prototyped and deployed on the preproduction server. New versions of the job monitoring applications were developed based on the common hBrowser framework. hBrowser framework offers common client-side implementation of the user interfaces with generic look and feel and with the advanced plotting, sorting and filtering functionality. After successful validation by the user community, new versions of the job monitoring applications were deployed in production. The Site Status Board (SSB) which provides a common solution for monitoring of the distributed sites and services was redesigned in order to improve performance of the data queries and the quality of the user interface. New SSB versions were successfully validated by the LHC VOs, were deployed in production and received a positive feedback of the user community.

The LHC VOs rely on the Service Availability Monitor (SAM) for remote testing of the distributed sites and services. Therefore one of the important WLCG monitoring tasks during 2011 and beginning of 2012 was migration to a new SAM architecture. The Experiment Dashboard provides a portal for navigating of the SAM test results and site availability calculated. The new portal compatible with the new SAM architecture was validated and deployed in production.

Further development of the Experiment Dashboard system is being aligned with the LHC computing needs and will follow the requests of the LHC virtual organizations.

3.5 WLCG Technical Evolution Groups

Following a number of discussions in the WLCG Management Board, Grid Deployment Board and Collaboration Workshop, a number of Technical Evolution Groups were established in the 2nd half of 2011 with the following mandate and goals:

To reassess the implementation of the grid infrastructures that we use in the light of the experience with LHC data, and technology evolution, but never forgetting the important successes and lessons, and ensuring that any evolution does not disrupt our successful operation.

The work should:

- Document a strategy for evolution of the technical implementation of the WLCG distributed computing infrastructure.
- This strategy should provide a clear statement of needs for WLCG, which can also be used to provide input to any external middleware and infrastructure projects.

The work should, in each technical area, take into account the current understanding of:

- Experiment and site needs in the light of experience with real data, operational needs (effort, functionality, security, etc.), and constraints;
- Lessons learned over several years in terms of deployability of software;
- Evolution of technology over the last several years;
- Partnership and capabilities of the various middleware providers.

It should also consider issues of:

- Long term support and sustainability of the solutions;
- Achieving commonalities between experiments where possible;
- Achieving commonalities across all WLCG supporting infrastructures (EGI-related, OSG, NDGF, etc).

Deliverables:

- Assessment of the current situation with middleware, operations, and support structure.
- Strategy document setting out a plan and needs for the next 2-5 years.

Groups were setup in the following areas (chairs in parentheses):

- Data Management.
  - Brian Bockelman (CMS,OSG), Graeme Stewart (ATLAS).
- Storage Management.
  - Daniele Bonacorsi (CMS), Wahmid Bhimji (Edinburgh).
- Workload Management.
  - Davide Salomoni (CNAF), Torre Wenaus (ATLAS).
- Databases.
  - Dario Barberis (ATLAS), Dave Dykstra (OSG).
- Security.
  - Romain Wartel (CERN), Steffen Schreiner (ALICE).
- Operations.
  - Maria Girone (CERN, EGI-InSPIRE TSA3.3), Jeff Templon (NIKHEF).

The reports from these groups were presented in February 2012 [TEG REPORTS]. A summary document covering all TEGs is being prepared by the TEG chairs together with a small team from WLCG Management and will be finalized prior to – and presented at – the WLCG Collaboration Workshop in May, NYC [WLCG 2012]. As the work on preparing the summary has not yet started, it is not possible to predict the outcomes. However, it is clear that a number of changes will be made even on a relatively short timescale, with more intrusive changes introduced (and heavily tested) according to the LHC accelerator schedule and in particular the long shutdown that is foreseen from the end of 2012 for at least 1 year.

4 SERVICES FOR LIFE SCIENCE

4.1 Introduction

The Life Science Grid Community (LSGC) was set up during PY1 with the central role to the coordination and support of the Life Science activity on EGI. In the context of the EGI-InSPIRE project, the Life Science HUC contributed to the LSGC effort by maintaining a production quality Grid environment for Life Sciences by providing technical skills and manpower for VRC operation, as well as some specific tools dedicated to the Life Science community. A significant effort is spent on the VO-wise monitoring and trouble-shooting of the EGI infrastructure. This effort is split between a daily trouble-solving activity to ensure the immediate usability of the Life Science resources for the community, and a longer term effort in VRC management tools that are being developed to simplify and lighten the VRC administrators' workload in the future. Moreover, the provision of additional services, mostly a Grid database interface (GRelC) and a data encryption service (Hydra) are being handled within the HUC.

4.2 LSGC user management tools

The Life Sciences HUC steers the LSGC ("Life Sciences Grid Community" VRC[4]) effort to organize the community and deliver new services. A particular effort is invested in assisting users to better exploit the Grid and rationalizing Grid usage. In this context, several services to the HUC users have been provisioned:

- Web gadgets listing Life Sciences applications and community requirements posted to the RT systems set up by User Communities Support Team have been added to the LSGC wiki;
- A HUC support service is delivered. A technical team of expert users has been set up. It addresses the difficulties reported by users on the VRC mailing lists or through GGUS. Bi-monthly phone conferences are scheduled and shifts are organized to ensure that there is always a team on duty tackling the problems. See technical team wiki for details[5];
- The technical team pro-actively monitors the infrastructure health at a VO level, to identify the problems occurring. The French NGI deploys a dedicated Nagios server[6] for that purpose. New probes to monitor all VO SEs, WMSs and CEs were developed;
- On-line reporting tools easing the monitoring of SEs space management have been added to the technical team wiki page.

Thanks to the tools set up, a large-scale action to clean up all obsolete files let behind by years of Grid usage from a heterogeneous community with 200+ registered users is currently being conducted.

4.2.1 Work Accomplished

PQ5:

- Setting-up of community communication channels and technical assistance to the end users.
- Provisioning of a redundant VOMS server for the biomed VO.

[4] LSGC wiki, http://wiki.healthgrid.org/LSVRC:Index

[5] Biomed technical team wiki, http://wiki.healthgrid.org/Biomed-Shifts:Index

[6] Biomed Nagios server, https://grid04.lal.in2p3.fr/nagios

- Investigation of technical means to similarly deliver a redundant LFC server.
- Strengthening of standard procedures to deal with storage resources decommissioning or failures.

PQ6:

- Dissemination of good practices as well as knowledge of tools used / developed among the VRC technical team.
- Improvements of the Nagios monitoring probes.

PQ7:

- Monitoring of new resources, in particular all WMSs and CEs accessible to the biomed VO users.
- Publication of a paper[7] describing the role, the organization and the impact of the Life Sciences technical team in experiments support.

The work completed is all part of the EGI infrastructure, the LS community does not intend to deploy an independent parallel infrastructure.

4.2.2 Current work and Outlook for 2012

The development of a HUC users database and management tools to assist VO administrators in their daily task is planned. This service will interface to Virtual Organization Membership Service (VOMS) servers as well as the EGI applications database, to avoid replicating existing information. It will complement the VOMS and applications database with extra-information on the users and their affiliations. It will be interfaced to the LSGC Dashboard.

4.3 GRelC data access interface

The Grid Relational Catalog (GRelC) access interface is a general purpose Grid database service that can be exploited in several Life Sciences use cases.

In this regard, two new use cases have been investigated during PY1. The first one concerns with the UNIPROT gene data bank (which has been completely imported into a new relational database and made available through the GRelC interface) whereas the second one (which is still being conducted) is a biodiversity use case, aiming at studying and sharing invasive species information. A detailed description about the two use cases is available in Section 7.

4.3.1 Work Accomplished in the LS area

PQ5:

- Availability of the UNIPROT proteins data bank through the GRelC interface.

PQ6:

- Improvement of the ETL (Extraction-Transformation-Loading) tool.

The ETL tool is comprised of a C library responsible for extracting the data from the flat file and for storing it into an XML file and the GRelC translator component responsible for translating the XML file entries into bulk queries for the new relational database.

[7] "Technical support for Life Sciences communities on a production grid infrastructure", HealthGrid 2012. http://hal.archives-ouvertes.fr/hal-00677839

PQ7:

- Definition of a new use case on biodiversity.

4.3.2 Current work and Outlook for 2012

The support to the Life Sciences domain will be a key activity also in Y3. New use cases will be defined jointly with LS groups and the existing ones (e.g. the one on biodiversity at the University of Salento) will be refined and finalized. A more detailed and cross-domain outlook about the GRelC service is discussed in Section 7.

4.4 Hydra encryption service

Hydra[8] is a file encryption/decryption tool developed by EMI to enable the protection of sensitive files stored on Grid storage resources. The service is composed by a distributed encryption key store (hence its name), and client command lines that can (i) upload/fetch keys to/from the key store and (ii) encrypt/decrypt data files using these keys.

An experimental Hydra service has been successfully deployed since the end of PQ4 (April 2011), first on a 32bits gLite version 3.1 and then on a 64bits gLite version 3.2 server. There are two difficulties identified for delivering a production service using the Hydra key store. First, many production sites are misconfigured, having deployed older version of Hydra in a former gLite release and sometime exposing hydra tags that are not valid anymore or older version of the hydra client. Second, the current version of the hydra client is developed as part of the EMI middleware and its dependencies are incompatible with the gLite 3.2 release deployed on the production infrastructure. The service delivered today is therefore only a test service mostly used for the validation of the functionality delivered.

4.4.1 Work Accomplished

PQ5:

- Deployment of a prototype Hydra service on a gLite release 3.1 UI;
- Work on the migration to gLite release 3.2.

PQ6:

- Provision of first beta-release of the Hydra server software compatible with gLite release 3.2;
- Fixes in the installation and configuration procedures.

PQ7:

- Work on Hydra client provision for sites providing resources to Life Sciences;
- Site configuration problems are being tackled. Work on Hydra CLI package with EMI to try installing it on the current release of gLite deployed on the production infrastructure.

4.4.2 Current work and Outlook for 2012

To deliver a production service, Hydra depends upon:

- Server's provision for Hydra key stores;
- Standard client packages installation on all EGI computing resources;
- Service monitoring and maintenance.

[8] Hydra service overview. https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS

A 3-server based Hydra key-store will be deployed as a service for the life sciences community. This task, due in the first year of the EGI-InSPIRE project, has been delayed until gLite 3.2 (64 bits) servers could be provisioned and further slowed down by the end of the HealthGrid association which was supposed to host one of the servers. An attempt will be made to produce client packages compatible with the current middleware release and deploy them. If it proves impossible, it will be needed to wait for the first deployment of the EMI release on the EGI infrastructure. The Life Science Grid Community (LSGC) technical team will handle the service client installation and monitoring for more sustainability. The service monitoring procedure will be integrated in the Nagios box-based monitoring service offered by the LSGC technical team.

4.5 Issues and mitigation

- **Status of the HealthGrid association**

The HealthGrid association was dissolved by court decision due to its inability to finance its debt. Reorganization is needed inside the CNRS partner to handle the tasks managed by HealthGrid. Despite this unexpected event, the LSGC VRC remains active. In the future, NGIs will be asked to host technical services, which will improve their sustainability.

- **Hydra service**

The end of the HealthGrid association prevents the deployment of a third Hydra key store on its resources. A new partner to host the service will be searched for.

- **Redundant LFC server**

Previous plans to deploy a redundant LFC server are currently on hold due to technical incompatibilities. The biomed LFC server is using an Oracle database back-end and its replication would mean finding a site to provide an Oracle licence for that purpose. The migration to MySQL back-ends is not recommended for performance reasons.

Two options remain:

1. Find a site that owns an Oracle license (like the LHC*b*), willing to host the redundant node;
2. Migrate the LFC node to MySQL: which is not recommended due to MySQL replication policy issues.

- **User assistance**

The work performed by the Life Sciences technical team is mostly around generic infrastructure monitoring, probing and troubleshooting activity, due to the complexity of middleware configuration and maintenance. Despite the automation of many monitoring procedures, the work remains dominated by the increasing number of incidents detected, the number of which follows the increased capability of the team to identify misconfigured or faulty services. The technical team longs for the time this background work to stabilize the infrastructure will be complete, but experience shows that this should not be expected in the near future. Hence the technical team cannot for now focus on its real, application community-oriented missions.

4.6 Summary

The LSGC set up during PY1 continues ramping up its work and now provides daily infrastructure quality of service improvements. It plays an important role for Life Science end users exploiting the Grid and it also contributes more widely to the set-up of a VRC administration environment that can be of interest of other VRCs in the future. The premature ending of the HealthGrid association activity delayed some of the tasks tackled by the Life Science Grid community but it did not seriously jeopardize the overall community activity.

5 SERVICES FOR ASTRONOMY AND ASTROPHYSICS

5.1 Overview

Activities carried out by the A&A community during PY2 have been focused on the following topics:

- Coordination of the A&A community focusing in particular on the long-term sustainability plan;
- Visualization tools and Services: VisIVO;
- HPC, parallel programming, and GPU computing;
- Access to databases from DCIs and interoperability with the VObs (Virtual Observatory) data infrastructure.

5.2 Coordination of the A&A Community

This activity, although not officially part of the work plan of task TSA3.5, is related to it given that one of the most important objectives of the task is the provision of requirements, use-cases and investigating: a) interactivity between e-Infrastructures based on different technologies (Grid, HCP and Cloud); b) access and management of astronomical databases from Grid Infrastructures. To fully achieve this, objective people working in task TSA3.5 (who are also in charge of the coordination of the European A&A community at large in EGI) need to interact with as many European A&A research groups and Institutes as possible.

During PY2 the coordination of activities within the A&A community concerning the usage of DCIs by both small-scale and large-scale projects, in particular by the ESFRI projects such as SKA (Square Kilometre Array) and CTA (Cherenkov Telescope Array) has been intensified. People have been encouraged to contribute to the process of gathering requirements to be fed to EGI; other initiatives aimed at enhancing the use of e-Infrastructures were also carried out. A coordinating workshop[9] was organized at the Astronomical Observatory of Paris on November 7th 2011, where the major astronomical projects and research areas were represented.

The workshop highlighted the need to activate multiple astronomical VRCs each aligned around a big project or Institution that would be representative of one of the major A&A research areas. A unique catch-all VRC, in fact, has proven to not be suitable as it is unable to effectively represent the whole A&A community and would impose strong limitations for a number of key aspects like the requirements gathering process and the preparation and implementation of a long-term sustainability plan. All participants agreed that the major effort within the A&A HUC should now be focused on the identification of big A&A projects and Institutions for which the adoption of DCIs could be beneficial, so they might be good candidates to lead new astronomical VRCs. The process aimed at the creation of new VRCs is currently in progress with regular checkpoints scheduled to verify its progress.

5.3 VisIVO

During PY2 significant results have been achieved for what concerns the porting to the Grid of VisIVO[10] (Visualization Interface for the Virtual Observatory), a visualization and analysis software for

[9] http://twiki.oats.inaf.it/twiki/bin/view/AstroVRC/AstroVRCWorkshop

[10] http://visivo.oact.inaf.it/index.php

astrophysical data. It consists in a suite of software tools aimed at creating customized views of 3D renderings from many types of datasets.

The activity mainly focused on the study and on the porting of the VisIVO MPI version to a gLite Grid. The relevance of this activity can be easily understood if one considers that, depending on the structure and size of datasets, the Importer and Filters components could take several hours of CPU to create customized views, and the production of movies could last several days. For this reason the MPI parallelized version of VisIVO plays a fundamental role.

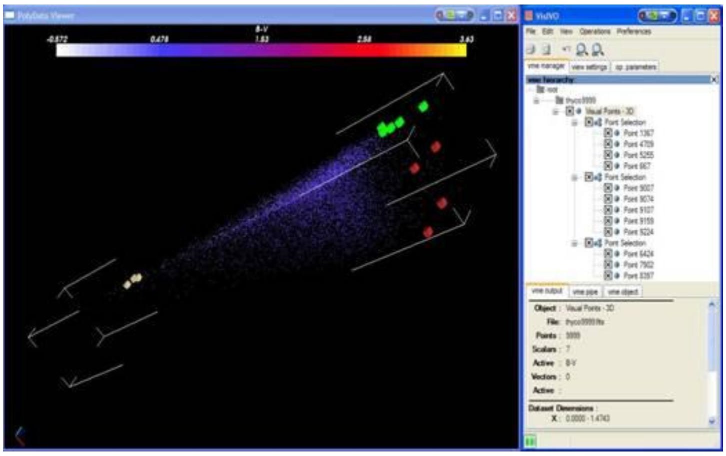

The preliminary study and the porting was mainly focused on the most important modifications of the code necessary to make VisIVO MPI fully compliant with the gLite Grid. The VisIVO MPI version, in fact, works assuming that the shared home directory and each process can directly work on the tables. Some classes of the VisIVO Filter component were modified to allow selections on a data table and preliminary tests were carried out.

Another important aspect that would enhance the performances, is the integration of VisIVO on Grid nodes where GPUs (Graphics Processing Units) are available. GPUs are emerging as important computing resources in Astronomy as they can be successfully used to effectively carry out data reduction and analysis. The option of using GPU computing resources offered by Grid sites to make visualization processing on VisIVO was then considered.

CUDA (Compute Unified Device Architecture) is the computing engine available in NVIDIA GPUs and accessible to software developers through variants of industry standard programming languages. Because VisIVO is developed in C++, the environment of CUDA is used to develop some computing-intensive modules of VisIVO. This activity started at the beginning of PY2 as a preliminary study on how to produce a CUDA-enabled version of VisIVO for gLite. This first study for the A&A community focused on the porting and optimization of the data transfer between the CPU and GPUs on worker nodes where a GPU is available. To provide a service able to take advantage of GPUs on the Grid, A&A acquired a new system (funded by the Astrophysical Observatory of Catania). It is an hybrid server CPU-GPU, with 2 quad-core processors Intel(R) Xeon(R) CP E5620 at 2.40GHz, 24 GB RAM DDR3-1333 NVIDIA TESLA C2070, 448 Cuda core and 6 GB of RAM. The server is configured as a Grid computing node.

The heaviest VisIVO Filter, the Multi-Layer Resolution Filter, was also analyzed. This filter makes possible the inspection of a very large user file (hundreds of gigabytes) to create data for the visualization of an entire dataset with different levels of resolution: starting from a fixed position, that represents the centre of the inner sphere, concentric spheres are considered. Different levels of randomization can be given, creating a more detailed table in the inner sphere and less detailed tables in the outer regions, or vice versa. The region external to the last sphere represents the background.

The performances of some additional VisIVO visualization filters such as randomizer, cut, select and swap operations on huge user data tables were also considered.

Now the first grid-enabled version of VisIVO (http://visivo.oact.inaf.it/index.php) service has been deployed. It is based on a specific grid-enabled library that allows users to interact with Grid computing and storage resources. The current version of VisIVO is also able to interface with and use the gLite Grid Catalogue.

Although VisIVO has been conceived and implemented as a visualization tool for astronomy, it has evolved to be a generic multi-disciplinary service that is capable of being used by any other community that needs 2D and 3D data visualization.

5.4 Grid and HPC

A working environment where Grid resources and HPC resources can be combined and used together is of utmost importance for many astronomical applications. Some typical A&A applications that could improve their performances in such an environment have been identified: FLY[11] (a cosmological code developed at INAF-OACt) and Gadget[12] + Flash[13], the most common cosmological codes in Astrophysics.

- The FLY code (a tree N-Body code) was executed on the gLite Grid. Some problems arise when running with a high number of N-Body particles. One of the most critical bottlenecks, in fact, is represented by the data transfer from the catalogue to the effective worker nodes. Unfortunately, the huge dimension of datasets makes tricky to have multiple replicas of them. The same problem happens during the data production phase, when a list of data files (each of them being tens of gigabytes in size) is produced. To overcome this problem workflows are being defined which can be adopted when the code is used in challenging scenarios where data files several tens of gigabytes in size have to be handled.
- Cosmological simulations we now being identified and use-cases defined. The goal of this activity consists in preparing the environment, the input file and the watchdog procedure to verify the run and the results.

Recently the activities related to Grid and HPC are carried out in close coordination with EGI and with IGI[14] (the Italian NGI) in Italy. A tight coordination with EGI and with NGIs, in fact, is crucial given that collaborations and agreements with PRACE and with other entities that operate and maintain HPC resources (typically supercomputing centers) are very important for this activity. One of the activities recently undertaken within the task TSA3.5 is to include small HPC resources into the A&A VO. This implies from one side to install and configure HPC clusters (based on low latency/high throughput networks, HPC libraries and tools, modules and compilers) and from the other side to make the grid middleware aware of these resources. The plan is to verify such small-size HPC resources and the related middleware aware version with the most popular cosmological applications mentioned previously, namely FLY and Gadget + Flash.

In the next period the A&A HPC activity in Italy will continue in close coordination with all the other communities involved in IGI that need HPC and MPI to efficiently run their applications.

[11] http://www.ct.astro.it/fly/

[12] http://www.mpa-garching.mpg.de/galform/gadget/index.shtml

[13] http://www.flash.uchicago.edu/site/

[14] http://www.italiangrid.it/

5.5 Access to Databases and interoperability with the VObs

Access to databases from DCIs and interoperability with the VObs (Virtual Observatory) data infrastructure remains one of the hot topics in astronomy. One of the most important related activities deals with the identification of use-cases and test-beds (both applications and complex workflows) that require simultaneous access to astronomical data and to computing resources. The data infrastructure of reference in astronomy is the VObs; end users should be able to access astronomical data in Grid through the VObs standards[15] and launch computational tasks on DCIs. In this context two key issues concern a SSO (Single Sign On) mechanism able to grant access to computing and data resources by means of a single authentication (users do not need to authenticate themselves multiple times) and tools/services to access astronomical databases federated in the VObs from DCIs.

For what concerns the way of accessing databases from DCIs, GRelC[16] is one of the tools under evaluation in order to verify its ability to meet the most important requirements of the A&A community.

Another important activity related to the integration of data and computing resources is the work in progress to integrate in Grid the BaSTI[17] (A Bag of Stellar Tracks and Isochrones) Astronomical Database and its feeding FARANEC code. A web portal was developed to make-easier the submission of the FARANEC code to the Grid. The portal allows the user to define a set of parameters and to simulate stellar evolutions in Grid, without worrying about the technical details concerning the underlying Grid Infrastructure. The portal is based on P-GRADE web portal. Since the portal is designed for the submission of an arbitrary job in Grid, it was necessary to customise it to make it able to fully satisfy our needs. The work aimed at implementing a new version of the portal based on the latest version of P-Grade is now in progress. It will be based on Web Services and this should enable a low level integration with the Virtual Observatory.

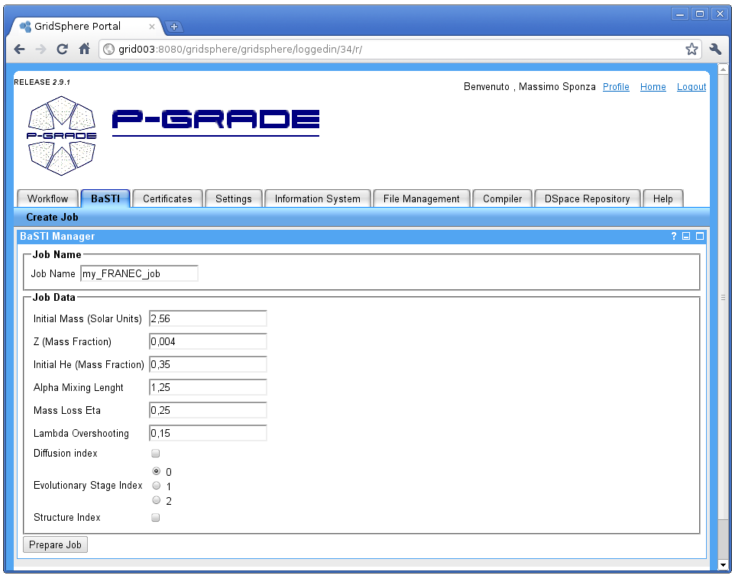

[15] http://www.ivoa.net/Documents/

[16] http://grelc.unile.it/home.php

[17] http://albione.oa-teramo.inaf.it/

6 SERVICES FOR EARTH SCIENCE

6.1 Overview

Earth Science (ES) applications cover various disciplines like seismology, atmospheric modelling, meteorological forecasting, flood forecasting, climate change and many others.

The presence of Earth Science in SA3 is centred in the implementation and maintenance of interfaces or tools to provide access to Earth Science specific resources from the Grid, in particular to large data infrastructures; for example resources within the infrastructure of the Ground European Network for Earth Science Interoperations - Digital Repositories (GENESI-DR), or climate data within the Earth System Grid (ESG). The community is supported independently by organisations and NGIs, and additional effort is put into fostering the community and to provide value-added services around EGI. The Services for Earth Science task covers the implementation of data access scenarios, to permit the utilization of Earth Science data resources in Grid jobs. The work can be divided into two efforts, separate but related in substance: discovery and access through GENESI-DR and ESG.

6.2 GENESI-DR

The GENESI-DR infrastructure provides a standardized data discovery interface based on OpenSearch and metadata standards for a federation of data repositories. While in the European project behind it (GENESI-DEC) focuses on a central portal as an interactive entrance point, the usage on EGI requires versatile clients such as a non-interactive, bulk oriented, tool. The work on such a command line client[18] that facilitates the usage in Grid jobs is ongoing. In recent months a data transfer component was included that downloads the bulk of resources found through the discovery process. This is required as the discovered data sets are not accessible through a unified channel, but depending on the data set instead through different access protocols such as HTTP, GridFTP or others. Because of changes to the underlying GENESI infrastructure, the tools had to be adapted. The command-line tool was extended with an ncurses based text user interface for interactive command-line usage over a shell, e.g. in the typical situation of a terminal SSH connection to a gLite User Interface. It now also has more sophisticated data handling and a few bugs were fixed to solve some rare crashes. In addition, a system for configuration over a dot file was added, which enables the user to permanently save and modify some variables like the preferred top request site. Documentation with a description of the functionality and a few standard use cases and examples to help users getting started with the utility has been prepared.

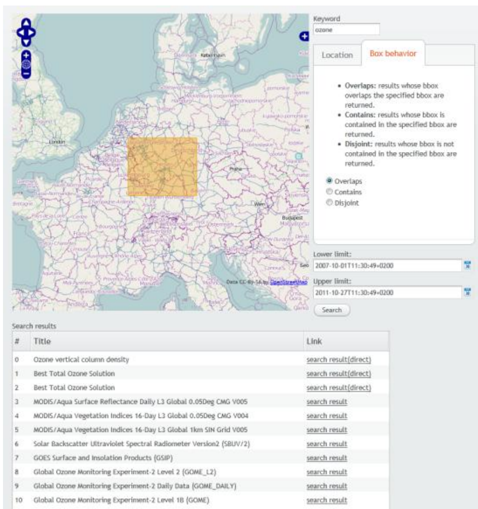

Additionally, a flexible web GUI was designed. It uses the OpenSearch interface provided by a running GI-cat instance (v.i.). The web interface accepts search parameters such as: a keyword, lower and upper time limits, a geographical bounding box (see figure 1) and more. The interface leverages the OpenLayers capabilities for defining the geographical bounding box. The list of results is subsequently being used to generate search queries for the respective datasets such as AIRS, MODIS, GOMOS or GOME

[18] https://appdb.egi.eu/?#!p=L2FwcHMvZGV0YWlscz9pZD03MDg

(mainly depends on the catalogues that are managed by GI-cat). Mostly, these datasets provide OpenSearch interfaces as well. The user defined search parameters from the first steps are then delegated to these second layer interfaces. The found data is presented to the user, who is able to select the required files. Based on this step, a job description file to be submitted to Grid resources is generated.

6.3 ESG

The Earth System Grid Federation (ESGF/ESG) is a distributed infrastructure developed to support CMIP5 (The Coupled Model Intercomparison Project, Phase 5), an internationally co-ordinated set of climate model experiments involving climate model centres from all over the world. Data access within ESGF is provided with two main services: OpeNDAP and GridFTP. A site that hosts these services is called a "Data node". The team that works on ESG interoperability is developing and testing a scenario based on an application from IPSL which uses CMIP5 data (climate model data stored on the ESG). The MPI code of the application is now running successfully on EGI. An important activity was the implementation of a multi-threaded climate data transfer program to download the data from the ESG data nodes. This smart data transfer tool, named '*synchro-data*'[19], facilitates the command line, bulk oriented access to ESG CMIP5 data. The tool can download files from the CMIP5 archive in an easy way, through a list of variables, experiments and ensemble members. The user defines one or many templates that describe the desired data, each of them listing variables, frequencies, experiments and ensemble members. The user separately defines a list of climate models to take into account. Using these templates, the tool explores the ESGF Grid and downloads all the corresponding available files. The program may be run regularly to download the possible new files. This tool has been improved considerably during this period. Major added features include fine grained priorities for transfers, support for CMIP5 "ensembles", an installer, an error watchdog, multi-threaded downloading, incremental downloads through keeping a history, caching to limit the stress on ESGF metadata server, additional actions (delete, cancel, retry), and a statistics module to report about remaining download volume, disk space requirements, etc. One of the team members regularly participates in ESGF developer meetings.

A critical point for the interoperability with ESG is the challenge of different authentication schemes for EGI and ESG, as they don't belong to the same federation. This situation, which is also relevant for the usage of the GENESI-DR infrastructure from EGI, was presented at the Workshop on "Federated identity system for scientific collaborations" in CERN on June 8-9 2011. Different solutions have been considered to make the use of ESG data with EGI easier. The software NDG Security Stack (http://ndg-security.ceda.ac.uk/wiki) has been set-up as a testing instance and has been reviewed. It was first developed for the NERC DataGrid and further on extended to be used as the Earth System Grid Federation security module. It includes support for OpenID, X509 and SAML; this security module is frequently used in climate data infrastructure. This instance is to be used for prototyping a solution for the authentication problem, which might be a new Credential Translation Service. The work is done in cooperation with Philip Kershaw (STFC, UK), who is working on the Earth System Grid Federation security model and software. The problem has been communicated to other EGI representatives in the form of a short internal report, describing the situation and possible approaches for a solution. Telephone conferences between representatives of EGI, SA3.6 and Earth System Grid Federation have been carried out, to clarify and discuss the situation and future plans. A

[19] http://dods.ipsl.jussieu.fr/jripsl/synchro-data

prototype modification of MyProxy is currently developed that will issue ESG certificates based on EGI certificate authentication.

While the GENESI-DR infrastructure can be used to discover many data sets of various data centres, and also includes now data from OGC web services, such as WMS instances, the data of the Earth System Grid is not discoverable through it. The search and discovery of ESG data sets is thus a separate issue, which is not a comfortable position for the users of EGI. It would be ideal to have single points of access for the search and discovery of data sets from the different sources, as well as a unified client to transfer the data, for easy inclusion in EGI job scripts or applications. The Earth System Grid, though, offers public access to THREDDS catalogues for the hosted data sets on a given Data Node (e.g. http://cmip-dn.badc.rl.ac.uk/thredds/esgcet/catalog.html for the BADC CMIP5 data). Therefore, solutions to unite the GENESI-DR search with a search capability for these dedicated catalogue instances have been searched. The applied solution is the GI-cat service developed by ESSI-Lab (http://essi-lab.eu/gi-cat), which uses a mediation approach to execute remote searches federating the results and remote catalogue crawling building up a local copy of the catalogues to be searched. As it evaluates and transforms the meta data of its sources to a common internal scheme, it can at the same time offer different catalog interfaces, allowing existing clients the access through different protocols (such as OGC CS/W, OpenSearch, GeoRSS and OAI-PMH). The software has been deployed at SCAI and the clients are tested against it. Unfortunately, the implementation of the protocol differs slightly, so the clients need to be adapted.

6.4 Current work and Outlook for 2012

The EGI Earth Science community is in contact with the FP7 project Virtual Earthquake and seismology Research Community e-science environment in Europe (VERCE), which aims at integrating a service-oriented architecture with an efficient communication layer between the Data and the Grid infrastructures, and HPC. A second novel aspect is the coupling between HTC data analysis and HPC data modelling applications through workflow and data sharing mechanisms. Discussions about effective and goal-oriented exploitation of Grid resources in the projects infrastructure have taken place. The projects members were suggested to register for the ESR general purpose Virtual Organisation to immediately be able to use computational resources. The setup of a dedicated VERCE VO or VRC is being discussed at present.

7 OTHER COMMUNITIES & SHARED SERVICES AND TOOLS

7.1 Grid Relational Catalog (GRelC)

The GRelC service is a Grid database management service aiming at providing access and management functionalities related to relational and non-relational databases in a Grid environment.

During the EGI-InSPIRE project, the GRelC software will support the HUC communities with a new set of functionalities available as web application through the GRelC Portal and the DashboardDB interface. The DashboardDB will represent the proper web access interface to the EGI Database of Databases. Such a “registry” will complement the functionalities provided by the EGI Application Database and will represent a distributed and multi-VO system supporting the HUC. It will be searchable and will allow people for instance to post comments and share opinions about the EGI data sources, ask to get access to a specific Grid-DB and join discussion groups.

7.1.1 Work accomplished

A major achievement is the first production release of the *DashboardDB registry*. It has been deployed and integrated into the official GRelC website. The functionalities tested, validated and now available to the end users in the official release are: user registration, user profile management, Grid-database registry, Grid-database search and discovery, discussion groups (including group subscription and notification), Grid-database rating and tagging. Moreover the *permalink* feature allows the registry to be easily “exported” (as a gadget) in other web contexts (the AppDB could be an example in the near future). This functionality has been exploited to integrate the DashboardDB registry into the official GRelC website, under the tab “Registry”. The back-end of the DashboardDB (developed during the first year) is also running in the production release.

Another important feature released in this period, is the *DashboardDB global monitoring* view which allows the monitoring of the GRelC instances deployed in EGI. Such a monitoring part complements the existing monitoring tools already available in the EGI environment, but provides a different perspective (which was missing before) focused more on “Grid-database management” services. By design and to address re-usability, also the DashboardDB global monitoring view is available as a permalink. This means it could be easily "exported like a gadget" in different web-contexts representing an add-on to more general EGI monitoring tools.

Another important progress is related to the *gLite 3.2 compliant release of the GRelC middleware*. The porting has been completed and the new RMPS are available on the GRelC website along with the needed documentation (both for end-users and administrators). An integration plan of the gLite 3.2 compliant release of the GRelC software into the Italian Grid Initiative (IGI) release has been defined and started. A new repository at the IGI level has been setup for this task.

Strong support has been provided to the HUC (in particular to ES and LS).

A use case (UNIPROT) related to the LS community has been jointly defined with the bioinformatics group working at the University of Salento and implemented over the past period. The UNIPROT data bank (about 13GBs flat file, now moved into 30 relational tables) is now available in Grid to this community through the GRelC service interface. A set of tools have been also designed and developed to support this use case and the management of the UNIPROT data bank. The porting of the UNIPROT data bank is an advantage because in a structured database the users are able to carry out complex queries that they were not able to perform on the flat file. Basically they are now able to run join queries across multiple information/concepts. Redundancy is also removed by definition, considering the intrinsic normalization process in building the relational data bank. Having the data bank into a relational database allows users to fully exploit the power of a declarative language like SQL. A key point is that we are porting several flat-file data banks into relational databases and this will allow researchers to cross-reference multiple (federated) data banks at the same time by running SQL queries.

The end users profiting from this are bioinformatics. In particular a group at the University of Salento is testing it for its own research activities. There are other potential users that could be interested in exploiting the new data bank and we'll contact them in the coming weeks, once we'll get some preliminary input from our beta-testers.

It is up to the end users to decide which version (relational or flat-file based) of the UNIPROT data bank they want to exploit for their purposes. If they are interested in extracting information by running complex join queries they would probably use the relational DB (as it is happening at the University of Salento), whereas if they want to run some existing data mining applications they will probably exploit the flat file. Of course, the two data banks (relational and flat-file based) can co-exist in the grid infrastructure.

Porting/adapting or developing new tools is up to the end users. SPACI is providing the ETL software to port the flat-file data banks into relational databases. It is comprised of a C library responsible for extracting the data from the flat file and for storing it into an XML file and the GRelC translator component responsible for translating the XML file entries into bulk queries for the new relational database.

Last but not least, the access to the UNIPROT data bank can take advantage of the features provided by the GRelC grid database interface (several kinds of delivery mechanisms, fine grained security control, etc.).

Another use case (IAS, *Invasive Alien Species*) in the Life Sciences domain and concerning biodiversity aspects has been defined in the last period jointly with the Biology group at the University of Salento taking into account their needs and requirements. This use case foresees the adoption of the GRelC middleware to support a platform for sharing invasive species information at a global level in the *Italian LifeWatch Virtual Laboratory*. Even though the use case is starting considering a single GRelC service instance at the University of Salento in Lecce (Italy) with some biodiversity databases, the long term plan (1 year) foresees more than ten nodes across the Italian country and hosting several databases exposed in Grid through the GRelC Grid-database interface. This interoperability exercise of data stored in distributed databases is meant to map the fragility of ecosystems to alien species invasion and estimating a long-term impact of alien species on biodiversity.

Concerning the Environmental domain, the GRelC service is exploited as a Grid metadata service at the Euro-Mediterranean Centre for Climate Change – (CMCC, Italy) to manage experiments/datasets metadata (compliant with ISO geospatial standards). Specific support for XML documents has been provided due to the complexity of the metadata hierarchy managed at the CMCC.

Visibility has been addressed updating the GRelC website and the wiki, giving oral presentations and tutorials, writing new papers and organizing sessions and workshops in ICT (e.g. International Conference on Computational Science - ICCS2011) and domain-related conferences (e.g. European Geosciences Union - EGU2011) as well documented in the Quarterly Reports.

Next year, to address sustainability, specific calls addressing HUC have been considered. For instance one of them related to the Italian LifeWatch Virtual Laboratory recently got approved. Another grant started in the second half of 2011 (for three years) concerns the extension of the GRelC software to include On-Line Analytical Processing (OLAP) functionalities for climate change data management. Other proposals including some extensions of the GRelC service to efficiently manage climate metadata have been also submitted too. Proposals submission represents a direct way to address sustainability and their acceptance can be seen as a concrete measure of the quality of the research activity related to the current work.

7.1.2 Current work and Outlook for 2012

The support in terms of management, monitoring and control of the GRelC services provided through the DashboardDB will be further extended and improved. The DashboardDB will provide new (service-specific) views about the GRelC deployment and the status of the registered services. A strong dissemination about the DashboardDB registry will be carried out over the next period to register the GRelC service instances already deployed in the EGI as well as the Grid-database resources they offer. Moreover, the questionnaire prepared during PY1 will be improved, extended and disseminated among the HUC to attract new users and make the existing ones aware of the new Grid-database registry tool. About the gLite 3.2 release, the integration of the GRelC middleware into the Italian Grid Initiative (IGI) release will be completed. A roadmap towards the *European Middleware Initiative Release* will also be defined.

Since the DashboardDB registry has been designed as a self-consistent component (it can be easily exported and embedded in other web contexts like a web gadgets) it will be publicized as a new *EGI gadget* at the European level (www.egi.eu/user-support/gadgets).

Additional support to the HUC will be provided to address new user needs and requirements. In particular the UNIPROT and IAS in the Life Sciences domain will be further supported. A major goal regarding the support for the LS community will be to create in Lecce (SPACI node) "*a Grid database node hosting several data banks addressing LS needs*". Concentrating at the same place many different biological databases will be relevant for the LS community and could be also crucial to attract new users.

Finally website, training events, tutorials, talks and papers will represent additional ways to disseminate the results of this activity.

The GRelC roadmap towards the European Middleware Initiative release can be considered an important milestone for the next period. Anyway, it can be considered a long term task since it will need continuous support to be implemented over the next years.

The DashboardDB registry as a whole can represent another long-term task due to the relevant role the registry can play in a production Grid environment. Moreover, even though it now contains Grid-database entries related to the network of GRelC services, in the future it could also address the needs coming from other tools in the same research area, becoming a more general and comprehensive Grid-database registry service.

For more information please refer to the GRelC section of the HUC Software Roadmap: https://documents.egi.eu/document/746.

7.2 SOMA2

SOMA2[20] is a versatile modeling environment for computational drug discovery and molecular modeling. SOMA2 is operated through a WWW-browser and it offers an easy access to third-party scientific applications. The SOMA2 environment offers a full scale modeling environment from

[20] http://www.csc.fi/soma

inputting molecular data to visualization and analysis of the results, and including a possibility to combine different applications into automatically processed application workflows.

During PQ5 CSC maintained and operated CSC's SOMA2 service. In addition, Autodock 4 integration work in SOMA2 continued. Also, we started investigations of how to setup a SOMA2 service which would also be provided to other user communities and not just for current CSC users. Program development effort was put in to upgrading the core UI library components and migration work was needed to be able to use the updated components. Work was also done to improve the basic Grid support in SOMA2 along with other minor fixes and improvements.

During PQ6 the focus was on a working Autodock 4 integration in SOMA2. Currently this provides a virtual screening service facilitating distributed resources. In addition, a new release of SOMA2, including the Grid support and more, was prepared. Basic setup for the SOMA2 service, which would be provided to other user communities as well and not only for current CSC users, also proceeded. However, this was slowed down by SOMA2 release not being ready yet. Also, CSC maintained and operated CSC's SOMA2 service.

The main outcome of PQ7 was a new version release of SOMA2. On 31.1.2012 SOMA2 version 1.4 (Aluminum) was released and made available on the SOMA2 web site. This release contains most of the EGI related development work so far consisting of the SOMA2 Grid integration and other improvements. Most of the work performed in PQ7 was put to finalizing and packaging the version release. Also during PQ7, CSC maintained and operated CSC's SOMA2 service.

The outlook for 2012 based on the achievements so far; we will enable a Grid-enabled application service for interested user communities. At the moment it seems that CSC is going to set up a separate pilot service for this. We will also operate and support the existing SOMA2 services. In addition, we will advertise the Grid enabled SOMA2 service to different user communities.

In a bit longer term our goal is to expand the scientific applications integrated in to the SOMA2 service, and integrate application services from different Grids hosted by different virtual organizations into SOMA2. This should be set as an important milestone because from the end users point of view, this would make using scientific applications in different Grids very easy and transparent. We will continue to maintain and operate the SOMA2 service and seek possible scientific applications to be added as part of the service. During all project years, we will support using SOMA2 services. Also development of SOMA2 gateway will continue according to feature roadmap including possible feature requests from the user community, bug fixes and other enhancements.

7.3 Workflow & Schedulers

During PY2 the work with Kepler (and Serpens suite for Kepler) has been focused on developing new template use cases basing on users requirements, improving the performance of the Kepler actors, and performing the tests of the different use cases. In addition thanks to the dissemination activities performed during the project the template use cases initially prepared for the Fusion community have been customised and reused by users from other communities (including astrophysics and computational chemistry). During the first year, several use cases were identified (see Table 2).

1. VMEC + DKES	2. VMEC + ISDEP	3. FAFNER + ISDEP
4. ASTRA + TRUBA	5. ASTRA + TRUBA + FAFNER	6. EUTERPE + ISDEP
7. GEM + ISDEP	8. ASTRA + GEM	9. ASTRA + GEM + TRUBA

Table 2: List of Fusion Use Cases

From this table, we have already built the template use cases that correspond with the following Fusion workflows: VMEC-DKES (1), FAFNER-ISDEP (3), and ASTRA-TRUBA (4). We also developed the CHEASE-MARSF use case (not in the previous list). This use case is the result of the success of the dissemination activities carried out.

The VMEC - DKES (1) case follows a parametric + parametric model. The number of different cases explored by the parameter scans can be particularly large in the case of the VMEC application, where many different parameters can be studied. DKES measures the transport of particles using an equilibrium previously calculated by VMEC. The initial work on this workflow has already started during PY1.

While testing this use case we realized that a large number of jobs were required for a single case. For example, just for the simplest case, each of the runs submits more than 2500 jobs (and the cases can only be increased). This number of jobs proved that an optimization in the actors handling jobs was required. Thus, all the statuses of the jobs use LB events level queries. This reduction improved the efficiency of the workflows, and increased the reliability of the whole run.

The FAFNER-ISDEP use case allows the movement of particles in the plasma to be studied. It follows a parametric+collector+parametric+collector model where the results from the first parametric scan are collected and put together so they can be used by the second parametric scan. The resulting workflow is as shown in Figure 2.

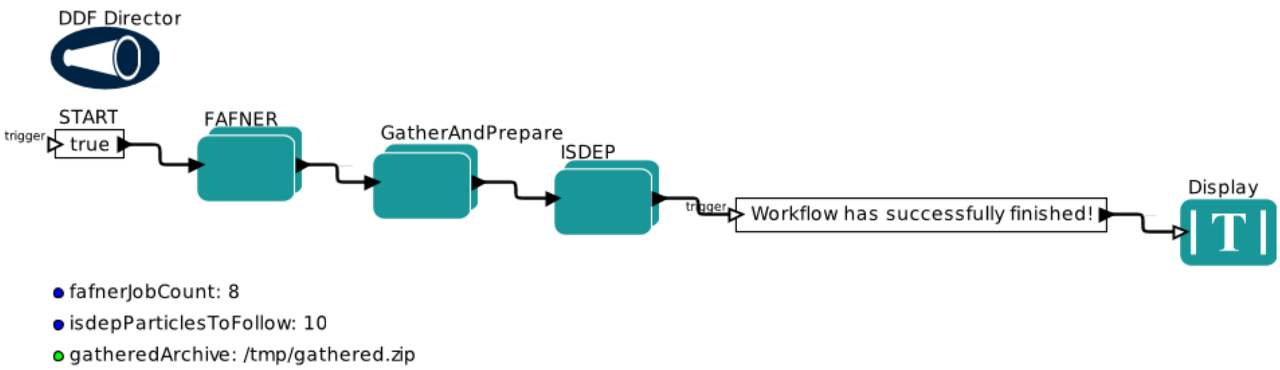

Figure 2: FAFNER-ISDEP workflow on Kepler

Both, FAFNER and ISDEP, use the actors previously developed that integrate the functionality to interact with the grid infrastructure. These actors can be seen in Figure 3. The actions of these actors change by using different values for the variables being considered for the different actors.

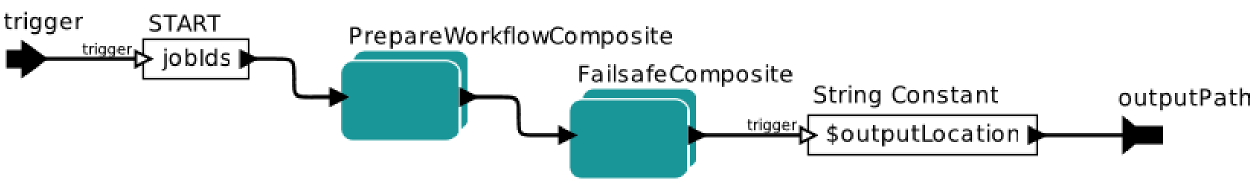

Figure 3: Grid Actors

The GatherAndPrepare actor, which retrieves the outputs from FAFNER and arranges these results so they can be used by ISDEP, differs from the actor used by FAFNER and ISDEP. This new actor can be seen in Figure 4. *$gatheredArchive* is the filename used for the file with all the results. This filename is also known by ISDEP.

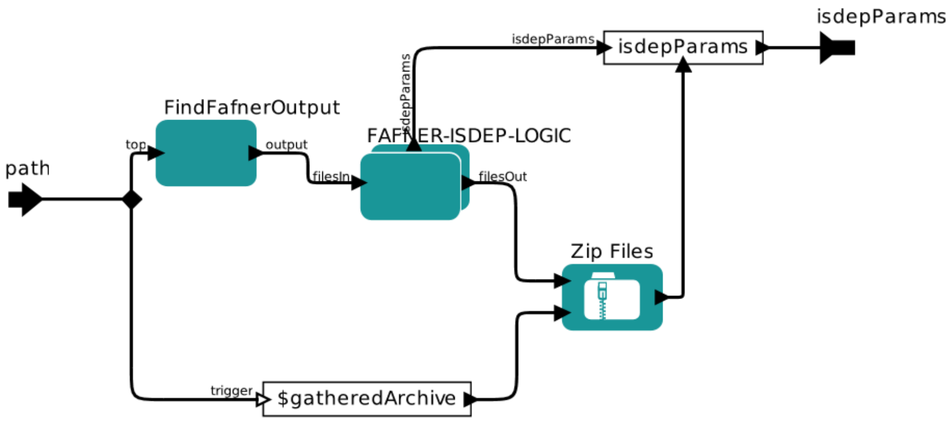

Figure 4: GatherAndPrepare actor

The first astrophysics workflow has been also customised and tested. It controls the production of realistic simulations of the anisotropies of the Cosmic Microwave Background (CMB) radiation. These simulations will be used to test the performance of various algorithms designed to detect the presence or absence of non-Gaussian features in the simulated CMB maps, before they are applied to actual data as observed by Planck, a satellite from the European Space Agency. In order to test the algorithms we need to produce large numbers of simulations. Each one of them is made of a combination of a Gaussian and non-Gaussian component plus realistic instrumental noise that takes into account the observing strategy of the satellite. The initial workflow follows the parameter study workflow template.

Also the first workflows from the CompChem area have been arranged and customised. This workflow follows the template of having application chain with the parameter study use case and the collector included. This use case is divided into two phases. First, the GFIT3C application fits the user provided ab-initio points of the potential energy surface. This is executed as a standalone task. The result is passed to the ABC, the quantum reactive scattering application, with a set of multiple alternative input files to be examined.

The dissemination activities carried out in different locations and meetings have allowed us to show the impact of the work to other communities. The improvements introduced in the actors to improve reliability and scalability, can be of interest for many other communities and can potentially extend the usage of this tool.

7.4 MPI and Parallel Computing

The MPI and Parallel Computing sub-task produces numerous MPI workbenches of increasing complexity with specific high impact on the Computational Chemistry, Earth Sciences, Fusion and Astronomy and Astrophysics (A&A) communities. These products are also intended to have an impact on other user communities. In addition, it focuses on ensuring that the user communities and site administrators benefit from several rudimentary improvements to the methodologies used and the available documentation. Many of these objectives are iterative, often requiring updates or fine-tuning. Other objectives, such as participation at the EGI Community Forum and the EGI Technical Forum, will be repeated at regular intervals. The core sub-task objectives (which bring definition to the tasks sustainability) are:

- Improved end-user documentation, addressing MPI application development and job submission in ARC, gLite and UNICORE;
- Quality controlled MPI site deployment documentation;
- Outreach and dissemination at major EGI events and workshops;
- User community, NGI and site engagement, gathering direct input;
- Participation in selected standardisation bodies.

To date, over 120 EGI Sites using the gLite/UMD middleware support MPI.

7.4.1 Work Accomplished

During the PQ5-PQ7 periods, TCD and UNIPG started independent work on exploiting GPGPU enabled devices. The proposed models have different implementations, however both focus on using hybrid grid/cloud solutions.

UNIPG worked on grid implementations of many common parallel scientific codes, including CHIMERE, RWAVEPR, ABC, DL_POLY, NAMD and VENU96. Many of these codes and their associated algorithms have direct re-use by other scientific domains. CSIC has contributed to the SIESTA computation chemistry application software.

CSIC has made continuous improvements to the MPI documentation wiki. CSIC has also tested MPI-START integration with SGE and PBS like job managers with newer versions of MPICH2.

CSIC is investigating incorporating support for MVAPICH into the MPI-START framework. MVAPICH has advanced Infiniband and fault-tolerance support.

TCD produced a proof-of-concept work using MPI-START (a script that encapsulates in a single command the different behaviours and interfaces to local resource management systems) to support generic parallel workloads. This will be demonstrated at the EGI Community Forum Parallel Computing Workshop. TCD will also demonstrate simple OpenMPI/OpenCL hybrid code running on the Grid.

The MPI SA3 team participates in the MPI-Virtual team. The formation of the team has help increase the visibility of the MPI task within the user communities, and a dedicated GGUS support unit has also been established to support sites and users. This has had a significantly good impact in ensuring that MPI related tickets are assigned to the proper expert team.

7.4.2 Current Work and Outlook for 2012

- TCD is investigating proof-of-concept execution of generic parallel codes using MPI-START. In particular using MPI-START to establish across the parallel computing resource the Hadoop "map/reduce" model, which is used in many "Big Data" problems and Charm++, an object oriented alternative to MPI.
- The current MPI-Virtual Team will come to an end in May 2012. The focused outcome of this groups work should be complete by then, including a better MPI testing infrastructure. The tests shall include better detection of job manager issues that can affect MPI/parallel jobs.
- TCD will propose a GPGPU virtual team. The aim of this group will be on how best to support GPGPU (and other generic resources) on the EGI grid infrastructure.
- An abstract entitled "Parallel Computing Workshop" was submitted to the EGI-CF 2012. The workshop is 90 minutes in duration and is led by Enol Fernández del Castillo (CSIC).

- CSIC is continuing to lead the MPI accounting sub-task to ensure correct accounting for MPI and parallel jobs.
- We will help produce better detection of PBS/Torque per-process limitations that affect parallel jobs at the EGI sites.

8 CONCLUSIONS

During the second year of EGI-InSPIRE, the SA3 work package built on and extended the achievements of PY1. Although domain-specific support is still required, further areas of commonality have been found and the communities understand that shared solutions are much more likely to be supported in the long-term than those that are highly-specific to a given VO. Further advances in this area are therefore possible for PY3 and beyond – some of these are highlighted in the Technical Evolution Group reports of the WLCG project.

Sustainability has been a key concern that has been addressed at the technical and strategic level. Sustainable does not, however, mean self-sustaining: all of the domains supported are dependent on external funding and this is committed in the short-term, planned in the medium and expected even in the (very) long-term. One of the youngest domains – HEP, which dates back some 60 years and is hence just an infant compared with e.g. astronomy – has a vision for the next 40 years, as presented in the CERN Director General's annual address in January 2012. Another positive is the trend towards highly distributed computing and "cloud-like" models is unlikely to be reversed.

The success of achieving a high quality of service with vast amounts of CPU delivered accompanied by hundreds of petabytes of storage and multi-gigabit networking across a heterogeneous federated infrastructure with loosely coupled management domains is now much more than an existence proof or even mere state-of-the-art: it is a reliable, performant and operational sustainable production system capable of meeting the needs of thousands of scientists worldwide 24 hours per day and close to 365 days per year.

9 REFERENCES

MS112	Quarterly Report 5:May 2011 – July 2011 https://documents.egi.eu/document/723
MS113	Quarterly Report 6:August 2011 – October 2011 https://documents.egi.eu/document/881
MS114	Quarterly Report 7:November 2011 – January 2012 https://documents.egi.eu/document/999
MS609	HUC Contact points and the support model: https://documents.egi.eu/document/419
MS610	Services for High Energy Physics: https://documents.egi.eu/document/540
MS611	Services for Life Sciences: https://documents.egi.eu/document/683
MS612	HUC Software Roadmap: https://documents.egi.eu/document/684
MS614	HUC Software Roadmap and Sustainability Plan: https://documents.egi.eu/document/746
D6.3	Annual Report of PY1: https://documents.egi.eu/document/312
D6.4	Capabilities offered by the HUCs to other communities: https://documents.egi.eu/document/472
TEG REPORTS	See http://indico.cern.ch/conferenceDisplay.py?confId=158775.
WLCG 2012	See https://indico.cern.ch/conferenceDisplay.py?confId=146547.

