5th EGEE User Forum

Monday 12 April 2010 - Friday 16 April 2010

Uppsala University

Book of Abstracts

Contents

Computer Science / **1**

Distributed parametric optimization using the Geneva library

Author: Ruediger Berlich[1]

[1] *Steinbuch Centre for Computing, Karlsruhe Institute of Technology*

The Geneva library implements parallel and distributed parametric optimization algorithms, capable of running on devices ranging from multi-core systems over clusters all the way to Grids and Clouds. The generic design makes Grid-resources available to user groups that have formerly not been exposed to such environments.

Detailed analysis:

Geneva ('Grid-enabled evolutionary algorithms') is a software library which enables users to solve large-scale parametric optimization problems on devices ranging from multi-core systems over clusters to Grids and Clouds. The generic design makes Geneva applicable to problem domains from a wide range of industrial and scientific application scenarios. From a user's perspective, parallel and multi-threaded execution can be achieved just as easily as serial execution on a single CPU-core, and does not require the user's evaluation functions to be aware of the parallel environment. Performance and extensibility are at the core of the C++-based, object-oriented design. The software has been shown to run in parallel with 1000 clients on a Linux cluster, each contributing a fraction of the overall solution. Given suitably complex optimization problems, scalability is almost linear. The code is available as Open Source, allowing customization under the terms of the Affero GPL v3.

Conclusions and Future Work:

The optimization environment will be further expanded to become a clearing house of different optimization algorithms, all based on the same data structures and using the existing framework for Grid-aware parallelization.

Impact:

The generic design makes Grid technologies accessible to users ranging from scientific to industrial application domains. It thus has a strong potential to bring new user groups to the Grid.

Keywords:

Optimization Grid Cloud

URL for further information:

http://www.gemfony.com

Justification for delivering demo and/or technical requirements (for demos):

The topic is sufficiently complex that it cannot be presented as a poster. Given sufficient network connectivity, a number of exa-catching demos can be presented.

Poster session / **2**

Eliminating and preventing Grid Security Vulnerabilities to reduce security risk

Author: Linda Ann Cornwall[1]

[1] *Particle Physics-Rutherford Appleton Laboratory-STFC - Science &*

Corresponding Author: linda.cornwall@stfc.ac.uk

The EGEE Grid Security Vulnerability Group was formed “to incrementally make the Grid more secure and thus provide better availability and sustainability of the deployed infrastructure”. The aim is to eliminate vulnerabilities from the Grid and prevent new ones from being introduced, thus reducing the risk of security incidents. This poster alerts users and developers to both the activities of the this group and problems that may be caused by vulnerabilities. It is also intended to inform what they should do to avoid introducing vulnerabilities and report any their find.

Impact:

At the time of writing, over 70 Grid Security Vulnerability bugs have been fixed since the activity started, and several pieces of less secure middleware have been taken out of use as more secure software comes into use.

As far as we are aware, no major Grid Security incident within EGEE or collaborating projects has occured due to a vulnerability in Grid Middleware. This could be due to the fact that we have successfully eliminated many vulnerabilities from the middleware, and successful in preventing new ones. It could also be due to hackers being less alert to Grids that other systems, or that the user community is particularly honest.

Keywords:

Grid Security Vulnerability Risk

URL for further information:

http://www.gridpp.ac.uk/gsvg/

Detailed analysis:

A system for handling Grid Security vulnerabilities was setup in 2006 at the beginning of EGEE-II. A Risk Assessment Team (RAT) was formed which currently has13 members from 10 different institutes to carry out this process, which is:
- anyone may report an issue
- the RAT investigates the issue
- if valid the RAT carries out a Risk assessment putting each issue into 1 of 4 risk categories - Extremely Critical, High, Moderate, or Low.
- a target date for resolution is set according to the risk
- an advisory is issued when the issue is fixed, or on the target date

Certain types of problem occur quite frequently - such as vulnerabilities resulting from incorrect file permissions, or failure to sanitize user input,. Suggestions for how to prevent new vulnerabilities entering the infrastructure will be made.

Some members of the group are currently participating in an ‘Overall Security Risk Assessment’ which looks at high level risks to the infrastructure, data, and various parties. A summary the reasons for doing this, the strategy, and the outcome (assuming it is available in time) will be included on the poster.

Conclusions and Future Work:

A good enthusiastic team has been handling Grid Security Vulnerabilities over the last 4 years, and many vulnerabilities have been eliminated from the deployed infrastructure. New types of vulnerabilities continue to be found and hackers get ever more ingenious in their quest to gain access to sites. This work therefore needs to continue into the future as grid technology increases it’s profile and hackers become more alert to this type of system and find new ways of exploiting systems.

High Energy Physics / 4

ATLAS Distributed Computing in Spain and Portugal: from data challenges to real data

Author: Xavier Espinal[1]

[1] *PIC/IFAE*

Corresponding Author: espinal@pic.es

The ATLAS PIC cloud is composed of the Iberian sites: PIC (Tier-1), IFAE, IFIC, UAM, LIP-Lisbon, LIP-Coimbra and NCG-INGRID-PT (Tier-2s) and is finalising preparations for the LHC data taking. To achieve readiness for data taking, all sites has been involved in the ATLAS Distributed Computing activities since early 2006: simulated event production, reprocessing, data and user analysis challenges and cosmic data taking. The evolution of the computing and operations activities from data challenges to the real data is described in this talk giving experiment and site experiences.

Detailed analysis:

After the deployment of the required services for the experiment, the Iberian sites now face the challenges associated with real data taking, and that is to provide a reliable service for data analysis for the world-wide scientific ATLAS community and particularity to the Iberian ATLAS scientists. The WLCG grid has been used so far by WLCG institute specialists and few scientists, but now it must be a daily common "tool" for every single LHC scientist with the first real data of the LHC appearing at the sites. The ATLAS institutes committed to ATLAS in the Iberian cloud have been preparing sites for the data taking since 2006, and achieving good results in all the wide activities that WLCG and ATLAS have been conducting. In this talk we will an overview of the ATLAS activities in Spain and Portugal since 2006, following the evolution of the different tools/frameworks that have been adopted by the experiment and the sites while awaiting the LHC start-up.

Conclusions and Future Work:

It has been a long road to have the first LHC data being analyzed at the sites. By the time of the 2010 Users Forum, sites should be flooded with new users willing to run their jobs. The real challenge is coming soon, when all of this people will be really grid-aware. The work until now has been exhaustive with the Iberian sites participated in all distributed computing activities, and we are now ready to face the data processing challenges associated with real LHC data taking.

Impact:

ATLAS distributed computing activities are reliant on the EGEE structure, the critical services abeing the catalogues (LFC), the transfers engine (FTS) and the Storage and Computing Elements that are crucial for the scientific community. Also ATLAS has important services running at every ATLAS cloud, such as the pilot factories and Frontier/squid servers which were recently deployed. All of these mechanisms have to be in place and in high availability to ensure the correct flow of data, and the execution of MonteCarlo/data processing jobs and user analysis jobs. In this talk we present the status and operations model of the ATLAS PIC cloud, following the evolution of ATLAS distributed computing and the improvement of the quality of service, which resulted from the increase of in robustness of the experiment distributed computing tools, and the continuous middleware improvements.

Keywords:

ATLAS, wLCG, LHC, Grid, PIC, Tier-1, IFAE, UAM, IFIC, LIP

URL for further information:

http://lhcatpic.blogspot.com/

Demo Session 1, Welcome Drink / 7

e-NMR: Computational NMR Infrastructure in life science and system biology

Authors: Marco Verlato[1]; Victor Jaravine[2]

Co-authors: Alexandre M.J.J. Bonvin [3]; Andrea Giachetti [4]; Antonio Rosato [4]; Anurag Bagaria [5]; David Galligani [4]; Eric Frizziero [6]; Gijs Van Der Schot [3]; Harald Schwalbe [5]; Henry Jonker [5]; Ivano Bertini [4]; Mirco Mazzucato [6]; Nuno Loureiro Ferreira [3]; Peter Guentert [5]; Rolf Boelens [3]; Sergio Traldi [6]; Simone Badoer [6]; Stefano Dal Pra [6]; Tsjerk Wassenaar [3]; Wim Vranken [7]

[1] *Istituto Nazionale di Fisica Nucleare Padua (INFN), Italy*

[2] *Goethe-University, Germany*

[3] *Bijvoet Center for Biomolecular Research, Universiteit Utrecht (BCBR), The Netherlands*

[4] *Interuniversity Consortium for Magnetic Resonance on Metalloproteins (CIRMMP), Italy*

[5] *Center for Biomolecular Magnetic Resonance (BMRZ), Goethe University, Frankfurt am Main, Germany*

[6] *National Institute of Nuclear Physics, Padua (INFN), Italy*

[7] *European Bioinformatics Institute (EBI-EMBL), Hinxton, UK*

Corresponding Author: zharavin@em.uni-frankfurt.de

In order to enable the life science community to make full use of the EGEE computing resources we have developed an e-infrastructure named 'eNMR' (EU 7th FP, Contract no. 213010).

eNMR deploys and integrates biomolecular NMR applications into a platform, so that EU scientists can easily access it via a standard browser interface and use it at every step of their research process.

The sequence of available web-portals covers all aspects of bio-NMR: data acquisition, processing, analysis, protein/DNA/RNA structure calculation, molecular docking, validation, deposition etc.

Detailed analysis:

Life-science researchers require the full range of advanced computational techniques in order to analyze complex biological data acquired experimentally at the EU NMR facilities. It becomes increasingly impossible to fulfill all of their requirements using the local computing resources.

eNMR is part of EGEE has developed solutions for performance computing, data access, authentication, security, accounting and usage statistics. Thus the e-NMR platform has all the strength of Grids, however, its focus is "science-driven". For example, in addition to already available applications, new ones are being identified and ported continuously, aiming to cover all aspects of bio-NMR; the enhancements to the grid middleware (gLite) are carried out to match the requirements of the applications.

Since the e-NMR portals have the same look and feel as bioinformatics web-gateways that have been available to the Life Science community for years, biologists can easily use the portals to execute their data analysis and structure calculation tasks without extensive knowledge of Grid technologies. The portals include documentation, tutorials and sets of use-cases from the worldwide labs.

Conclusions and Future Work:

An NMR e-Infrastructure has been successfully deployed at the various EU partner sites: the Universities of Utrecht, Frankfurt, Padova, Florence, and EBI-EMBL. The e-Infrastructure tests, compares, integrates and provides access to the computational tools for NMR structural biology community. Streamlined protocols and efficient workflows are being developed that enable the researcher to run in parallel their calculations, and to perform all the steps from the basic data to the structures of biomolecules and to a better understanding of biological systems and living organisms.

Impact:

Biomolecular NMR is truly a multi-disciplinary science comprised of all bio – physics/chemistry/informatics/tech/nano/om
etc. This has resulted in the fragmentation of research methods, preventing full interoperability among different labs, and thereby limiting the scientific impact of European bio-NMR research.

The eNMR by integrating applications enables scientists to address complex inter-disciplinary problems. Comparing the efficiency of different methods (CASD) leads to a more thorough and consistent analysis of experimental data, and thus a higher reliability of structural information.
The eNMR includes both established and emerging applications, so that scientists can depend on it to achieve their research and education goals and eventually could conduct potentially all of their computational work using it.
To use eNMR e-infrastructure capabilities via a user-friendly web-browser interface does not require extensive knowledge of Grid and IT technologies, thus increasing the potential user base from just bioinformatics to across life sciences. The number of users is growing and currently eNMR is the second largest VO in EU life sciences as measured by the volume of computation.

Keywords:

Life sciences, NMR, biomolecules, structure calculation, pipeline, web portals, structural biology

URL for further information:

http://www.enmr.eu

Justification for delivering demo and/or technical requirements (for demos):

We'd like to demonstrate the sequence of the web-portals implementing the applications for biomolecular-NMR; we shall need a big screen (preferably on a stand; to be connected to our computer).

High Energy Physics / **8**

Response of the ATLAS Spanish Tier2 for the collisions collected in the first run at LHC

Author: Santiago Gonzalez De La Hoz[1]

Co-authors: Alejandro Lamas [2]; Alvaro Fernandez [2]; Andreu Pacheco [3]; Carlos Borrego [3]; Elena Oliver [2]; Fco. Javier Sanchez Martinez [2]; Gabriel Amoros [2]; Jordi Nadal [3]; Jose Salt [2]; Jose del peso [4]; Juanjo Pardo [4]; Manuel Gila [4]; Marc Campos [3]; Miguel Villaplana [2]; Mohamed Kaci [2]; Pablo Fernandez [4]; Xavier Espinal [5]

[1] *Instituto de Física Corpuscular (IFIC)-Universitat de València-CSIC*

[2] *IFIC-Valencia*

[3] *IFAE-Barcelona*

[4] *UAM-Madrid*

[5] *PIC-Barcelona*

Corresponding Author: santiago.gonzalezdelahoz@cern.ch

The distributed analysis tests during the STEP09 and UAT exercises were a success for ATLAS and for the sites involved in the ATLAS computing model. The services were exercised at records level with good efficiencies and real users continued to get real work done at the same time.

Sample problems found were :

data access, was troublesome under heavy loads;
pilot factories need to be better organised;
monitoring could be improved

Solutions have ben developed for these problems

URL for further information:

http://ific.uv.es/grid/e-science/activeprojects.htm

Impact:

STEP09 and UAT exercises involved all major offline activities done in conjunction with other LHC experiments: Monte Carlo Production, Full Chain Data Distribution, Reprocessing at Tier1s, User Analysis Challenge and ATLAS Central Services Infrastructure. Those tests were a successful exercises for ATLAS and for the sites because they showed that data distribution was generally good, ATLAS central services worked well, analysis tested at very high rates at Tier2s and the reprocessing from tapes works well, with CMS concurrently active at shared Tier1s (for instance the Spanish Tier1).

STEP09 and UAT were a useful exercise for our Tier2 in order to solve the problems highlighted as soon as possible. It was the first time we had that level of feedback and information. Storage resources are sometimes undersized but it will not be a long-term problem; data transfer showed a timeout problem that may not be related to storageware, and Intra-VO fairshare (50% production, 50% analysis) was tested.

Keywords:

EGEE, CLOUD, LHC, ATLAS, TIER, GRID

Detailed analysis:

A Scale Test of Experiment Production was executed in June 2009 (STEP09) and in October 2009 (User Analysis Tests). The STEP09 full production activity stressed a number of critical areas, including tape writing/reading at Tier1 as well as analysis. Tier2 participated in Monte Carlo simulation and in user analysis components of the challenge. User analysis jobs in the EGEE cloud occurred through both WMS job submission and pilot jobs. User Analysis Tests (UAT) was a follow-on test to the STEP09 exercise and the last one before data taking. The goal was to get many user analysis jobs running over worldwide resources. This has the advantage of including potential problem jobs that might be missed in a more controlled test like STEP09.

The goal was to cover the major Tier2 activities: Monte Carlo production, data distribution and the user analysis challenge during the STEP09 and UAT exercises from a site point of view (in this case the ATLAS Spanish Tier2). The outcome of these exercises was that there were a number of areas where limitations were found. Improvements were made, defining the "final" WLCG operation environment that will be used for the first pp run of the LHC.

Conclusions and Future Work:

Spanish ATLAS Tier2 sites are ready, showed robustness, stability and good performance ready for the data taking.
The ATLAS computing system is ready as well. The Distributed Data Management (DDM) system improved during the last year and the PanDa Monte Carlo and User Analysis system increased global efficiencies and running stability.
The last updates will be made well in advance to have the sites ready for the LHC data taking, and to avoid big computing system interventions.

Poster session / 9

jGridstart: request, obtain, and install grid user certificates in a friendly way

Author: Willem van Engen[1]

[1] *Nikhef*

Corresponding Author: wvengen@nikhef.nl

The use of X.509 certificates gives flexibility in authentication and authorisation on the grid. The associated key is usually stored on the user's computer. While this is good practice from a security standpoint, managing keys and certificates is far from trivial. jGridstart attempts to bridge this gap by providing a friendly user-interface to guide the user in requesting, renewing, and installing user certificates. It is currently in use at the DutchGrid certification authority.

Detailed analysis:

Grid applications running on the user's desktop typically use Globus certificates to obtain access to the grid. In addition to this, the certificate needs to be present in the user's web browser to access online grid services. One of the first experiences with the grid are often the generation of a private key and certificate signing request,
submission of the request to the certificate authority, downloading the certificate, and installing it into the web browser. This generally involves file-manipulation and invocation of command-line utilities (openssl), which can be quite daunting to less seasoned grid users.

To make requesting and managing grid certificates straightforward even for new users, jGridstart has been developed with the following things in mind: (1) Friendly user-interface that detects the current state of affairs, and presents sensible options to the user; (2) Easy deployment using Java web start, which

enables low-level system access; (3) Single application (3.5MB) that has no external dependencies, apart from Java.

jGridstart takes care of the whole cycle from new request, retrieval, installation into the Globus certificate store and web browser, to renewal.

Conclusions and Future Work:

jGridstart is currently tailored towards the DutchGrid certificate authority's request process, and some work needs to be done to easily support other procedures as well.

The current approach is a single program that covers all aspects, providing a consistent user-interface. It would be beneficial to be able to use an adaptable web portal as user-interface, with jGridstart just covering the portions that require low-level system access.

Impact:

A user-friendly interface for managing grid certificates lowers the entry barrier for grid users on the desktop. Less time and effort is needed to get started with the grid (certificate requests), and to remain an active grid user (certificate renewals). jGridstart also makes the grid more accessible for less technical disciplines, where the command-line may be unknown altogether.

Keywords:

x509 certificate ca grid user interface usability ui java client

URL for further information:

http://jgridstart.nikhef.nl/

Infrastructure Tools and Services / 10

A solution to distribute software packages at Grid sites using Squid

Authors: Arnau Bria[1]; Elisa Lanciotti[2]; Ricardo Graciani[3]

Co-authors: Adria Casajus [3]; Gonzalo Merino [1]; Manuel Delfino [1]

[1] *PIC*

[2] *PIC (Spain)*

[3] *Universitat Barcelona*

Corresponding Author: elisa.lanciotti@cern.ch

In the current model of Grid computing the software needed by virtual organizations (VO) is stored at the sites in a shared area mounted on NFS and jobs running on the worker nodes access the necessary applications through NFS protocol. An alternative model, based on Squid, is proposed in this article. With this approach, the software packages are pulled from the central VO repository to the sites using http protocol. Squid also takes care of managing the transfers of packages from a central repository at the site to the worker nodes, and optimising network throughput over the LAN.

Impact:

This project has impact on both the virtual organization (VO) and the site infrastructure. From the VO perspective this solution provides an alternative to the NFS protocol which presents several limitations, mainly due to its not so good reliability and some compatibility issues with other applications, which result in a non negligible source of job failures. For the site infrastructure, this is an interesting alternative to the NFS file system, which can require a quite complicated deployment, whereas Squid relies on the very standard http protocol. Thank to its flexibility, Squid can be configured to optimise the data flow between client and server to improve performance and caches frequently-used content to save bandwidth. Squid also ensures scalability since more servers can be easily configured in load balance.

Keywords:

squid cache http computing farm software distribution

URL for further information:

http://www.pic.es/index.gsp

Detailed analysis:

A prototype has been set up to test the possibility of distributing the LHCb specific software packages to the worker nodes (WN) of the site through http protocol using Squid servers. In this framework Squid acts as http proxy and, in addition to that, provides caching functionality reducing the network traffic from the remote software repository of the VO to a central repository at the site. The test setup has a two level hierarchy of squid servers: The first level of the hierarchy consists of one central cache for all the site. The secondary level runs at the worker nodes where a local cache is provided. When a job lands on a WN it issues an http query to download the software tarball from the VO remote repository. If the request is already cached in the local Squid cache of the WN, the package will be immediately available with a disk to disk copy. Otherwise, the Squid server of the WN will escalate the request to its parent cache at the site. The central Squid server receives the request and, if cached, will send the response back through the site LAN, otherwise will forward the http request to the remote VO web server.

Conclusions and Future Work:

The feasibility of the proposed model of software delivery has been proved with a prototype setup at PIC Tier1 for the LHCb VO. Future work will aim to put this setup in production, first with a subset of nodes of the local computing farm, and tune Squid parameters in order to optimise the performance. In case the model proves to be more advantageous than the current one, based on NFS, we will consider to extend it to the whole farm.

Security / 11

An Active Security Infrastructure for Grids

Author: Stuart Kenny[1]

Co-author: Coghlan Brian [1]

[1] *TCD*

Corresponding Author: stuart.kenny@cs.tcd.ie

To date, grid security activities have largely focused on prevention mechanisms, i.e., authorization,authentication, and secure communications. Here we present an Active Security Infrastructure (ASI) for grids, the design of which focuses on the areas of detection (e.g. intrusion detection), and reaction, i.e., taking action to prevent, or to recover from, a security incident. The infrastructure is composed of a distributed monitoring and control layer and an analysis layer. Communication between layers is via a grid information system.

Detailed analysis:

Monitoring the current state of security is a vital task for those administering a grid. To secure a grid it is important that security information be available to site administrators in a timely and efficient way. ASI comprises a distributed monitoring and control layer and an analysis layer. The monitoring is performed using standard security monitoring tools (e.g. Snort) that have been enabled to allow any information gathered to be introduced into the grid information system. This information is analysed using an alert correlation approach based on the definition of attack scenarios. These scenarios model attacks as a sequence of steps, where each step takes the system from an initial secure state to a final compromised state. The output of the analysis is a single high priority correlated alert, possibly comprised from several lower priority alerts, and optionally an new grid policy. The grid policies generated by the analysis layer are distributed to the control layer, again through the grid information system, where they are enforced and any actions that should be taken to mitigate a possible security incident are applied.

Conclusions and Future Work:

ASI delivers an end-to-end monitoring and control solution. The monitoring component combines standard NIDS and HIDS, and so is able to detect a broad range of intrusion types. It also allows for the dynamic addition of extra tools. The analysis component provides aggregation and and correlation services as well as dynamic grid policy generation. The policies automatically generated by the analysis can be used to provide active control on grid resources. Our experience deploying ASI on the Grid-Ireland infrastructure has shown it to be a useful addition to current grid security.

Impact:

Grid-Ireland was established in 1999 to develop and coordinate the provision of a national grid service for the academic research community in Ireland. Grid-Ireland currently has a point-of-presence (Gateway) at 18 institutions. The day-to-day running of the Gateways is centrally managed by the Grid Operations Centre (OpsCentre) based in Trinity College Dublin. As the first stage in the full deployment of ASI on the Grid-Ireland infrastructure the monitoring component has been installed at 10 of the 18 Grid-Ireland sites. The analysis component is hosted at the OpsCentre. ASI has been continuously monitoring the Grid-Ireland infrastructure since June 2008, although several previous prototypes were deployed prior to this. At time of writing some 14 million alerts have been collected from the monitoring component to a security alert repository, the vast bulk of which are from Snort monitoring at the sites. Although the alert repository provides a useful auditing tool in its own right, clearly it would not be practical for a human operator to detect a potential attack from amongst such a large number of alerts. It is for this reason that the analysis component is required.

Keywords:

Grid information system, intrusion detection, IDS, alert correlation

URL for further information:

https://www.cs.tcd.ie/Stuart.Kenny/index.php?id=activesecurity

Fusion / **12**

Montera: a framework for efficient executions of Monte Carlo codes on the Grid

Author: Manuel Rodriguez-Pascual[1]

Co-authors: Ignacio Martín Llorente [2]; Rafael Mayo García [1]

[1] *CIEMAT*

[2] *DSA-Research, UCM*

Corresponding Author: manuel.rodriguez@ciemat.es

Monte Carlo codes constitute a powerful tool for scientific computing. Because of their architecture their parallelization is straightforward, and they have been successfully ported to the Grid in multiple occasions. However, there is still a lack of a deep analysis on their optimization for being executed on a distributed environment. To solve this issue we present Montera, a framework that efficiently executes this kind of codes on Grid infrastructures making the most of their particularities.

Detailed analysis:

A characterization of Monte Carlo codes and Grid sites are proposed, so their behaviour can be modelled. Then we have implemented Montera (MONTE carlo RApido - Fast Monte Carlo from its Spanish acronym), a framework that implements the aforementioned modellings with a 2 step dynamic scheduling to efficiently execute Monte Carlo codes on Grid. It is coded in Java, and employs DRMAA API to manage the execution of tasks. GridWay is the chosen metascheduler to control the Grid execution of tasks and provide Montera additional information about the status of the Grid infrastructure at any given moment.

Conclusions and Future Work:

Although the aforementioned framework already fits its requirements, there is still room for further development. An additional degree of parallelization can be added in order to reduce the time employed on analyzing the Monte Carlo code to be run and the sites that constitute the Grid infrastructure. Also, the proposed modellings can be modified to fit other kind of problems, thus becoming a useful solution for an efficient execution of a wider set of codes.

Impact:

Monte Carlo codes are widely employed in many different areas of knowledge. By creating a tool that greatly simplifies their execution on Grid infrastructures and at the same time improves their performance, we expect the scientific community to benefit from it. By developing the code in a tight collaboration with the users, we are able to also guarantee its correct functionality and usability. The establishment of a direct feedback among the developing team and the final users ensures the fitting of the proposed application to the specific needs and problems of the non-expert Grid users.

Keywords:

Scheduling; Monte Carlo; Grid.

URL for further information:

http://www.ciemat.es/portal.do?IDR=1481&TR=C

Poster session / 13

Checking Grid Certificate Profile Compliance

Author: David O'Callaghan[1]

Co-authors: Brian Coghlan [1]; Louise Doran [1]

[1] *Trinity College Dublin*

Corresponding Author: david.ocallaghan@cs.tcd.ie

Digital certificates are used to secure international computation and data storage grids for e-Science projects in EGEE. The International Grid Trust Federation has defined a set of guidelines for digital certificates used for grid authentication. We have designed and implemented a program and test suites to check X.509 certificates against profiles and policies relevant for use on the Grid to assist implementers and users of PKI to reach appropriate trust decisions.

Detailed analysis:

The IGTF has defined the Grid Certificate Profile for X.509 digital certificates used for grid authentication. As certificates are machine-readable and the Grid Certificate Profile is relatively well-defined, there is clearly scope for automation in checking compliance of certificates to the profile, which we believe is important for grid PKI scalability. We designed a practical tool for CA operators and grid sys-admins based on Perl and OpenSSL, both widely used in the target community. We investigated several approaches to defining tests and different certificate access libraries. We made use of the Perl Test / TAP framework to define the certificate test suites and added necessary functionality to Crypt::OpenSSL::X509 to allow us to implement the tests, which was a challenge due to poor OpenSSL documentation. We encountered some parts of the Grid Certificate Profile that pose difficulties for automatic testing.

Conclusions and Future Work:

Some bug fixes and complete coverage of the Grid Certificate Profile are necessary. Test suites for other requirements are desirable, as is an online certificate status or validation service. As the grid CA certs are easily available a web application could present current test results for all known CAs. In closing, the current version is a useful practical tool for grid PKI implementers and users, and indeed IGTF reviewers now make use of it when assessing CAs for accreditation

Impact:

The Grid Certificate Profile contains 88 distinct provisions, 8 of which were not implementable as automatic tests as they require comparing multiple certs, online checks, or subjective judgement. 69% of Grid Certificate Profile provisions were implemented as tests. Only 22 out of 91 IGTF CAs fully passed the test suite, although most failures were against recommendations rather than requirements. Compliance may be increasing over time but even some recently accredited CAs are not strictly compliant. This work is of considerable interest in IGTF where it is seeing use in the accreditation process. There exists complementary work in evaluating policies against requirements but there does not appear to be any well-known alternative to our work, which bridges the gap between policy and the practice of issuing certificates.

Keywords:

PKI, X.509, Trust, Authentication

URL for further information:

http://grid.ie/wiki/CheckCerts

Poster session / 14

Modelling aerosol pollution over Athens, Greece using grid technology

Author: Eleni Athanasopoulou[1]

Co-authors: Aggeliki Dandou [1]; Maria Tombrou [1]

[1] *NKUA, GREECE*

Corresponding Author: eathana@phys.uoa.gr

This scientific work aims at studying aerosol pollution over Athens, through an on-going modeling effort using grid technology. Aerosol predictions by the eulerian model CAMx, will be obtained for different emission scenarios. Focus is given on the role of conventional anthropogenic emissions versus natural emissions (sea-salt and Aeolian dust) on aerosol pollution. Predictions will be evaluated against aerosol measurements by the National Ministry of the Environment.

Detailed analysis:

This study involves air quality modeling, applied on HellasGrid CA through SEE user interface. Two domains are simulated for 5 days of 2007, which results in 45,375,120 spatial-temporal "points", where input-output data are treated. The computational time for this job could reach 30 days on a conventional local computational unit. The motive for using grid technology was a more efficient system for such simulations.
CAMx has no software requirements. Thus, the executable file was locally, statically and serially compiled and then imported into the grid system along with meteorological and pollution inputs, through gLite middleware.
Accidental run abortion results in job submission from scratch, which leads to delay of output retrieval. Corruption problems are better manipulated on local units, since the model can be restarted from where it stopped.
A local, static but parallel CAMx compilation is a further step, so as to take advantage of more computational units of the grid system.

Conclusions and Future Work:

A first attempt to use grid technology for aerosol simulations over Greece and Attica is presented. Anticipated results will offer scientific knowledge of the impact of different emission scenarios on aerosol chemistry and concentrations.
No substantial effort was needed to incorporate a CAMx model run into the grid. On-going simulations are based on a local, static and serial compilation. Parallel job compilation and execution will definitely give a boost in job evolution and output retrieval.

Impact:

Aerosol is an important contributor of pollution over the highly populated Attica peninsula, according to measurements. Aerosol sources are not only transportation and industry, but also Aeolian erosion, sea and re-suspension. Greece is an area with extended Archipelago and wind produces sea-salt particles. Additionally, the land surface is dry enough to offer dust bursts by the action of wind.
The contribution of each source to aerosol concentrations necessitates the modeling approach. In the current study performed in a PhD research, CAMx is applied over Greece, with a fine grid nested over Attica. Different emission scenarios show the importance of each source to aerosol predictions.
The on-going simulations are one with all emissions, and one without Aeolian erosion emissions. Their difference shows the role of Aeolian erosion on the atmospheric chemistry and on aerosol concentrations. Further simulations will be treated in analogy, so as to study sea and re-suspension contribution.

Keywords:

aerosol, modeling, Athens

URL for further information:

http://env.mg.uoa.gr/index.php?option=com_content&view=article&id=57&id=57&Itemid=72

Security / 15

Encrypted Data Storage

Author: John White White[1]

Co-authors: Akos Frohner [2]; Joni Hahkala [1]; Kalle Happonen [1]

[1] *Helsinki Institute of Physics HIP*

[2] *CERN*

Corresponding Author: john.white@cern.ch

Biomedical researchers are required to encrypt sensistive patient data before use/storage on Grid Data Management services. The services described enable this operation.

Detailed analysis:

Security and Data Management are two cornerstones of the
distributed production computing environment that the EGEE project is providing as a general e-Science infrastructure. An important requirement on the Data Management services
is the provision for securely storing sensitive data.
This privacy requirement has been given by the general Biomedical research community in conjunction with various national and international regulations and is met
through encrypting data and distributing the encryption keys.

The Encrypted Data Storage system is comprised of:

Hydra, the split encryption key storage and retrieval system;
one or more metadata catalogues such as the gLite LFC or AMGA;
a set of clients to communicate with
any GFAL-enabled storage element such as DPM.

The Biomedical research community
typically works with a far lower volume of data than, for instance, the High Energy Physics collaborations. Therefore, the components of the encrypted data storage have been designed and implemented with security rather than data throughput in mind.

Conclusions and Future Work:

From the experience of test services, we will describe areas of work that are ongoing or planned as dictated by the user feedback and project technical direction.
In addition, some upgrades and important bug fixes for more general Grid users will be presented.
Future work would inlcude extending this type of service to a Cloud infrastructure.

Impact:

This collection of services has allowed parts of the Biomedical community to conduct their research. Further upgrade and generalization of the service for other Grid users will enable others to protect data. As Cloud services become more prevalent, an encryption scheme such as this would add privacy.

Keywords:

Data Managment, encryption, keys

URL for further information:

https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS

Poster session / 16

gLite Porting to the Play Station 3 using ETICS for electronic High Throughput Screening (eHiTS)

Author: Eamonn Martin Kenny[1]

Co-authors: Brian Coghlan [1]; John Walsh [1]; Peter Lavin [1]

[1] *Trinity College Dublin*

Corresponding Author: eamonn.kenny@cs.tcd.ie

Since 2003 TCD has invested heavily in middleware porting, constantly engaging with the middlware development groups of the EU DataGrid (EDG), LHC computing Grid (LCG) and EGEE projects. In 2008, TCD ported gLite worker node to Yellow Dog Linux 6 on the Play Station 3, without data management. The node, built in ETICS, and tested in a production environment with gLite WMS job submissions were successfully submitted to it. The decrease in computational performance of porting the eHiTS software to the Grid is too expensive, so porting the Grid to other architectures such as the PS3 is important.

Detailed analysis:

As a result of the accumulated expertise, TCD was invited in 2006 to join a new EGEE-II integration and testing activity (SA3) as the main portability partner for further development in this area, and then for EGEE-III, Trinity College Dublin handled the porting and multi-platform coordination.

It was found that applications running on hardware such as the Graphical Processor Unit (GPU) and the Synergistic Processing Element (SPE) will incur very large performance hits if ported to the Grid. In particular, electric drug screening software such as eHiTS is one such application.

The meta-package generation facility of ETICS was proposed as a method to quickly generate a minimal gLite worker node for the PS3 that would accept glite WMS job submission but would not support data movement. Producing such a solution first involved porting VDT globus to the PS3, something not done before. The addition of a number of no architecture RPMs and VOMS then allows the node to be contacted by a standard gLite computational element (CE). However, a separate queue is required per

platform type to avoid confusion with other Linux platforms. A small cluster of 7 nodes is now running production jobs.

Conclusions and Future Work:

An integrated build and test environment (including external projects such as dCache, VDT and gridsite) on a very large set of platform types will allow EGEE to produce sustainable gLite clients for worker nodes and user interfaces. This is ongoing work at TCD, PSNC and CERN.

The exploitation and support of MPI on multi-core technology has proven useful in recent years. However, the GPU and SPE are relatively untouched at present, for highly parallelisable tasks in EGEE.

Impact:

The impetus to move the VDT globus source code stack into ETICS to allow it to be ported to exotic platforms such as the PS3 and openSUSE has resulted in the ability to build and maintain a version of VDT globus in EGEE. This in itself is a useful achievement.

The more portable the EGEE Grid is, to platforms such as RedHat Linux, Debian Linux, Mac OS X, Windows and PPC Linux (such as the PS3), the more likely it is that users will use the Grid, since they will not have to port their applications to the Grid.

eHiTS Lightning's computational speed is 26-fold to 60-fold faster on the PowerPC Cell Broadband Engine (Cell B./E.) compared with the equivalent application on an Intel-based processor. Therefore porting the software to the Grid makes little sense from a purely economical point of view.

Moreover, this immediately points to the fact that there is a whole application space in existence for which the EGEE Grid could be used, but isn't exploiting due to its lack of portability.

The use of GPU's is becoming much more cost effective. EGEE should therefore pay close attention to the exploitation of such technologies in the future.

Keywords:

PS3, portability, eHiTS, EGEE, SPE, Cell-Broadband, VDT, TCD, Grid

URL for further information:

http://grid.ie/autobuild/etics/org.glite/WN-3.2.0/yellowdog6_ppc64_gcc411/

Demo Session 2 / 17

Complex Scientific Workflows exploiting Grid and HPC

Authors: Antonio Gómez-Iglesias[1]; Francisco Castejón[1]; Marcin Plociennik[2]; Michal Owsiak[2]; Tomasz Zok[2]

Co-authors: Daniel López Bruna [1]; José Luis Vázquez-Poletti [3]; Álvaro Cappa [1]

[1] *CIEMAT*

[2] *PSNC*

[3] *Universidad Complutense*

Corresponding Author: antonio.gomez@ciemat.es

Simulating a full fusion Tokamak requires a range of codes and applications which address different aspects of the plasma and at specific ranges of space and time scales. The computational complexity of all these tools is so high that only paradigms like grid computing or HPC allow to carry out all the simulations. The EUFORIA project enables the European fusion modelling community to exploit Grid and HPC.

In this work we show a complex workflow that uses two applications that run in different computer architectures: grid and HPC.

Detailed analysis:

As previously shown by the fusion community, the usage of grid infrastructures for fusion research has provided interesting and relevant results that have created a large number of new possibilities to explore in the future. The Fusion community has been traditionally focused on HPC and it still reluctant to use grid computing, but with this work we show the feasibility to join both technologies and take the maximum advantages from these computational technologies. Our experience shows how to create and execute complex workflows to simulate plasma physics is a critical point nowadays and the key point for the future of the simulation in our field.

EUFORIA supports fusion modellers in this simulation work by providing programming expertise for porting and optimising codes on a range of computing platforms, providing grid and HPC resources, and providing the tools required to orchestrate the range of simulation codes necessary to simulate the full fusion reactor.

Conclusions and Future Work:

With the modifications introduced to Kepler, we can demonstrate the technical feasibility of launching jobs from this workflow orchestration tool to a mixed DEISA and EGEE environment. We show not only the technical feasibility, but also the scientific relevance of the results that can be obtained and the importance of this development in future researches.

Impact:

Kepler, a workflow orchestration tool, has been adapted by EUFORIA to enable fusion modellers to submit simulations to both Grid and HPC resources from their desktops and to visualise the results they obtain.

The project partners collaborate with DEISA and EGEE in order to ensure a wide adoption of the tools developed and deployed by EUFORIA in the infrastructure of DEISA and EGEE, by using the Fusion VO resources.

The possibility to create and run complex workflows, with some components running in HPC and some components running in Grid, with the use of the results obtained from these components by another component that can also run on Grid or HPC becomes a reality with these developments.

Keywords:

Fusion, Workflow, Grid, HPC, EGEE, DEISA

URL for further information:

http://scilla.man.poznan.pl:8080/confluence/x/mQAY

Justification for delivering demo and/or technical requirements (for demos):

The main feature of this work is the hability to execute complex workflows in mixed HPC-Grid environments. This is perfect for a demo where we can show the functionality of the developments.

Demo Session 2 / **18**

User-friendly access to computational services with FARO portal for EGEE and ENEA-GRID

Authors: Alessio Rocchi[1]; Andrea Santoro[1]

Co-authors: Andrea Quintiliani [1]; Carlo Sciò [2]; Francesco Beone [1]; Giovanni Bracco [1]; Salvatore Podda [1]; Samuele Pierattini [1]; Silvio Migliori [1]

[1] *ENEA*

[2] *S3S*

Corresponding Author: andrea.santoro@enea.it

A simple and flexible user access to applications is the final target for any system providing access to computational resources, even more for highly abstract infrastructures as grids and clouds. In this work we present a portal to the EGEE user that allows to submit widely used scientific softwares as jobs to EGEE. The main target of this interface is to hide from the end user low-level details connected with job submission. The expected result of this work is to make the GRID more palatable to other scientific communities, which at the moment are not interested in using this technology.

Detailed analysis:

FARO [Fast Access to Remote Objects] is a portal based on the integration of Open Source Software elements with a XML based Java application whose aim is to provide a simple, flexible and customizable Web access to user applications in grid/clouds context. It has been originally developed to provide user access to ENEA-GRID, a GRID-like infrastructure used to interconnect computational and data storage resources of ENEA. Many researchers from several different fields (nuclear fission and fusion, climate and ocean simulation, computational chemistry, CFD) use this vast array of computational resources to run their own simulations. Interoperability between EGEE and ENEA-GRID is in production since the beginning of the EGEE project. This work is a new FARO implementation that allows to submit to the EGEE grid jobs based on software widely used by the computational chemistry community (e.g. CPMD), all with few mouse clicks. The key feature of this interface is that the job submission happens in a very transparent and easy-to-use way through a dedicated portal that hides most of the low-level details to the user so that application access becomes really the access to a service.

Conclusions and Future Work:

Many scientific communities are not yet involved into the GRID, despite the potential benefits they might obtain from it. Part of the reason can be easily found in the intrinsic complexity of submitting jobs to the grid itself. You have to obtain certificates, write jdl files, often manage data transfers, etc. By the means of FARO, we aim at making specific softwares of those communities available to users through a common web browser (IE7, Firefox, Opera) and run them with a single mouse click.

Impact:

One of the main reasons why the GRID has difficulties in reaching a large part of the scientific community that might benefit from its resources, is because of its inability to be transparent to the end-user. The end-user (i.e. the “generic” scientist) is interested only in launching the code, possibly giving it one or more input files, and in retrieving the output. Ideally he would like to click on a “submit job” button, without concerning himself with certificates, jdl files, data transfers to SEs, and other low-level issues, and focus on the result of the job. By developing a user interface that 1) allows users to authenticate only once and 2) submit to EGEE some widely used softwares by using a graphical interface which hides most of the technicalities involved, we aim at interest the “average” researcher more into the GRID by reducing the prerequisite knowledge he should have about the EGEE inner workings, allowing him to focus more on his research.

Keywords:

User-friendly, interface, portal, open source, computational chemistry

URL for further information:

http://www.afs.enea.it/project/eneaegee/

Justification for delivering demo and/or technical requirements (for demos):

The aim of this work is improving the ability of users to interface with the GRID. Thus it is important for potential users in the interested communities to see and try the interface first-hand.

High Energy Physics / 19

LHCb operations: organizations, procedures, tools and critical services.

Author: roberto santinelli[1]

[1] *CERN/IT/GD*

Corresponding Author: roberto.santinelli@cern.ch

The proliferation of tools for monitoring both activities and the status of the critical services, together with the pressing need for prompt reactions to problems impacting data taking, user analysis and scheduled activities (e.g. MC simulation) brings the need of better organizing the huge amount of information available. The monitoring system for the LHCb Grid Computing relies on many heterogeneous and independent sources of information offering different views for a better understanding of problems while an operations team and defined procedures have been put in place to handle them.

Detailed analysis:

With LHC taking real data the "expert" user of the Grid left his place to rather many "normal" users in turn eagerly using the LHCb infrastructure to access data on the Grid. This suddenly introduced an impelling need to minimize service unavailability by proactively monitoring them on one hand and made immediately available all relevant information to debug problems either to the first line user support or to the shifter crew on the other hand. The definition of such a monitoring system is by far trickier than its implementation and has necessarily to start with an attentive analysis of the whole system in its complexity in order to isolate the relevant aspects. This abstract summarizes what is the state-of-art about LHCb Grid operations emphasizing the ultimate reasons that brought to various choices and what are the tools currently in use to run our daily activities: the most common problems experienced across years of activities on the WLCG infrastructure, the services with their criticality, the procedures in place some of them so well exercised and trustable to be made fully automatic, the relevant metrics to watch at, the tools available and what is still missing

Conclusions and Future Work:

This is a continuously working in progress activity. The evolution of various components necessarily brings to new unpredicted situations that have to be sorted out with new procedures or new monitoring tools. Nonetheless the system is getting more and more stable and the learning curve achieving its plateau. We ca not exclude that - far to be an ideal system - DIRAC and/or the grid middle-wares it is interfacing to will not evolve, perhaps dramatically. The organizations of LHCb operations has however to guarantee - during the LHC life period - the right quality of the access to the Grid.

Impact:

All above aspects can really be sorted out only after years of experience running a such complex system as the DIRAC system interacting with EGEE (or other) middle-ware stacks is. The learning process bringing the LHCb Grid Operations up to its current organization is by itself a valuable and certainly sharable with other EGEE communities. This is even more important if one considers that this is achieved also after numerous interactions with other HEP experiments having to face similar problematic situations too. The impact in LHCb of such a well trained and organized operations system for grid computing is crucial to permit the whole community smoothly accessing the LHC data on the distributed computing infrastructures and to run their analysis on top of that. The procedures defined so far, the alarming mechanisms in place, the multi-layered support structure, the various tracking systems adopted, the organizations of internal meetings, the adequate channeling of the information to/from services and resources providers, all of that might offer a useful operations-oriented platform for even more general use.

Keywords:

LHCb operations monitoring

URL for further information:

http://lhcbweb.pic.es/DIRAC/info/general/diracOverview

Earth Science / 20

Climate data storage in e-INIS

Author: Geoff Quigley[1]

Co-authors: Alastair McKinstry [2]; Brian Coghlan [1]; John Ryan [1]; Keith Rochford [3]

[1] *Trinity College Dublin*

[2] *Irish Centre For High End Computing*

[3] *Dublin Institute of Advanced Studies*

Corresponding Author: gquigle@cs.tcd.ie

We describe the federated national datastore activity of the e-INIS project, aimed at building a sustainable national e-Infrastructure for the Irish academic research community and how the CMIP5 project is using the datastore to meet their storage requirements. The datastore builds upon existing infrastructure and services, including Grid-Ireland, the National Grid Initiiative. Read access to the data is to be offered to international researchers using GeoNetwork and OPeNDAP, requiring that Grid technology be interfaced with community technologies using e-INIS's bridge servers.

Detailed analysis:

Using Grid technologies such as LFC and DPM allows distributed storage and replication of data on inexpensive hardware. The CMIP5 project requires the storage of approx. 100,000 netCDF files on 198 TB of storage,from mid-2009 to 2011 for the IPCC AR5 project. This data is being generated by EC-Earth climate model runs over November 2009 - December 2010. Data will be made available to other scientists and the public via http and OPeNDAP in 2010-2011, using a GeoNetwork catalog server. OPeNDAP enables optimised access to netCDF datasets, making it possible to download parts of files and access metadata for files. GeoNetwork provides a federated catalog service to ISO 19115 standards, fulfulling the INSPIRE directive requirements for public access, and enables data discovery via a web portal and clients such as Google Earth, NASA's Worldwind and its own GeoNetwork utilities. The metadata is then compliant with METAFOR conventions, for intercomparison with other climate models in the CMIP5 intercomparison as part of the IPCC Assessment Report 5. A subset of the data generated in Ireland is forwarded to other centres for comparison, but the majority is available only via e-INIS.

Conclusions and Future Work:

Bridge servers are being used to interface grid technology with servers that comply with user-community specified standards. The CMIP5 work is testing this with a reasonably large quantity of data over a period of time. As the project progresses it is envisaged that the various institutions involved will acquire 10Gb/s paths. The main technical challenges are establishing an integration in the bridge layer and scaling in step with both the size of dataset and speed at which it is being accessed.

Impact:

Integration of gLite and OPeNDAP in the e-INIS datastore is a large-scale demonstration of use of the e-INIS bridge servers to interface the existing back-end storage with community defined protocols. OPeNDAP is a very widely used standard but is better suited to access of subsets of large data-sets than to the bulk transport of data. The gLite middleware is better suited to secure transport of large quantities of data but does not provide the introspection features of OPeNDAP or comply with the relevant international standards for this climate modelling community of users. Other groups are looking at integrating gLite and OPeNDAP but we present a generic architecture that can be applied to other services, illustrated by the specific example of CMIP5 using this architecture to manage scientific data of international interest. The resultant system uses grid technology for back-end storage and to manage writing of data while providing a read-only front-end with a different security model that is compliant with international standards.

Keywords:

Grid, Storage, Data, Climate Model,

URL for further information:

http://www.ichec.ie/research/met_eireann#cmip5

Infrastructure Tools and Services / 21

Application Domain Accounting for EGI

Authors: Alvaro Simon[1]; Andres Gomez[2]; Antonio Laganà[3]; Aurelio Rodriguez[1]; Carlo Manuali[3]; Carlos Fernandez[1]; Esteban Freire[1]; Javier Lopez[1]; Ruben Diez[1]; Sergio Diaz[1]

[1] *CESGA*

[2] *Unknown*

[3] *University of Perugia*

Corresponding Authors: javier.lopez@cern.ch, esfreire@cesga.es, aurelio@cesga.es

Accounting is a powerful tool for users and VOs to obtain information on grid resources usage. Currently, they can access the accounting portal to display executed jobs. However, grid users and VOs prompt further improvements to enhance the EGEE accounting system and foster cooperative endeavors. One of the most requested features is application-level accounting. Together with COMPCHEM VO the initial requirements are being analyzed. It has suggested to gather extra information on the most executed programs, on the amount of retrieved results (per user/program) and on the used resources.

Detailed analysis:

The current EGEE III accounting system is oriented to gather job-level information from the batch system and merge it with VO and user-level information. This kind of information, however, is not detailed enough although there is demand on filling such a gap with application-level accounting. At present VOs have no information whatsoever on the software utilized and on the results gathered by the users. The development of application-level accounting tools will require the joint effort of different development teams along the next years. COMPCHEM and CESGA are beginning to design a preliminary draft which contains specific requirements about users and services parameters. To satisfy this new application-level accounting request, gLite middleware will require some improvements, to collect and store the new information in the EGI accounting repository database. On the other hand the new fields to be stored should be agreed and validated by Open Grid Forum (OGF) to become part of the UR standard. UF will be a good opportunity to present the work done allowing other VOs and SSCs to comment about it and add their requirements.

Conclusions and Future Work:

Application-level accounting was originally requested by COMPCHEM VO to improve its own portal, but this work can be of benefit for most of the EGI SSCs. It would be useful for users and VOs to establish in EGI a common science gateway able to record some information that is not collected by the current grid middleware. This is not an easy task, and involves the development of new sensors, the agreement between different EGI groups like ROSCOE and EMI, and the installation of these new sensors on grid sites. This should be discussed and developed over the next fewyears.

Impact:

The principal aim of this proposal is to further improve the EGEE accounting system. If these improvements are implemented in next years, they will have an immediate benefit for grid users and SSCs. One of the most visible improvements will be reflected in the EGEE accounting portal, based on the current internal application used at CESGA since 2003, a new method to gather and export the application usage will be used. This information will be aggregated, filtered and published for each site and later it be integrated in the accounting portal. Several new reports will be produced going from which application consumes more CPU time to which scientific area is using more computational resources. Other reports can be used to assist administrators, VO managers and users discovering application performance loss or application failures. Additionally, it is possible to detect in advance problems with the applications (because exit codes can be recorded) and which applications deserve more attention by the managers of the site, as well as, bad usage of the resources. One of our goals is to show and discuss this proposal to other SSCs to contribute accounting enhancement.

Keywords:

Application domain accounting,application-level resource usage,VOs and SSCs management optimization

URL for further information:

http://www.egee.cesga.es/Accounting

Astronomy and Astrophysics / 22

Dark energy in the Grid: the PAU survey project

Authors: Christian Neissner[1]; Nadia Tonello[1]; Santiago Serrano[2]

[1] *PIC*

[2] *IEEC*

Corresponding Author: neissner@pic.es

The PAU survey is intended to study the dark energy and its evidences such as the accelerated expansion of the Universe. With this aim, a large number of galaxies will be catalogued forming a map of the universe in large scale. High volumes of data need to be managed, processed and stored. The facilities of the PIC computing centre and the GRID environment will allow the challenging data management that the PAU survey requires.

Detailed analysis:

The PAU survey camera will scan a large area in the sky with an extensive set of filters. The data consists of astronomical images in FITS format. The process of obtaining the calibrated galaxy catalogues from the raw astronomical images is called data reduction. Simulated images are used as input for the validation of the process, since in the simulation the exact properties from the input galaxies and stars are known. The matching software used for validation is written in ROOT framework.
The reduction pipeline, written in python, consists of initializing the system, producing master calibration images and obtaining the catalogues from the science images.
As the code needs to access to a large number of files during the execution, the raw images, permanently stored in tapes, are temporally copied in a nfs disk, providing a significantly quicker and more efficient access.
A grid environment is the optimal location to carry out these computationally intense and loaded tasks.

Conclusions and Future Work:

The PAU survey will provide data of a huge number of galaxies that will be catalogued to form a new map of the universe in large scale. A high volume of data will be managed, processed and stored using the facilities of the PIC computing center and the grid framework. The data are intended to give new information for the understanding of the physics of the accelerated universe and the dark energy.

Impact:

We expect about 350GB of data coming from the telescope each night of observation so it would be likely to ingest, process and store this volume within 24 hours.
To reach the 24 hours time constraint, each science image reduction will be distributed to an independent job, so multiple cores will run in parallel at this stage. Such a technical challenge will be only possible into the grid framework.
The raw images and the final output catalogues are planned to be stored in tapes. The grid environment will make them available to all the VO members who, from their own institutions will be able to perform data analysis and cosmological studies using the unique quantity and quality of data that the PAU survey will provide.
The experience of PIC to deal with astronomical data, coming from the support it is giving to the MAGIC experiment, will be taken as a starting point and further optimized, for advantage of the two scientific groups.

Keywords:

Cosmology and Astrophysics, Data Management

URL for further information:

http://www.ice.csic.es/research/PAU/PAU-welcome.html

Bioinformatics / **23**

A Grid Implementation For Genetic Linkage Analysis

Authors: Andrea Calabria[1]; Davide Di Pasquale[1]; Gabriele Trombetti[1]; Luciano Milanesi[1]; Matteo Gnocchi[1]; Paolo Cozzi[1]

[1] *National Research Council - Institute of Biomedical Technologies*

Corresponding Author: luciano.milanesi@itb.cnr.it

The Genetic Linkage Analysis of SNP markers aims to discover the genetic correlation in monogenic diseases by following their inheritance in families through the generations. The computational cost and memory requirements of the major algorithms in literature make large data sets very hard to be analyzed on a single CPU. The work here presented is a Grid implementation of a data pipeline application for Linkage Analysis, a web based tool for Grid submission, monitoring and retrieval of Linkage challenges with large pedigrees and big markers datasets derived from genotyping chips up to 1M SNPs.

Detailed analysis:

Enabling the use of the EGEE Grid Infrastructure for the execution of linkage analysis on very large SNPs datasets is achieved through the implementation of a web application conceived in 3 main layers: the User Interface is designed to visually make up the pipeline for the linkage process, to set up the linkage challenges and the related otpions, and to monitor the workflow execution masking the complexity of low level interactions with the Grid middleware; the Application Layer executes data pre-processing operations like the retrieval of the SNPs data from the genotype database and the opportune formatting and splitting of the computational pipeline into Grid compliant jobs; the Submission layer, built on top of the grid middleware, monitors each single grid process and ensures the success of its computation by managing the automatic resubmission of failed jobs, bypassing certain Grid limitations such as the size of the input/output sandbox, while simultaneously optimizing the Grid storage and transfer overheads. When all tasks are computed, the results are retrieved, merged, showed to the user and made available for downloading through the web interface.

Conclusions and Future Work:

This Linkage Application enables the exploitation of the Grid infrastructure to set up and run analysis without a dedicated cluster and without the need of mastering complex informatics skills, achieving ease of use through the implementation of a dynamic and visually interactive User Interface. The data splitting approach has been tested as a valid approximation and the distributed implementation is mostly useful in high-end challenges, where Grid overheads are negligible with respect to parallelization benefits, making very large analyses accessible even without dedicated hardware.

Impact:

The Application is aimed at biologists with minor informatics skills, all the user interactions being implemented through an intuitive and interactive web interface where visual techniques are adopted to set up the system parameters and options. The system was tested with real case analyses to evaluate both the efficiency on computational time and the functionality of the data splitting approach, comparing its performances with a single 2 GHz CPU workstation and with a cluster composed by 280 CPU cores. The analysis, run using the Merlin software, involved different size datasets (10k - 1M of SNPs) and logged all phases duration of jobs life span. Results show that distributed analysis pipelines with big datasets can achieve a speedup up to more than 70x compared to a mid range dual-core 2 GHz CPU execution; the average performance shows a scaling trend comparable with the cluster due to the similar workload distribution technique adopted (not MPI). With the increase of the challenges size in terms of number of jobs, the variability in the total lasting time increases, due to the higher job resubmission rate, nevertheless the benefit in the distribuited approach gets even higher.

Keywords:

grid, distributed computing, linkage analysis

URL for further information:

www.itb.cnr.it/linkage

Medical Imaging / 24

Medical image indexing and retrieval in a distributed computing environment

Authors: Ihar Safonau[1]; Kanstantsin Kurachka[2]

Co-author: Vassili Kovalev [1]

[1] *United Institute of Informatics Problems of the National Academy of Sciences of Belarus, Minsk, Belarus*

[2] *Gomel State Technical University, Gomel, Belarus*

Corresponding Author: igorsafonov@gmail.com

The problem of indexing and content-based retrieval of medical images in a distributed computing infrastructure is discussed in the paper. It is considered in the context of a national-wide lung screening and diagnosis programme which is currently under development. High-resolution images are drawn from a test database containing results of X-ray chest examination of about 250000 people. The image content indexing and retrieval software is developed using GRID middleware UNICORE and a communication protocol MPI.

Detailed analysis:

Because of computational capacity of image processing algorithms and very large amount of visual data, it makes sense to use an approach capitalizing on a distributed computing architecture. A number of tasks were stood out including data communication, efficient allocation of computing resources, synchronization and optimization of computing process.
We tested the suggested software in a local network (100 Mbit/s) with 15 computing nodes. The cluster node configuration included computers with an 1.6GHz CPU, 2Gb RAM, Windows XP Professional. The Firebird was used as a database server. Experimentation revealed that the critical bottleneck in the environment was the data communication because of the large volume of transmitted visual data. Each compute node should receive a chunk of data, to compute them and send the result to the database. For minimization of data transfers, a special computational procedure was developed and implemented. In order to secure data safety and provide an authorized access to the private medical information we plan to utilize the security mechanism of the UNICORE GRID middleware.

Conclusions and Future Work:

The computational efficiency of calculation of image descriptors increases with increasing computation and it increases with decreasing transmitted data. The implemented approach allows the processing algorithms to be isolated from the specific elements of computing environment which provides a certain flexibility for software implementation and testing.
We plan to develop a computing cluster on the basis of a tuberculosis dispensary and then to build a dedicated computational GRID infrastructure for supporting large-scale image computing.

Impact:

An experimental study of indexing software on the test image sets containing 1000, 2000, 3000 and 4000 images has demonstrated interesting results. For the indexing of small number of medical images (up to 4000), it is best to use a dedicated computer without the organization of the distributed computing environment. In the GRID infrastructure, time is also spend on the exchange of information between the service components. A substantial contribution is the time for interaction with the database, which is located on a dedicated computer. The processes of the reception and transmission of the images or their descriptors also have an effect on the overall computational time too .
But with increasing medical images databases, using distributed computing becomes viable, and the benefit from the proposed scheme of calculation is becoming more noticeable.

Keywords:

Lung image, Content-based Image Retrieval, Indexing, Distributed computing

URL for further information:

http://safonov.by/grid/index.html

Data Management / **26**

Deployment and Management of Grid Services for a data intensive complex infrastructure

Author: Alvaro Fernandez Casani[1]

Co-authors: Alejandro Lamas Daviña [1]; Gabriel Amoros Vicente [1]; Javier Nadal Dura [1]; Javier Sanchez Martinez [1]; Jose Salt Cairols [1]; Victor Muñoz Mendez [1]

[1] *Instituto de Fisica Corpuscular (IFIC) UV-CSIC*

Services needed to access scientific data and computational resources in a distributed computing infrastructure are getting more and more complex. Our Atlas Tier2/Tier3 users, and National e-Science and GRID-CSIC initiatives researchers, have different requirements but all need ubiquitous and efficient access to these resources. We will provide details on how we use WLCG and Glite middleware, Storm, Lustre filesystem and our developments within this software to meet our specific necessities, to allow users direct access to the data, and to reach high availability in the provided services.

Detailed analysis:

We are currently providing 590 TB of online data and potential capacity of 560 TB in offline data tapes, available to several communities which users in some cases intersect. For example we have Atlas Tier2 users that are also members of our Tier3, and we need to apply group and user quotas for the different spaces. National Grid Initiative and Grid-CSIC communities have also common members and applications, in some cases parallel computational jobs that also need direct access to data. Authorization for these data is done at the filesystem level with the usage of ACLs, and to be respected by higher levels of middleware we had to develop custom plugins for Storm for example. Lustre is mounted in the WNs to directly access with posix file calls, and local users have direct read-only access from the UI. We also provide X.509 secured web access for commodity. The data access patterns for newer applications outside the HEP scope is being analyzed, as well as optimal parameters like the read ahead for different types of workloads. As we are reaching higher levels of served data, we are scaling the number of Lustre OSTs, and distributing grid middleware services among several machines.

Conclusions and Future Work:

The presented approach has proven to work fine and meets the required quality levels for the operation of the different supported virtual organizations, but in order to constantly improve, we are analyzing the usage of Lustre capabilities to access the geographically distributed data stored at the 3 sites of our distributed Tier-2. We are also planning to support 3rd level data tape storage, using the future developments of the Lustre/CEA with HSM when mature. We also intend also to continue supporting our developments and making it available to the community.

Impact:

Recent tests as the STEP09 and UAT tests successfully exercised our infrastructure for the Atlas Tier2 in conjunction with other LHC experiments. The usage of Lustre as the backend cluster filesystem has proven to reach excellent performance levels, in addition to the configuration that allow jobs at the WN to directly access data in the servers. These tests also showed some problems mainly related with underestimation of the amount of stored data, which are being easily solved adding more storage servers thanks to the scalability of the services. The inclusion of several newer e-science Vos in the context on NGI and GRID-CSIC projects have reinforced the need for group and user quotas, and it has been adopted the new pool paradigm of Lustre 1.8 to group several OST, proving to work fine. Development

of newer software plugins for Storm to allow local authorization policy based on ACLs for the different spaces, was also adopted in the main source code, making it available for other Grid Sites to use it and watched as a proof of concept. Also Grid services are improving, and storm team is aware of the need to support group and user quotas at the srm level.

Keywords:

Grid, Lustre, Data Access, SRM, Middleware development

URL for further information:

http://ific.uv.es/grid/e-science/activeprojects.htm

Novel Architectures and Technologies / **27**

Offering GridWay users more power by means of the Amazon public cloud

Authors: Alejandro Lorca[1]; Jose Luis Vazquez-Poletti[1]

Co-authors: Eduardo Huedo [1]; Ignacio Llorente [1]

[1] *Universidad Complutense de Madrid*

Corresponding Author: alejandro.lorca@fdi.ucm.es

We present a mechanism to easily provision public cloud resources for grid users. The extension of the underlying grid infrastructure benefits demanding situations coming from a single user, a group belonging to a Virtual Organization, or even from a institutional requirement. A set of very simple tools allows the GridWay administrator to deploy arbitrary instances and monitor how the enrolment is performed, guaranteeing the usability of the new resources to the specific target community.

Detailed analysis:

Major work is being carried out nowadays in order to integrate cloud and grid, into a flexible, on-demand, and heterogeneus computing infrastructure. Such approach could overcome the stringent requirements to be met at grid computing. Moreover, the latencies coming from long awaited queues, diminish the operational efficiency of time-critical studies. Those experiments therefore, would not consider the trade-off from porting the application to the grid. By means of deploying more machines coming from a public cloud provider as Amazon EC2, we show how users at different grouping levels can profit from the instantaneous provisioning of such services. Indeed, it is left to the GridWay administrator (i.e the administrator of the grid user interface), the responsability of safeguarding the public and private keys coming from the Amazon EC2 billing account. A new set of commands, similar to the ec2-tools, let the administrator determine, launch, monitor and enrol the new machines into the GridWay information system. The connectivity between GridWay and the virutal new instances is handle with the ssh middleware access driver, ensuring therefore safe and encrypted data transmission.

Conclusions and Future Work:

We show how the launching, provisioning and addition of public resources might help users with specific and highly-demanding requirements to speed up their jobs. The extra work needed to set up such infrastructure extension is assumed by the administrator, who gets at his disposal new tools that make straightforward the whole process. Future work is oriented towards the automation of the mechanism, to establish a clear protocol for getting access to the Amazon EC2 account data, and new policies for self-acquiring resources while monitoring the GridWay jobs' queue.

Impact:

This study provides grid communities more arguments to consider the possibility of including public cloud as a new actor, besides traditional computing centres, for extending their existing infrastructure. This extension is not free, requiring a shift from the static scientific model of budget allocation. At this point, we consider necessary a debate whether on-demand resources are required as utility computing

and should be addressed as running costs depending on the relevant sciencitic analysis.
From a technical point, we consider a new communication mechanism to be integrated into the GridWay metascheduler, allowing much faster using of grid resources. We believe this is one of the key points where grid computing is still far off the desirable standards.

Keywords:

metascheduler, grid, public cloud, ssh

URL for further information:

http://gridway.org, http://aws.amazon.com

EGI Session / 28

EGI InSPIRE HUC SA3 - Services for Heavy User Communities of DCIs

Authors: Jamie Shiers[1]; Maria Girone[1]

[1] *CERN*

Corresponding Author: jamie.shiers@cern.ch

At this last User Forum of the EGEE project, it is appropriate to consider how services for "heavy users" of Distributed Computing Infrastructures will continue to be provided. The role of the European Grid Initiative (EGI) has been discussed for some time now and a specific activity regarding Service Deployment is foreseen as part of the EGI InSPIRE ("Integrated Sustainable Pan-European Infrastructure for Researchers Everywhere") proposal. This talk explains which communities are targeted by the work that is foreseen and outlines mechanisms whereby its progress can be tracked. Moreover, it examines how the potential benefits – within individual communities, between them (e.g. common tools and services) as well as to the more general DCI community – can be achieved.

Detailed analysis:

The communities identified as Heavy Users Communities (HUCs) within the EGI InSPIRE proposal 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)
After briefly summarizing the main areas of work that are foreseen, this talk examines potential interactions between the communities and with the various "Virtual Research Communities". Using well-known case studies, it attempts to quantify the value of previous spin-offs – such as the CERN Program Library as well as more recent and ubiquitous examples – and argues the case for long-term international and inter-disciplinary collaboration, either between related domains and/or based on shared technology.

Conclusions and Future Work:

By highlighting the value of interdisciplinary and international collaboration, this talk will emphasize both the short-term benefits of services for user communities that are heavy users of distributed computing infrastructures, as well as the much broader socio-economic spin-offs that clearly justify the associated investment.

Impact:

The planned work addresses first and foremost the short to medium term priorities of the supported communities. This includes efficient exploitation of international distributed computing infrastructures which in turn leads to improved return on investment for the agencies funding the associated scientific work. However, the expected impact goes much further, both in terms of number of people affected,

as well as the timescale of the project itself and the costs will be examined in terms of both direct and indirect benefits.

Keywords:

EGI INSPIRE, HUC, SA3

URL for further information:

http://sites.google.com/site/egiinspiresa3/

Data Management / **29**

StoRM, a flexible Storage Resource Manager solution

Authors: Alberto Forti[1]; Riccardo Zappi[1]

[1] *INFN-CNAF*

Corresponding Author: riccardo.zappi@cnaf.infn.it

StoRM is a high performance and flexible SRM solution for general disk based storage resources that bring in Grid the advantages of high performance storage systems based on cluster file system, such as GPFS from IBM and Lustre from SUN. StoRM can be used in small data centre with low rates of storage management requests and, at the same time, it is capable to grow in terms of storage and workload managed to fulfil requirements from large-scale production sites. The latest version of StoRM satisfies the wish list built with the desiderata of current and future users.

Conclusions and Future Work:

StoRM project will be part of EMI,and so it will be supported as a software solution in the European grid middleware. Moreover, a special focus will be placed in ensuring the fulfillment of any new requirements coming from the many users who are using, or will adopt, StoRM as a solution.

Impact:

Being able to use the file protocol allows to use the now increasingly popular cluster file systems like GPFS and Luster. The jobs can directly access the files needed without having to transfer them to the local disk space. The advantage in terms of performance and network load is very great. StoRM's architecture also allows a simple scaling model, as it is possible to add Front-End components if needed, making StoRM usable even in large centers such as the Tier-1. One of the last feature is developed to support hierarchical storage solution when they are configured as GPFS and TSM. This is the solution adopted by the Italian Tier-1.

Keywords:

SRM, Storage Element

URL for further information:

http://storm.forge.cnaf.infn.it

Detailed analysis:

StoRM is a storage resource manager implementing the SRM interface version 2.2. It is designed to work on generic disk based storage systems separating the data management layer from the underlying storage systems. StoRM provides a flexible, configurable, scalable and high performance SRM solution. It supports standard Grid access protocols as well as direct access (native POSIX I/O call, that is "file://" protocol) on data fostering the integration of non Grid aware application providing local access on shared storage. Another important characteristic of StoRM is the it doesn't use a database to store the namespace, as done by other implementations as dCache and DPM, but uses an algorithm working only in memory. The physical location of a requested data is computed without querying any database service but evaluating a configuration file, an XML schema that describes the storage namespace and input parameters as the logical identifier and SRM attributes. Indeed, StoRM relies on the underlying

file system structure to identify the physical data position. For this reason we consider StoRM a solution which is efficient and reliable.

Poster session / 30

The HELIO project.

Author: Gabriele Pierantoni[1]

Co-authors: Brian Coghlan [1]; Eamonn Kenny [1]; Geoff Quigley [1]; John Ryan [1]

[1] *Trinity College Dublin*

Corresponding Author: pierang@cs.tcd.ie

HELIO is a project of the 7th Framework Program that aims at creating a collaborative environment in Heliophysics. As it includes different services that can be computationally intensive and involve large, dispersed volumes of data, HELIO will use a Grid-based system for processing and storage for its most intensive analysis.
The architecture must balance conflicting requirements: to have a powerful, yet easy to use, instrument, to comply with the policies of the back-ends, and to be flexible in orchestrating, locating and co-locating the services.

Keywords:

Heliophysics, grid service, storage, workflows

URL for further information:

http://www.helio-vo.eu/

Detailed analysis:

HELIO will provide a set of instruments with which scientists can discover, understand and model the connections between solar phenomena, interplanetary disturbances and their effects on the planets. It does so by orchestrating services that help to identify and analyze events and observations across the solar system, and to correlate these using propagation models.
HELIO must be simple to use, yet ideally its architecture must support different workflow engines to orchestrate its services, it must encompass Grid services fully complying with their policies and, must allow the different services to be deployed flexibly.
HELIO strives to tackle all this requirements with an architecture encompassing different layers:

• The user access layer, an abstraction of all the possible ways in which the different services can be accessed, such as a local workflow engine, the HELIO web site or programs written by the users.
• The service access layer, an abstraction of all the possible deployments of the services that also allows the independent use of some of the services
• The services themselves, with interfaces for remote and local access.
• The back-ends upon which some services rely.

The back ends encompass data sources; web sites that provide specific computational facilities and, more importantly for this paper, grid middleware based back ends.
Grid Middlewares will provide three main functionalities: security, computation and storage.

• A Security Service, called Community Interaction Service, will provide voms-proxy certificates that will be used for authentication and authorization.
• A Processing Service, partially based on Grid Middleware, will offer a unified service for execution. The Grid-based back end of this service is based on the gLite job submission system.
• A Storage Service, based on Grid Middleware, will offer storage for the HELIO system. It will be based, at first, on the gLite storage services but, potentially, will be later connected to an iRods service.

Conclusions and Future Work:

HELIO started in June 2009 and, as of now, the overall architecture is just being finalized while prototypes of the most important services are being developed.
A first official release of the project that encompassed the prototypes of several services without the Grid-based back ends was released at the end of January 2010.
A second release, due at the end of May 2010, will encompass prototypes of the Community Interaction Service and the Processing service that will be completed and joined by the Storage Service in the third release of HELIO that is due at the end of September 2010.

Impact:

HELIO's impact will cover three main fields. It will have an impact on Heliophysics helping to foster cooperation among different scientific communities, it will have an impact in the Computer Science community as the problems it is trying to tackle are common in projects that must balance similar constraints, and finally it will provide a useful testing case for the Grid back-ends.
This paper focuses on the technological aspect of the project.
If successful, the HELIO architecture could be adapted as a solution to similar problems that deal with large, distributed data sets arranged in unforeseen topologies, intensive computational tasks and the need to offer the functionalities of a complex system in different ways, e.g. with a workflow or used independently.
It could also prove useful in implementing systems that expose a simple interface to users that want to use complex instruments but do not want to be bothered with technical details.

HELIO will also use the Grid facilities of Grid Ireland. Of particular interest will be the assessment of how well the Storage Facility being developed in the national e-Inis project will act as a back-end to the Storage Services of the HELIO project

Scientific Gateways and Portals / 31

An advanced web portal for accessing Grid resources with Virtual Collaboration features.

Authors: Andrea Del Linz[1]; George Kourousias[1]; Milan Prica[1]; Roberto Pugliese[1]

Co-authors: Daniele Favretto [1]; Fabio Bonaccorso [1]

[1] *Sincrotrone Trieste S.C.p.A.*

Corresponding Author: milan.prica@elettra.trieste.it

The Virtual Control Room (VCR) is an open source web portal that puts together a rich collaborative environment with the simplified access to the gLite Grid resources. The latest version of the VCR is based on the Gridsphere 3 and Google Web Toolkit (GWT). It uses the DORII Java Common Library for accessing Grid resources, integrates DORII Workflow Management System, presents a much improved Application Manager, introduces tags for user-application mapping and support for robot certificates.

Detailed analysis:

Registered portal users may access grid resources from the VCR 3.0 using their personal certificates, or the portal's robot certificate, the latter an approach that proved to be most useful for occasional users of the infrastructure. Users are linked to the various projects through the VCR tags so each user is presented with the correct set of resources that he is entitled to use, and his proxy certificate has the correct VOMS attributes set automatically. Integration of the scientific instrumentation is provided through a graphical Instrument Element client. The VCR's tunneling allows for remote access to the legacy control system and supports interactive application through visualization of the i2glogin client in a user's browser.

Conclusions and Future Work:

Future enhancements will be based on the feedback provided by the users of the DORII and its follow-up projects (DORII+). One of the first additions will be support for the visualization using Gvid. Other

improvements will come from the developments of a pure Java client for the gLite middleware that is the DORII common library.

Impact:

VCR is the main user interface adopted by the DORII project. DORII focuses on application from three different fields of science (Experimental, Environmental, Seismic) which provide for a wide and diverse user base. The VCR portal has already been successfully applied for applications like the on-line and batch data analysis in experimental science, oceanographic and coastal observation and modeling (using imaging or through Mediterranean Ocean Observing Network) and network-centric seismic simulations. VCR is endorsed by EGEE's RESPECT program.

Keywords:

Grid, Portal, e-Infrastructure, Scientific Applications, Virtual Collaboratory, Workflow

URL for further information:

http://www.dorii.eu/

High Energy Physics / **32**

The Grid as an Extended Application Service Provider in an Synchrotron Radiation Facility

Authors: Andrea Del Linz[1]; George Kourousias[1]; Milan Prica[1]; Roberto Pugliese[1]

Co-authors: Alessio Curri [1]; Daniele Favretto [1]; Fabio Bonaccorso [1]

[1] *Sincrotrone Trieste S.C.p.A.*

Corresponding Author: milan.prica@elettra.trieste.it

The Synchrotron Radiation Facilities (SRF), as large research establishments, have a very important role in Science and a great impact in the community. Due to the data-parallelism of the computational problems in the general field of physical sciences and the very high data volumes in terms of storage, the Grid Computing has been a successful paradigm. But other than their Computational requirements there are additional Grid characteristics and features that can be highly useful. We discuss the Grid utilisation in the SRF ELETTRA.

Detailed analysis:

The first section of this paper presents the use of suitable Grid technologies for beamlines and labs of the Synchrotron Radiation Facility ELETTRA. The second part is focused on novel technologies (Instrument Element) and Grid characteristics (on-line), other than that of HTC, that can be very useful especially for SRFs. The general modus operandi of these facilities is that they host laboratories/beamlines that have resident scientific personnel that assists visiting scientists to perform experiments. The visiting scientists often have to use software that is preconfigured for the tasks of the beamline. We demonstrate how the Grid can be used as an Application Service Provider to serve the ELETTRA in-house and visiting users with an extended Software as a Service model that includes the Computation and secure user Authentication. The beamlines that will be discussed are in the field of Computed Tomography, Medical Physics, Small Angle X-ray Scattering, and X-ray Diffraction. The Grid infrastructure that is utilised in based on gLite and the user front-end to the Grid is through an advanced web portal with Virtual Collaboration features.

Conclusions and Future Work:

The future plans aim to the improvement and the extension of the above-mentioned Grid technologies. Particularly there are on-going efforts for the advancement of the remote instrumentation control technologies. Additional beamlines and labs in ELETTRA will adopt these technologies. The web portal for Grid access (VCR v.3) is in active development. We expect additional results in user management especially for the case of visiting users.

Impact:

By deploying a suitable and customised set of Grid technologies for Synchrotron Radiation Facilities there may be a direct positive impact to the hosted scientist. In the described deployment in ELETTRA, the interaction with the computational resources has been greatly improved. The user can use a specialised form of Grid Certificate, belonging to the beamline, thus additional authorisation overheads are minimised. The Grid as Application Service Provider offers to the user the required scientific application in addition to computational and storage resources. The advances of the middleware in the form of an Instrument Control technology can greatly improve the on-line elaboration of the scientific data. Finally the VCR, a web portal for secure and user-friendly access to Grid resources, can offer a better user experience. Such a set of technologies can enable more users to access the services of an SRF.

Keywords:

Synchrotron,Remote Instrumentation,On-line processing,e-Infrastructure,Scientific Applications,Grid

URL for further information:

http://www.dorii.eu/ http://www.elettra.trieste.it/

DCI programming / **33**

PROOF on Demand

Author: Anar Manafov[1]

Co-author: Peter Malzacher [1]

[1] *GSI Helmholtzzentrum für Schwerionenforschung GmbH, Germany*

Corresponding Author: a.manafov@gsi.de

PROOF on Demand (PoD) is a set of utilities, which allows starting a PROOF cluster at a user request, on any resource management system. It provides a plug-in based system, to use different job submission frontends, such as LSF or gLite WMS. PoD is fully automated, and no special knowledge is required to start to use it.

Impact:

Since November 2009, ALICE users at GSI have switched completely to PoD, instead of a static PROOF cluster, which we had before. With the help of PoD, each user at any time can request a desired number of PROOF workers. And his/her private PROOF cluster is set up with just a few seconds of startup time. The ALICE group uses LSF plug-in.
The ATLAS team on IN2P3-CPPM has successfully tested PoD as well, where the gLite plug-in was utilized. And a DPM node with xrootd door was used as a storage element.

Keywords:

PROOF, gLite WMS, PoD, Distributed Analysis, Interactive analysis

URL for further information:

http://pod.gsi.de

Detailed analysis:

Main components of PoD are pod-agent and pod-console. pod-agent provides the communication layer between a PROOF master on a local machine and PROOF workers on remote resources, possibly behind a firewall. pod-console provides a user-friendly GUI, which is used to setup, manage, and shutdown the dynamic PROOF cluster. Installation is simple and doesn't require administrator privileges, and all the processes run in the user space. PoD gives users, who don't have a centrally-administrated static PROOF cluster at their institution, the possibility to enjoy the full power of interactive analysis with PROOF. It is also a very good alternative to static PROOF clusters.
PoD is a specially designed solution to provide a PROOF cluster on the fly.

Conclusions and Future Work:

Currently a dynamic proxy machinery is in development. It will give athe possibility for PoD to automatically make a decision on whether to make a direct connection between PROOF workers and servers or to proxy it. In the actual version of PoD, it always proxies the connection.
Also PoD is going to be extended to cover more resource management systems. We therefore plan to develop AliEn Grid plug-in, SGE plug-in, Condor plug-in and SSH plug-in which will help to setup a dynamic PROOF cluster on any machine where SSH connection is available.

Computational Chemistry / 34

A priori modeling of chemical reactions on a grid-based virtual laboratory: towards standard representations of data for molecular chemistry

Authors: Antonio Laganà[1]; Sergio Rampino[1]

Co-authors: Antonio Monari [2]; Elda Rossi [3]; Kenneth Ruud [4]; Stefano Evangelisti [2]

[1] *Università degli Studi di Perugia*

[2] *Université Paul Sabatier Toulouse III*

[3] *CINECA*

[4] *University of Tromsø*

Corresponding Author: ser_ram@dyn.unipg.it

We have assembled on the COMPCHEM segment of the EGEE Grid Infrastructure the core of the Grid Empowered Molecular Simulator (GEMS) queueing in a common workflow a suite of - adequately "gridified" - codes for the a priori modeling of elementary chemical processes. The communication between applications from different scientific domains is fostered by the use of common data formats. A test calculation is shown for the benchmark H + H2 reaction.

Detailed analysis:

GEMS is a grid-based molecular simulator that can act as a molecular science computational engine in complex multiscale chemical contexts. The general procedure for the modeling of chemical reactions is articulated as follows:

A) evaluate the potential energy between the atoms involved at a set of geometries concurrently;
B) get an analytical representation of the Potential Energy Surface (PES) by fitting the obtained values;
C) run the dynamics on the resulting PES at several initial conditions concurrently and extract the reactive properties.

The intermediate B) step is a single, fast task. On the contrary, the A) and C) steps are extremely Grid-empowerable. In fact, in line with the grid philosophy, larger sets of geometries (or of initials conditions) can be dealt with at the time of a single geometry (initial condition) calculation just adding more computer power to the Grid.

We confine our attention here to the simple gas-phase prototype exchange reaction A + BC (v,j) -> AB (v',j') + c. This is the simplest model of a chemical reaction were a bond is broken and a new one is formed and can be treated at the highest level of theory.

Conclusions and Future Work:

Due to the use of common data formats, we have been able to assemble the core of GEMS for a grid execution aimed at accurately modeling chemical reactions. The complete workflow has been executed for the benchmark H + H2. Grid performances are highlighted and results compared to previous calculations.

We are currently working at enlarging the sets of codes integrated, further improve their interoperability, and extend the treatment to the more complex cases to which the three-step procedure sketched in a previous section applies. For instance, in the near future, to the four-atom problem.

Impact:

Progress in the capability to simulate chemical processes at the molecular level is an important component of the advance in modeling natural phenomena, designing new materials and products, mastering new technologies and carrying out innovative experiments. This progress requires assembling various pieces of software, convergence of the competences of different experts, concurrence of the elaboration on several processors. The difficulty of gathering in the same place all the necessary hardware, software and human resources, makes computational grid platforms the ideal environment for the exploitation of collaborative computing and interoperability.

The past few years have seen significant progress in the grid-porting of scientific applications by the COMPCHEM Virtual Organization on the EGEE grid. Moreover, the problem of interoperability among a large class of codes pertaining to the Quantum Chemistry domain singled out the need for the definition of common data formats, such as the Q5Cost standard.

This work treasures the efforts made in both directions and represents a solid ground for future advances towards a universal molecular simulator.

Keywords:

elementary reactions, virtual laboratory, grid-workflow, interoperability

URL for further information:

http://www.hpc.unipg.it/srampino/urgems.html

Demo Session 2 / 35

On demand Grid services for training in Earth Observation

Author: Dana Petcu[1]

Co-authors: Daniela Zaharie [1]; Dorian Gorgan [2]; Marc Frincu [1]; Marian Neagul [1]; Silviu Panica [1]; Teodor Stefanut [3]; Victor Bacu [3]

[1] *West University of Timisoara*

[2] *Technical University of Cluj Napoca*

[3] *Technical University of Cluj-Napoca*

Corresponding Author: petcu@info.uvt.ro

Current applications involving satellite data need huge computational power and storage capacities. Grid computing technologies that have evolved in the last decade promise to make feasible the creation of an environment, for these kinds of applications, which can handle hundreds of distributed databases, heterogeneous computing resources, and simultaneous users. In this context the recent Grid-based platform GiSHEO offers, for training purposes, satellite image processing services, workflow-based service composition, and user interaction combined with e-learning facilities.

Detailed analysis:

GiSHEO (On Demand Grid Services for Training and High Education in Earth Observation) addresess the issue of specialized services for training in Earth observation. Special solutions were proposed for data management, image processing service deployment, workflow-based service composition, and user interaction. A particular attention is given to the basic services for image processing that are reusing free image processing tools, like GDAL. The platform has distributed data repositories. A special feature is the connection with the ESA's GENESI-DR catalog. The physical platform is based on four clusters that are geographically distributed at four academic institutions. GiSHEO uses the Grid platform for

near-real time applications for short-time data-intensive tasks. The data sets that are used for each application are of several tens of GBs, and the tasks are specific for image processing. In this particular case a scheme of instantiating a service where the data are located is required in order to obtain a response in near-real time.
A particular component is eGLE, the eLearning environment. It uses templates to allow teachers specialized in Earth observation to develop new lesson

Conclusions and Future Work:

The GiSHEO platform includes at this moment several basic services of the new training facility for Earth Observation. Advanced lessons and tutorials, that involve more than just remote sensing data, and related to disaster management and archeology, are planned to be developed in the next future.

Impact:

Grid-based platforms were recently build all over the world to satisfy the huge needs of computational power and storage capacities for Earth observation activities. Training activities in the field are not following these developing activities, resulting a big gap between the field request for specialists and the labor market offer. Currently there are only a few number of resources involved in educational activities in Earth Observation, one of the most complex one being EduSpace. Contrary to the current existing platforms that provides tutorials and training materials, GiSHEO intends to be a living platform where experimentation and extensibility are the key words. The feature of service composition support the creativity of the trainees.

Keywords:

Grid services, Earth observation, workflows

URL for further information:

http://gisheo.info.uvt.ro

Justification for delivering demo and/or technical requirements (for demos):

The platform will be available for public use starting from 2010. Its impact on the community depends on the dissemination and the feedback. EGEE UF is an adequate forum to provide them.

Poster session / 36

Migration of the MAGIC data transfers to a Grid infrastructure

Author: Roger Firpo Curcoll[1]

[1] *Port d'Informació Científica*

Corresponding Author: rfirpo@pic.es

The MAGIC collaboration is moving from a computing model based on local computer farms to a Grid based model, including the migration of all the services of the official data center. Here we present the recent progress in the adoption of the Grid infrastructure in the MAGIC data center, which relates to the data transfer from the observation site.

Detailed analysis:

The MAGIC (Major Atmospheric Gamma-ray Imaging Cherenkov) telescopes, run by an international collaboration of institutes from 9 different countries, is a two 17-meter Cherenkov telescope system located on La Palma (Canary Islands) dedicated to the study of the universe in very high energy gamma-rays. The addition of the second telescope caused the data output of the system to rise up to 300 TB per year of data. The MAGIC data center hosted by the Port d'Informaci— Cient'fica (PIC) in Barcelona is migrating its services to Grid as part of an upgrade needed to deal with the increased data volume. After migrating the data to a Grid filesystem, we have ported the data transfer service. In the last months we have set up an SRM endpoint in the observation site and ported all the data transfer tools to use Grid

file transfers.

Conclusions and Future Work:

The MAGIC data transfer from the observatory to the data center has been migrated to Grid. This is a further step of the data center into Grid after the migration of the storage, computing and data access services. The next step in this process will be to develop high level analysis tools for users to better use the datacenter services.

Impact:

The computing system in the MAGIC observation site was not initially designed to interact with Grid, and data transfer and publishing to Grid were made through ssh connections and an intermediate system in the data center. In the last months we have installed a new server to publish the data in the observation site through a SRM endpoint, and also ported all our data transfer tools to use this infrastructure. This allowed us to save resources and reduce the complexity of the system. Our initial tests with the new infrastructure also point to a better use of the available bandwidth, which can lead to deprecating the tape-based transfer of raw data.

Keywords:

MAGIC, Gamma-Ray Astronomy, Cherenkov Telescope, data center, data transfer

URL for further information:

http://magic.pic.es

Poster session / 37

Running Bag-of-Tasks Applications in EGEE with MaWo

Author: Oleg Sukhoroslov[1]

[1] *Centre for Grid Technologies and Distributed Computing, ISA RAS*

Corresponding Author: oleg.sukhoroslov@gmail.com

MaWo is a distributed computing framework which implements a well-known master-worker pattern. It provides a programming interface as well as easy-to-use tools for running Bag-of-Tasks applications across heterogeneous computing resources. MaWo allows the user to seamlessly utilize all available resources including local workstations, clusters and grids among which EGEE is a primary target. The framework significantly reduces the time and effort needed to port an application to grid.

Detailed analysis:

The grid is an ideal platform for Bag-of-Tasks (BoT) applications composed of many independent tasks. The efficiency and run time of such applications in a grid strongly depend on a strategy used for scheduling of tasks. A well-known "master-worker" strategy proved to be efficient for heterogeneous and unreliable distributed resources. It can also improve the run time of an application in a grid by bypassing the central grid scheduler. In order to streamline the porting of BoT applications to grids a generic application framework MaWo was developed. MaWo implements basic parts ofa master-worker pattern such as workers allocation, communication with master, task scheduling, failure recovery, etc. Since many users have access to several computing resources and infrastructures MaWo supports the simultaneous use of various types of resources by means of pluggable adaptors. The current implementation supports the EGEE infrastructure as a primary target and was successfully used to run several applications in EGEE.

Conclusions and Future Work:

The MaWo framework provides an easy-to-use solution for the development and running of BoT applications across heterogeneous computing resources. The applicability of the framework was demonstrated by successfully running several real applications both on local and EGEE resources. The planned future

work includes the optimization of the current implementation, measurement of the total efficiency of computations, support for pluggable schedulers and the implementation of adapters for other grid infrastructures.

Impact:

MaWo is a generic framework which can be used to streamline the development of BoT applications. The framework provides a programming interface for the implementation of problem-specific parts of application. MaWo also supports the declarative description of BoT applications which enables the quick porting of applications without using MaWo API. The current implementation supports the simultaneous allocation of workers onthe local machine, cluster and EGEE infrastructure. MaWo also provides a built-in Web interface which can be used to monitor the status of running applications. In contrast to existing master-worker frameworks for EGEE, such as DIANE, MaWo requires less effort to develop or port applications because it doesn't require the installation of gLite User Interface and supports simple declarative description of an application in addition to a programming interface. MaWo was used to implement and run several real BoT applications in EGEE including ray tracing and atomic cluster conformation problem.

Keywords:

Bag-of-Tasks Applications, Master-Worker Framework, Application Development

URL for further information:

http://dcs.isa.ru/os/mawo/

Fusion / **39**

Fusion metadata navigation and data delivery services

Authors: Nikolay Marusov[1]; Sergey Portone[1]; Vladimir Voznesensky[1]

Co-author: Igor Semenov [1]

[1] *RRC "Kurchatov institute"*

Corresponding Author: nikolay.marusov@grid.kiae.ru

The services for user-friendly metadata navigation and data delivery have been developed. The metadata service enables exploring of fusion-specific data storage via a web browser. The data delivery service allows to download a required piece of data from Storage Elements using a link obtained from the metadata service.

Detailed analysis:

In the course of the conceptual design of the ITER data acquisition and control system (CODAC) existing solutions for fusion data acquisition, storage and delivering were considered. However, it turned out that the set of existing data management systems used by fusion community does not satisfy the requirements. Solutions designed forthe Large Hadron Collider are not able to cover CODAC requirements entirely. In this work we propose a service for navigation over semantically linked metadata which describe the structure of ITER data sets. The service provides an intuitively clear GUI assisting the user to understand interrelations between entities in CODAC system knowledge domain. The GUI application uses Stellaris information service to store metadata. Stellaris allows to store ontologies and knowledge bases described in RDF and OWL languages. Semantic nets describing CODAC system knowledge domain provide URL's of related web-pages (similar to Wikipedia ones) and URL's to make requests to data delivery service. The data delivery service selects data subsets in the form of Universal Access Layer's Consistent Physical Objects and other formats from Storage Elements emulating future ITER storages.

Conclusions and Future Work:

The proposed services are intended to demonstrate data management on the basis of the semantic description of the CODAC system domain. They could serve as the complete components for fusion re-

search social networks for joint exploitation of data from ITER and other fusion facilities. The navigation service could also be used as part of the CODAC human-machine interface to assist ITER personnel

Impact:

The services for GUI metadata navigation, data selection from Storage Elements and data delivery to an end user have been developed. The proposed solution should simplify the deployment of the ITER IT-infrastructure, and the involvement of existing facilities.

Keywords:

fusion, data management

URL for further information:

http://vo.nfi.kiae.ru/fuworld

Poster session / 40

The APEL CPU Usage Accounting Infrastructure for NGI Grids

Authors: Cristina Del Cano Novales[1]; Gilles Mathieu[1]; John Casson[1]; John Gordon[1]; Ming Jiang[1]; William Rogers[1]

[1] *e-Science Centre, Science and Technology Facilities Council, United Kingdom*

In order to implement and deploy a scalable and flexible distributed CPU usage accounting infrastructure for the NGI Grids, the accounting records transport mechanism of the APEL (Accounting Processor for Event Logs) tool is modified and extended to integrate with ActiveMQ message broker network. The new APEL CPU usage accounting infrastructure supports a robust accounting capability at an NGI level and flexible across VOs accounting records queries.

Detailed analysis:

A general topic publication and subscription messaging model enables distributed components in a system to publish and subscribe messages to/from a well defined topic that can be viewed as a virtual destination and source of messages. The definition of topics and low level reliable delivery of messages among components can be achieved by using a concrete message broker implementation to this model. The work reported here investigates the feasibility of adopting such a messaging model to implement a distributed accounting infrastructure and utilises the Apache ActiveMQ message broker network as the accounting records transport layer of APEL for robust delivery of accounting record messages. In the distributed infrastructure, while the ActiveMQ message brokers will manage the delivery of accounting records messages, at a NGI level and between NGI accounting instances and the central records cache, the original user interfaces for existing APEL clients will remain consistent.

Conclusions and Future Work:

The implementation and testing with production CPU usage records publication of the new APEL based accounting infrastructure demonstrate that the transport mechanism of a distributed accounting infrastructure can be implemented based on topic publication and subscription messaging model in an ActiveMQ message broker network. Further investigations on the scalability and fault tolerance features of ActiveMQ broker will be conducted.

Impact:

Within the distributed accounting infrastructure, accounting records are transported from APEL Publishers at Grid sites to either a regionalised accounting system or the central one by choice via a common ActiveMQ message broker network. This provides an open transport layer for other accounting systems to publish relevant accounting data to a central accounting repository via a unified interface provided an APEL Publisher and also will give regional/NGI Grids the flexibility in their choice of accounting system. The robust and secure delivery of accounting record messages at an NGI level and between NGI accounting instances and the central one are achieved by using configurable APEL Publishers and an

ActiveMQ message broker network.

Keywords:

APEL, CPU Usage Accounting, Message Broker, NGI

URL for further information:

http://goc.grid.sinica.edu.tw/gocwiki/ApelHome

Poster session / 41

Integration of Nagios plug-in into a data model to improve the Grid stability from the user point of view

Author: Silvio Pardi[1]

Co-authors: Andrea Apicella [2]; Francesco Palmieri [2]; Guido Russo [2]; Leonardo Merola [2]; Simone Celestino [2]

[1] *INFN - Naples Unit*

[2] *University of Naples Federico II*

Corresponding Author: spardi@na.infn.it

In this work, we propose a new approach to publish and consume monitoring information about Grid sites within a gLite based infrastructure. Starting from a set of tests that are crucial for a Grid site or for a Virtual Organization, we created a data model to represent them in the standard gLite information system. Through the Nagios tool, we have periodically performed sanity checks publishing the results concerning resource monitoring or resource discovery. Preliminary tests have shown the effective benefits provided by this approach in term of successful jobs.

Detailed analysis:

In the EGEE framework, we have many monitoring tools that provide a large amount of site administrator oriented information: ie. GridIce, GSTAT, Gridmap, Service Availability Monitoring SAM and the Site-Nagios.
Our approach aims at using monitoring information at a user level, in order to minimize the effects of grid services failure and improve the user's perception. Therefore, we have selected a set of tests crucial for the success of Grid jobs and data-management activities. Then we have created a data model able to represent this information by defining attribute-value pair. In particular, each result has been represented by a string containing the test status, followed by the death line of check validity. We implemented a set of plug-in for Nagios, able to publish the results of the test on the site-BDII. The plug-in have been designed to test the crucial aspects within the pilot-virtual organization matisse, and they are periodically scheduled by the nagios. In this scenario, each site publishes in real-time the status of grid services to the top level Information system. Then the users can specify the monitoring metrics as job requirements at submission time.

Conclusions and Future Work:

Preliminary tests show interesting improvements from the user's point of view. In future works we plan to deploy this approach in a large scale Grid and obtain feedbacks in order to detect the best metrics to publish, to improve the data model and to maximize the positive impact on the Grid stability. Finally, we plan to investigate the introduction of site reputation concept as a new metric to use during the resource discovery process.

Impact:

The introduction of monitoring metrics on the information system opens up new interesting scenarios to improve the grid stability and facilitate operations in a production grid infrastructure.
We obtain interesting stability improvements from the user's point of view. This effect emerged through a set of preliminary tests that showed the positive impact of the proposed approach in term of number

of jobs successfully completed. By adding the requirement tag in the JDL at submission time, the users are able to avoid in a transparent way, the instable resources still present in the informative system. Finally regarding the operation, the expected impact consists in supporting administrators that can take advantage from a more stable system during the recovery of incoming problems.

Keywords:

Monitoring, Information System, User interface

URL for further information:

http://people.na.infn.it/spardi

Poster session / **42**

Data awareness in gqsub

Authors: Anthony Doyle[1]; Stuart Purdie[1]

[1] *University of Glasgow*

Corresponding Author: stuart.purdie@cern.ch

gqsub is a user (command line) interface for submitting and monitoring Grid jobs that conforms to the IEEE standard for qsub (and friends).

Recent development has focused around data and data awareness; both in terms of data local to the submission machine, and data elsewhere in a Storage Element.

Detailed analysis:

The gLite and Arc computing infrastructures are both distinguished by their data awareness - that is, the location of data matters to job scheduling.

Firstly, specifing the request. For staging in of data, this is not too tricky; by allowing the files to be specified as a url, it's straightforward to handle this. The staging out of data, to push it to a Storage Element at the end of the job, is more complex, as it requires two distinct bits of information. At the time of writing, there are a few alternatives in development, pending user trials before final selection.

The second part is to handle the requests. In the cases where the computing infrastructure supports the desired behaviour, it's a matter of compiling the request into the appropriate form. Where it's not, then the requisite code needs to be supplied, typically via the use of a wrapper script.

One interesting case of this is local use. A design goal of gqsub was to have job scripts that could be submitted either locally or on the Grid without modification. Data access changes this, as most local systems are not aware of storage elements. At the time of writing, no clear solution is available.

Conclusions and Future Work:

This work takes gqsub from being able to treat the Grid as a cluster extension to being able to leverage the full power of the the Grid model. As noted, there are a couple of outstanding points to be resolved, which is fully expected by February. Once complete, this will enable the Grid to come to any cluster.

As a further work, the opposite is planned - taking a cluster to the Grid. By capturing a sub directory tree, from the users machine, and staging that remotely, te setup of complex environments can be simplified.

Impact:

The data awareness increases the scope of suitable application scales, from using the Grid as a cluster replacement to being able to manage tera scale applications. It has lost none of the simple user interface in the progress, meaning it is able to scale with the users needs, from the simplest use case, out to multi-institutional data-analysis projects.

gqsub has always stood out from other attempts to provide a 'simple' way to use the Grid, because it offers an interface that is *already* known to many e-Scientists, rather than expecting them to learn a new one. For users that have previously used a cluster of some sort, the time spent on learning gqsub is trivial - less, and often significantly less, time than it takes to sort out the X.509 certificates.

It is designed to facilitate having a cluster head node also serve as a submission system for Grid jobs, so that it is natural and straightforward for a user to switch between the two.

Keywords:

gLite, data management, HCI

URL for further information:

http://www.scotgrid.ac.uk/gqsub/

Justification for delivering demo and/or technical requirements (for demos):

(see comments)

A demo would need just an internet connection, and be run from any connected system (i.e. my laptop). An external monitor would be handy, to increase visability.

Medical Imaging / **43**

Gwendia workflows for cardiac image analysis

Authors: Benjamin Isnard[1]; Christopher Casta[2]; Eddy Caron[3]; Johan Montagnat[4]; Ketan Maheshwari[4]; Patrick Clarysse[5]; Tristan Glatard[5]

[1] *INRIA - LIP*

[2] *University of Lyon - CREATIS*

[3] *University of Lyon - LIP*

[4] *University of Nice Sophia-Antipolis - I3S*

[5] *CNRS - CREATIS*

Corresponding Author: glatard@creatis.insa-lyon.fr

We will present the description of a cardiac image analysis workflow with the new Gwendia language and its enactment on two different grid platforms, namely EGEE (using the MOTEUR workflow engine) and Grid'5000 (using the DIET workflow engine). Technical achievements enabling the execution of the very same workflow on those two grids will be presented. Finally, early application results (both performance and medical imaging) will be reported.

Detailed analysis:

This presentation will detail achievements of the Gwendia French national project to describe and execute medical imaging grid workflows on various platforms. First, specific features of the Gwendia workflow language will be highlighted on a cardiac segmentation pipeline. In particular, the use of list manipulations and iteration strategies to handle multidimensional medical images will be exemplified. Second, experiments using two workflow engines will be reported. DIET MA-DAG is used to run the cardiac workflow on the Grid'5000 while MOTEUR is used for the deployment on EGEE. Thanks to significant technical effort, the very same workflow can be run on those two grids. Based on that, performance comparisons between those two grids will be presented.

The resulting implementation offers a unique tool for large-scale processing of cardiac images. Results will be shown highlighting how parameter estimation and myocardium segmentation can be improved using the grid. In particular, sweeps on physical and imaging parameters will be detailed.

Conclusions and Future Work:

The implementation of a cardiac image analysis pipeline with the Gwendia workflow language has been discussed. Based on that, performance comparisons between Grid'5000 controlled conditions and EGEE production ones were presented. DIET was used to run the workflow on Grid'5000 while MOTEUR was the EGEE enactor. Medical imaging results illustrating how a particular use-case can benefit from the grid computing power have been shown.

Impact:

Although workflow technology has been acknowledged many times to be a very powerful tool to support grid porting of a variety of applications, it has to be recognized that writing grid workflows is still an art only mastered by a few. Because it is dedicated to the description of applications handling a lot of data, the Gwendia language is expected to reduce such a usability gap faced by domain scientists. Moreover, since DIET and MOTEUR can both interpret this language, the very same workflows can be executed on the Grid'5000 and EGEE grids, which somehow bridges research and production grid communities.

Regarding medical image analysis, the Gwendia project offers a concrete tooling for the large-scale processing of cardiac images. In particular, variability among cardiac sequences (due to pathologies, acquisition conditions, etc.) make robust myocardium segmentation very challenging. Workflows such as the one built in the scope of Gwendia allows to rapidly identify suitable parameter sets, benefiting from the computing power offered by the grid.

Keywords:

Workflows, cardiac imaging, Grid'5000, EGEE, DIET, MOTEUR

URL for further information:

http://gwendia.polytech.unice.fr

Earth Science / **44**

Experiences on porting and running a Climate - Air Quality Modeling System on the Grid

Authors: Christos Kanellopoulos[1]; Dimitrios Melas[2]; Eleni Katragkou[2]; Ioannis Tegoulias[3]; Paschalis Korosoglou[1]; Prodromos Zanis[3]

[1] *Grid Operations and HPC Centre, Aristotle University of Thessaloniki*

[2] *Laboratory of Atmospheric Physics, School of Physics, Aristotle University of Thessaloniki*

[3] *Department of Meteorology-Climatology, School of Geology, Aristotle University of Thessaloniki*

We present results of regional climate-air quality simulations performed over Europe for the future decade 2091-2100 and the control decade 1991-2000 and briefly discuss the impact of climate change on air quality metrics over Europe throughout the 21st century. In order to meet the increased demands on computational resources the high resolution simulations were performed on the EGEE Grid. Special tools to allow for easy coupling of RegCM3 and CAMx on the underlying resources have been developed and we briefly discuss the benefits and drawbacks of using the Grid for such simulation campaigns.

Detailed analysis:

The RegCM3 regional climate model has been applied for the climate simulations whereas the air quality simulations have been performed using the CAMx model. CAMx has been coupled to RegCM with a FORTRAN-based code interface, which reads the basic meteorological parameters from RegCM and converts them into the format that is expected by CAMx. The spatial resolution was set to 50km x 50km. The vertical profile of CAMx was split into 12 layers of varying thickness extending to almost 6.5km with the lower layer being 36m deep and the uppermost layer 1.2km thick. To execute our simulations on the Grid, we have implemented a DAG job. With respect to the two application models themselves, we have used their parallel versions in order to leverage the distributed resources of the Grid infrastructure.

RegCM3 was compiled and linked against the MPI library, while CAMx was compiled with OpenMP. In order to facilitate the execution of the OpenMP enabled CAMx model on the Grid, we have developed a modified version of the GRAM module, which has been installed on the CE at the GR-01-AUTH cluster and which allows the scheduler to properly allocate resources for shared memory parallel applications.

Conclusions and Future Work:

Our model simulations yielded average decadal changes of tropospheric ozone up to 1.5 ppb for the mid decade of the century and higher ozone increase (1-5 ppb) in the 2090s. The response of ozone to climate change varied in magnitude depending on season the and the geographical location. Future work includes application of newly developed online coupled climate-chemistry models, which are more suitable for the study of the complicated particle interactions. One of the technical challenges is the production of large datasets, which can be efficiently managed with tools, such as AMGA and GrelC.

Impact:

The result of our work is a database of air quality data corresponding to two different time slices: a present decade 1991-2000 and a future decade (2091-2100). The air quality information concerns concentrations of gaseous species and particulate matter (O3, NOx, CO, SO2, PM10) for the present time and provides an estimation of how anthropogenic climate change may affect these concentrations at the end of the century.
This information is of great importance to the atmospheric science community and, more specifically, to the researchers involved in climate change studies and the impact on air quality metrics on regional scale.
Most of the climate change studies available in the literature focus on a global scale, where the model resolution is too coarse to resolve atmospheric phenomena over a complex terrain such as the European. In this work, we introduce regional models in the study of climate change, and work with a resolution of 50 km, which can be easily further downscaled to even finer spatial resolutions, enabling us to simulate climate-air quality phenomena on a national level.

Keywords:

Atmospheric Science, MPI, OpenMP

URL for further information:

http://www.cecilia-eu.org/meetings.htm

Justification for delivering demo and/or technical requirements (for demos):

N/A

Medical Imaging / 45

Experience feedback on some scientific gateway components for medical imaging applications

Authors: Hugues Benoit-Cattin[1]; Sorina Camarasu-Pop[2]; Tristan Glatard[2]

[1] *University of Lyon - CREATIS*

[2] *CNRS - CREATIS*

Corresponding Author: glatard@creatis.insa-lyon.fr

This presentation summarizes experience acquired at Creatis with some Scientific Gateway components to deploy medical imaging applications on the EGEE grid during the last years. We will detail feedback about various tools that we tested, including pilot-job systems, workflow managers, user front-ends, monitoring components and data management systems. Based on application use-cases ported at the lab, we will discuss advantages and drawbacks of these gateway components and highlight remaining challenges. Finally, an integration of some components in a coherent grid environment will be shown.

Conclusions and Future Work:

This presentation summarizes user feedback regarding medical image simulation obtained at Creatis on several scientific gateway components during EGEE-III. The resulting system covers the complete application-middleware stack and provides useful application monitoring tools for production control. However, effort is still needed on certain aspects like MPI and data management issues. We are confident that this work contributes to the improvement of gridified applications performance as well as to the interoperability and re-usability of existing tools in the context of the EGEE-EGI transition.

Detailed analysis:

The discussion will be based mainly on three use-cases with different characteristics: the GATE Monte-Carlo simulation code, the Simri MPI image simulator and the FIELD-II Matlab ultrasonic simulator. We will show how workflow systems (in particular Taverna and MOTEUR) can facilitate the porting of those applications on infrastructures operated by gLite and ARC. Some remaining interoperability issues (in particular regarding service developments) will be highlighted as well as performance road blocks. Besides, we will illustrate to what extent pilot-job systems (in particular the DIANE tool) improve performance and reliability in some cases. Missing features regarding, e.g., MPI codes, control of pilot submission or the handling of GLUE requirements will be summarized. Various user front-ends (application-specific portals, generic workflow interface) will be shown and their features will be discussed as well as their current limitations regarding for instance the handling of grid certificates. Finally, data management issues related to SE(s) selection and secure handling of medical data will be summarized and potential solutions will be reviewed.

Impact:

We integrated RESPECT tools in a complete middleware stack offering a powerful and easy to use system for executing applications on the grid. Pilot jobs improve QoS in terms of latency reduction and fault-tolerance. They allow to easily reach 100% success, while the gLite-only approach success rate is often between 70% and 80%. The workflow implementation provides a generic framework for the integration of new applications with different computing models and their execution on other systems. Results show that this generic workflow approach is not penalizing in terms of performance. The user-friendly frond-ends are essential for opening the access to the grid to larger research communities. The monitoring allows on the one hand for production control and on the other hand for gathering logs that feed grid research models. Last but not least, we hope that our work and feedback could contribute to identify some common bricks among existing scientific gateway components and to facilitate their interoperability.

Keywords:

Medical imaging, pilot jobs, workflows, front-ends, data management

URL for further information:

http://www.creatis.insa-lyon.fr

Poster session / 46

Experiences with a lightweight GRID infrastructure using AFS

Author: Roger Barlow[1]

Co-authors: Alessandra Forti [1]; Andrew McNab [1]; Michael Jones [1]

[1] *Manchester University*

Corresponding Author: roger.barlow@cern.ch

The Andrews File System (AFS) with certificate-based authentication can be used to provide a system of input and output sandboxes which is simpler and easier to use than the usual tarball. We describe how such a system was set up at Manchester, the experiences of various different users, and the implications for design of Tier3 facilities.

Detailed analysis:

Development was carried out using the 1800 CPU Tier2 centre at Manchester and the local user community. The worker nodes run an afs client system, and the gssklog program is used to connect running jobs to the local afs server.
Several user projects have been and are being successfully undertaken using this, in the areas of accelerator simulation, reactor studies, and particle physics data analysis. We describe these cases and the lessons learned from them.
Using this system users can run jobs on the Tier 2 centre in exactly the same environment as when running test jobs on the local cluster. This makes the path the online test and debugging environment to the offline production environment extremely smooth.

Conclusions and Future Work:

This is a successful lightweight system which makes access by users considerably easier and encourages use of the centre, and speeds up their computing tasks

Impact:

This system provides a way of opening up the facilities of a large Grid centre to a user with minimal adaptation on
their part, and has encouraged our transition to grid computing. Other centres and other users will be interested in adopting this system.

Keywords:

infrastructure; user adoption; sandbox

URL for further information:

http://www.hep.man.ac.uk/u/roger/gssklog

Demo Session 1, Welcome Drink / **47**

DrugScreener-G: Towards Grid-enabled Large-Scale Virtual Screening Coming into Handy

Author: Soonwook Hwang[1]

Co-authors: Ana Lucia Da Costa [2]; Doman Kim [3]; Dongsoo Han [4]; Jean Salzemann [5]; Jincheol Kim [1]; Sangdo Lee [1]; Sehoon Lee [1]; Vincent Breton [5]

[1] *KISTI*

[2] *Healthgrid*

[3] *CNU*

[4] *KAIST*

[5] *LPC-IN2P3*

Corresponding Author: hwang@kisti.re.kr

DrugScreener-G (DSG) is an intuitive and easy-to-use grid-enabled in silico screening tool, aiming to help scientists in drug discovery including biologists and biochemists to conduct a large-scale deployment of virtual screening process on the Grid. With the help of DSG, scientists can easily have access to the PDB database, download and view the 3D structure of target proteins of interest, launch and manage millions of protein-ligand docking simulations on the Grid. To facilitate the analysis of docking results, visualization tools like Jmol and Chimera have been integrated into DSG as well.

Detailed analysis:

DSG was developed based on client/server architecture. The DSG server takes care of details of the complexity of Grid-related operations including maintaining thousands of agents distributed on the Biomed

VO of EGEE that actually perform the docking simulations. Indeed, it is quite challenging to get millions of in silico dockings done on the public Grid infrastructure like EGEE as it requires the grid-relevant issues of scalability, performance and fault tolerance to be seriously addressed. The DSG server was implemented as a web service that has adopted, as its core grid engine, the WISDOM Production Environment, a widely accepted grid-enabled virtual screening technology. The DSG client was developed as an Eclipse RCP-based application, providing an easy-to-use GUI interface for (1) preparing a set of ligands and target proteins by allowing accessing well-known public biological databases such as ChemBridge, ZINC and PDB database, (2) submitting docking simulation jobs to the Grid via the DSG server, (3) monitoring the progress of the simulation jobs, and (4) the post-processing of the docking results to help uncover their biological meaning.

Conclusions and Future Work:

DSG is an easy-to-use integrated tool for scientists, who usually have a resistance to learning complicated yet powerful IT technologies like Grid computing, to harness the power of its production infrastructure to carry out large-scale in silico experimentation of their own. DSG has implemented some basic ideas of the virtual screening processes. In future versions, it will support (1) consensus docking by integrating multiple in silico docking tools (e.g., Flexx, Autodock, etc) and (2) multistep docking by supporting the execution of Molecular Dynamics (MD) simulations.

Impact:

The WISDOM initiative, a large-scale deployment of virtual docking on production Grid infrastructures, has been seen in the EGEE community with great success from the scientific and grid deployment perspective. However, from the usability point of view, there seems to be a lot to be done in order to be able to get the grid-enabled in silico approach easier to use for non-experts of grid computing. DGS targets researchers and scientists in drug discovery, biology and biochemistry to help them have access to the Grid infrastructure and exploit its full power in drug discovery and virtual screening without having to learn details of grid middleware services and tools. DSG is expected to help them capture more insight into the grid-enabled approach to large-scale virtual screening and thus enhance their research productivity in drug discovery. The DSG client tool is now being used and tested by biologists in Chonnam National University, one of long-time partners of the WISDOM collaboration. With the help of tools like DSG, their long-time dream of conducting large-scale virtual screening on their own at any time without any help from grid experts is to be made realized in the near future.

Keywords:

virtual screeing, in-silico docking, grid-enabled, production grid, drug discovery

URL for further information:

http://anakin.kisti.re.kr/DrugScreenerG

Fusion / **48**

Numerical reconstruction of fusion-origin carbon structures using refactored GIF portal

Authors: Vladimir Voznesensky[1]; Vladislav Neverov[1]

Co-authors: Aleksander Kukushkin [1]; Igor Semenov [1]; Konstantin Dyabilin [1]; Nikolay Marusov [1]; Nina Kuvshinova [1]

[1] *Kurchatov Institute*

Corresponding Author: vovic@nfi.kiae.ru

Two developed technologies are shown. The top one is a method for the massive modelling of X-ray propagation in carbon films found in fusion devices as well as numerical reconstruction of such films composition. The basic technology is the strongly re-factored GIF portal that allows the development of high-level logics and user web interface for such grid applications.

Detailed analysis:

In this paper we present a grid application for the numerical modelling of X-ray beams diffraction on various targets made of nanostructures like nanotubes, graphens, e.t.c., to determine the structural constitution of the films found in plasma confinement devices. Every variant of the target is defined by a set of scalar parameters and is computed in a single grid job. Every variadic parameter is computed for a specific node in a user-defined mesh and by a user-defined formula that depends on that node. A re-factored GIF web portal is used for running this application. The portal's core operations are re-written in C++. Their data is stored in a fast embedded database. A special Lua interpreter was developed (instead of the previously used Python) to fulfill security, performance and language simplicity requirements of the UI declaring and jobs driving scripts. The interpreter stores all it's objects as DB records, and so there is zero overhead on referencing large objects like file bodies. It also allows reading and writing jobs sandbox files as language objects concurrently with interpreting.

Conclusions and Future Work:

The application to be presented is a building block for a scientific conveyor for structural analysis of films found in fusion devices as well as other targets. It could be extended to a comprehensive tool for numerical reconstruction of films constitution by different optimisation methods. A re-factored GIF portal is used for running application scripts. It provides a minimalistic development and execution environment for abstract user interface and grid jobs manipulation scripts.

Impact:

Carbon films found in plasma confinement devices are a subject of intensive research. These films consist of different carbon structures. The determination of the presence and quantity of such structures may help to reveal these structures formation mechanisms. The film could be X-rayed in a special source of tightly focused radiation. Such X-raying would give a photograph of photon rays diffracted by the film's atoms. In the simplest model, the diffraction processes on different structures are independent, so this photograph could be linearly decomposed to a set of pictures, each of them caused by the appropriate structure, by methods of linear programming. Each picture could be obtained by numerical simulation of photons propagation and diffraction on the appropriate structure, which is described in the detailed analysis above.

Keywords:

fusion, portal, workflow, user interface

URL for further information:

http://vo.nfi.kiae.ru

Poster session / 49

Supporting 60 e-Science applications from Europe and Latin-America: Best practices and analysis extracted from the EELA-2 project

Author: Leandro Ciuffo[1]

Co-authors: Riccardo Bruno [1]; Roberto Barbera [2]

[1] *Italian National Institute of Nuclear Physics (INFN)*

[2] *Italian National Institute of Nuclear Physics (INFN) and University of Catania*

Corresponding Author: leandro.ciuffo@ct.infn.it

Investments to promote e-infrastructures within new scientific communities in several regions of the world have been attracting new research groups interested in porting their applications on the Grid. Such an expansion across many institutions/countries facing different maturity levels of IT infrastructures, network connections and e-science awareness represents new challenges to support a diversity of users/applications. This work aims at presenting a complete picture of all EELA-2 applications as well as to share some best practices and methodologies used to support them.

Detailed analysis:

Based on our experience supporting 60 applications in the framework of the EELA-2 project, we have noticed that Grid users may be broadly divided into three groups: (i) those participating in collaborative experiments which requires High-Throughput Computing across many computing and storage clusters; (ii) those that have true computational and storage demands that cannot be handled by their local resources in a reasonable time; and (iii) those with modest computational needs that could be easily handled by a local cluster or storage server.
The advantages of Grid Computing is clear for the first two groups of users, but what about the third group? Is Grid computing not suitable for them?
We surveyed all EELA-2 application groups in order to identify how they are using the Grid. We focused our analysis on four common aspects of applications that are ported on the Grid: (a) frequency of runs, (b) number of jobs submitted on each run, (c) average input file size and (d) job execution time. We have also investigated the use of MPI. The statistics of all these aspects – separated between Latin-American and European applications - will be presented in the paper.

Conclusions and Future Work:

In this paper we present the profile of 60 applications ported on the EELA-2 infrastructure and describe a successful methodology to support e-Science applications, starting from the experience of EELA and EELA-2 projects.
From the Latin-American side we have noticed that the affiliation of small laboratories with large Grid projects may opens exciting new opportunities to collaborate with other institutions belonging to the same Virtual Organization, and thus realise significant gains at institutional, national and international levels.

Impact:

On one hand, EELA-2 supports applications that runs thousands of jobs per week, lasting for many hours and handling Gigabytes of data, but on the other hand there is also "bag-of-task" applications that use to run 1 single job on an occasional basis and consumes a very few computing resources. Even so, the use of a Grid infrastructure to run this kind of application can bring some added value to them. The analysis of the results can be a valuable study for several EGEE-like Grid projects collaborating with institutions from developing countries. Our report will discuss what are the potential advantages of using Grids for each kind of users/applications.
In addition, the best practices and tools used in EELA-2 can be shared with others application's support centers.
It's worth mentioning that some EELA-2 supported applications are also running on the EGEE infrastructure. The implications of having applications playing in both sides are also discussed.

Keywords:

Application support EELA methodology

URL for further information:

http://applications.eu-eela.eu

Bioinformatics / 50

Distributed System Based Strategies for Search, Design and Evaluation of COMBO-FISH probe sets

Authors: Eberhard Schmitt[1]; Michael Hausmann[1]

[1] *University of Heidelberg, 69120 Heidelberg, Im Neuenheimer Feld 227*

Corresponding Author: eschmitt@kip.uni-heidelberg.de

A prominent means to detect genetic aberrations is the method of fluorescence in situ hybridization FISH. To avoid labelling large genomic regions by one polynucleotide like in standard FISH, for COMBO-FISH we search for a set of about 30 colocalizing short sequences with the requirement that no more than 4 of these stretches colocalize within 250 kb anywhere else in the genome. The

exact search is parallelized and applied to certain subsets motivated by kinetics and stability considerations.

URL for further information:

http://services.medigrid.de/5thegee

Impact:

As opposed to algorithms like BLAST, we have to perform an exact search. In the pU and pY cases, the whole genome of the respective species is scanned for pU and pY sequences of minimally 15 bases by a primitive comparison algorithm, which takes one day on a desktop computer. The detected pU and complementary pY sequences are stored in an ASCII data base for each chromosome with a total of 120 Mb ASCII strings. For the design of a specifically labelling COMBO-set, the pU sequences of the gene are extracted. In a second parallelized step, the sequences are located in the whole genome, transferring the partial data bases to the distributed system. The newly developed cluster detection algorithm and deletion process regarding the side conditions runs on one processor using the collected location information to reduce the set. Improving a refined search algorithm, the time for the automatic construction of one COMBO-set was reduced to 10 minutes on a 8 processor system.

Keywords:

Services@MediGRID, personalized medicine, COMBO-FISH, gene labelling, genome cluster search

Detailed analysis:

The design of an oligonucleotide set colocalizing within a genomic region of typically 80 to 250 kb and respecting the requirement that no further clusters of more than 4 oligos exist calls for a parallel search on a distributed system. Taking binding and stability considerations into account, three biochemically different oligo systems can be distinguished: a) polypurines pU, b) polypyrimidines pY, c) mixed sequences. Cases a) and b) are dual in the sense that they both also allow triple helical conglomerates and the complementary sequences have to be located, too. About 2 to 4 percent of the genome (species specific) consist of pU and pY sequences of more than 15 bases, which we all collect into a special data base. The genetic region of interest is investigated for pU and pY sequences and these oligos are located within the whole genome in a parallelized distributed search. In an automatized deletion process, the set is reduced until no further clusters of more than 4 oligos exist.

Conclusions and Future Work:

The design of specifically labelling COMBO-sets has been extended to 150 locations with 3 sublocations of genes at one time. The combined search is being accelerated by cross-referencing of search results, which calls for additional run time correlation algorithms. Furthermore, the search for mixed sequences must use the whole genome as data base. The work presented here and its extension can be a valuable contribution to personalized medicine using individual genetic profiles.

Earth Science / 51

Gridifying the Soil and Water Assessment Tool (SWAT) for the sustainable development of the Black Sea region

Author: Lukasz Kokoszkiewicz[1]

Co-authors: Andrew Maier [1]; Ann van Griensven [2]; Anthony Lehmann [3]; Dorian Gorgan [4]; Gregory Giuliani [3]; Nicolas Ray [3]; Seleshi Yalew [2]

[1] *CERN*

[2] *UNESCO-IHE*

[3] *University of Geneva*

[4] *Technical University of Cluj-Napoca*

Corresponding Author: lukasz.kokoszkiewicz@cern.ch

The Soil and Water Assessment Tool (SWAT) is one of the main tools used in the hydrology community. In the EU/FP7 EnviroGRIDS project it is used to study the Black See catchment and to simulate complex scenarios. In several cases, such as the study of the interplay of global and local scenarios, the availability of substantial computing resources in the Grid is important.

Detailed analysis:

We have ported the SWAT application to the Grid (EGEE infrastructure) and we are preparing to offer the software infrastructure to allow users a seamless access. Ease of use is provided by the availability of a simple yet complete monitoring (based on the Dashboard), and on a low- latency robust job service (Ganga/Diane). In order to further improve performances and the parallelization options within the EnviroGRIDS project, existing tools and activities will be integrated from within the Ganga/DIANE infrastructure, depending on performance and other user requirements.

Conclusions and Future Work:

Using grid technology in hydrology is a great opportunity to push it to another level, enabling people to simulate more and more complex and much larger models of our environment. It also brings a new, interesting use case to the grid itself, and provides important experience in solving problems like the parallelisation of applications that were not originally designed to run on the grid.

Impact:

This work aims to offer a service to the whole hydrology community. The utilization of grid massive computing power makes problems like the calibration and uncertainty analysis of large SWAT models solvable in reasonable periods of time. Thus it will enable the community to undertake much bigger challenges.

Keywords:

envirogrids, grid, swat, ganga, diane, black sea, hydrology

URL for further information:

http://www.envirogrids.net/

Workflow Management / **52**

A novel approach to Workflow Management in Grid Environments

Authors: Frank Berretz[1]; Sascha Skorupa[1]

Co-authors: Adam Belloum [2]; Bernd Schuller [3]; Volker Sander [1]

[1] *FH Aachen University of Applied Sciences*

[2] *University of Amsterdam*

[3] *Research Center Jülich*

Corresponding Author: berretz@fh-aachen.de

Most workflow management systems in Grid environments provide a push-oriented job distribution strategy, where jobs are explicitly delegated to resources. In those scenarios dedicated resources execute jobs according to the request of a workflow engine or Grid wide scheduler. This approach has various limitations, particularly if human interactions should be integrated in workflows. To support such interactions with the benefit of enabling volunteering approaches, this presentation proposes a pull-based approach, where resources can actively select tasks for execution from a repository.

Detailed analysis:

So far, workflow management systems in e-Science have focused on automatic job-execution. The actual schedule of a concrete workflow is basically a result of a decision process that relies on the workflow's history and the availability of resources in VOs. Human interactions, either within the decision process or with respect to tasks in a workflow, are not considered yet or rather cumbersome to implement. The established push-oriented job distribution in Grids imposes execution patterns that also limit scalability due to central decision points and that do not provide flexibility in the above discussed sense. The presentation proposes a pull-based approach by implementing a task repository mediating between resources and concrete workflows. This concept is based on the emerging WS-HumanTask standard and allows resources to request jobs according to their individual workload and capabilities. We implemented a prototype by extending the existing jBPM engine, so that tasks are sent to an external repository. These tasks are accessible via a Java client using Web Service Interfaces. Further challenges consist of open security issues, UNICORE integration and the proof of scalability.

Conclusions and Future Work:

The proposed solution is an evolutionary step towards the support of pull-based approaches in Grid environments resulting in better scalability and flexibility of current workflow executions. By integrating humans and special resources into workflows, this approach provides mentioned significant benefit. The presentation summarizes first results of the HiX4AGWS project that is gracefully supported by the BMBF. Several challenges still have to be solved. Particularly the prove of scalability, the UNICORE integration, and the implementation of a security infrastructure are work in progress.

Impact:

Instead of selecting resources by schedulers, in a pull-based approach the autonomous resources decide about the execution of a particular job. SLAs between resources and VOs as well as market-based approaches are viewed as support technologies for steering this selection process efficiently. By avoiding a central brokering strategy performed by a single scheduler, improved scalability is feasible. Also, the traditional role of resources is rather passive, since the enactment is done by a central entity, leading to the requirement of stringent VO management, which has to be continuously transferred to the user mapping lists of site administrators. Among the high administration complexity, this approach is not transferable to a more flexible workflow solution that relies on human interactions. Pilot-based infrastructures already propose a pull-oriented approach to realise a late job binding. However, pilot-jobs are still bound to a push-based distribution strategy, even though they provide a vehicle for implementing a pull-oriented concept for batch-oriented jobs. But neither human interaction, e.g. required in e-Health, nor resource-driven contributions are appropriately supported.

Keywords:

workflow management, human interaction, pull-based, resource scheduling, UNICORE

URL for further information:

http://www.fh-aachen.de/17458.html

Poster session / 53

Job management in gLite

Author: Marco Cecchi[1]

[1] *INFN*

Corresponding Author: marco.cecchi@cnaf.infn.it

The gLite WMS has been designed and implemented to provide a dependable, robust and reliable service for efficient and transparent distribution and management of end-user requests to high-end resources shared across a production quality Grid.
The WMS comes with a fully-fledged set of added-value features that hide to end users the complexity of such a heterogeneous and ever growing infrastructure and enable, thanks to a flexible, service oriented and general architecture, applications coming from largely different domains

Detailed analysis:

The WMS is responsible to translate users' requirements and preferences into concrete operations, interactions and decisions, in order to bring the execution of a request for computation, storage and the like (also known as 'job') to a successful completion. This is done transparently, while acting on behalf of the user.
Several types of jobs are supported: simple, intra-cluster MPI, interactive, collections, parametric and workflows in the form of directed acyclic graphs.
The Grid is a complex system and errors can occur at various stages throughout the so called submission chain. The WMS has the ability to automatically recover from infrastructure failures by implementing resilient strategies which include resubmission and retry policies. Additional benefits concern sandbox management - with support for multiple transfer protocols, compression and remote access - data-driven match-making, automatic credential renewal, service discovery and optimisations for collections such as bulk-submission and matchmaking.
Job tracking information in terms of relevant events, milestones and overall status can be retrieved and used by the WMS via the so called Logging & Bookeeping service.

Conclusions and Future Work:

After all these years operating in the EGEE infrastructure, the latest WMS releases have reached unprecedented stability and a performance which can smoothly accomodate for the current needs. By the end of EGEE-III, the WMS will have extended its support to more architectures and platforms.
Nevertheless, a new and challenging era is coming which will require the whole gLite stack to deal with other middleware distributions and an expanded use base. Consequently, the WMS will have to be deeply involved in managing different computing paradigms, standards, services and emerging technologies.

Impact:

Managing a grid job, from submission to completion, typically involves coordinating and interacting with a number of different services: computing elements, storage elements, information systems, data catalogues, authorization, policy and accounting frameworks, credential renewal. In this respect, the WMS, especially by virtue of his central, mediating role, has to deal with a wide variety of people, services, protocols and interfaces. Interoperability with other Grids must also be taken into account in this scenario.
On the user's side, the WMS exposes a Web Service based interface in accordance to the WS-I profile, which defines a set of Web Services specifications to promote interoperability. Access to the WMS is also granted by a dedicated User Inferface and APIs which are available in C/C++, Java and Python bindings.
Furthermore, the WMS fully endorses the Job Submission Description Language, an emerging standard which aims at facilitating interoperability in heterogeneous environments, through the use of an XML based job description language that is free of platform and language bindings.
On the resource's side, both legacy and OGSA\BES based interfaces are supported.

Keywords:

Job Submission and Management, Resource brokering, Interoperability, Grid Computing, Metascheduling

URL for further information:

http://web.infn.it/gLiteWMS

Astronomy and Astrophysics / 54

LOFAR Archive Information System

Authors: Andrey Belikov[1]; Fokke Dijkstra[2]; Willem-Jan Vriend[1]; Zheng Meyer-Zhao[2]

Co-authors: Adriaan Renting [3]; Danny Boxhoorn [1]; Dennis Van Dok [4]; Edwin Valentijn [1]; Floris Sluiter [5]; Hanno Holties [3]; Kor Begeman [1]

[1] *Kapteyn Astronomical Institute, University of Groningen, The Netherlands*

[2] *Donald Smits Center for Information Technology, University of Groningen, The Netherlands*

[3] *Astron, The Netherlands*

[4] *Nikhef, The Netherlands*

[5] *SARA, The Netherlands*

Corresponding Author: f.dijkstra@rug.nl

The LOFAR radio telescope, consisting of several stations with antennas, will become operational in spring 2010. The signals are collected and correlated by a central supercomputer. The resulting data, several petabytes a year, will be stored in an archive, distributed over a number of partners. The LOFAR Archive Information System plays a crucial role in this long-term archive. It will keep track of the data stored in the archive, and allows for further processing of that data. The information system supports different computing and storage technologies, including the EGEE Grid.

Detailed analysis:

The LOFAR Archive Information System is a cooperation between OmegaCEN/RUG, Big Grid and Target, and built on the basis of the Astro-WISE Information System. It supports different storage technologies, like the Astro-WISE data-server, SRM storage and the Target storage system. It not only keeps track of where the data is stored but also keeps track of the processing workflow, which was used for the data processing. Both the processing parameters, including software versions, and relevant intermediate results are stored. Processing jobs can be submitted to the EGEE Grid.
Much work has been performed on the integration of Astro-WISE with the Grid. Special attention has been paid to the integration of 3rd party java versions of Grid client software with Astro-WISE. This, to make the environment as user friendly as possible, without requiring a full gLite installation for all clients. For transferring proxy credentials to the job submission engine a MyProxy server is used.
The different storage interfaces are hidden from the user by making use of the object oriented approach of Astro-WISE. The proper storage interface to be used for the data is therefore instantiated automatically.

Conclusions and Future Work:

The LOFAR Archive Information System is a key component of the LOFAR long-term archive, enabling researchers to find and process LOFAR data, while keeping track of the data lineage.
Future work will focus on further refinement of the processing and storage capabilities, using input obtained during processing and service challenges. To allow general use of the archive, without requiring end-user certificates, the use of robot certificates is planned for services that allow retrieving data and starting up processing. This also allows immediate coupling with the Virtual Observatory.

Impact:

The construction of the LOFAR long-term archive is fundamentally important for astronomers who want to work with the LOFAR data. This archive will store petabytes of data, and will include processing capabilities. The LOFAR Archive Information System is completely metadata oriented. The full data lineage can be stored for the data that is processed, and data can be processed on the fly. Furthermore the Astro-WISE system is very flexible and gives the astronomers powerful tools to analyse their data, using a common approach and a common data model. All of this has been connected to the EGEE Grid as one of the suppliers of storage and computing.
What differentiates the LOFAR information system from other Storage Grid environments is the key role for the metadata, the full data lineage, and the target-processing capability, where processing is started when necessary.

Keywords:

Astronomy, LOFAR, Data management, metadata, processing, Grid

URL for further information:

http://www.astro-wise.org/Lofar.html

Demo Session 1, Welcome Drink / 55

NeuroLOG: a federated environment for neurosciences supporting multi-centric studies

Author: Alban Gaignard[1]

Co-authors: David Godard [2]; Diane Lingrand [3]; Franck Michel [4]; Johan Montagnat [1]; Rojas Balderrama Javier [3]

[1] *CNRS/I3S*

[2] *Visioscopie*

[3] *UNS/I3S*

[4] *INRIA*

Corresponding Author: johan@i3s.unice.fr

The NeuroLOG project develops a distributed environment dedicated to the support of collaborative and multi-centric studies in neurosciences. The NeuroLOG environment leverages grid technologies to deploy a highly secure, large scale data federation and to deliver high throughput neuroimages analysis. It is interfaced to the gLite middleware and takes advantage of the EGEE grid. It integrates in existing site-specific neuro-informatics environments. NeuroLOG clients, deployed on user desktops, provide neurosciences services through a coordinated federation of site servers and grid services.

Detailed analysis:

The NeuroLOG middleware has reached maturity and is deployed over a federation of 4 collaborating neuroscience centers. Its data management layer integrates neuro-images, associated neuro-pathological test results and other metadata stored in the databases. Advanced data representation technics, a data mediation and a data federation layer are used to represent in a unified view, manipulate and integrate the data source. A specific access control policy was designed to support multi-centric collaborations while ensuring that the local privacy policy of each provider prevails.

In addition, the middleware provides neuro-data analysis tools representation and deployment functionality. The collaborating centers can expose and share their data analysis algorithms for integration in neuro-image analysis pipelines. The MOTEUR workflow designer and enactor is used to describe neurosciences studies, taking advantage of the distributed computing resources provided by EGEE. The workflow engine is seamlessly integrated to the data management layer to facilitate the manipulation of experimental data sets. All facets of the NeuroLOG middleware are integrated in a unique graphical environment.

Conclusions and Future Work:

This demonstration will show the use of the NeuroLOG middleware in concrete neuroscience data analysis use cases: adding users to the federation, granting access to the resources, browsing the distributed data sets, selecting data sets relevant for a given study and executing a neurosciences pipeline. It will emphasise on the use of distributed resources and the exploitation of the EGEE grid for the addressing the computation needs. It will discuss the usability of grid resources for neurosciences and the perspectives to better support that community in the future.

Impact:

Neurosciences are increasingly relying on computerized analysis of large, coherent data sets including images and associated information on the clinical and environmental context. Indeed, epidemiological, therapy and drug impact studies require the analysis of large population of patient images over long periods. Furthermore, large data sets are required to build neuro-atlases characterizing the anatomy and physiology of the normal or pathological brain. Federating domain-specific resources is increasingly important to assemble the data sets required, or to compare data analysis procedures developed in neuroscience centers. Data sets can be assembled for pathologies with low occurrence rates. Specialized atlases can be composed for specific populations. Variations of pathologies can be studied over large geographical areas. Best practices can be identified and exchanged.

The neuro-scientists are standing half way between clinical neurologists and computer scientists. They are accustomed to the use of computing environments for supporting their experiments and they are a vector to demonstrate the use of computerized models and tools to the clinical world.

Keywords:

Neurosciences, Workflows, Data Representation, Data Federation

URL for further information:

http://neurolog.polytech.unice.fr

Justification for delivering demo and/or technical requirements (for demos):

The NeuroLOG environment is accessible through a high level, portable and integrated GUI. The client has a strong visual impact. It provides convincing grid-enabled use cases from a sensitive domain.

Computational Chemistry / 56

pKa Calculations of Key Ionizable Protein Residues in Acetylcholinesterase

Author: Jiri Wiesner[1]

Co-authors: Jaroslav Koča [1]; Zdeněk Kříž [1]

[1] *Masaryk University*

Corresponding Author: wiesner@chemi.muni.cz

One of the most abundant quantities characterising proteins is its isoelectric point, which is directly dependent on the number of charged ionizable residues, on the pKa of all ionizable residues more exactly. In this work, the pKa of buried protein aminoacids are estimated using method based on molecular dynamics called thermodynamic integration. The main scope is to compute pKa in acetylcholinesterase, but the computations on a small protein, thioredoxin, will be also presented as benchmark calculations.

Detailed analysis:

The pKa computations will be performed as a thermodynamic integration by Sander program from the Amber program suite. The thermodynamic integration method runs molecular dynamics with a mixed potential in different integration points with variable value of a mixing parameter. The mixed potential is combined from the potentials of two distinct states of the molecule differing by van der Waals or electrostatic properties of a few of its atoms. The result of thermodynamic integration is a deprotonation free energy difference between deprotonation in the solution and in the protein. The planned simulation length is about 10ns with 1fs timestep for each of the selected protein ionizable residues, which constitutes, even these days, considerable amount of computational time. In order to perform a thermodynamic integration calculation, Sander has to run in parallel on at least two processors. The whole run will divided into small pieces computed separately in the grid environment.

Conclusions and Future Work:

Even such sophisticated and computationally expensive methods as thermodynamic integration using force fields still provide only rough and qualitative results in the field of pKa computation. The future work will include further attempts to compute more precise free energy values using polarizable force fields.

Impact:

There are no reported computational studies on the pKa computations on acetylcholinesterase. The results will provide better picture of electrostatic interactions in the enzyme and valuable input for setting up the ionizable residues for force field molecular dynamics simulations. So far, the benchmark computations on thioredoxin are nearly finished, acetylcholinesterase will be computed in the grid environment. The main purpose of the benchmark was to reproduce thioredoxin pKa results already reported (Simonson et al., J.Am.Chem.Soc.,2004,126,4167.) and evaluate following: our model for the electrostatic change connected with deprotonation, alternative usage of new Amber force field 03, Self-Guided Langevin dynamics used to encrease sampling. We found the Amber force field 03 less suitable for this

task giving even higher overestimation of the free energy values than Amber force field 99. We were also able to successfully reproduce the reported free energy value for thioredoxin, 9.1 using Amber force field 99, 10.5. Our value is within the error estimated in the original work. The other evaluations and acetylcholinesterase computations are under way.

Keywords:

thermodynamic integration, acetylcholinesterase, thioredoxin, ionizable residues, pKa, ff03, ff99

URL for further information:

http://www.ncbr.chemi.muni.cz/group/lcc/acetylcholinesterase.html

High Energy Physics / **57**

LHC ATLAS users analysing data on the Grid

Author: Massimo Lamanna[1]

Co-authors: Dan Van Der Ster [1]; Johannes Elmsheuser [2]

[1] *CERN*

[2] *LMU Munich*

Corresponding Author: massimo.lamanna@cern.ch

In this talk we will describe the experience of the experiment ATLAS in commissioning the system (large scale user-analysis exercises in late 2009) and the first experience with real data.
First of all, we will describe the tools to commission and control the sites contributing to user analysis.
We will describe the support structure and the user experience (study of monitoring data and direct feedback from users). We will also show examples of analysis running on new data collected after the first collisions starting December 2009.

Conclusions and Future Work:

The effort ATLAS dedicated in preparing for distributed analysis is demonstrating to be a necessary investment in the preparation for the data taking and the analysis of the first LHC data. This experience is also an useful example for the entire Grid community.
The future direction is to extend the system in view of the increasing user activity (and more sophisticated use cases). Since most of the development are of general interest we are considering to team up with other VO to support and further develop our tools.

Detailed analysis:

Analysis in High-Energy Physics is characterised by a very large number of users. This is true both in absolute terms, due to the size of the HEP experiments (2000+ physicists in the case of ATLAS) and as a percentage of the research community.

The most important feature is that each user submits custom applications (normally built on top of the experiment-specific framework) instead of using pre-loaded services like in case of common portal applications.

In general user analysis tend to be I/O limited and iterative (sub data samples read many times by sets of users), which poses new constrains on site performance. The need to minimise the latency for users (compared to maximizing the throughput as in detector simulation and data reconstruction) is fundamental.

To prepare for the data taking, we have contributed (and we will report) in the following areas:
- Development and support of end-user tools (e.g. Ganga)
- Integration in the experiments framework (e.g. ATLAS Distr. Data Management (DDM) and in the experiment scheduling system (gLite WMS on EGEE, PanDA on EGEE, OSG and NDGF, ARC in NDGF)

- Setting up a user support and tutorials
- Organising the commissioning of sites

Impact:

Sharing the ATLAS experience in the field of analysis is beneficial for all the grid communities (users, operations and middleware):

-1- New user communities can consider to the ATLAS experience in order to assess the global impact of a pervasive grid adoption. Solutions like the distributed analysis shift system for supporting users will be illustrated.

-2- ATLAS put in production a sophisticated system to commission and control grid sites for analysis. The specific characteristics of this activity (short jobs, I/O contention, multiple users) should be taken into account by the sites. Some of the tools developed have an interest beyond the ATLAS community.

-3- The commissioning of sites for analysis produces performance data which can be used to compare different set up and different usages mode (in particular from a data access point of view). These data (like read efficiency, error rates) are collected and made available also to the middleware community. We plan to extend the error analysis in order to help to pin down common error sources. Similar data are important since they are extracted from the live system and can track the system evolution across HW and SW upgrades.

Keywords:

ATLAS analysis

URL for further information:

https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasComputing

Computer Science / 58

An investigation of the effect of clustering-based initialization on Learning Classifier Systems' effectiveness: leveraging the Grid infrastructure

Authors: Fani Tzima[1]; Fotis Psomopoulos[1]

Co-author: Pericles Mitkas [1]

[1] *Aristotle University of Thessaloniki*

Corresponding Author: fpsom@issel.ee.auth.gr

Strength-based Learning Classifier Systems (LCS) are machine learning systems designed to tackle both sequential and single-step decision tasks by coupling a gradually evolving population of rules with a reinforcement component. ZCS-DM, a Zeroth-level Classifier System for Data Mining, is a novel algorithm in this field, recently shown to be very effective in several benchmark classification problems. In this paper, we evaluate the effect of clustering-based initialization on the algorithm's performance, utilizing the EGEE infrastructure as a robust framework for an efficient parameter sweep.

Detailed analysis:

Clustering-based initialization is based on the idea that starting from a non-random set of rules may help the evolutionary process focus on the search-space optima (the optimal ruleset for the given classification task in our case) more effectively and quickly. Intuitively, this non-random set of rules should be based on the given dataset and provide an effective summary of the knowledge available in it. Our solution tries to leverage the potential of clustering algorithms to provide a representative set of centroids for a given dataset, that we then try to transform into rules suitable for the initialization of ZCS-DM.

The ultimate goal is to boost the algorithm's performance, both in terms of predictive accuracy and in terms of training times, through the reduction of the evolutionary process' execution time. In our current investigation, after detailing the proposed initialization process, we report the results of deploying the algorithm on the Grid infrastructure by means of a DAG workflow process. The conducted series of experiments evaluates alternative initialization parameter sets, aiming towards the optimization of the algorithm in terms of both efficiency and accuracy.

Conclusions and Future Work:

Our studies so far have proven ZCS-DM to be a robust and accurate data mining tool, which can outperform its rival algorithms in most of the benchmark datasets used and to achieve a prediction accuracy well above the baseline on all of them. However, given the evolutionary nature of the algorithm, further optimization in terms of time efficiency is necessary. In this direction, we have employed a clustering-based initialization phase and evaluated its effect on algorithm performance through an extensive set of experiments conducted by leveraging the Grid infrastructure.

Impact:

Among the various methods used to tackle classification problems, rule-based (or tree-structured) classifiers are particularly popular, because they combine: i) an intuitive representation that allows for easy interpretation of the resulting classification model; ii) a nonparametric nature that is especially suited for exploring datasets where there is no prior knowledge of the attributes' probability distributions; iii) fast, computationally inexpensive construction methods that produce models storable in a compact form; and iv) fast classification of new observations, once the model has been constructed. Inspecting the above list, one can easily conclude that LCS share most of the advantages of these methods, with the exception of the third point, as genetic algorithm-based search is an arguably slow and computationally expensive search method. Towards this end, the optimization of the ZCS-DM algorithm using Grid resources may provide researchers with an invaluable tool for performing data-mining tasks, and end-users with an efficient application for enhancing decision making tasks.

Keywords:

Classification, Learning Classifier Systems, Parameter Sweep, Algorithm Optimization

URL for further information:

http://issel.ee.auth.gr/

Computer Science / 59

Grid-enabled parameter initialization for high performance machine learning tasks

Authors: Fotis Psomopoulos[1]; Kyriakos Chatzidimitriou[1]

Co-author: Pericles Mitkas [1]

[1] *Aristotle University of Thessaloniki*

Corresponding Author: fpsom@issel.ee.auth.gr

In this work we use the NeuroEvolution of Augmented Topologies (NEAT) methodology, for optimising Echo State Networks (ESNs), in order to achieve high performance in machine learning tasks. The large parameter space of NEAT, the many variations of ESNs and the stochastic nature of evolutionary computation, requiring many evaluations for statistically valid conclusions, promotes the Grid as a viable solution for robustly evaluating the alternatives and deriving significant conclusions.

Detailed analysis:

Both NEAT and ESN involve a large parameter space of around 23 continuous parameters for the former and about a dozen for the latter. Even though our joint approach slightly reduces these numbers, the order of complexity remains the same. Our aim is to discover parameter areas that drive the algorithm

optimally while preserving its sensitivity to parameter variation. We apply Grid-enabled brute force search to parameter setups, with values chosen from small expert-created sets containing 2-4 parameters at a time. Performance is measured on three time-series benchmarks widely used in the community of ESNs: a) the Mackey-Glass system, b) the Multiple Superimposed Oscillator and c) the Lorentz attractor as well as on Reinforcement Learning testbeds including a) the single and double pole balancing and b) the 2D and 3D mountain car tasks. An additional search is made on the vicinity of the optimal parameter values and for the most influential ones. Statistical significance is also measured.

Conclusions and Future Work:

Our goal is to demonstrate the use of Grid resources for a) the optimisation of our algorithm for the topologically and weighted evolution of ESNs in terms of parameter fine-tuning and b) the macroscopic observation of the algorithm behaviour to different setups in order to eventually improve the underlying mechanisms. Our ongoing efforts include improving the algorithm in terms of speed, accuracy, generalisation and robustness and augmenting our pool of test cases for better performance evaluation.

Impact:

NEAT and ESN represent two of the most influential algorithms in the areas of neuroevolution and reservoir computing, respectively. To the best of our knowledge, there is no reported research on parameter initialization for NEAT, while for ESNs, reservoir optimisation is an active area of research. Their explicit combination has not been tried before and, in order to avoid local optimum behaviour, parameter fine-tuning is routed through Grid resources. Additionally, we believe that this study will provide a better intuition on how to improve such algorithms by studying macroscopic features, since microscopic observation is extremely difficult when evolutionary computation and neural networks are involved. The Grid enables such a strategy. The optimised results can be compared to recent efforts of optimising reservoirs and networks in the same domains. The overall goal is to make the algorithm more computationally efficient and thus applicable to real-life tasks.

Keywords:

Neuroevolution, Echo State Networks, Parameter optimisation

URL for further information:

http://issel.ee.auth.gr/

Computational Chemistry / 60

Porting of Computational Chemistry legacy applications on the EGEE Grid platform: computational aspects and visualization tools

Author: Alessandro Costantini[1]

Co-authors: Antonio Laganà [2]; Osvaldo Gervasi [2]

[1] *COMPCHEM-University of Perugia*

[2] *University of Perugia, Itlay*

The work carried out to implement complex computational chemistry suites of codes on distributed systems and, at the same time, to develop appropriate graphical tools for the visual rendering of the outcomes of the calculations on the production EGEE Grid infrastructure available to the COMPCHEM VO, is here presented and discussed with some examples.

Detailed analysis:

To this end the ABC code has been ported into the EGEE Grid environment using the P-GRADE Grid Portal. It has also been complemented with relevant visualization tools.
ABC is a quantum mechanical atom-diatom reactive scattering program that carries out accurate calculations of the quantum S matrix elements to evaluate reaction probabilities as well as state-to-state

integral and differential cross sections.
Typically, ABC can be executed several times for different sets of input parameter values, consuming a large amount of CPU hours and collecting large amounts of output data. The visual analysis of the output files carried out by the procedure has been greatly improved with respect to that of the standard P-GRADE version using a set of java based graphical tools implemented as standard pluggable user interface components called Portlets. The final user can in this way compare the outcomes, with no need to download all the output files which remain on the server, and evaluate the possible strategies for a new calculation.

Conclusions and Future Work:

The porting of legacy computational chemistry applications onto the Grid infrastructure, together with the development of the related visualization tools, is being carried out as part of a more general effort to build a solid platform for assembling accurate multi scale realistic simulations and for establishing an advanced molecular and material science research environment.

Impact:

The increasing availability of computer power on Grid platforms has become a strong incentive to implement complex Computational Chemistry suites of codes on distributed systems and to develop appropriate distribution models. On this ground the Virtual Organization (VO) COMPCHEM assembled out of a group of molecular and material sciences laboratories committed to implement their computer codes on the EGEE production Grid infrastructure is building a library of molecular dynamics codes to be offered to its users as services including not only the concurrent production of the numerical results but also related graphical rendering.
The implemented case study demonstrates the possibility to reuse the present porting process in order to provide a reusable example for other groups which are interested in porting their applications to production Grid systems.

Keywords:

Computational Chemistry, application porting, visualization tools

URL for further information:

http://compchem.unipg.it

Computer Science / 61

Phenomenology of Minority Games in Efficient Regime

Author: Karol Wawrzyniak[1]

Co-author: Wojciech Wiślicki [2]

[1] *Interdisciplinary Centre for Mathematical and Computational Modelling, University of Warsaw*

[2] *A. Soltan Institute for Nuclear Studies*

Corresponding Author: kargulek@gmail.com

We present a comprehensive study of the utility function of the minority game in its efficient regime. We develop an effective description of the state of the game. For both the step-like payoff function g(x) = sgn(x) and the proportional function g(x)=x, we explicitly represent the game as the Markov process and prove the finiteness of number of states. We also demonstrate boundedness of the utility function. Using these facts we can explain all interesting observable features of the aggregated demand: appearance of strong fluctuations, their periodicity, and existence of preferred levels.

Detailed analysis:

Minority game (MG) was designed as a microscopic model of adaptive behavior observed in multi-agent systems. The MG is a typical bottom-up construct and therefore the usual definitions of the game first specify rules of behavior for individuals. Then, piecing together microscopic variables, one defines higher-order quantities characterizing grander systems. Despite the simplicity of the basic rules of

taking decisions by agents, adaptive abilities and phenomenology of populations playing MGs appear to be surprisingly interesting and their properties are nontrivial. It was shown that the MG exhibits different modes of behavior, depending on the game parameters: the random, cooperation, and herd. The latter case is characterized by a small strategy space compared to the overall number of agents. Our study of this regime is motivated by the interesting phenomenology observed in numerical simulations and the lack of satisfactory interpretations of them. In our previous work we found, in a different context, that the crucial role in the explanation of observable behavior in the MG is played by the utility function. Therefore, we further exploit the utility to study the phenomenology of MGs in their efficient regime.

Conclusions and Future Work:

MG is a stochastic, multi-agent model where some interesting phenomena are observed only for large populations. The analysis of each set of parameters requires the processing of copious sets of realizations of the game but no synchronicity, and therefore no heavy data traffic during execution, is needed between them. Using the EGEE infrastructure we were able to study the MG in the deep efficient mode. Our results are based on more than 7000 runs, where each lasted from few minutes to few hours. We observed interesting collectivity in agent behavior and provided the detailed mathematical explanation.

Impact:

Depending on the payoff function g(x), the game is driven by different dynamics which requires different methods of the analysis. For the step-like payoff function g(x) = sgn(x), we explicitly represent the game as the Markov process and prove the finiteness of number of states. Since the MG represents a system with many degrees of freedom, the dimensionality of states is expected to be large. We demonstrate the boundedness of the utility function which allows for the substantial reduction of the number of state parameters and simplification of the state description. For the step-like payoff the state is reduced to (i) a history of m minority decisions and (ii) utilities of all pairwise different strategies. For the proportional payoff, g(x) = x, the number of states is still finite and the utility remains bounded, but effective analysis requires a different concept of state. In such a case, the state is based on the order of utilities in the ordered list instead of their values. Using these representations we can explain all interesting observable features of the aggregated demand: the appearance of strong fluctuations, their periodicity, and the existence of preferred levels.

Keywords:

Minority game; adaptive system; Markov process; de Bruijn graph; EGEE infrastructure

URL for further information:

http://kargul.polgrid.pl

User Support and Services / 62

User Support for Distributed Computing Infrastructures in Europe

Author: Torsten Antoni[1]

[1] *Karlsruhe Institute of Technology*

Corresponding Author: torsten.antoni@kit.edu

In a worldwide distributed computing infrastructure such as EGEE (Enabling Grids for E-sciencE) one of the challenging tasks is to build and maintain a reliable user support infrastructure. During the last years the GGUS (Global Grid User Support) system has been developed as EGEE's answer to that challenge. GGUS is designed as a centrally coordinated distributed infrastructure that integrates existing tools instead of trying to replace them.
The next challenge will be to adapt the user support infrastructure to the new operations model on which the EGI/NGI infrastructure will be based.

Detailed analysis:

With the series of EGEE project reaching its end in 2010 and the move towards building a sustainable infrastructure based on the national grid initiatives from a large number of European countries and governed by EGI (European Grid Initiative), it will become necessary to adapt the user support infrastructure to fit this new operations model for grid computing in Europe. The big challenge here will be scaling the infrastructure up to 40 NGIs instead of 12 ROCs and at the same time ensuring the production quality of the service provided for the user communities by the established partners. The new partners will need assistance from EGI in reaching the same high level of service.
The EGI infrastructure will consist of a large number of independent projects (EGI-InSPIRE, NGIs, EMI,...). A frictionless and fluent communication between all these partners will become even more important than it is already now. It has to be based on generally accepted tools and has to adhere to agreed formalisms. GGUS will be one of the main components in this communication.

Conclusions and Future Work:

This presentation will describe the status of the user support infrastructure close to the transition from EGEE to EGI and present user support processes and integration plans for EGI. The interplay between the involved actors, like EGU.eu, NGIs, EMI and SCCs, will be presented to give a full picture of the EGI user support spanning various projects.
Additionally an outlook will be given on how a unified user support infrastructure for the major European DCIs could be realised.

Impact:

The number of user communities making use of different types of computing resources is constantly growing. Therefore one of the mid-term goals of the European e-Infrastructure strategy is to provide the users with seamless access to various distributed computing infrastructures (DCIs). Achieving this will also be a key element in reaching sustainability for these projects. An area of utmost importance for this is to provide a single point of contact for a user experiencing problems with one of the infrastructures. The user should not be forced to acquaint himself with the details of the various DCI's support tools.
GGUS could be used as the integration platform providing this single point of contact as well as providing the means of communication between the experts from the various infrastructures , like for example EGI, PRACE and GEANT. Creating this could be achieved without replacing tools currently in use in the projects, as experience from EGEE has shown.

Keywords:

user support, help desk, EGI, DCI

URL for further information:

www.ggus.org

Workflow Management / **63**

OpenMOLE: a grid enabled workflow platform

Author: Romain Reuillon[1]

Co-authors: Mathieu Leclaire [1]; Nicolas Dumoulin [2]

[1] *ISCPIF*

[2] *CEMAGREF*

Corresponding Author: romain.reuillon@iscpif.fr

OpenMOLE is a free and open source workflow engine providing distributed computing facilities, especially suited for scientific research in complex systems. Third-party software packages can be embedded in a workflow that
automatically transfers and processes their input and output (files and data). Embedded software packages are called "tasks". Any task within a workflow can be either executed locally or delegated to distributed computing environments, including the EGEE grid.

Detailed analysis:

The main purpose of OpenMOLE is to provide a high-level workflow platform for designing scientific experiments. OpenMOLE decouples the scientific business logic from the resources used to execute it. It enables the definition of scientific workflows and the delegation of workflow tasks to high-performance computing environments in a declarative way.

In a first phase, we have conceived workflow tasks in such a way that the platform is able to migrate them on demand to another execution environment than the end-user local personal computer (PC). Thus we have defined what is an execution context and the consumed resources of a task. The implementation of these concepts allows remote execution of a given task.

We also faced the challenge of establishing direct links between a user PC and external distributed computing environments. Although OpenMOLE does not assume that the user PC owns a public IP address, no particular third-party server (i.e., other than the user PC and the execution environment servers) is ever required. Furthermore, no preliminary installation step on a remote environment is required: OpenMOLE transports everything it needs to be run remotely.

Conclusions and Future Work:

The core functionalities of the workflow engine have been implemented. The greatest part of our effort will be spent next on the development of a user-friendly GUI to take full advantage of OpenMOLE. A component is also under construction to keep track of experiments carried out in OpenMOLE and allow collaborative design of workflows among scientific communities. In parallel to platform development, more specific applications and scientific methods will be tested to ensure the broadest coverage of the diverse and interdisciplinary scientific domains composing complex systems.

Impact:

Taking advantage of grid computing remains a tricky issue for the non-expert user. Software and hardware are heterogeneous, bad workload management decisions happen and, generally, failure rate is higher than other computing environments. In this context, a certain amount of technical and methodological knowledge must be acquired before making efficient use of grid computing. Fortunately, many types of applications and methods contain inherent parallel aspects. For these problems, we claim that a platform can completely hide the intricacies of execution environments to the user.

Software tools such as g-Eclipse, JSAGA or Ganga propose a high-level object layer to abstract the execution environment. Yet, they only partially hide the technical details and overall heterogeneity. OpenMOLE goes further and hides the whole distributed environment of the business layer.

Other software tools such as Taverna offer similar features. The OpenMOLE project, however, follows a different approach in which everything runs on the user desktop by default and no third-party server is ever called. Tasks
are delegated on demand to distributed execution environments by the user.

Keywords:

Workflow, generic platform, model exploration

URL for further information:

http://www.openmole.org/

Computer Science / 64

Modeling Grid Job Time Properties

Author: Lovro Ilijasic[1]

Co-author: Lorenza Saitta [2]

[1] *University of Turin*

[2] *University of Eastern Piedmont*

Corresponding Author: lovro@di.unito.it

Large-scale analysis of grid job data, consisting of more than 28 million jobs gathered during 20 months from all major EGEE Resource Brokers, provides interesting results regarding the properties of job time parameters. The main parts in a job's life cycle are observed to have different distributions, having also different origins that cause these behaviors. While 'match time' parameter corresponds better to power law distribution,the total job length is well modeled with log-normal distribution. Here we attempt to fit the distributions and give explanation through generative models.

Detailed analysis:

We can recognize three main parts in the life cycle of a job: time from registration to Resource Broker dispatching job to a suitable Computing Element (match time); time the job waits in CE's queue (wait time); and time it was actually being executed (run time). We can also observe the total lifetime of the job, the sum of the three, as a separate parameter, exploring also the stake of the parts in the total length. When analyzing distributions of these parameters, we note the straight line signature in log-log scale diagram, particularly for the 'match time' parameter, pointing to power law. However, the 'total time' parameter can be fitted more accurately with log-normal distribution, which comes as a consequence of the generative model considering the dependencies between consecutive jobs. Multiplicative processes, where in every step the size of the event (here, job length) grows or shrinks according to a random variable multiplier, can be applied here to give the explanation for the possible log-normal distribution.

Conclusions and Future Work:

Close relations between power-law and log-normal distributions have already been noted in the literature. Very small variations in generative models are shown to decide between the two distributions of event sizes. As in the case of job length parameters, other generative models can be tested and applied, possibly bridging the gap between the two with the double Pareto, or double Pareto log-normal distributions.

Impact:

The results of the large scale analysis of job length parameters give us insight into global behavior of the grid network. Power laws and log-normal distributions are often associated to natural processes and are related to emergent behavior of complex systems. In this case, understanding the distributions of time parameters can be used in network simulation, optimization, scheduling and self-management. Understanding how jobs are correlated between themselves and what kind of behavior it causes on global level is valuable information in many different contexts, proving it usable even for predicting future job parameters. Besides, more analysis is performed regarding the correlation between different time parameters of the same job (like efficiency, for example, which is defined as run time over total time), bearing useful results as well.

Keywords:

Job length, distribution, run time, generative models, efficiency

URL for further information:

www.grid-observatory.org

Poster session / 65

Harnessing Heterogeneous Computing Resources for Solving Hard Computational Problems

Author: Mikhail Posypkin[1]

[1] *ISA RAS*

Corresponding Author: mposypkin@gmail.com

The talk presents the BNB-Grid tool aimed at solving hard combinatorial, discrete and global optimization problems in a distributed heterogeneous computing environment. The BNB-Grid can run applications across different service and desktop grid platforms as well as individual workstations and clusters. The BNB-Grid efficiently copes with difficulties arising in such systems: the software diversity, unreliability of nodes and problems with batch (queuing) system. The talk discusses the implementation details and computational results for two challenging problems.

Detailed analysis:

The BNB-Grid tool can harness the consolidated power of computing elements collected from service grids, desktop grids and standalone resources to solve hard optimization and combinatorial problems. Adding new type of computational resource is done in an easy and transparent way via shell scripts incapsulating details of the concrete middleware. Currently the tool supports pure SSH, Unicore and gLite service grids, BOINC and XWHEP desktop grids. The BNB-Grid hypervisor submits applications to different computing elements and organizes their interaction via specially designed protocol. Interaction includes management commands to start and stop computations and exchanging algorithmic information. Two ways of interaction are supported – via TCP/IP sockets and via files where establishing the socket connection is problematic. The BNB-Grid has been successfully applied to molecular conformation problem that plays an important role in computational chemistry and to cryptanalysis of A5/1 cryptosystem. For both problems new results were obtained.

Conclusions and Future Work:

In the talk we described the BNB-Grid tool for harnessing heterogeneous computational resources to solve large scale combinatorial and optimization problems. We also discussed the implementation issues of two hard optimization problems from different domains. In future we plan to improve load management part in BNB-Grid and consider new classes of optimization problems.

Impact:

The BNB-Grid tool provides an application-level interoperability among different types of computational resources. The flexible scripting mechanism simplifies adding new resources to the system. Different ways (via TCP/IP or file system) of interaction between the hypervisor and the running application make it possible to easily connect different types of resources. The efficiency of the proposed approach has been demonstrated on two hard optimization problems.

Keywords:

Global and Combinatorial Optimization, Application Level Interoperability, Grid Computing

URL for further information:

http://dcs.isa.ru/posypkin

Computational Chemistry / 66

Protein Molecular Dynamics and Free Energy Calculations on the EGEE Production Grid

Author: Stavros Farantos[1]

Co-authors: Manos Giatromanolakis [1]; Massimiliano Porrini [1]; Osvaldo Gervasi [2]; Vangelis Daskalakis [1]

[1] *Foundation for Research and Technology-Hellas/ Institute of Electronic Structure and Laser*

[2] *Department of Mathematics and Computer Science, University of Perugia*

Corresponding Author: farantos@iesl.forth.gr

Atomistic simulations of large biomolecules such as proteins require extensive computational resources. Dynamical and thermodynamical properties can be obtained either by averaging over very long time trajectories or sampling the phase space by running hundreds of short time trajectories.

The latter methods are the most appropriate for high throughput computers such as the Grid distributed computers. Algorithms are presented for calculating vibrational spectra of the active site of cytochrome c oxidase as well as free energy landscapes based on thermodynamic perturbation theory.

Detailed analysis:

Proteins are large molecules with thousands of atoms whose motions cover a broad range of time intervals, from a few tens of femtoseconds, the oscillation periods of strong chemical bonds, to milliseconds, the period of large scale conformational changes. The study of dynamics of such complex systems in a broad range of spacial and temporal scales as well as thermodynamic quantities and structural properties remain a challenge for atomistic simulations. Here, we present an algorithm that assists us to harness the current computational Grid infrastructure for carrying out extended samplings of phase space and integrating the classical mechanical equations of motion for long times. A bundle of shell scripts has been written to automatically submit and propagate trajectories in the Grid and to check and store large amounts of intermediate results. We report our last years experience in employing the Enabling Grids for E-sciencE production infrastructure via the CompChem and SEE virtual organizations in investigating the dynamics of enzymes such as Cytochrome c Oxidases.

Conclusions and Future Work:

The positive experience gained from running classical dynamics of proteins on the EGEE Grid will be transfered to projects involving Quantum Molecular Dynamics. A highly parallelized Fortran code written for solving the time dependent Schroedinger equation by formulating the Hamiltonian in Cartesian coordinates will be deployed into the Grid [1].

[1] Jaime Suarez, Stavros C. Farantos, Stamatis Stamatiadis, and Lucas Lathouwers, A method for solving the molecular Schroedinger Equation in Cartesian coordinates via angular momentum projection operators, Comp. Phys. Comm., 180:2025-2033, 2009.

Impact:

It is demonstrated that by using the thousand of cpus available in EGEE important biophysical/biochemical problems can be solved in a realistic way. Converged thermodynamic quantities can be obtained. Also the availability of flexible scripts can cope with drawbacks of the Grid and make runs stable and long running.
Results for Cytochrome c Oxidases have been published [1-3].

[1] V. Daskalakis, S. C. Farantos, C. Varotsis, Assigning vibrational spectra of ferryl-oxo intermediates of Cytochrome c Oxidase by periodic orbits and Molecular Dynamics, J. Am. Chem. Soc. 130(37), 12385-12393, 2008.
[2] Massimiliano Porrini, Vangelis Daskalakis, Stavros C. Farantos, and Constantinos Varotsis, Heme Cavity Dynamics of Photodissociated CO from ba3-Cytochrome c Oxidase: the Role of Ring-D Propionate, J. Phys. Chem. B, 113(35), 12129-12135, 2009.
[3] Vangelis Daskalakis, Stavros C. Farantos, Victor Guallar, and Constantinos Varotsis,
Vibrational Resonances and CuB displacement controlled by proton motion in Cytochrome c Oxidase,
J. Phys. Chem. B, in press, 2009.

Keywords:

Atomistic Simulations, Classical and Quantum Molecular Dynamics, Distributed computing

URL for further information:

http://tccc.iesl.forth.gr/general/intro/node1.html

Bioinformatics / 67

Estimating the Performance of BLAST runs on the EGEE Grid

Authors: Abel Carrión-Collado[1]; Ignacio Blanquer-Espert[1]; Vicente Hernández-García[1]

[1] *UPV*

Corresponding Author: iblanque@dsic.upv.es

Estimating the response time of large experiments is a key issue for achieving an efficient load balancing and minimizing the failure rate. This requires having a good knowledge of both the application and the infrastructure. This work describes a set of experiments that have conducted to the definition of a performance model that can be used to estimate the response time of the selected resources and to adapt the partition of the load to the dynamic status of the resources.

Detailed analysis:

The work has concentrated on two lines: First, the analysis and selection of the parameters that provide information about the performance of a resource, using the GLUE schema (SPECint, SPECfp, average and maximum queuing time among others). Second, the analysis of the factors affecting the performance of BLAST searches (reference database size, input sequence size, sequence length, operation parameters and resemblance of input and reference data). Therefore, several experiments have been completed fixed one or several of the variables. The experiments were repeated fixed the same conditions to reduce variability and increase significance. Input data has been based on the UniProtdatabase, introducing random modifications to control the execution conditions. Submission was performed explicitly selecting the CEs and measuring the information provided by the BDII. The publication delay is also considered as an error factor. Results were obtained in different resources at different times, and a performance model is being adjusted using this information.

Conclusions and Future Work:

The work describes a model to estimate response time of BLAST runs in the EGEE grid. The work will be completed with experiments validating the model and introducing new parameters from other components, such as the Workload Manager and the LRMS. This will provide a more complete model to characterise the execution of BLAST jobs in the grid, and the improvement of scheduling policies. Dynamic adaptation of load is feasible, especially on pilot submission schemas, which are also considered.

Impact:

The results obtained have revealed that database size and input data size have a direct linear impact on the response time, thus leading to a simple prediction model. Variability due to similarities on the input and reference data, as well on the request for larger or shorter output has a minor impact on the performance and can be ignored. Results on the effect of the performance capabilities of the results are on the way. Static values will be easier to fix, but dynamic parameters will need to be feed directly on the model.
Although there are other studies on the literature about the estimation of performance of BLAST, in the knowledge of the authors this is the first study on a production Grid infrastructure. The results will be relevant, not only for the users of the biomed community, but also for the developers of QoS schedulers. This is an important research line, especially considering the sustainability of Grid infrastructures, and the consideration of external providers.

Keywords:

Performance estimation, BLAST, dynamic scheduling

URL for further information:

www.grycap.upv.es

Infrastructure Tools and Services / 68

Site Status Board: WLCG monitoring from the experiment perspective

Authors: Jacobo Tarragón Cros[1]; Julia Andreeva[1]; Marco Ferreira Devesas Campos[1]; Pablo Saiz[1]; Ricardo Rocha[1]

[1] *CERN*

Corresponding Author: jacobo.tarragon.cros@cern.ch

Now that the LHC has started, the experiments require a high standard of reliability and performance on their computing activities. Monitoring these activities is not a trivial task mainly due to two reasons: first of all, asserting the proper behavior of a site depends heavily on the software model of each experiment; secondly, the number of sites taking part in WLCG has increased drastically compared to previous HEP experiments.

Detailed analysis:

The Site Status Board (SSB) web application, developed under the Dashboard Experiment framework, has been designed to provide an overall view of the sites performance from the experiment perspective. Designed originally for the LHC VOs, it allows the experiments to define a set of activities, also known as views. For each view, the experiment administrator can define the metrics that have to be collected. For instance, CMS has currently five different views ('computing shifters', 'site commissioning', 'space monitoring', ...). For the first view, the metrics include the number of running jobs, transfer status, availability of software on the site, etc.

The SSB collects the status of the metrics over time and presents it in several formats.
SSB will also include pointers describing the possible errors and solutions if this information is provided. Thanks to the SSB, the organizations can analyse site statuses easily, and at the same time, they keep track of the evolving metric results.

Conclusions and Future Work:

Production level services are being built using the Site Status Board. At the same time, the application is constantly evolving since it needs to adapt to the experiments growing needs. Future changes will focus on a performance boost for the historical data browsing, improvements on the reliability of information gathering, and extending the flexibility of the metric definitions.

Impact:

The SSB is being widely used by CMS and LHCb for several activities: computing shifts, site commissioning and space monitoring in the case of CMS, and job and space monitoring in LHCb. ATLAS and ALICE are also evaluating the SSB.

Keywords:

grid, monitoring, site, status

URL for further information:

http://dashb-ssb.cern.ch/ssb.html

Poster session / 69

Gustav : a lightweight approach to CPU accounting

Authors: Emidio Giorgio[1]; Gianni Mario Ricciardi[2]

Co-authors: Roberto Barbera [3]; Soonwook Hwang [4]

[1] *Unknown*

[2] *KISTI institute*

[3] *University of Catania & INFN*

[4] *KiSTI institute*

Corresponding Author: emidio.giorgio@ct.infn.it

Gustav is a CPU usage accounting tool developed by INFN, COMETA and the KISTI institute. Gustav collects accounting records from resources and publishes them to a centralised relational database, that can be queried through a web interface. Gustav is characterised by a lightweight architecture, that makes its usage ideal for small sized infrastructures. However, interoperability with more powerful tools, like DGAS or APEL, is accomplished by the adoption of standard format for accounting records.

Detailed analysis:

Gustav publishers, implemented in Python, run on resources and analyse periodically batch systems logs to produce statistics about CPU usage by grid users. Records are published to a central collector, implemented through a MySQL database; the central collector runs also a web interface that make possible to perform queries about collected data. The query mechanism is very flexible, allowing to focus from single user/VO/resource/day to any combination of these parameters. Supported batch systems are currently Torque and LSF. The publisher implementation, which relies more on the batch system rather than the middleware, makes Gustav usable for CPU accounting of any grid middleware, like gLite, UNICORE or ARC, as long as it supports one the above mentioned batch systems. Beside support to most widely used batch systems, it allows also interoperability with more widely used tools, making Gustav records processable by them through the adoption of standard formats for records.

Conclusions and Future Work:

In a distributed computing infrastructure, based on shared resources, an accounting system is a fundamental requirement. Gustav aims to reach this goal in a simple and reliable way, with a simple installation and an easy to use web interface. It's independent of the middleware used to run the Grid infrastructure, and supports two of the most common LRMS. Efforts will be made in the near future to enhance its capability to show collected data in a graphic way, and plugins for other LRMS such as Condor and SGE are being considered. The deployment of Gustav on Asian sites has been also planned.

Impact:

Gustav is currently deployed on the regional Grid of the COMETA Consortium in Italy, where it proved to work well on a gLite-based infrastructure made up of a dozen of sites, tens of VOs and some hundreds of users. The installation performed by local system administrators ran smoothly, and the whole system has been deployed very quickly. Several tools were already available for accounting purposes, like DGAS/HLRmon or APEL. They have a clean and robust design, suitable for larger scale infrastructures with some thousands of sites and users. These tools provide also exhaustive aggregated usage information. However their deployment is not trivial, and can require a large effort for regional Grids or single sites requiring a much simpler accounting system. Furthermore, given the small sizes of infrastructures whom Gustav has been designed for, it is not required to aggregate records, which are stored with their full details, allowing in this way a finer granularity on queries.

Keywords:

CPU usage accounting, gLite, UNICORE, ARC, grid, PBS, Torque, LSF

URL for further information:

https://gustav.consorzio-cometa.it

Poster session / 70

An integrated monitoring system for Grid Data Centers

Authors: Alessandra Doria[1]; Gianpaolo Carlino[1]; Giovanni Battista Barone[2]; Giuseppe Vitagliano[2]; Guido Russo[3]; Leonardo Merola[3]; Maurizio Pollio[2]; Silvio Pardi[1]; Vania Boccia[2]

Co-authors: Alessandro Russo [2]; Armando Cosimo [2]; Enrico Maria Rossi [2]; Giuseppe Petrizzo [2]; Luigi Ilardi [2]; Salvatore Naddeo [2]

[1] *INFN, Napoli*

[2] *Universita' Federico II, Napoli*

[3] *Universita' Federico II & INFN, Napoli*

Corresponding Author: guido.russo@cern.ch

We present a monitoring system developed for the Data Centers used for the SCoPE and ATLAS projects in Napoli, Italy. The system is based on a portlet container which gives an integrated view of the Data Center, and allows a graphical-based, hiearchically organized navigation for all the equipments, from the racks to the active components. The system allows monitoring of the whole infrastructure (UPS, cooling, electrical power consumption at single socket level), but also of the network (Gigabit, 10 Gigabit, Infiniband, Fibre Channel) and, of course, of storage and server nodes.

Detailed analysis:

We are approaching the integration of applications for monitoring, realizing a portal-integration (PI) and an enterprise-information-integration (EII).
The work we have done has followed this path: (i) identification of the functionalities that need to be monitored, along with their characteristics; (ii) identification of the already available commercial or self-made software products, if any; (iii) identification of the available open–source software; and (iv) development of dedicated plug -in and java web – applications.
The work has been carried out because of the need of having a dedicated troubleshooting and monitoring system for the newly built Data Center (we opened on february 2009), for which we needed a robust, flexible and easily adaptable system, without the complexity of functionalities often present in commercial systems but of rare usage, if any. But soon the job moved to the construction of a true portal system for the Data Center, which shall soon become the unique point-of-entry for all kinds of access, from the novice user, the expert user, the management team, and the referee group which has its role in future funding.

Conclusions and Future Work:

More modules and more functionalities are being implemented, as well as an optimization of existing ones, initially integrated off-the-shelf, e.g. by using struts as a development framework for java web application. As an example of a tool which we are implementing, we are actively working on a tool for automatic emergency shutdown and restart, and for programmed shutdown and restart for maintenance downtimes, on a rack basis.

Impact:

The system we realized represents an integrated system, giving to most users a unique approach in accessing the grid Data Center; in particular, for the management team, it allows a thorough integration of all monitoring subsystems, and thus able to easily accommodate new hardware and software tools, thus following the evolution of the Data Center. For first-level monitoring, the approach we followed guarantees a simplified view of the infrastructure, so that operators shifts do not need very qualified personnel. We realized our own modules for a lot of equipment, but we also integrated existing applications such as GRIDICE – GANGLIA – POWER FARM. Such an integration strategy is in our opinion very useful in a situation, like the one in EGEE, in which a single tool is not enough for an exhaustive management of all the aspects of the Data Center.

Keywords:

monitoring - grid - data center

URL for further information:

www.scope.unina.it

Poster session / 71

MEG - MyProxy Enabled GSISSHD

Author: Kevin Haines[1]

[1] *STFC*

Corresponding Author: kevin.haines@stfc.ac.uk

MEG allows users to logon to a grid-resource using any SSH-enabled client, so long as they have uploaded a credential to a MyProxy server. MEG accepts the username and password to the credential from the SSH client and retrieves the proxy on behalf of the user, and the uses the proxy to determine if the login can proceed. Resource providers benefit too, as only grid-based authentication mechanisms are needed, so a UI box can be implemented without adding an additional layer of user management.

Detailed analysis:

The user starts by using his certificate to generate a credential that is stored in a MyProxy server. The Myproxy Upload Tool [1] can be used to accomplish this task. During this operation the user will select a username and password that allows the credential to be retrieved at a later time.
At any point during the lifetime of the delegated credential, the user can run an SSH (or SFTP) client of choice to connect to the MEG resource. To login, the user supplies the username and password of his MyProxy credential. Then, MEG uses these to retrieve the credential from the MyProxy server, and uses that credential to authenticate the user against the resource. Assuming this succeeds, the user is logged into the resource, where a proxy credential will be waiting in the environment for further use.

Conclusions and Future Work:

The system has proved very popular with users at STFC and on the UK-NGS. Further extensions to this system should be simple due to the very modular nature of the solution.

Further benefits have been proved within STFC, using a MyProxy-SSO (Single Sign On) server, and the UK NGS is looking at providing a Shibboleth-based extension.

Impact:

MEG is a benefit for all users of X509-based grid resources, by enabling them to choose which SSH client they want to use. It is lightweight, small, and easy to maintain and understand, and removes the need to maintain SSH-based portals (which require further user account management).

Keywords:

gsisshd ssh myproxy meg

URL for further information:

http://wiki.ngs.ac.uk/index.php?title=KGSISSHD

Justification for delivering demo and/or technical requirements (for demos):

Laptop (own supplied), and if possible a large screen.

High Energy Physics / 72

Client/Server Grid applications to manage complex workflows

Authors: Daniele Spiga[1]; Fabio Farina[2]; Filippo Spiga[2]

[1] *CERN*

[2] *Dipartimento di Fisica "G. Occhialini"-Universita degli Studi Milano-Bicocca*

Corresponding Author: filippo.spiga@cern.ch

Grid Applications adopting a client/server paradigm allow easier enforcement, improvement and optimization of the middlewares. A server allows us to enact specific and complex workflows, in order to centralize application management and hiding Grid complexities to the end users. These features aim to enable Grid to a large and heterogeneous community not only at the infrastructure

level, but also at application level. A successful example for the grid server paradigm is presented for CMS and its analysis tool, CRAB.

Detailed analysis:

HEP computing models require the coexistence of specific and complex workflows granting high scalability as their main QoS. The Grid provides an effective solution to scalability problems thanks to infrastructures that allow building distributed applications coping with specific use cases. We propose the adoption of a server layer between the users and the Grid middleware based on CRAB. This intermediate service guarantees modularity and overcomes the heterogeneity of the Grid, but also represents the key point to implement a flexible and reliable application by isolating further the end users from their specific application workflow details (e.g. data location catalogs interactions and data movements duties). The client/server paradigm lets in addition to build a manageable environment with the chance to handle centrally policies and deployment of new features.

Conclusions and Future Work:

Significant experience has been achieved in CMS during past years through many computing challenges. A solid methodology to design flexible and scalable Grid servers has been acquired. Future activities are aimed at generalizing and spreading the proposed approach by defining a common core of libraries to be shared within various CMS workload management projects.

Impact:

The proposed solution as been enforced by the CMS experiment with the implementation of a service dedicated to end user analysis workflows execution. There are more than thousand users from different countries who run more than 40K jobs every day. A deployment activity has been performed successfully and our large community is using the service both for Monte Carlo simulations and for the analysis of the data collected during first operational phase of LHC until the end of 2009.

Keywords:

distributed analysis, HEP, Client/Server

URL for further information:

https://twiki.cern.ch/twiki/bin/view/CMS/SWGuideCrab

73

Constraints on primordial non-Gaussianity using Planck simulated data: First Results

Author: Marcos Lopez Caniego Alcarria[1]

Co-authors: Andres Curto [1]; Enrique Martínez-González [1]

[1] *INSTITUTO DE FISICA DE CANTABRIA (CSIC-UC)*

After the successful porting of the point source detection code and the SZ clusters detection code to the EGEE GRID, we have ported and tested a new application. This application is composed of two codes, one that produces simulations of Planck and another one that looks for non-Gaussianity signatures in these maps using spherical wavelets. These applications are part of an ongoing project being carried out by the Observational Cosmology and Instrumentation Group at the Instituto de Física de Cantabria (CSIC-UC) on different analyses of the Cosmic Microwave Background (CMB).

Conclusions and Future Work:

References
1.Martínez-González E., Gallegos J. E., Argüeso F., Cayón L., Sanz J. L., 2002, MNRAS, 336, 22
2.Vielva P., Martínez-González E., Barreiro R. B., Sanz J. L., Cayón L., 2004, ApJ, 609, 22 3.
3.Cruz M., Turok N., Vielva P., Martinez-Gonzalez E., Hobson M., 2007, Science, 318, 1612
4.Curto, A., Martínez-González, E., Barreiro, R. B., 2009, ApJ, 706, 399

5.Komatsu E., Spergel D. N., 2001, Phys. Rev. D, 63, 063002
6.Liguori M., Matarrese S., Moscardini L., 2003, ApJ, 597, 57
7.Liguori M., Yadav A., Hansen F. K., Komatsu E., Matarrese S.,Wandelt B., 2007, Phys. Rev. D, 76, 105016

Detailed analysis:

This application consists of: 1) 11.000 Gaussian CMB maps are simulated and processed through the Planck simulation pipeline for 4 frequency bands, 70, 100, 143 and 217 GHz, producing four maps per simulation. 2) these maps are combined and convolved in harmonic space with the spherical mexican hat wavelet, a tool used in the past to do several Gaussianity analyses of CMB maps (see for example 1,2,3,4). The convolved map, the wavelet coefficients map, will contain information about those structures of the initial map with a characteristic size of R and will be used to compute third order statistics in a similar way as the bispectrum (see 5). We will repeat this process for several angular scales R between 2.9 arc min - 170 degrees. This process is going to take approximately 70K CPU hours and 550 GB of disk. Finally these statistics are used to constrain the levels of non-Gaussianity of the local type (see 6,7,8) which could be present in Planck data. In particular, we have focused on the uncertainties of a parameter known as the non-linear coupling parameter, fnl. The constraints have been obtained using both analytical and numerical methods. We will present the first results.

Impact:

The uncertainties on fnl, sigma_fnl, depend on the cosmological model, the instrumental properties and the available fraction of the sky to be analyzed. For instance, a better estimation on fnl can be achieved with a low instrumental noise and a high sky coverage. The Planck mission will provide CMB maps with unprecedented resolution and quality at several frequencies. Therefore, we expect to obtain very competitive constraints on the primordial non-Gaussianity with this experiment. The study of this kind of non-Gaussianity has become a question of considerable interest as it can be used to discriminate different possible scenarios of the early Universe and also to study other sources of non-Gaussianity non-intrinsic to the CMB.

Keywords:

Application Porting, Planck, Non-Gaussianity

URL for further information:

http://max.ifca.unican.es/webcmb/research/research.html

Computational Chemistry / 74

A multiconfiguration time-dependent Hartree algorithm for non-Born-Oppenheimer calculations

Author: Dimitrios Skouteris[1]

Co-authors: Antonio Lagana' [1]; Osvaldo Gervasi [1]

[1] *University of Perugia*

Corresponding Author: dimitris@dyn.unipg.it

We have embarked upon a project of performing dynamical calculations treating electrons and nuclei on exactly the same footing,
considering all particles as degrees of freedom of a dynamical system. In that way, the Born-Oppenheimer approximation is circumvented. The very concept of a potential energy surface is abolished and the problems associated with kinetic energy couplings or Berry phase simply vanish. Since an exact treatment of this kind would be prohibitive, we split the various degrees of freedom according to a MCTDH scheme.

Detailed analysis:

According to the MCTDH scheme, in our code, the overall wavefunction of the system is expanded as a sum of configurations. Each configuration is a product of single degree of freedom (DOF) wavefunctions (orbitals) and has its own coefficient in the overall representation.

The single configuration limit is analogous to the Hartree-Fock limit in electronic structure calculations. The wavefunction is propagated according to a variational principle analogous to the time-dependent Schroedinger equation, both coefficients and orbitals being time-dependent.

All orbitals can be visualised throughout the calculation, thus providing a qualitative image of the dynamical progress. At each point in the calculation, the norm and the energy of the wavepacket are calculated in order to ensure their conservation. Moreover, the error introduced (into both orbitals and coefficients) by time discretisation into steps is estimated and the time step is accordingly modified. If necessary, care is also taken of Pauli antisymmetrisation.

Quantities of experimental interest are obtained throughout the calculation, such as autocorrelation functions and time-dependent flux functions.

Conclusions and Future Work:

We are currently utilising our code for simple, prototype three-particle systems such as the collision between a H+ ion and a H atom to study the probabilities of inelastic scattering of the ion and/or transfer of the electron between atoms. Various initial angular momenta are considered and their effect on the overall process is evaluated.

In the future, we plan to extend our work to systems with more than one electrons. Such an extension is expected to scale very favorably with the number of degrees of freedom and the Pauli principle antisymmetrisation is very easy to implement.

Impact:

The MCTDH method lends itself very conveniently to use in a grid environment. The decomposition of the overall problem into several single degree of freedom problems, where intercommunication is only necessary at the beginning of each time step, suggests the handling of each degree of freedom by a single processor and the consequent facilitation of the computation, both in terms of memory and computer time.

Moreover, going beyond the Born-Oppenheimer approximation entails many advantages. Problems traditionally associated with this approximation include nonadiabatic couplings and Berry phase effects, which simply cease to exist. Furthermore, the user can, at each point in time, visualise the orbitals for each degree of freedom separately. This provides unique visualisation opportunities which would be harder to implement in a more traditional
'exact' calculation.

As with any dynamics calculations, the quantities calculated throughout the process (autocorrelation function, time-dependent flux etc.) can be used to predict optical spectra and scattering cross sections. Comparison with the result of more traditional calculations can help to evaluate the effect of the approximation.

Keywords:

wavepacket, MCTDH, wavefunction, orbital, flux, autocorrelation, algorithm

URL for further information:

http://www.chm.unipg.it

Novel Architectures and Technologies / 75

Enabling the use of e-Infrastructures with Microsoft HPC and the Matlab distributed compute server

Author: zeeshan ali shah[1]

Co-authors: Erwin Laure [1]; Mikael Sterner [2]; Saeed Yahyanejad [1]

[1] *PDC-Centre of High Performance Computing, KTH*

[2] *Microsystem Technology Lab, School of Electrical Engineering, KTH*

Corresponding Author: zashah@pdc.kth.se

Many scientists with smaller scale problems could benefit from e-Infrastructures but are often discouraged by their complexity. These users have little experiences with shell based Linux environments typically offered; instead, they
often use Windows-based platforms and higher level packages like Matlab. In this talk we present a case study from microsystems research, focusing on how the usage of recent technologies like the Microsoft HPC server and the Matlab distributed computing toolbox can improve the productivity of the researchers and allows them to exploit existing e-Infrastructures.

Detailed analysis:

While portals are an excellent means of shielding the user from the complexity of e-Infrastructures, this approach has limitations, particularly, it typically doesn't allow the user to run their own programs. Many users still require direct access to run their computation but have little experiences with shell based Linux environments typically offered on today's e-Infrastructures.
Instead, these users often use Windows-based platforms and higher level packages like Matlab. This has also been recognized by EGEE and prototypes that allow the integration of resources enabled by Microsoft HPC and Matlab's distributed computing toolbox have been developed, the latter has even been used for a tutorial at SC'08.

Microsystems research is one examples of users that could benefit significantly by having these tools offered as a regular service on today's e-Infrastructures. These users often require complex simulations of e.g. fluidic or Electromagnetic systems, which are on the limit of a standard PC. These users depend on e-Infrastructure providers like local data centers or campus Grids to enlarge their available resources without having the burden of administering these systems.

Conclusions and Future Work:

We have shown how the entrance barrier to exploit e-Infrastructures can be lowered by using recent technologies like the Microsoft HPC server and the Matlab distributed compute toolbox thus making the infrastructures accessible to new user communities.
For future work we intend to expand the prototype locally by increasing the number of resources available to it and integrating it with the existing EGEE infrastructure making it accessible worldwide. We will also explore different federated identity management solutions like Shibboleth.

Impact:

By providing access to a Windows HPC and Matlab distributed compute server based platform at PDC, KTH, the Microsystems researchers have seamless access to an enlarged resources pool and benefit from professional support available at the data center. These resources act as an extension of their local desktop and relying on the active directory based user management already established at KTH the usage of these resources is actually completely transparent to the user.

This prototype shows the potential e-Infrastructures have supporting also users with small to medium sized problems and how the entrance barrier for this class of users can be lowered such that e-Infrastructures become a viable tool for them. Using the techniques prototyped by other groups in integrating Microsoft HPC with the EGEE infrastructure we intend to widen the scope of this service from local campus e-Infrastructure to an international scale. This will also require the use of federated identity management system as experience has shown
that these kind of users have great difficulties in dealing with Grid certificates.

Keywords:

MS HPC, Matlab

URL for further information:

http://www.pdc.kth.se/projects/mshpc/

76

Computational chemistry experiments over a web service-oriented grid

Author: Vincenzo Barone[1]

Co-authors: Gianni Morelli [2]; Julien Bloino [3]; Nico Sanna [2]; Tiziana Castrignano [2]

[1] *Scuola Normale Superiore*

[2] *CASPUR*

[3] *Scuola Normale superiore*

Corresponding Author: vincenzo.barone@sns.it

We will describe the implementation of effective numerical differentiation techniques aimed to obtaining different energy derivatives of quantum chemical energies and properties on a service-oriented grid where the computational workflow is spawned over multiple, geographically distributed, sites. The application porting over the grid and its extension as a web service over local and wide area networks is fully outlined.
A description of the procedures developed for these experiments is given in full detail and the applicability of the methodology to similar problems pointed out.

Detailed analysis:

The application is composed of three layers each performing a given task: (1) a pre-processing tier; (2) a task execution tier and, finally (3) a data collector
tier for the post-processing phase.
The first step of the pre-processing is implemented via a proper web
interface where a WS interacts with the end-user for the authentication, data input and job configuration.
In a second step the pre-processing WS builds up an execution table containing all the input files that are generated taking into account the user data of the first step. The domain of this phase is the http server running the WS and all the data are locally stored and managed.
The second tier of the application is a java class responsible for the
task generation for each of the input files created by the pre-processing WS. The java class, dispatches over the grid nodes the scripts for
the execution of the QM code. While the load balancing algorithm
is able to cope with various execution queues to obtain the maximum node loads, QM jobs depend also on the specific input data.
The post-processing stage is the most customizable one depending on the specific results to be extracted from QM outputs.

Conclusions and Future Work:

We have presented an innovative solution
to the execution of top class applications such those currently
in use in QM calculations.We have
setup a grid environment and developed the necessary
code sections to port largely used QM packages
over a distributed, wide area network grid. This
computational experiment has proven to provide a
stable environment when built up with advanced
technologies such as service-oriented grid architectures
and web services.
The computational benefits of this solution have been
proven to be even better than locally executed parallel applications

Impact:

Computational chemistry requires huge computing resources in order to
solve quantum mechanical (QM) Hamiltonians describing the properties of molecular species by means of variational and/or perturbative many-body approaches.
As a consequence, several studies are in progress with the aim of setting up innovative strategies able to deal in an effective way with the increasing

complexity of the molecular systems of current interest, which are still in the domain of classical simulations.
However, in the foreseeable future software developments will not change the general situation that leading calculations using quantum chemistry methods
require a huge computing power either in terms is of CPUand/or I/O resources and top-level computer architectures are often used. Those computing machines are often built on top of commodity components and this
trend seems well consolidated even for the upcoming parallel architectures.
As an alternative we propose a new Web Service which has several new appealing features, i.e. (1) it is architecture independent,(2) it is modular, that is, configurable for user or site requirements,(3) it is dynamical, that is,
executable on available nodes.

Keywords:

Computational chemistry, energy and property derivatives, Web Service

URL for further information:

http://idea.sns.it

Regional Activities, International Projects and Collaborations / **77**

Disseminating the Grid: An essential guide

Authors: Catherine Gater[1]; Neasan O'Neill[2]; Sy Holsinger[3]

[1] *EGEE*

[2] *QMUL*

[3] *Trust-IT*

Corresponding Author: catherine.gater@cern.ch

When it comes to dissemination, what works and what doesn't? During EGEE-III, NA2 has coordinated a rich programme of activities through an increasingly wide range of media – press releases, booths, blogs, Twitter and websites to name just a few. This session gathers together best practices based on the lessons learnt with input from collaborating dissemination project, GridTalk. Practical examples and case studies highlight how to deal with the media, make an impact on policy makers, get business leaders interested in the grid and bring new user communities on board.

Detailed analysis:

EGEE's NA2 brings together 27 partners from 22 countries; this session examines the advantages and disadvantages of a geographically distributed team – how much effort is needed to execute an effective dissemination plan and how much does it all cost? Regardless of a project's size, collaborating with other projects and organisations can reach a wider range of audiences. The session explores ideas for building up working relationships and networks, including MoUs and media sponsorship deals.

Events are a highly effective means for engaging new and existing audiences in your work, but which should you attend? With booth fees sometimes running to thousands of Euros, it pays to target these carefully. The session reviews some of the highlights from the dissemination team's travels, including the must-attend events, how to build a better booth and most importantly, how to let everyone know that you are there.

Getting businesses excited about grid technologies is a particular challenge and this session also presents the lessons learnt from the EGEE business outreach activities – through the Business Forum, targeted business events and the successful Business Associates programme.

Conclusions and Future Work:

Once the EGEE-III project closes, the future of grid dissemination, like the grid itself, will be highly distributed – dissemination will be shared by a host of communications teams, including the EGI.eu communications team, the Specialised Support Centre dissemination activities and the National Grid Initiatives. This session concludes by introducing some of the expected new players in the field of grid dissemination including EGI.eu and the National Grid Initiatives and outlines how the user community will continue to benefit from highly professional and dedicated dissemination teams.

Impact:

The impact of dissemination activities is potentially very large, with some events attracting ten thousand or more delegates. Press releases are circulated to thousands of journalists across dozens of countries, triggering articles and interviews that can reach a whole range of audiences, from the specialist press to the general public. The LHC GridFest event in October 2008, which EGEE contributed to, generated 160 international press clippings from TV, radio and press and significantly raised the profile of grid computing in the minds of the general public and policy makers.

Policy makers and funders can also be targeted highly effectively by publishing articles in key high circulation journals and by sending booths to international policy events. By outlining the benefits that grids offer to science and society, this kind of outreach can help to ensure essential funding for future e-Infrastructures. The grid only thrives through its active user communities – by reaching out to new and existing users through publications, events, social media and community-building websites, dissemination activities help to bring these communities together.

Keywords:

dissemination, communications, outreach, business, media

URL for further information:

http://press.eu-egee.org/index.php?id=228; http://www.gridtalk.org

User Support and Services / 78

Supporting diverse local virtual organisations

Author: Douglas Mcnab[1]

Co-authors: Mike Kenyon [1]; Sam Skipsey [1]; Stuart Purdie [1]

[1] *University of Glasgow*

Within ScotGrid-Glasgow, primarily an ATLAS Tier-2 grid site, we are involved in supporting a number of local virtual organisations (VO)'s, including Electrical Engineers, Solid State Physicists (SSP) and Optics all with very specific non-high energy physics requirements (HEP). An account of the main issues and achievements that resulted from working with other groups at our site who sought direct access to a batch system rather than a grid site will be presented. This includes trialling various submission mechanisms and middleware components to achieve the goals of each of the specific VO's

Detailed analysis:

At ScotGrid-Glasgow primarily a High-Energy Physics (HEP) grid cluster we support a number of local virtual organisations (VO)'s through our grid site. These include Electrical Engineers, Optics and a Solid State Physics Group (SSP). Each group has its own unique requirements that differ from the usual HEP usage of large VO's with many people to support and run the VO, produce and install software and VO operated submission systems that we see at ScotGrid. The lessons learned from this will be presented. The Electrical Engineers had not run on the EGEE grid and were used to a local batch system. The issues from transitioning from batch to grid are presented alongside the measures taken to increase adoption rates within their community. Optics had a requirement to run licensed software on the grid called FDTD by Lumerical Inc. A solution to this licensing issue using functionality within CREAM will be presented. Optics and SSP had a requirement to run MPI codes, FDTD and CASTEP respectively. Something new for ScotGrid not necessarily EGEE. A discussion of issues surrounding this implementation are presented.

Conclusions and Future Work:

There are lessons to be learned from the experiences of other user communities using ScotGrid and from sites administrators working directly with them to accomplish their goals. This work has described three local VO's and their use cases on ScotGrid. Issues encountered and their solutions have been presented with specific examples. Future work will include continuing to support non-HEP communities and easing their adoption of grid technologies.

Impact:

This research has allowed ScotGrid to understand the requirements across many diverse virtual organisations (VO)'s and learn from this work to ease the transition of new users to the Grid world. More importantly it has enabled researchers in their various fields to carry out their work more efficiently. At ScotGrid this has meant helping our electrical engineering group run their work on the EGEE grid and enabling users of Lumerical's FDTD package to perform large-scale design of devices across a diverse range of applications in biophotonics, display technologies, solar energy, optical communications, sensing and imaging. Running MPI codes has brought up some suggestions for new functionality within the current EGEE MPI implementation. This work has lead to the creation of a unified batch/grid submission framework titled gqsub (submitted under a different abstract) which has applications beyond ScotGrid.

Keywords:

ScotGrid Glasgow Lumerical FDTD CASTEP MPI CREAM

URL for further information:

https://www.scotgrid.ac.uk/wiki/index.php/Users

Justification for delivering demo and/or technical requirements (for demos):

no demo

Data Management / 79

Enabling Large Data Transfers Between Web Services

Author: Spiros Koulouzis[1]

Co-authors: Adam Belloum [1]; Edgar Meij [1]

[1] *Informatics Institute, University of Amsterdam, the Netherlands*

Corresponding Author: s.koulouzis@uva.nl

Despite numerous benefits, many Web Services (WS) face problems with respect to data transport, either because SOAP doesn't offer a scalable way of transporting large data-sets or because orchestration workflows (WF) don't move data around efficiently. In this paper we address both problems with the development of the ProxyWS. This is a WS utilizing protocols offered by the Virtual Resource System(VRS)(Java API used by the VBrowser), to enable other WS to transfer and access large datasets without modifying WS nor the underlying environment.

Detailed analysis:

There is currently an abundance of deployed (legacy) WS using SOAP, which fail to produce access and return large datasets. Moreover, orchestration WF causes WS to pass messages containing data back through the WF engine. To address these problems we introduce the ProxyWS: a WS that is able to access data from remote resources (GridFTP, LFC, etc.), thanks to the VRS, and also transport larger data produced by WS, both legacy and new. For the ProxyWS to be able to provide larger data transfers to legacy WS, it has to be deployed on the same Axis-based container, just like a normal WS. This enables clients to make proxy calls to the ProxyWS instead of a legacy WS. As a consequence the ProxyWS

returns a SOAP message containing a URI referring to the data location. For new implementations the ProxyWS is used as an API that can create data streams from remote data resources and other WS using the ProxyWS. This approach proved to be the most scalable since WS can process data as they are generated from producing WS. Thus with the introduction of the ProxyWS we are able to provide a separate channel for data transfers, that allows for more scalable SOA-based applications.

Conclusions and Future Work:

We have presented the ProxyWS, which may be used to support large data transfers for legacy and new WS. We have verified its performance to deliver large datasets on two real-life tasks: Indexing using WS in a distributed environment and annotating documents from an index. From our experiments we have found that ProxyWS is able to facilitate data transports where normal SOAP messages would have failed. We have also demonstrated that with the use of the ProxyWS legacy WS can scale further, by avoiding data delivery via SOAP and by delivering data directly from the producing to the consuming WS.

Impact:

Many different approaches have been introduced in an attempt to address the problems mentioned earlier. Examples of these include Styx Grid Services, Data Proxy Web services for Taverna and Flex-SwA. Some noteworthy features of these approaches are: Direct streaming between WS, Usage of alternative protocols for data transports, and larger data delivery to legacy WS. However, each of these examples only addresses one part of the problem and, furthermore, do not include any means of allowing access to remote data resources. Leveraging these existing proposals and combining them with the VRS we implemented a ProxyWS. To validate it, we have tested its performance using 2 data-intensive WF. The first is a distributed indexing application that uses a set of WS to speedup the indexing of a large set of documents, while the second relies on the creation of that index for retrieving and recognizing protein names contained in results coming from a query. With the use of the ProxyWS we are able to retrieve data from remote locations (8.4 GB of documents for indexing), as well as to obtain more results relative to a query (8300 documents using the ProxyWS versus 1100 using SOAP).

Keywords:

web services, data transfers

URL for further information:

http://staff.science.uva.nl/~skoulouz/pmwiki/index.php/ProxyWS/ProxyWS

Regional Activities, International Projects and Collaborations / 80

NDGF approach to serve the non-LHC communities needs: Nordic Material Science VO

Author: Ivan Degtyarenko[1]

[1] *CSC - IT Center for Science Ltd.*

Corresponding Author: ivan.degtyarenko@csc.fi

We present the joint NDGF and CSC (Finland) project on setting up and supporting a new Nordic Material Science Virtual Organization inside NDGF (Nordic MS VO). The activities are focused on making the flexible and Nordic MS VO users friendly grid environment available on ARC enabled resources.

Conclusions and Future Work:

The work is in progress. Within the next few months the applications to be checked on grid usability, mainly the licensing issue. Up to 10 scientific application will be chosen to be at the VO support level. VO support will prepare Runtime environment configs, and end-user documentation on the NDGF wiki. After the installation and preparation phase is over the ongoing user support work will be continued.

Detailed analysis:

The Nordic material science community - physicists and chemists and other computational scientists, but not limited to a particular field of science - is spread over all Nordic countries. Therefore NDGF support for the community is a natural addition to its list of supported virtual organizations (VOs).

Tha main target of the VO supporters is non-LHC computational scientists. The main goal of the project is to get those scientists running their jobs on ARC enabled resources. We identified several factors which make a grid service attractive and usable for the computational scientists. Namely (i) availability of the pre-installed scientific applications on the grid resources, (ii) solid MPI environment support, (iii) possibility to run long time jobs and (iv) working environments for the interactive use.

In this project we address the first two issues.

The project has been started in Oct 2009 and will last till the end of the 2010.

Impact:

The approach is based on the solid numbers. We have collected the statistical information from the VO participants identifying their needs in application software. We have also found out type of the most popular MPI flavour and collected statistics about CPU number usage by applications.

The impact is that the Grid must become a user-friendly environment for the former cluster and main-frame users. Nordic MS VO to be easily expandable. Starting with the 10-15 participants from the Nordic countries we can easily add more. Having 5-10 top applications pre-installed from very beginning we can consider other applications pre-installation, based on the upcoming end-users needs.

Keywords:

Nordic MS VO, NDGF, end-user environment, ARC, runtime environment

URL for further information:

http://ndfg.org/

Computer Science / 81

Distributed evolutionary optimization framework on the Grid and HEP use case

Author: Adam Padee[1]

Co-author: Krzysztof Zaremba [1]

[1] *Warsaw University of Technology*

Corresponding Author: adam.padee@cern.ch

This paper describes an optimization framework which utilizes a distributed evolutionary algorithm with the ability of adaptation to the complex physical structure of the grid fabric. The algorithm is divided into several partitions (demes) which are located on physical clusters. This architecture allows simultaneous use of all the available resources, regardless of their geographical location, within one application. The paper also presents an example use case of this technique, which is track reconstruction optimization for COMPASS experiment.

Detailed analysis:

Numerical optimization by means of EA (Evolutionary Algorithms) became popular in many disciplines of technology and science because it is more resistant to local optima than classical, gradient-based methods. On the other hand, EA require much more computational power, and this limits their usage. The grid environment, with its huge number of CPUs,seems to be a perfect environment for this kind of applications. The main difficulty here is how to use Grid resources, coming from many clusters, in a coordinated way within one program, without huge performance losses on communication. In this

application we adopted the multiple-deme architecture of EA in which every cluster is occupied by separate master-slave populations. These populations (called demes) occasionally exchange their best solutions, but most of the communication takes place within the cluster. For internal communication each population use MPI messages, and for external (between demes), files on the Storage Elements accessed via GFAL. This way, there is no need for non-standard extensions to grid middleware, and the performance penalty caused by slow access to files is marginal (due to the low intensity of this kind of communication)

Conclusions and Future Work:

The tests performed on the example application show that even big differences between processing speeds of the clusters hosting individual demes have little influence on the final solution quality, making the method very useful. However, slow communication between demes may be the limiting factor when using the program to solve problems with less computationally expensive objective functions (and thus faster swapping of generations). To extend the functionality we plan to introduce faster communication methods to the migration operator, like PACX-MPI, or DIANE messages.

Impact:

In the scientific world there is a strong need for general-purpose optimization frameworks. This is confirmed by a vast number of tools of this kind available in the Internet, offering broad variety of optimization algorithms, including EA. These tools usually work in abstraction from the hardware layer, making it possible to run on many platforms. This comes at a price. It is virtually impossible to run such a general-purpose tool in a highly heterogeneous grid environment and utilize all the available resources in a coordinated way. Our program is aware of the structure of the underlying fabric and is specifically designed to minimize the losses caused by communication bottlenecks. This allows to use EA-based optimization for the problems too big to fit on one supercomputer. Development of our program was initially driven by one application, which is track reconstruction optimization in COMPASS experiment. Still, a considerable effort has been put into maintaining the modular structure of the program. It is thus possible to easily adapt it to other problems, changing only a couple of options in the configuration file. For this reason the framework can be useful for other users.

Keywords:

Evolutionary Algorithms, Computational Grids

URL for further information:

http://wiki.polgrid.pl/index.php/DEAG

Poster session / 82

Using the EELA-2 grid Infrastructure to Perform HeModynamics Computer Simulations: The HeMoLab Experience

Author: Paulo G. P. Ziemer[1]

Co-authors: Bruno Schulze [1]; Pablo J. Blanco [1]; Raúl A. Feijóo [1]

[1] *LNCC - National Lab for Scientific Computing*

The main goal of the HeMoLab project is the development of computational models used in the simulation of the Human Cardiovascular System, performed by a general purpose numerical solver that apply distributed computing techniques through the use of the MPI paradigm. This paper describes the main steps of the porting process of SolverGP to gLite environment, covers the main characteristics of gLite, grid tools that have been used in the gridification process, and examples of performance results of computer simulation that were obtained from the application already ported to Grid of the EELA-2.

Detailed analysis:

The simulation of the human cardiovascular system is performed using 1D and 3D models. The computer simulation is performed through the solution (using Finite Element Method) of a Navier Stokes

equations system that describe the motion of fluids, such as blood inside of an artery. A set of numerical methods were implemented using Fortran 90 and the libraries PETSc and Parmetis in a framework called SolverGP, that uses MPICH2 as the MPI standard implementation.
We created a module between HeMoLab and EELA-2 infrastructure, in order to launch the remote script for execution of the SolverGP, encrypt the input files, trigger the capabilities for jobs monitoring and receive the output files generated remotely.
Before submitting a job, the encripted inputFiles, the SolverGP executable, and the HeMoLab libraries must be submitted. A JDL file is configured to define the files that will be sent directly to the Worker Nodes: the script of solver, the Watchdog (to enable the online monitoring) and the Secure Store (libraries required to decrypt the files).
After the calculation of the simulation, the files are moved from Grid Node to User Interface, in order to be sent for display in HeMoLab.

Conclusions and Future Work:

This work covers the main characteristics of gLite, the Grid tools that have been used in the gridification process and examples of performance results of computer simulation that were obtained from the application already ported to EELA-2. We will use the C++ API from the HeMoLab interface in order to create a more friendly way to submit jobs, run some larger cases (that take more than a week of computing time), make tests to discover the speedup of SolverGP, explore DIRAC capabilities (and others tools) and integrate our cluster in the EELA-2 infrastructure.

Impact:

Some of our calculations need many days or a month for processing. Thus the HeMoLab project, along with the infrastructure of EELA-2, will provide simulating computational tools that can be used in the training of physicians and speeding up surgical planning. For instance, the computer model of an aorta, needs approximately 500,000 points to represent the artery geometry and then originates a system with 3,500,000 equations. The computing of just one heart beat period in such a model takes approximately 35 days (with 640 time steps) when using 8 processors in a 2X quadcore Xeon 3GHz 64GB RAM.
Despite the delay in the allocation of processors for the execution of jobs (in tests with 60 processors), we have noticed a decrease in processing time with respect to the tests used in our local cluster of 8 machines. Through the creation of the interoperability layer between HeMoLab and the services of the EELA-2, the processing power had a great increase. The increase from 36 to 715 processors will permit the speedup of SolverGP. Another advantage is the ability to run 8 jobs in a parallel way (number of EELA-2 sites that satisfy requirements to run SolverGP).

Keywords:

Scientific Computing; Grid-based Simulation and Computing; Complex Systems Simulation

URL for further information:

http://www.lncc.br/prjhemo/

Demo Session 2 / 83

Interoperabilty among gLite and ARC middleware in training grid infrastructures: GRIDSEED testcase

Authors: Sergio Maffioletti[1]; Stefano Cozzini[2]

Co-authors: Iztok Gregory [3]; Riccardo Murri [4]; Tyanko Aleksievic [5]

[1] *Grid computing Competence Center (GC3), university of Zurich.*

[2] *CNR-INFM Democritos and ICTP*

[3] *CNR-INFM Democritos and Sissa/eLab*

[4] *CSCS*

[5] *Sissa/elab*

Corresponding Author: cozzini@democritos.it

GRIDSEED provides a simple tool to setup up a portable fully fledged gLite grid infrastructure based on virtual machines. It exploits the concept of "grid in a box" providing a self-containing grid that could be easily deployed in any existing infrastructure for training and dissemination purposes. In this work we present a recent effort to include ARC middleware in the original GRIDSEED infrastructure
The purpose of the exercise is to provide a training facility where users could learn how two similar grid middleware can interact and have a first knowledge on interoperability issues.

Detailed analysis:

Gridseed now incorporates a complete minimal and functioning GLite installation as well as an ARC installation. The novelty of this approach is the exploit of the commonalities of both infrastructures with an emphasis on the functional interoperability between the two.
The UI from both middleware have been unified in a single gridseeed_UI where users as enabled to use same credential as well as same VOMS proxy to interact with both CE frontends.
Both LCG_CE and ARC_CE are configured to access the same computing resources.
Both GLite and ARC are using the same LFC file catalogue as well as the same storage element.
The above services are complemented by a fake Certification Authority. Any user can generate her personal certificate and/or a host certificate to add for instance a new VM host to the system through simple web interfaces.
The information system interoperability has not yet been reached GLite and ARC are using different and non-compatible schema. New release of ARC is supposed to use Glue-2 schema.
On top of the middleware components, GRIDSEED provides high level tools like GANGA and DIANE and other demo application developed within euindia grid project.

Conclusions and Future Work:

This work presents a first attempt to setup a training facility that would provide to end users the possibility of facing grid interoperability features and issues when using quite widely deployed European grid middleware solutions.
While using GRIDSEED, users are able to experience different levels of interoperability: Middleware level, scientific gateways level, application access level.
GRIDSEED will be expanded to include support for other grid middleware like GARUDA used in the Indian Grid infrastructure, as well as other high level tools like web-based portals or workflow engines

Impact:

GRIDSEED tool was developed to easily deploy a training grid infrastructure almost everywhere in the world with the only requirement a set of machine (simple PC are) locally connected among them. It uses widely available and standard virtualization tool like VMware.
It has been used several times in Training events and it is a key training resource for EU-IndiaGrid project.
In this work GRIDSEED has been enhanced in order to be a completed training environment formed by a virtual infrastructure where different grid middleware are inter-operating. High-level tools are also installed to help users in porting, deploying and accessing their application across the different available middleware.
The purpose of the exercise is to provide a training facility where users could learn how to use ARC and GLite middleware and have a first knowledge on interoperability issues at different levels.
At the best of our knowledge GRIDSEED is, at the moment, the only grid virtual training infrastructure that provides an easy configurable environment where end users can experiment with interoperability issues on tightly integrated middleware.

Keywords:

Virtual training infrastructure, GLite, ARC, interoperability

URL for further information:

gridseed.escience-lab.org

Justification for delivering demo and/or technical requirements (for demos):

A live demo of the features of GRIDSEEDwould result in a more direct impact for the end users.
We are currently working on it but we would like to take a final decision on live demo at a later stage

Bioinformatics / 84

Grid assisted structure calculation of large protein systems in solid-state NMR context

Authors: Alexandre Joseph[1]; Anja Bockmann[1]; Antoine Loquet[1]; Christophe Blanchet[1]; Fabien Mareuil[2]; Michael Nilges[3]; Therese E. Malliavin[3]

[1] *CNRS IBCP*

[2] *Unite de Bioinformatique Structurale, URA CNRS 2185 et Institut Pasteur*

[3] *Institut Pasteur*

Corresponding Author: christophe.blanchet@ibcp.fr

Due to the present experimental limitations of solid-state NMR (ssNMR), 3D structure calculations of proteins using these NMR spectra as a source for structural constraints are demanding in terms of computing power. The application ARIA solid-state NMR is dealing with the automated assignment and structure calculation of large protein
systems in ssNMR context. We are currently adapting and evaluating that structural-biology application to grid computing within the EGEE infrastructure, and to cloud computing infrastructure using the HIPCAL and Eucalyptus systems.

Detailed analysis:

The solid state NMR (ssNMR) is a valuable alternative to X-ray crystallography and solution NMR for structural analysis of insoluble proteins. Iterative assignment methods like ARIA (Ambiguous Restraints for Iterative Assignment), based on successive simulated annealing (SA) procedures, reliably assign NMR cross-peaks and calculate protein structures by using ambiguous distance restraints derived from the NMR cross-peaks intensities. The application of ARIA to ssNMR spectra is demanding in term of computing resources. Indeed, because of the low spectral resolution, a large number of possible assignments has to be processed. Also, it is difficult to derive precise distances from the peak intensities and that impairs the convergence. In order to face the higher complexity due to assignment ambiguity and the poor intensity quantification, several parameters have to be substantially increased, as the number of integration steps in the SA procedure, the number of protein conformations generated, and the number of possible assignments explored. These increases rely on a more intensive use of molecular dynamics softwares performing the SA, as CNS (Crystallography & NMR System).

Conclusions and Future Work:

Solid-state NMR (ssNMR) spectroscopy can address questions on structure, dynamics and interactions of insoluble proteins. The large CPU resources of the grid computing allows a straightforward parallelisation of the conformer generations at each step of the iterative procedure. Adapting ARIA to grid computing infrastructure will thus help structural biologists to process highly complex NMR experiments. Future work includes the porting and integration of ARIA on both grid and cloud infrastructures, in collaboration between CNRS IBCP and Institut Pasteur Paris, and also through a collaboration with the eNMR EU project.

Impact:

Adapting the ARIA method to the grid computing infrastructure will help structural biologists to study and assign large biological systems, as membrane proteins and protein fibrils using ssNMR, with more efficiency and reliability. Indeed, with the help of the grid it will be possible to enhance the convergence capability and the assignment reliability of the software by substantially increasing the complexity through the number of integration steps in the SA procedure, the number of protein conformations generated, and the number of possible assignments explored. A typical iterative procedure requires 8 steps. Each step consists of the calculation of several conformations, ranging from 20 to 100 instances, using the CNS software. Between two steps, the ARIA software analyses the calculated
structures from the previous step and defines the new restraints used in the conformation calculations during the next step. Having access to a large scale grid platform such as EGEE opens the way to run in parallel several structure determination by several users, and/or increases the capabilities of structure and assignment procedures on large systems.

Keywords:

Applications porting; Structural Biology; Grid Computing

URL for further information:

http://aria.pasteur.fr

Poster session / 85

grid-CICADA – Efficient Explorer of Conformational Space of Biomolecular Systems

Authors: Jakub Stepan[1]; Jan Kmunicek[2]; Jaroslav Koca[1]; Ludek Matyska[2]; Petr Kulhanek[1]; Zdenek Kriz[1]

[1] *National Centre for Biochemical Research, Masaryk University, Kotlarska 2, CZ-61137 Brno, Czech Republic*

[2] *CESNET, z. s. p. o., Zikova 4, 160 00 Praha 6, Czech Republic*

Corresponding Author: kmunicek@ics.muni.cz

Knowledge of conformational space is essential in the assessment of dynamical behavior of biomolecules, which is important, for example, in drug design and folding studies. Unfortunately, systematic exploration of conformational space is impossible due to its high complexity. To circumvent this problem, we have developed the program CICADA that tries to rationalize the search in such a space. In this work, we will demonstrate the utilization of a newversion of CICADA in the grid environment during the study of selected middle size biomolecules.

URL for further information:

http://ncbr.chemi.muni.cz; http://egee.cesnet.cz

Impact:

The main purpose of our contribution is to transfer the knowledge from local solutions into production utilization using available grid research area built around the EGEE infrastructure, and thus permitting scaling up the performed computations. The faster exploration of conformational space achieved by the Grid utilization forms an important step towards answering more challenging research questions. Moreover, our implementation serves as a show case of high-throughput type of computational chemistry calculations running within the worldwide distributed infrastructure, thus demonstrating its suitability for such kind of computations.

Detailed analysis:

The program grid-CICADA is composed of two main components: a) control unit and b) explorers. The control unit collects information about alreadythe searched conformational space. It optimizes found interconversion pathways among conformers and rationalizes the strategy of further space exploration. It also manages explorers, which perform a simple search in a limited conformational subspace. Since individual explorers are independent to each other they can be executed independently, which makes a basis for efficient utilization in the Grid. The program CICADA was used for the exploration of conformational space of single amino acids, small peptides (Leu- and Met-enkephalin), DNA and RNA fragments and also mono- and oligosaccharides. We would like to show the efficiency of grid-CICADA on Leu-enkephalin and on 19 amino acids long C-end of Caseine kinase I in native and hyperphosphorylated form.

Conclusions and Future Work:

Here we present an advanced way how to perform efficient exploration of conformational space of biomolecules within the worldwide EGEE/EGI environment. Our extension to the approach, already tested using single CPU and MPI-parallelized implementations, allows routine utilization within the worldwide Grid, thus boosting the research in the field of flexibility and folding of biomolecular systems.

Keywords:

CICADA, PES, conformational space, folding, grid, high-throughput computing, distributed computing

Earth Science / **86**

Grid implementation and application of the WRF-ARW prognostic model

Authors: Aleksandar Pajović[1]; Davor Davidovic[2]; Ibrahim Hadžismailović[3]; Maja Telišman-Prtenjak[4]

Co-authors: Danijel Belušić [4]; Karolj Skala [2]

[1] *Republic Hydrometeorological Institute, Banja Luka, Bosna and Herzegovina*

[2] *Ruđer Bošković Institute, Zagreb, Croatia*

[3] *Federal Hydrometeorological Institute, Sarajevo, Bosna and Herzegovina*

[4] *Faculty of Science, Dept. of Geophysics, Zagreb, Croatia*

Corresponding Author: davor.davidovic@cern.ch

In this work the implementation and deploying of the Advanced Weather Research and Forecasting (WRF-ARW) model on the SEEGRID-SCI grid infrastructure is presented. The goal of porting the model to the grid is to get more accurate and detailed forecast with operational speed-up on high-resolution model grids. The results and the application of the WRF model will be shown through the examples of the operational weather forecast, foehn effects over Banja Luka in Bosna and Hercegovina (BIH) and bora winds simulations over the Adriatic coast in Croatia.

Detailed analysis:

The numerical weather prediction models, like WRF-ARW model, demand large execution time and resource allocation; therefore we need large parallel computation facilities, like the grid systems, in order to execute them. The idea of deploying the WRF-ARW model to the grid is that all of the model data, binaries and results, are stored on the grid storage elements while the submission and data manipulation scripts are installed on the grid user interface node. In order to store huge amount of data on the grid, we have created the storage scheme to store files and folders required or produced by the model. The storage scheme is implemented in the LFC catalogue, so that the data retrieving is independent of the storage where the data physically resides. All grid execution is controlled by the scripts responsible for model data collection, generating jdl files, submitting jobs and results collecting. These scripts integrate the glite tools for job submission and output retrieving and hide the grid complexity from the end users as much as possible. The biggest issues are to avoid jobs remaining in the queuing system and the copying time for large model data.

Conclusions and Future Work:

We have achieved a great success in porting WRF-ARW model to the grid and improvements gained through deployment of the model increases the support and feasibility of new projects in the area of the SEE-GRID-SCI infrastructure and Earth Science community. Implementation of the WRF-ARW model on the grid infrastructure, increasing scientific research in regional scale particularly in the field of ecology. Some improvements in the model execution on the grid is to be done. In the plan are some new projects on this subject which will extend the model functionality and increase efficiency.

Impact:

We have achieved significant progress in accelerating the model execution and simplicity of starting jobs on the grid. The weather forecast is now more accurate due to the increased grid resolution, while the execution time is significantly smaller than before. In our examples we have found that the benefits are not only in the reduced execution time (up to 30% for large enough number of CPUs) but also in the simultaneous runs of the same model. This is especially important in research activities when many similar jobs can be executed simultaneously on different grid CE that results in enormous speedup. For research purposes, the WRF model was used to show the increase of air temperature due to the Foehn effect, and to study the reasons of the failure of the weather forecast for a local extreme situation, like a thunderstorm or hailstorm, which occurred in Banja Luka this summer without a clear weather forecast outputs. Also, the model was used for very accurate and detailed simulations of the bora wind over the north Adriatic. For the forecast purposes, the model is used for the daily weather prediction over BIH.

Keywords:

Grid, meteorology, WRF/ARW, foehn effect, bora winds, weather research, SEEGRID-SCI

URL for further information:

http://wiki.egee-see.org/index.php/WRF-ARW

Computer Science / 87

Towards an interoperable representation of the Grid Observatory information: an experiment with Common Base Event

Authors: Cecile Germain-Renaud[1]; Philippe Gauron[1]

[1] *LRI*

Corresponding Author: cecile.germain@lri.fr

Reaping the full benefit of the Grid Observatory (GO) initiative requires providing end-users with convenient representation of the traces. The present lack of standardization creates considerable difficulties for developing automated analysis and situation handling solutions. We report an experiment on the representation of internal logs of the gLite WMS with the IBM Common Base Event (CBE) format, categorizing the information of hundred of thousand events to a few generic type situations, with potential applications of the open-source supporting technologies to job monitoring.

Detailed analysis:

The CBE XML schema defines the format of an event, which at its core is a 3-tuple (reporting component, impacted component, situation). Component may be hardware of software. Twelve generic situations are available, e.g. Start, Stop, and the very important Feature and Dependency describing component availability. As a test case, we considered the CondorG logs of GRIF-LAL site from 2008-09-16 to 2009-03-24, with 883,701 events amongst which 118,191 are associated to one or several identifiers. The main challenge is to type each logged event into a situation. The CondorG log format being un-documented, we developed a software suite for elucidating its syntax and to some extent semantics. The syntax is "class::fonction:message"; we identified 7 different classes, and 21 functions. We are currently in the process of segmenting the messages (finding and deleting url or process identifiers) in order to type them. The next step will be interpreting each type as a suitable CBE situation type. For instance, the event "ControllerLoop::run(): Aborting daemon..." will be classified as a CBE "StopSituation" with arguments "successDisposition=SUCCESSFUL" and "situationQualifier=ABORT INITIATED".

Conclusions and Future Work:

One year of GO experience shows that, despite documentation efforts, the simplest traces are by far the most used. The GO will exploit our work with CBE for the presentation of the traces broadening the scope of their utilization by the computer science community.
The a posteriori conversion of the operational logs is also an opportunity for interacting with system administrators about the possible usage of the advanced tools from IBM to improve the QoS to the end-users. These interactions will help assessing the interest of the technology for the future evolution of the EGI middleware.

Impact:

The second priority of the GO, after data collection, is to provide parsimonious and informative representations of the traces. A first step is converting the traces towards user-friendly formats, in other words taking the burden of the "80% preprocessing" of data mining. While some traces are natively organized along standards (e.g. the information system –IS - is Glue compliant), other ones use proprietary formats (e.g. the Logging and Bookkeeping –LB -), and are in the worst case fully undocumented (internal logs of the WMS). To maximize the added value of the costly conversion process, the target format should 1) be a (de-facto) standard and 2) come with an exploitation framework. We choose CBE, which as a format and associated technologies (automatic analysis engine, visualization tool) is the result of IBM's extensive experience with autonomic management. CBE is not suitable for all gLite logs, e.g. the IS, or in the WMS scope jobmap, are not event-oriented. However, CBE adequately covers many of them: in the WMS scope, wmproxy, jobcontroller, condorG, logmonitor and workload manager; outside, the LB. Our work will be a first step towards consolidating these disparate sources of information.

Keywords:

Monitoring, standardization, events

URL for further information:

www.grid-observatory.org

Poster session / 88

Stellar energy flux modeling under SYNTSPEC application

Authors: Gražina Tautvaišienė[1]; Šarūnas Mikolaitis[1]

[1] *ITPA VU*

Corresponding Author: sarunas@itpa.lt

We present the stellar energy flux modeling tool under the SYNTSPEC application. SYNTSPEC is the gridified tool for stellar spectra analysis. It is an example of a data- and compute-intensive application running on the testbed of the EU BalticGrid-II Project (http://www.balticgrid.org), which brings new quality to the research in astrophysics. The multi job application is run within the Gridcom system – the user friendly interface that allows a common (virtual) work of the physically spread scientific group.

Conclusions and Future Work:

The stellar flux and normalized spectra modeling in the GRID environment is a tool that increases the abilities of the SYNTSPEC application and creates a new powerful structure for astronomers. It is a good example ofa user friendly tool with a significant science potential. With the imminent start of the GAIA mission, the future development of the application is planned, but the main tasks are achieved. It is still important to make it more autonomous and time saving for scientists.

Detailed analysis:

The stellar energy flux modeling is very important for stellar analysis. The stellar interior isa dense and hot plasma environment where energy is produced by the fusion process. It irradiates a specific flux of energy, that is spread within all range of electromagnetic wavelengths. The photosphere of spectra is the part of star which modifies the energy flux distribution over wavelength range. The atomic and molecular structure of photosphere redistributes the initial energy flux through all the spectrum employing absorption, reemission, scattering processes and paints the unique shape of flux image of the specific star. The energy flux calculation under SYNTSPEC is a very good example of application that benefits from usage of the BalticGrid-II testbed because of the need of powerful computing resources. It calculates the energy flux normalized to the continuum stellar spectrum, that is applied for stellar classification and determinations of e.g. chemical compositions, effective temperatures and surface gravities of stars. The specific energy flux modeling is the important tool for analysis of data, which will be produced by the European Space Agency's GAIA space observatory (to be launched in 2011).

The gridification of the application was performed in the Institute of Theoretical Physics and Astronomy of Vilnius university. It is set of Fortran and C++ coded programs joined together by scripts. Parametric submission is performed by the GRIDCOM interface. Single job runs for 8 – 10 hours. About 50 jobs are required to derive the main parameters of the star and about 15 jobs for every other chemical element. In general it takes more then 400 jobs for one star. The output is not very huge, but a bigger amount of temporary store in every single node is required. Depending on the initial atomic database the program stores more then 100 GB of temporary data. The output follows VOTABLE standards to be compatible within the Virtual Observatory infrastructure.

Impact:

The SYNTSPEC application benefits from the usage of large project resources, which makes possible the calculation of synthetic stellar spectra for significant wavelength ranges which is essential for the

galactic and stellar research studies. The special added value is the implementation of the energy flux modeling, which makes the application ready for very specific data processing of data for GAIA and other modern observatories.
Running of the application within the Gridcom graphical interface on the BalticGrid- II Special Interest Groups site (http://sig.balticgrid.org) improved the job submission procedures and the possibilities of user interaction. The application can be submitted and analyzed by a dispersed group, which is especially important in remote teaching or analysis of data.

The BalticGrid-II project established a production-level infrastructure, interoperable and complementary with the EGEE grid. It provides the scientists in the Baltic States and Belarus access to critical resources, supports effective research collaborations and made possible the efficient sharing of unique instruments and data.

Keywords:

Stellar spectra, energy flux, stellar astrophysics, BalticGrid-II, Gridcom, SYNTSPEC

URL for further information:

https://sig.balticgrid.org/SIGs/stellar-spectra/

Poster session / **89**

Performance evaluation of the Estimated Response Time strategy: tools, methods and an experiment

Authors: Alain Cady[1]; Cecile Germain-Renaud[2]

[1] *LAL*

[2] *LRI*

Corresponding Author: cecile.germain@lri.fr

An extensive body of research focuses on economic and intelligent scheduling models. Conversely, the gLite matchmaking engine adopts an agnostic approach for estimating the waiting time of incoming jobs, derived from the Copernican principle: “this job is not special”. An open question is the evaluation of this minimalist strategy. This work reports on the creation of the software tools required for this evaluation, the methodology and criteria, and presents preliminary results using the Grid Observatory logs of the Information System (IS) and PBS scheduler of GRIF/LAL.

Detailed analysis:

gLite maps jobs to Computing Elements (CEs) by considering the categorical requests and breaking ties along the Estimated Response Time (ERT), an estimation of the queuing delay published by the CEs in the IS. Evaluating the accuracy of the ERT computation requires comparing it along and extended period with the actual response time (ART). The ART is readily available from the logs of the local schedulers. We created a logging system of the IS for the Grid Observatory (GO). To cope with the massive redundancy, IS logs are stored in compressed diff format, with a reference snapshot each day, and diff at 15 mn period. For the evaluation, Perl scripts use regular expressions to extract ERT from the patched (reciprocal of diff) IS logs, and ART from the scheduler logs. Next, we use the Object capability of Matlab to create an efficient data representation (CE, VoView, Site, with matched ART and ERT), with flexible functionalities (subsampling, elementary or advanced statistics, plotting, outliers management). A first usage of this representation is to exhibit rare but repetitive abnormal patterns where relatively long-lasting spikes of load are correlated with low ERT.

Conclusions and Future Work:

The evaluation criteria should integrate both the statistics about the reliability of the estimation (BQP concept), and a parameter measuring the operational impact of the prediction. We are exploring a description by the receiver operating characteristic (ROC) curve that segments the ERT along the load regimes. In this work, the analysis is aggregated over the whole available history. Future work will

combine the segmentation method developed in the GO and this description in order to refine the performance evaluations along time-differentiated activity regimes.

Impact:

The process described so far is not fully satisfactory:1) the fixed frequency of IS logs is not adapted to the varying job arrival rate, 2) the IS ERT might not be the value actually used by the WMS, and 3) nor of the best one that the CE could provide. Recording the used ERT in the LB would provide much more accurate information with respect to 1) and 2). With the information we had, we decided to estimate the ERT of a job by the last published ERT for its target CE. Thus, we get irregular time-series, with a (ERT, ART) pair at each job arrival. The tool allows easy search for correlation at different lags. We observed a strongly differentiated behavior between the two most heavily loaded queues: Atlas has a nearly flat correlation landscape, while Biomed shows a peak at lag 20. The main issue for evaluation lies in defining the objective function. The Batch Queue Predictor (BQP) initiative accurately pointed that synthetic indicators e.g. Mean Squared Error, or the correlation coefficient, are misleading: for instance, the (accurate) estimation of a null ERT for an empty CE is not very informative, and should not mask the above-mentioned spikes.

Keywords:

scheduling, estimation

URL for further information:

www.grid-observatory.org

Poster session / 90

Visualizing the dynamics of e-science social networks

Authors: Alain Cady[1]; Asterios Katsifodimos[2]; Cecile Germain-Renaud[3]

[1] *LAL*

[2] *INRIA*

[3] *LRI*

Corresponding Author: cecile.germain@lri.fr

Computational grids provide new natural examples of large-scale complex networks emerging from collective behavior. An interesting question is thus whether these networks exhibit properties similar to those of other social networks (SN), or original ones, which would be the specific signature of e-science. As a first step towards answering these questions, we build a scalable interactive visual exploration tool targeted at the spatiotemporal structure of data-access and data-sharing. The primary goal is to help computer science researchers getting intuition about the dynamics of the system.

Detailed analysis:

This work exploits the gsiftp and gridftp logs for providing six relevant fields: identifiers for the user and requesting machine, direction of the request (get or put), file name, transfer size, outcome (success or failure) and timestamps. As such, the planned tool will be able to represent the output of any monitoring providing similar information. Two graphs capture the data-user relationship: the bi-partite data access graph, with users and file as nodes, and edges connect a user with the file she uses; the data-sharing graph [1], in which nodes are users and edges connect users with similar interests in data. As in most SNs, our data require true multi-scale visualization, but with a strong requirement for smoothness and continuity: for instance, the files popularity is much more heavy tailed than a classical Zipf's law. Moreover, useful attributes can be extracted from the file names (e.g. output error files versus data files), contributing to the need for interactive selection of the representation. Finally, the system should allow easy manipulation of the time frame and create representations with the appropriate level of detail.

Conclusions and Future Work:

Resource sharing is the specificity of Grids. Identifying and characterizing the sharing patterns amongst data and users gives a partial view of the more elusive e-science network, which is ultimately about the research the users are actually conducting with that data. Such characterization is also a requirement for addressing both long-term dimensioning and short-term allocation of resources. Building on state of the art research, we propose to contribute to these goals through information visualization. The planned tool exploits only generic information and will thus be of general usage.

[1] A. Iamnitchi, M. Ripeanu and I. Foster. Small-World File-Sharing Communities., Infocom 2004, Hong Kong, March 2004
[2] A. Aris, B. Shneiderman. Designing semantic substrates for visual network exploration. Information Visualization 6, 4 (2007), 281–300.
[3] N. Henry and J-D Fekete. MatrixExplorer: a dual-representation system to explore social networks. IEEE Transactions on Visualization and Computer Graphics 12, 5 (2006), 677–684.

Impact:

The Grid Observatory experience concurs to the general urgent requirement for information visualization. The challenges of visualization for SN exploratory analysis have been recently formalized from users' requirements [2,3]. One of the most important is that scientists do not want just to draw a graph, but 1) to build several representations according to the different attributes and 2) to be able to compare and identify a consensus among actor clusters across representations. In order to build a state of the art framework, this work is realized as collaboration between the INRIA AVIZ group and the Grid Observatory, with support from the Paris-Sud University. The AVIZ group has developed the GraphDice [4] environment, a multivariate network visualization system for exploring the attribute space of edges and actors. This environment has proved to be both user-friendly and information rich. Moreover, the capacity to deal with a high number of attributes is essential for future integration with the other source of information provided by the Grid Observatory, for instance the actor's VO, their usage of computing resources and the failure events.

Keywords:

interactive visualization, social networks

URL for further information:

www.grid-observatory.org

Earth Science / 91

Bridging the gap between applications geospatial data and the Grid

Authors: Alberto Proenca[1]; António Esteves[1]; António Pina[1]; Bruno Oliveira[1]; Joel Puga[1]

[1] *Universidade do Minho*

Corresponding Author: pina@di.uminho.pt

CROSS-Fire is a Portuguese NGI funded project aiming to develop a grid-based risk management decision support system for the Civil Protection (CP), using forest fires as the main case study and FireStation (FS) as an application that simulates the fire spread over complex topography.
CROSS-Fire uses EGEE to provide raw technological capability provision, including data management and storage, access to meta-data data bases and HPC and a Geospatial Information Infrastructure based on OCG-SWE Web Services to provide the access and management of remote geospatial data and meteorological data.

Detailed analysis:

The CROSS-Fire overall software approach take advantage of the adoption of OGC-WS standards proposal and INSPIRE directives on the development of : i) a standard-based Spatial Data Infrastructure layer, based on Geoserver to exploit and enable geospatial services for data access and processing ii)

a 52N's implementation of a OGC-SWE compatible layer, to address sensors CP data sources, such as meteorological stations data and satellite images and iii) a graphical user interface to access the platform facilities.

The core of the CROSS-Fire is a WPS 52North OGC standard layer divided into three interoperable components: the Application Business Logic, the Grid Services and Geospatial Services.
The WPS hides from the user the complexity of the access to a wide range of distributed computing resources, by providing: i) the mechanism to access the grid facilities for processing and data management, ii) all the algorithms, calculations, and execution models that operate on spatially referenced data, iii) an interface to the SOS server that provides dynamic meteorological data and iv) the services requested from the portal and other GUI clients.

Conclusions and Future Work:

Current work is centred on the design and implementation of a web portal where many players can connect, to request services through the WPS core layer.
The portal will allow to update the input data required to estimate the risk of the natural hazard, or to access the past simulations to validate the predictions with actual field data.
To provide FS with the dynamic data, we are using the 52º North implementation of the OGC-SWE layer.The dynamic data includes information coming from sensors in weather stations (such as DAVIS Vantage Pro2) and sensors MODIS from satellites such as Terra/Aqua.

Impact:

Cross-Fire provides a general approach for the development of a Civil Protection application that requires not only the availability of high-performance computing resources and data management at remote Grid sites, but also the ability to access, to integrate, to analyze and present geospatial, available data repositories and sensor networks data across a distributed computing environment.

The approach allows different components to interact with each other in a standardised mode, acting each one as clients of the OGC WPS web services layer.
The use of WPS offering a modular addition of new facilities, algorithms and services to clients and users of an application as long as the requests conforms to OCG-SWE and EU/INSPIRE standards.
The platform may also be used to support new brands of CP applications, such as Flash floods, implemented as a new Application Business Logic component that responds to the specificities of the new application domain.

Keywords:

civil protection, geospatial services, interoperability, forest fires, OCG

URL for further information:

http://pop.cp.di.uminho.pt/crossfire/

Parallel Processing / **92**

MPI support and improvements in gLite/EGEE - a report from the MPI Task Force

Authors: Isabel Campos Plasencia[1]; John Walsh[2]

Co-authors: Brian Coghlan [3]; Dennis van Dok [4]; Enol Fernández del Castillo [1]; Jeroen Engelberts [5]

[1] *IFCA*

[2] *Trinity College Dublin/Grid-Ireland*

[3] *Trinity College Dublin*

[4] *NIKHEF*

[5] *SARA*

The MPI Task Force (MPI-TF) was established in response to calls from the EGEE user communities and site administrators during the MPI session at EGEE09. Improvements in the overall deployment

and support of MPI on the EGEE infrastructure were demanded. The objectives of the MPI-TF are to provide solutions which will improve the testing, installation and configuration of MPI-enabled middleware, and to support greater MPI application usage on the Grid. We report on the work performed by the MPI-TF and report on the changes and improvements in MPI support delivered since its establishment.

Detailed analysis:

The reported high failure rates for MPI jobs on the EGEE grid is a source of frustration for many of its users. At the EGEE09 conference, the Earth Science community reported that only 7 sites from 26 supporting both MPI and its virtual organisation ran their applications without errors . Computational Chemistry reported that only 53% of jobs ran successfully at the CompChem enabled sites when requiring 8 processors, and this dropped to 21% when requesting 16 to 64 processors. Similar issues were reported by the Astronomy and Astrophysics community.
The site administrators reported on problems associated with installing and configuring MPI at their institutes.
This was backed up with a report from the MPI Working Group (MPI-WG), who gathered extensive feedback from a large sample of users and site administrators.
The provision of a single MPI solution for gLite is non-trivial. This is in part due to lack of good solid support from upstream OS providers, and in part due to the diverse site system setups, such as batch systems, multi-core worker nodes, and shared versus non-shared filesystems. We present a pragmatic approach to systems support.

Conclusions and Future Work:

The MPI-WG delivered a set of recommendations to improve MPI support on EGEE based on user and site administrator feedback. The MPI-TF was established to deliver many of the short term improvements which could be dealt with within the lifetime of EGEE-III. However, some issues cannot be easily concluded within this timeframe. Firstly, full interoperability with other grid infrastructures and middlewares such as ARC and UNICORE. Secondly, the emergence of multi-core and hybrid systems needs to be evaluated - this require a generic “Parallel” job job type. These will be examined in the EGI era

Impact:

In December 2009, over 100 sites were publishing support for MPI on the EGEE infrastructure. Some 50 of these sites pass the MPI validations tests; these probe the range of basic MPI distributions available at each site. Most sites successfully support at least one version of MPI correctly. The introduction of critical SAM validation tests is expected to help isolate problems at the sites and contribute to the development of a “trouble-shooting” guide. Updated user documentation covering numerous use cases and recipes will help users and developers avoid unforeseen issues and in turn help get the best out of the infrastructure. Similarly, much needed updates to the MPI/gLiite related middleware will accommodate the diverse systems setups currently support for non-MPI job execution.

The MPI-TF consists of representatives from multiple EGEE-III activities, including JRA1, NA4, SA1 and SA3, with oversight from the EGEE TMB. We shall provide an evaluation of the impact of the MPI-TF by measuring improvements to the number of sites actively supporting MPI, sites successfully passing MPI SAM tests, and by seeking feedback from the user community through their virtual organisations.

Keywords:

MPI, User Support

URL for further information:

https://twiki.cern.ch/twiki/bin/view/EGEE/MpiTools

Computational Chemistry / **93**

Ab initio grid chemical software ports – transferring knowledge from EGEE to Polish NGI

Author: Mariusz Sterzel[1]

[1] *ACC Cyfronet AGH*

Corresponding Author: mariusz.sterzel@cern.ch

Chemical software, especially with ab initio methods, have been developed over several years of research in the field of numerical methods in chemistry. As the research was performed by many groups of scientists this resulted in a variety of software suites. From this set the commercial packages are of particular interest among the community due to the availability of many computational methods, the fast development of new ones and better user support. To attractthe chemical community to grid computing a set of frequently used software suites have been ported to the grid.

Conclusions and Future Work:

The present solution has been applied only to serial versions of chemical codes . Future work will mainly focus on the parallel execution of program suites on the PL-Grid infrastructure, as the scientific problems chemists are dealing with are involving larger and larger molecules. Taking into consideration EGEE achievements in that field we expect this task to be done with similar ease. Further plans beyond this include the integration of all grid software ports within the Virtual Laboratory, to allow the execution of complicated chemical experiments by non-specialists.

Detailed analysis:

The majority of most commonly used chemical software has been ported to the EGEE Grid during the second and third phases of EGEE projects. The most spectacular achievements include the grid ports of commercial packages like Gaussian, Turbomole and Wien2k (both serial and parallel versions), the most popular among chemical community. The major difficulty with grid ports of commercial packages is related to their strict license requirements. Those include not only a code protection against unauthorized use but, more importantly, the check of usage patterns compliance. Unfortunately the solutions worked out do not allow simultaneous use of several commercial packages during one job run, as in most cases commercial packages belong to different VOs. The model applied in PL-Grid is different. In contrast to the EGEE solution all the packages have been ported to a single VO. Such a model allows not only the simultaneous execution of several instances of chemical packages at the same time, but it also eases the execution of complicated experiments where the next step of an experiment depends on the result(s) obtained previously.

Impact:

The EGEE infrastructure together with the grid ports of chemical software has proven its usefulness in research in the field of computational chemistry. Direct availability of easy to use chemical packages increased Grid attractiveness, ad resulted in better grid utilization by the community and enhanced user satisfaction, making the chemical communitythe second largest Grid user. The same results are expected for Pl-Grid and others National Grid Initiatives (NGI). It is important to point out that only a very little additional effort related to porting procedures was required, as the main obstacles for porting have been solved by the members of Computational Chemistry Scientific Discipline Cluster. On the same basis other NGI's can take advantage of EGEE acquired knowledge and expertise in several scientific areas.

Keywords:

ab initio, chemical software, porting, NGI, PL-Grid

URL for further information:

www.plgrid.pl (NGI www page)

Demo Session 1, Welcome Drink / 94

Grid Web Portal with software plug-ins to chemical software

Author: Mariusz Sterzel[1]

Co-authors: Daniel Harezlak [1]; Tomasz Szepieniec [1]

[1] *ACC Cyfronet AGH*

Corresponding Author: mariusz.sterzel@cern.ch

New users, especially those accustomed to Graphic User Interfaces (GUI), often face difficulties during the adoption of grid computing for the research. The main source of issues is related to command line interfaces which are hard to adopt by non-experts. To avoid users' disappointment a new web based interface has been proposed. In contrast to existing web-based tools the portal developed facilitates not only job management but primarily can serve as a work environment for chemists. Here, we present the current status of the portal development including several plug-ins to first principles chemistry packages.

Detailed analysis:

Existing web portals focus mainly on tasks, which simplify job management on the Grid. Although this is very important, this part of portal functionality is not sufficient for scientists accustomed to GUI environments supporting their research conduction. Our solution, in contrast to other existing web portals, is "user centric" instead of "grid centric". Therefore, it provides tools facilitating research conduction and "hides" the existing grid infrastructure from the user as much as possible. The tools adopted in the portal, based on ViroLab Project technology together with Google Web Toolkit library, allow for easy planning, development and execution of complicated computational experiments without being distracted by the grid technology. Currently users can execute and modify the existing experiments allowing execution of Gaussian, Turbomole and GAMESS US packages including a variety of computational chemistry methods. Other tools integrated in the portal allow an analysis of the results of computational experiment(s) (energy for example) and their visualization. Particular experiment results can be displayed via external software.

Conclusions and Future Work:

The Grid web portal for chemists has been developed to fulfill community needs especially coming from users newly adapting to the Grid. Although the number of software packages supported by the portal is still limited, the collection of computational methods provided by the above software suites allows studies of a variety of molecular properties at various levels of accuracy. Nevertheless, our future work will focus on the adoption of other packages to extend the available portfolio, as well asthe integration of the portal with national grid infrastructures for the benefit of their users.

Impact:

The development of grid web portal for chemists is a next step towards better community satisfaction. As it focuses on computational experiments in chemistry rather then job management only, we expect high interest in it, coming mainly from new users adopting the EGEE Grid. The availability of the portal will not only avoid new users' disappointment but what is more important will drastically shorten the time needed for the adoption of grid computing to particular computational problems in chemistry. Although the current version is limited to the several first principles codes only, most of the important chemistry tasks like chemical reaction paths, potential energy surfaces, or even ab initio molecular dynamics studies, can be easily performed onthe EGEE infrastructure.

Keywords:

web portal, software plug-ins, computational chemistry software

URL for further information:

Under construction. Probable address: http://chempo.grid.cyfronet.pl/chempo

Justification for delivering demo and/or technical requirements (for demos):

A demo offers possibility to reach much broader audience than during ordinary session. No special requirements is required – a tv screen and fast internet connection for live demonstration

Poster session / 95

Trading Computing Resources across gLite and XtreemOS Platforms

Author: Pushpinder Chouhan[1]

Co-authors: Alvaro Arenas [1]; Andrew McDermott [1]; Nick Trigg [1]

[1] *STFC*

Corresponding Author: pushpinder.chouhan@stfc.ac.uk

The objective of this work is to present a system that facilitates the commercialization of Grid resources through a Virtual Marketplace of computational resources, where a seller is capable of listing the Grid resources, and buyer can request/bid dynamically for required computing resources for their applications.
This model exploits the benefits of Grid computing, especially the inter-operability and scalability of Grid platforms. Interoperability is achieved by using the OGF SAGA standard on two Grid platforms: XtreemOS, a Linux-based Grid OS, and gLite, a Grid middleware.

Detailed analysis:

Our trading system is a virtual marketplace of resources that comprises three layers:
The user-application layer, the top layer, allows end-users to submit applications to the deployed resources after analysing economical and performance related information.
The virtual market place layer is the middle layer that implements trading services anddynamically schedule applications on Grid resources. Scheduling depends on criteria such as cost, processing power, execution time or resource availability.
The resource management layer is the bottom layer that manages the Grid resources used by the end users.
Our implementation leverages on existing technologies such as trading algorithms, grid middlewares/OS and API for Grid applications.
Monitoring services inform the trading service about cost and performance of Grid resources. Monitoring is achieved by either direct or indirect capture of resource status and pre-defined events. The indirect interface uses logs generated at run-time by the Grid infrastructure. The direct interface is a portal collecting dynamically events generated by monitoring services associated to the Grid infrastructure.

Conclusions and Future Work:

Our work answers questions such as “which Grid should be used that will minimize cost along with achieving efficient applications’ execution time?”, “how end-user can select Grid resources according to pre-defined policies, including cost policies?” and “how to achieve interoperability when using gLite and XtreemOS platforms?”.
Our trading system provides a portal for end users to avail the computing power of Grid resources, depending on economical and performance parameters. In future, we plan to extend our system considering ecological parameters such as power consumption.

Impact:

The main benefit of this system is to integrate in a single framework four key features: cost saving for end-user, dynamic scheduling, interoperability, and elasticity.
Cost saving for end users is guaranteed by allocating the applications to the more economical resource(s), following policies defined by the end users.

Dynamic scheduling is achieved through the virtual marketplace, which implements scheduling algorithms that allocates applications following classical performance parameter as well as the cost of resource usage.

Interoperability is achieved by using a standard programmable interface, the Simple API for Grid Application (SAGA), to bridge the gap between existing grid middleware and application level needs. The same system could run on XtreemOS or gLite Grids, or interoperate on Grids using resources from both platforms.

Elasticity is provided within the Grid as our system enables adding and deleting resources dynamically. In addition, resource prices can change on the fly, and such changes are taken into account by the virtual marketplace.

Keywords:

Grid Middleware, Virtual Resource Market Place, Grid Applications, Workload Management

URL for further information:

http://www.xtreemos.eu/

Poster session / 96

Monitoring CREAM Jobs in RTM via L&B

Authors: Aleš Křenek[1]; František Dvořák[1]; Janusz Martyniak[2]; Jiří Filipovič[1]; Jiří Sitera[1]; Luigi Zangrando[3]; Massimo Sgaravatto[3]; Miloš Mulač[1]

[1] *CESNET*

[2] *Imperial College*

[3] *INFN*

Corresponding Author: ljocha@ics.muni.cz

The Real Time Monitor (RTM) is a high-level monitoring tool which aggregates information on grid jobs and presents it in a suitable form. The success of this tool depends on the accuracy of information it is able to receive from lower layers. However, withthe recent increase of the number of jobs submitted directly to Computing Elements (CE) the fraction of jobs seen by RTM decreases. gLite Logging and Bookkeeping service (L&B), with its recent extension to CREAM jobs, provides the necessary level of abstraction and a glue between different job flavours to allow high-level tools to see full range of grid jobs.

Detailed analysis:

Recently we introduced extensions to the CREAM Computing Element and the L&B service, which unify the view on jobs executed by CREAM regardless of their submission path (via gLite WMS or directly to CREAM). CREAM is able to distinguish direct submission; in this case the incoming job is registered with L&B first. In both scenarios CREAM logs events on progress of job execution as well as possible failures to L&B. This information is processed at L&B into a view on overall job state which is consistent between WMS and CREAM-only jobs.

On the other hand, the RTM was modified to receive notifications on job state changes from L&B, rather than extracting job state information from raw data in L&B database as before. Besides improving overall reliability, this binding is done on the level of L&B job state which is already common to both WMS and CREAM-only jobs. Therefore the RTM needn't make any further distinction between different job types.

Conclusions and Future Work:

The described work is mostly integration. With a relatively small effort we were able to put recently developed pieces of code together to bring additional considerable benefits. Besides the practical desirable results of a more accurate RTM view on the EGEE grid, this is a positive improvement on the overall architecture, with L&B as the glue monitoring service between various job types and the massively exposed high-level tools. Therefore in the near future we will concentrate on hardening the existing prototype towards production quality, and on its wide-scale deployment on the infrastructure.

Impact:

Some of the grid users prefer using their workload management systems (e.g. Atlas Panda) bypassing gLite WMS and submitting jobs directly to CEs. The amount of workload distributed to grid sites in this way is not negligible.

Our work, by unifying the view on all grid jobs (going through WMS or directly to CE) at the level of L&B, enables the uniform monitoring of al lgrid jobs with high-level tools like RTM . Consequently, the

real-time view on the grid state is considerably improved. Additional benefit is extending the time span of CE job data (CREAM purges them soon after job completion), enabling better post-mortem analysis of problems etc.

The RTM is to date the only system which has access to distributed L&B servers worldwide. This makes in an important tool not only for dissemination purposes and individual users, but also for large experimental communities delivering a single monitoring entry point. Data collected by the RTM may also be analysed off-line providing an opportunity to study performance of the GRID in greater detail. Adding direct submissions to the RTM monitoring system will make this tool more attractive also to communities not using WMS resources in their work.

Keywords:

CREAM, L&B, RTM

URL for further information:

http://egee.cesnet.cz/cms/export/sites/egee/cs/info/{UF4-poster-rtm.pdf,CREAM-poster.pdf}

Regional Activities, International Projects and Collaborations / **97**

Improvements of the grid infrastructure and services within SEE-GRID

Author: Misev Anastas[1]

[1] *Faculty of Natural Science and Mathematics*

Corresponding Author: anastas@ii.edu.mk

Enlargement of the grid user community is probably the most important challenge of the grid infrastructure development community. To attract new users, with new applications and needs, the infrastructure should be flexible enough to satisfy their needs and requirements. We present the experience of the development community from the SEE-GRID-SCI project, organized in a Joint Research Activity. The main focus if this JRA was to improve the usability of the infrastructure and Grid services for the end-users from target communities, as well as manageability of the underlying infrastructure.

Detailed analysis:

Based on the analysis of the user's needs, taken from the 3 main SEE-GRID user communities, with special emphasis on the commonalities between them, several applications services and operational tool were successfully implemented. Also, some of the preexisting operational tools were enhanced with application-specific extensions, enabling their deeper integration with the current and future grid applications.
The Operational Tools developed or enhanced in this framework were mostly oriented toward better monitoring(BBmSAMeX), alerting (AMS) and job tracking and analysis (JTS, Logwatch-G), but also security (GSSVA), ticketing (NMTT) and portal enhancements (USGIME).
The Application Services aim was to support the current and future applications with a richer set of services, needed for more complex and robust grid applications. They ranged from the advanced usage of workflows (AWT, CWRE), better data and file management (DM-Web, FM-J-API), applications platforms and services for specific areas of grid usage (ESIP, MEWS, RAS, SDSASS), event logging (Event Logger) and advanced, user based monitoring and analysis (UMON).

Conclusions and Future Work:

Further development in this (JRA1) and other future frameworks will enable strengthening of the infrastructure in the SEE region and stronger involvement of the scientist from this regions in the European scientific and research activities

Impact:

The newly developed tools and services are already helping the SEE-GRID users and site administrators to achieve more productivity, to have more robust and stable sites, to have higher job throughput, to identify and categorize the underlying resources and to have overall higher grid utilisation. Using the applications services, more complex applications in all 3 user communities are already running. With the help of the new tools, better monitoring and overall understanding and utilization of the platform is achieved.

Keywords:

grid, tools, services, monitoring, alerting, workflows, portals, security, data management

URL for further information:

www.see-grid-sci.eu

Poster session / 98

A new middleware component for job monitoring in interactive mode: The SCoPE Experience

Author: Silvio Pardi[1]

Co-authors: Francesco Palmieri [2]; Guido Russo [2]; Jonathan Cacace [2]; Leonardo Merola [2]

[1] *INFN - Naples Unit*

[2] *University of Naples Federico II*

Corresponding Author: spardi@na.infn.it

This work aims to extend the standard gLite middleware with a new functionality that enables the final users to interact with their submitted grid-jobs, for monitoring purposes.
The new service is composed by a worker nodes Proxy server, with web-services interface VOMS compliant, easy to deploy and not invasive to respect the standard gLite solution. Through the Proxy, a generic user can connect, from the user interface, to the Worker Node that is taken his job. The tool has been implemented on the SCoPE infrastructure, to support the e-science community of the University of Naples Federico II

Detailed analysis:

The application is based on a client-server infrastructure, with the core installed on the Proxy machine that works as server. The Proxy is composed of a web service interface using the https protocol, that pilots a set of local applications, responsible to find information about the user-jobs and to open and close GSI secure connections from the User Interface to the Worker Nodes.
The interaction with the user is done through a CLI, that take in input the Grid-Job ID of the process that we want to control. The client contacts the web-service via https and sends the JobID plus the use proxy certificate. The server verifies the user credential and the VOMS extensions, then interacts with the CE, through the gsiftp protocol, in order to understand which WN hosts the user-job and to find other information about the job. Finally the Proxy server opens a GSISSH secure connection with the User Interface and gives a shell on the WN that runs the user-job with the right local account, in the execution folder. The shell is closed by the Proxy as soon as the job will finish.
The adoption of this software, does not involve any alterations of the standard glite roles (i.e. CE or WN).

Conclusions and Future Work:

The interaction with the job during the running is a very important activity during the software preparation, job tuning and pre-production phase, and is a requirement of many scientific communities involved in computational science. In this work we proposed a solution, secure, simple to deploy and to use, and transparent with a minimum impact on the standard gLite middleware. Currently our implementation can work on the PBS queue system, and in the next months we plan to optimize the use of this application for parallel jobs, and also to port the software on different batch systems.

Impact:

The impact of our work is very interesting in term of simplification of the interaction with the Grid, from the user point of view.
The proposed application, allows the final user to bypass the standard restrictions of a batch system, by using the standard gLite interaction. The user can control the effective state of his job and is enabled to control in realtime the job outputs, the memory occupation and many other parameters
The interactive monitor gives a very important instrument to the final user during the software development, by allowing to follow the flow of the jobs, and to do troubleshooting with a strong impact in term of time saving during the application tuning.
Finally the simplicity of the system allows the user to interact with his job in a transparent way, without any additional concerns on the use of computer grid.

Keywords:

Interactive session, Grid Shell, monitor

URL for further information:

http://people.na.infn.it/spardi

Regional Activities, International Projects and Collaborations / 99

EUAsiaGrid – e-Infrastructure Platform for Asian-Pacific Region

Authors: Alex Voss[1]; Jan Kmunicek[2]; Ludek Matyska[2]; Marco Paganoni[3]

[1] *School of Computer Science, University of St Andrews, United Kingdom*

[2] *CESNET, z. s. p. o., Zikova 4, 160 00 Praha 6, Czech Republic*

[3] *University of Milano-Bicocca and INFN, Piazza della Scienza 3, 20126 Milano, Italy*

Corresponding Author: kmunicek@ics.muni.cz

EUAsiaGrid is the EU-funded project aiming to bring together researchers from Asia-Pacific region around the distributed infrastructure being built within in the region. The primary purpose of the project is to support researchers willing to perform scientific discovery though advanced computing. In addition, the project is expected to further facilitate the uptake of e-Infrastructure approach.

Detailed analysis:

Internally, the EUAsiaGrid project is divided into several work packages responsible for distinct aspects of the overall tasks (dissemination, training, application support). A set of application domains as disaster mitigation, social science, computational chemistry, biomedicine and bioinformatics and high energy physics has been recognized from the beginning of the project thus forming the critical mass of end users who could profit from the EUAsiaGrid. As the regional researchers represent the key element gluing together the fabric infrastructure, underlying services and targeted applications within EUAsiaGrid project, the main focus is to attract all possible users through active search for new application areas that could benefit from the run in distributed environment.

Conclusions and Future Work:

Here we present the current state of EUAsiaGrid. Based on the 18 months of existence of the project the current achievements will be presented and discussed. Moreover, the planned advances in the field of sustainability and further globalization of cyber infrastructure will be demonstrated.

Impact:

The main purpose of our contribution is to inform about current advances concerning building, maintenance and continuous development of e-Infrastructure within the Asia-Pacific region. To fully utilize

the momentum obtained through EUAsiaGrid project, the current activities are pointed towards a transformation into a completely sustainable and persistent Asia-Pacific e-Infrastructure, fully integrated into the global grid environment. In short term view, a EUAsiaGrid Roadmap document describing the policies and organizational steps necessary for seamless formation of uniform, multinational digital platform is being formulated. The long term sustainability of the grid infrastructure in the region will be further supported and achieved through the incorporation of EUAsiaGrid into emerging European global grid infrastructure.

Keywords:

EUAsiaGrid, grid, distributed computing; roadmap

URL for further information:

http://www.euasiagrid.org/

Regional Activities, International Projects and Collaborations / 100

Developing e-Infrastructure services for e-Science applications: the EELA-2 experience

Authors: Diego Scardaci[1]; Francisco Brasileiro[2]; Roberto Barbera[1]

[1] *INFN Catania*

[2] *Universidade Federal de Campina Grande, Brazil*

Corresponding Author: diego.scardaci@ct.infn.it

This work aims at presenting the results obtained in the joint research activity of the EELA-2 project. In order to satisfy all heterogeneous applications requirements and to simplify the access to the Grid infrastructure, in the context of this activity, a set of special services has been developed to enhance the functionality of the gLite middleware and provide users with a richer platform. The middleware services developed have widened the number of potential applications taking benefit of the grid e-infrastructure and, moreover, have speeded up the porting process.

Detailed analysis:

To meet all heterogeneous requirements of the applications involved in the EELA-2 project, a set of special services has been designed and developed. These services enhance the functionality of the gLite middleware giving the EELA-2 users access to a richer middleware, reducing the amount of application development, and generally accelerating the adoption of Grid technologies. The identification and the development of these services has been conducted with the goal of increasing the reach and the usability of e-Infrastructure by assembling/reengineering existing technologies and developing new ones that facilitates the installation, management and use of the grid infrastructure. By increased reach, we mean more sites belonging to the grid infrastructure, with more resources being shared, more users being served and a more diverse range of applications being supported. By increased usability we mean the addition of services that can ease the tasks of developing and deploying new applications, as well as installing, managing and deploying the core infrastructure.

Conclusions and Future Work:

Several sites and applications involved in EELA-2 benefited of the additional services developed by the joint research activity. The usability improvement allowed more applications to be supported and attracted more users, increasing either the number of resources available or the institutions joining the grid. Considered the success of this methodology, we will continue to apply it, developing new ad-hoc services, to attract more e-Science applications and users to the EELA-2 e-Infrastructure.

Impact:

The following infrastructure-oriented services have been developed to provide alternatives to ease the installation, management and use of the e-Infrastructure: (1) a gateway between gLite and Ourgrid, a simpler peer-to-peer technology; (2) the porting of the gLite User Interface and Computing Element

to the Microsoft Windows platform; (3) the Storage Accounting for Grid Environments, a system to measure the usage of gLite storage resources. The following application-oriented services have been developed according to the feedback of the users and site administrators: (1) the Grid Storage Access Framework, an Object Oriented Framework designed to access and manage Data Grid via APIs; (2) the Secure Storage, a service for the gLite middleware providing users with a set of tools to store in a secure way confidential data; (3) OPeNDAP Meta-Finder, a tool to search geographical information data sets available at the Web through OPeNDAP servers; (4) the Watchdog, a tool that allow users to watch the status of a running job tracing the evolution of produced files.

Keywords:

services, security, monitoring, accounting, application, peer-tp-peer, windows

URL for further information:

https://grid.ct.infn.it/twiki/bin/view/EELA2/JRA1Services

Scientific Gateways and Portals / 101

SOMA2 – Open Source Gateway to Molecular Modelling Workflows

Author: Tapani Kinnunen[1]

Co-authors: Pekka Lehtovuori [1]; Tommi Nyrönen [1]

[1] *CSC - IT Center for Science Ltd.*

Corresponding Author: takinnun@csc.fi

SOMA2 gateway is a WWW-browser operated molecular modeling workflow environment developed and deployed by CSC. The SOMA2 environment allows users to control and combine scientific applications available in the computing system into unique application workflows, which are automatically executed. SOMA2 offers a flexible framework for integrating and executing molecular modeling applications, including molecular data exchange. SOMA2 source code is distributed under the GPL license.

Detailed analysis:

For end users, SOMA2 offers a secure and personalized environment for inputting molecular data, submitting and controlling jobs and analyzing the results. In SOMA2, scientific applications are presented and configured via web forms, which guide users to correctly configure a program by supplying default values, thresholds, runtime help and validation. For experts, SOMA2 offers a framework to make virtually any molecular modeling application accessible to the end users.

SOMA2 has a modular design where third-party scientific applications are described as pluggable capsules with interfaces to manage the data. A capsule consists of an XML description, used e.g. to generate an application web form, and scripts & file templates to enable program execution and processing of the program output. The XML descriptions are based on an original schema. SOMA2 enables communication and data exchange between applications by employing a common data exchange format, CML (Chemical Markup Language).

Conclusions and Future Work:

At CSC, SOMA2 is available as a service for CSC's academic users providing access to 14 different molecular modeling applications, which are seamlessly integrated within the system. System is fully integrated with the local computing infrastructure.

In addition to UI enhancements, SOMA2 development plans include DCI-integration by making use of grid middleware. SOMA2 framework is fully compatible for this approach. Technical challenges mostly concern with grid authentication and authorization components.

Impact:

SOMA2 gateway provides an easy to use and intuitive single user interface to scientific applications and it hides all technicalities from end users. The system automates repeating tasks and eliminates redundant work so that end users can focus in the actual scientific task instead of dealing with technical issues requiring manual work. SOMA2 does not only integrate applications but also different computing platforms connected to the system making the whole infrastructure easily reachable for all users.

The framework for describing scientific applications in SOMA2 facilitates transfer of technical know-how from experts to service users so that everything remains in machine readable form. In addition, core SOMA2 software is separated from the program descriptions and basically no real programming skills are required to create a SOMA2 capsule. Flexibility, application oriented approach for abstract level reusable workflows and open source license make SOMA2 system unique in its domain.

Keywords:

gateway, user environment, open source, molecular modeling, workflow, CML, XML, knowledge transfer

URL for further information:

http://www.csc.fi/soma

Bioinformatics / 102

Setup and usage of an e-BioScience infrastructure for high throughput DNA sequence analysis

Authors: Angela CM Luyf[1]; Barbera DC van Schaik[1]; Silvia D Olabarriaga[1]

Co-authors: Antoine HC van Kampen [1]; Frank Baas [1]; Katja Ritz [1]; Michel de Vries [1]

[1] *Academic Medical Center*

Corresponding Author: b.d.vanschaik@amc.uva.nl

Current DNA sequencers produce a large amount of data. Because the amount of data is growing and the computation time for analysis is increasing, we initiated a pilot to run applications on the Dutch Grid (Big Grid, part of EGEE). We used the software platform that was developed for medical imaging in the VL-e project (e-BioInfra), and applied it to DNA sequence analysis. With the knowledge gained in this pilot, we now use this platform routinely for computational intensive analyses.

Detailed analysis:

We ported sequence analysis applications to the grid with the e-BioInfra platform. The analysis tools (executables) are wrapped as workflow components with the Generic Application Service Wrapper (GASW) service. The workflows are described in the Scufl language of the Taverna workbench and executed on the grid with the MOTEUR workflow engine. Computation and data management are based on the gLite middleware. The analysis tools, GASW descriptions, workflow descriptions and data are located on grid storage. For the analysis of sequence data the user transfers files from local resources to grid storage with the Virtual Resource Browser (VBrowser). Thereafter, the user can start workflows that are enacted on the Dutch Grid by the MOTEUR Web Service from the VBrowser. The user monitors workflows and grid jobs from the same front-end, using specialized grid services and web interfaces. When the analysis is complete, the user recovers or browses results with the VBrowser.

Conclusions and Future Work:

A generalized layer on top of the gLite middleware (MOTEUR, GASW) can help end-users to interact with grid resources from a user friendly interface (VBrowser). We observed that besides the hard- and software tools, in-house expertise about workflows and grid infrastructures is needed, because porting applications to the grid is not a trivial task. In the future we need to improve data security and error handling, which are known challenges for the deployment of grids in practice.

Impact:

The platform has been used for two different sequencing projects. In a pilot study, a metagenomics study of viruses in human samples, we identified the experimental sequences using BLAST, a popular sequence alignment program in the bioinformatics community. For the second project, identification of alternative splice variants, we ported in-house developed Perl and R scripts to the grid to perform all versus all comparisons between 400,000 sequences per experiment. This has become a routine tool in the Bioinformatics Laboratory, facilitating the analysis of sequence data.

We expect that the number of applications and the throughput of the DNA sequencers will increase fast in the next years. The workflow technology can help us to reuse earlier developed software components and we need grid infrastructure to cope with the increasing data flow and analysis. The workflows are available for members of the same Virtual Organization and via myExperiment.org

Keywords:

Next generation sequencing, e-BioInfra, Bioinformatics, Grid workflows, Porting, MOTEUR, VBrowser

URL for further information:

http://www.bioinformaticslaboratory.nl/ (EBioScience infrastructure)

Poster session / 103

ASC: an Adaptive Scheduling Controller

Authors: Alessandra Doria[1]; Catello Di Martino[2]; Davide Bottalico[2]; Giovanni Battista Barone[2]; Giovanni d'Angelo[2]; Luisa Carracciuolo[3]; Vania Boccia[2]

Co-authors: Christian Esposito [2]; Gianluca Busiello [4]; Giovanni Cantele [3]; Giuseppe Vitagliano [2]; Maurizio Pollio [2]; Mauro Petrillo [4]; Silvio Pardi [1]

[1] *INFN*

[2] *University of Naples Federico II*

[3] *CNR*

[4] *CEINGE*

The deployment, management and TCO of large computing environments always involve huge investments. These systems, once in production, have to meet the needs of users belonging to large and heterogeneous communities: only an efficient and effective use of these systems can repay the investment made. In this context, we report the experience made to design, implement and validate an adaptive scheduling system able to reconfigure itself when a lack in utilization efficiency occurs.

Detailed analysis:

The adaptive scheduling controller (ASC) relies on top of a Maui-Torque scheduling system and has been developed following the following steps: (i) we have identified a set of Maui key-parameters, related to a combination of fairshare, reservation, preemption and backfill mechanisms, used to achieve an efficient and effective use of the system; (ii) we have evaluated the system behavior with respect to some key-statistics (queue waiting time, jobs throughput, resource usage, and so on); (iii) we have developed a control loop that uses information about the key-statistics and the desired performance profile in order to dynamically define a new set of Maui key-parameters values.
The default profile of the ASC control loop, based on automated log analysis and neural network techniques, can be chosen among a set of available profiles, each one identifies a target class of applications/users (e.g., parallel jobs, multi-thread jobs, concurrent jobs, and so on).

Conclusions and Future Work:

Here we describe the work made to devise an adaptive scheduling controller (ASC), which aims to gain a balanced, efficient and effective use of the computing resources by heterogeneous communities.

Actually ASC has been deployed for Maui-Torque scheduling system, but it can be easily implemented on the top of other scheduling systems (e.g., LSF, PBS/Moab, and so on).

Impact:

This work has been deployed and validated on computational resources of the University of Naples Federico II, acquired in the context of PON "S.Co.P.E." Italian National project. The resources are shared among three different contexts all based on gLite middleware: EGEE, Southern Italian and metropolitan GRIDs.
Due to the heterogeneity of the user community, the computational resources are used both for traditional GRID jobs and for HPC applications.
The adaptive scheduling represents an appealing solution to gain the needed trade-off among the needs of these, usually contrasting, class of applications.
We have validated the system by tests both with "driven load" and real production load (i.e., 100 Kjobs/month on about 2000 CPU and HPC jobs required about a 10% of resource usage).
The whole user community has experienced a good level of satisfaction. In particular, HPC community, usually penalized by a general-purpose scheduler configuration, registered an improvement.

Keywords:

adaptive systems, job scheduler, resource management systems, log analysis, neural networks.

URL for further information:

http://www.scope.unina.it/C2/scheduling/default.aspx

DCI programming / 104

ETICS a success story from research to production

Author: isabel matranga[1]

Co-author: andrea manieri [2]

[1] *Engineering Ingegneria Informatica SpA and ETICS2 team member*

[2] *Engineering Ingegneria Informatica SpA*

Corresponding Author: isabel.matranga@eng.it

ETICS is a software build, test and quality certification system which, after four years of development activities and use by the research environment, is now mature to start its way into the commercial software development environment. The experimentation of the use of ETICS also in the commercial environment was part of the challenge since the very beginning. This experimentation has now become a reality thanks also to the project commercial partners like Engineering which is going to offer ETICS as a service to the project managers from its production units.

Detailed analysis:

gLite and gCube, and more recently Unicore have exploited the ETICS services and contributed to their improvement. On the 28th February 2010 the ETICS 2 project will end and ETICS services - per se - will no longer be provided by the ETICS 2 Consortium (the adoption of ETICS has been considered in several project proposals e.g EMI, SGI). The ETICS open source SW (Apache 2.0 license) will be hosted on a public forge managed on a voluntary bases by members of the ETICS 2 consortium. In the view of a fully exploitable ETICS system, Engineering has started an experimentation by using ETICS in commercial projects aiming at creating higher consensus around the ETICS service and promoting it throughout the company to all its Project Managers. This activity will also allow the identification of further improvements in terms of performance and usability of the entire system, as well as the usage patterns and operational costs of running such service. The ongoing experimentation aims at installing ETICS at Engineering premises on a Virtual infrastructure, and will include the recent developments of ETICS submission engine towards the Amazon Web Services to allow higher capabilities on demand

Conclusions and Future Work:

The success of the ETICS experimentation within Engineering company will mean also a success for the innovation cycle by witnessing the adoption of research results in the more pragmatic commercial domain justifying investments. According to the results of the experimentation (in terms of percentage of adoption within the company) Engineering will analyse the opportunity of operating the ETICS service from its data centre. In the future, the same service may be opened to external users, mainly SMEs cooperating with Engineering or Research Projects in which Engineering is involved.

Impact:

During the presentation Engineering team will present experiences and results from the above described experimentation. Engineering can be considered the right example of adopter of ETICS in a context of software production as it is a large company with 7000 people essentially working on software development projects. It runs about 100 development projects per year, from small ones (3months - few developers) to large ones (spanning over several years - up to 70 developers - several development units). Engineering is also running a data centre which currently serves more than one hundred customers with an average of three thousand users. Moreover the company is certified ISO 9001 and CMMI lv 3 and since 1996 has devoted resources to improve the quality of its results and internal processes. Engineering aims at fully exploiting the ETICS system within the company and for this reason it will support the use of the ETICS service after the end of the ETICS 2 project (the activity will only partially be supported by the ETICS 2 EU co-funding) for the D4Science II development team, the Easy-Rider project (a national research project) and two of its commercial projects.

Keywords:

Build, Test, Software quality certification, Cloud services, commercial exploitation

URL for further information:

www.eng.it; www.eticsproject.eu ; cloud.eng.it (under construction)

Poster session / 105

Programming gLite without an UI - a JavaGAT adaptor for gLite

Author: Maximilian Berger[1]

Co-author: Thomas Zangerl [2]

[1] *University of Innbruck*

[2] *KTH*

Corresponding Authors: tzangerl@pdc.kth.se, maximilian.berger@uibk.ac.at

Programming in the EGEE Grid with gLite currently requires a large toolchain of software, which is installed on UI machines. It is often difficult, if not impossible, to install the full gLite software stack on development machines or on deployment machines, e.g. for web portals.

We present a generic adapter for programming gLite which is independent from the gLite software stack and from UI machines. It can be used on any machine providing the Java Runtime Environment. It provides an API compatible with JavaGAT, the standard for Grid programming.

Detailed analysis:

When trying to program software for the Grid, application developers are usually faced with the problem that there exist several APIs for Grid programming. Each Grid middleware demands its own software stack, which is incompatible with the software stack and API of other Grids. This problem will increase in the Future, as the current EGI proposal defines Unicore, ARC, and gLite as possible middlewares to be used.

To overcome this problem, a standard API for Grid Programming, GAT, has been introduced. In particular, bindings for Java are published under the term JavaGAT. Until recently, however, there was no functioning implementation for gLite available in JavaGAT.

The main reason for the lack of gLite support was the dependency of the full gLite software stack in the public available API. Software accessing gLite would have to be run on a machine which has gLite installed, which in most cases meant a special User Interface machine.

We present a working JavaGAT adaptor for gLite which is independent of the gLite software stack. It only uses the public APIs of the gLite services, and as such can be run on any machine providing the Java Runtime Environment.

Conclusions and Future Work:

The adapter presented here is included in the JavaGAT distribution. It is currently in use in several projects, ranging from workflow management systems to Grid portals. It has been tested and verified in the EGEE Grid, and it is ready to be used by more Grid application developers.

In the future, we plan to maintain and enhance the gLite adaptor. We plan further bug fixes and enhancements, and hope to provide the best independent API for programming the gLite Grid.

Impact:

With the adapter presented here, we broaden the access to gLite. Programming gLite no longer means having to use a specific operating system with special software installed. Application development can now happen on laptops, desktops, and existing servers. We also support different operating systems. Whereas the original gLite software API is only available for Linux, our adaptor also works on other OSes, such as Mac OS X or Windows. This also means that developers can write Grid applications which work under these OSes, and therefore broaden the range of available users.

Second, the adapter confirms to a standard interface currently widely in use for Grid programming. This means that applications currently written for this standard can be adapted for gLite easily. This enables Grid application developers to support multi-Grid operation, and also to switch from other Grid systems to gLite at their leisure.

Keywords:

gLite programming JavaGAT

URL for further information:

http://gforge.cs.vu.nl/gf/project/javagat

Demo Session 2 / **106**

XtreemOS Grid Operating System

Author: Yvon Jegou[1]

Co-authors: Christine Morin [1]; Peter Linnell [1]

[1] *INRIA*

Corresponding Author: yvon.jegou@irisa.fr

The XtreemOS operating system provides for Grids what a traditional operating system offers for a single computer: abstraction from the hardware and secure resource sharing between different users. When a user runs an application on XtreemOS, the operating system automatically finds all resources necessary for the execution, configures user's credentials on the selected resources and starts the application. It simplifies the work of users by giving them the illusion of using a traditional computer.
The XtreemOS technology is being developed in the frame of the XtreemOS European IP.

Detailed analysis:

The XtreemOS operating system provides three major distributed services to users: application execution management, data management and virtual organization management.

XtreemOS supports legacy Linux applications as well as MPI and SAGA parallel applications. Applications can be run in the background or interactively. This last possibility allows to exploit numerical computation platforms such as Scilab or Mathlab on the Grid. The application execution manager provides scalable resource discovery
through a peer-to-peer overlay which dynamically connects all resources.

XtreemOS provides location-independent access to user data through XtreemFS, a POSIX-like file system spanning the Grid.

User management in XtreemOS is delegated to Virtual Organization managers. Access rights to resources in XtreemOS are based on policies. Policy rules are defined by Virtual Organizations as well as by administration domain, are checked at reservation-time and are enforced on resources during execution.

This system is available for PCs, clusters and mobile devices. The cluster flavour exploits the Kerrighed single system image Linux operating system.

Conclusions and Future Work:

The XtreemOS project has released the second public release of its Linux-based Grid operating system in November 2009 under the motto "Making Grid Computing Easier". The consortium has conceived and integrated a platform of open source technologies to enable easier deployment, management, usage and programming on top of large-scale computing Grids. Supporting more volatile resources such as those provided by Cloud infrastructures is currently being investigated.

Impact:

Many enterprises are operating in a distributed fashion. Thus, the whole company is divided into several administrative domains. Further, many joint research and development programmes exploit resources spanning multiple administrative domains. In order to run the overall business effectively, the different locations must cooperate and dynamically adapt as a whole during changes. One of the main goals during this operation is the minimization of administration tasks. Furthermore, it is essential for enterprise to be able to execute legacy software within these environments without the need to modify or recompile the various system components.

Users logged into an XtreemOS box transparently exploit VO-managed resources through the standard POSIX interface. In contrast to middleware approaches, XtreemOS is an operating system able to execute any kind of application on the Grid, including unmodified existing applications as well as interactive applications.

XtreemOS can be seen as an alternative to traditional Grid middleware, facilitating the use of federated resources for scientific as well as business communities.

Keywords:

Grid operating system, Virtual Organization, Grid file system

URL for further information:

www.xtreemos.eu

Justification for delivering demo and/or technical requirements (for demos):

XtreemOS will be demonstrated either on a single PC running multiple virtual machines or on external testbeds if Internet connectivity is provided by the conference organization.

High Energy Physics / 107

CMS Remote Analysis Builder (CRAB) with the ARC grid middleware

Authors: Erik Edelmann[1]; Kalle Happonen[2]; Tomas Lindén[2]

Co-authors: Jesper Koivumäki [2]; Joni Välimaa [2]

[1] *CSC - Finnish IT Centre for Science*

[2] *Helsinki Institute of Physics HIP*

Corresponding Author: kalle.happonen@cern.ch

The Compact Muon Solenoid (CMS) is one of the general purpose experiments at the CERN Large Hadron Collider (LHC). For distributing analysis jobs to computational resources scattered across the world, the CMS project has developed the CMS Remote Analysis Builder software (CRAB). Until now CRAB has only supported the gLite and OSG middlewares, but with the help of a new plugin, the CMS analysis jobs can be run on sites running the ARC middleware.

Detailed analysis:

CRAB uses the underlying grid middleware to find a computing element (CE) close to an storage element (SE) where the data to be analyzed resides, and submit the jobs there. CRAB uses the middleware with the help of plugins. These plugins have been available for the gLite and OSG middlewares. The Tier-2 center hosted by Helsinki Institute of Physics (HIP) uses the ARC grid middleware. To enable CRAB jobs on HIP, two methods are being used. The first is based on grid level interoperability between ARC and gLite grid middlewares, together with the gLite plugin for CRAB. Recently, also an ARC plugin for CRAB has been developed, to enable jobsubmission using an ARC UI and native interaction between CRAB and ARC. This plugin allows CRAB to better support ARC, and enables CRAB to use all its features for jobhandling.

Conclusions and Future Work:

CRAB can used in a standalone mode or through a server managing the users job on the grid. In standalone mode the CRAB ARC plugin has been proven to handle the CRAB analysis jobs succesfully. This method has the drawback of only being able to use one middleware type at the time. The CRAB ARC plugin will be integrated with the CRAB server to enable transparent job submission to all middleware flavors in use in CMS. This will removes the last large obstacle of completely transparent usage across middlewares.

Impact:

The sites supporting the CMS experiment have predominantly run only on gLite and OSG middlewares. When the CMS tools better support different middlewares, the sites have more freedom choosing how to run their site. By making the CMS tools more middleware agnostic this also helps the users use the largest amount of resources, transparently and efficiently.

Keywords:

ARC, interoperability, CMS

URL for further information:

http://wiki.hip.fi/twiki/bin/view/Extranet/Tier/CRABModificatioN

Scientific Gateways and Portals / 108

A Scientific Gateway for Distributed Mining of Biomedical Image Data utilizing Web Services

Author: Charalampos Doukas[1]

Co-authors: Aristotle Chatziioannou [2]; Ilias Maglogiannis [3]

[1] *University of the Aegean*

[2] *National Hellenic Research Foundation*

[3] *University of Central Greece*

Corresponding Author: doukas@aegean.gr

Recent advances in biomedical applications like DNA sequencing, microarray data generation, high-throughput, gene-function studies, medical imaging, and electronic medical records, have resulted in the automatic generation of new and vast data repositories. Mining and managing such biomedical data is a complex procedure that requires several processing phases. Especially in the case of biomedical images, proper preprocessing, (e.g., image enhancement, color processing), feature extraction and classification are required.

Conclusions and Future Work:

Tools and workflows to support the different phases of the image mining process will be accessible through a simple semantic search in a web repository administered by open source workflow tools. Proper authentication and encryption mechanisms have been utilized in order to guarantee the appropriate security.

Impact:

The proposed framework supports the knowledge discovery cycle downstream of image analysis tasks, i.e., from exploratory image preprocessing and transformation to pattern discovery, hypothesis induction and image recognition. In cases where the available resources do not match the requirements of the task at a researcher's mind, the platform can make far easier the development of new solutions by reusing and combining available partial solutions. An ultimate goal of the proposed approach is to reduce time from innovation to deployment. A researcher's work for a specific image processing task will be assessed by the image processing community, and outstanding research results will be picked up and tested on other type of images.

Keywords:

distributed mining gateway, biomedical images, Web Services

URL for further information:

http://www.e-lico.eu/

Detailed analysis:

This work in the context of e-LICO (e-Laboratory for Interdisciplinary Collaborative Research in Data Mining and Data-Intensive Sciences) project funded by Information Society Technology program of the European Commission, provides an open framework based on Web Services that provides access to complete tools for distributed data mining of biomedical image data. Web Services are emerging as a promising technology to build distributed applications. It is an implementation of Service Oriented Architecture (SOA) that supports the concept of loosely-coupled, open-standard, language - and platform-independent systems. The latter allow service providers to modify backend functions while maintaining the same interface to clients. Web Services are accessed through the HTTP/HTTPS protocols and utilize XML (eXtendible Markup Language) for data exchange. This in turn implies that Web Services are independent of platform, programming language, tool and network infrastructure. Services can be assembled and composed in such a way to foster the reuse of existing back-end infrastructure.

Evolution of density of interacting agents for the two species irreversible reaction A+B->2A in complex networks

Authors: Aristotelis Kittas[1]; Panos Argyrakis[1]; Paschalis Korosoglou[2]

[1] *Physics Department, Aristotle University of Thessaloniki*

[2] *Grid Operations and HPC Centre, Aristotle University of Thessaloniki*

We use Monte-Carlo simulations to study the dynamics of the irreversible A+B->2A two species reaction on regular lattices, Erdos-Renyi (ER) and scale-free (SF) networks. The problem we study is an analogue to the spread of a virus in computer networks and other epidemiological models, such as information diffusion (i.e. rumor spreading) in social networks, and information propagation among stations in ad hoc mobile networks. We monitor the density of the species with respect to time, which we find to be an exponential function of time in the long time regime.

Detailed analysis:

We study the dynamics of the irreversible A+B->2A reaction where A and B represent agents of two species that move randomly on a lattice or a network topology. Besides the one dimensional lattice case we show that the agent density follows an exponential function of time in the long time regime. We also study the transition mechanism from the short to the long time regimes and estimate the crossover time as a function of initial densities and network topology. In order to perform our numerical simulations in a robust manner and to leverage the underlying Grid resources we have implemented a three phase workflow. In the first phase we have constructed the networks (both ER and SF) and stored them on SEs. In the second phase we dispatched a series of Monte Carlo simulations that used the networks designed in the first phase along with the distinct initial parameters. Our results were joined together in the third phase with a program that accumulated the data produced in the second phase. Post processing of the results was performed on local workstations.

Conclusions and Future Work:

Our results show that with the right immunization strategies we can limit infection spreading on normal lattices and ER networks where a transition mechanism from a short time to a long time regime is evident. We plan to further continue our studies by examining cases of reversible species interaction and recovery phenomena.

Impact:

We have developed a model to study the A+B->2A two species irreversible reaction, which is related to epidemiological modes, such as the spread of a virus in networks, the spread of rumors in social networks and information propagation in ad-hoc mobile networks. While the 1D lattice presents the most trivial case with a linear dependence of the density from time, 2D and 3D lattices as well as ER networks can be described with an exponential decay in the long time regime. The crossover time from short to long time scales as a power law in these systems, with ER networks characterized by a different exponent than 2D and 3D lattices. The connectivity of ER networks was shown to have little or no influence to the infection process. In SF networks the crossover is almost completely absent and the infection is spread immediately, with the highly connected nodes acting as the infection centers. The behavior of density in SF networks is more complex than their ER counterparts. For small initial density in SF networks, the connectivity of the network is shown to severely influence the infection process, which proceeds much faster in well connected networks.

Keywords:

SI model, Monte Carlo

URL for further information:

http://kelifos.physics.auth.gr

Poster session / 110

GILDA training infrastructure: present successes and future triumphs

Author: Valeria Ardizzone[1]

Co-authors: Diego Scardaci [1]; Emidio Giorgio [1]; Giuseppe La Rocca [1]; Leandro Neumann Ciuffo [1]; Marco Fargetta [1]; Riccardo Bruno [1]; Roberto Barbera [1]

[1] *Istituto Nazionale di Fisica Nucleare sez. catania*

Corresponding Author: valeria.ardizzone@ct.infn.it

GILDA is a very successful initiative, led by INFN, aiming to provide a special production-quality Grid infrastructrure (t-Infrastructure) for higher education, training and dissemination purposes. Since its earlier phases, EGEE adopted GILDA as its t-Infrastructure, providing many scientific and humanistic communities with a fully fledged virtual laboratory where to "learn and try". GILDA has officially been adopted as the default t-Infrastructure by other projects such as EELA and EELA-2, EPIKH,EUAsiaGrid, EUChinaGrid, EUMEDGRID/EUMEDGRID-Support, SAGrid, the UNESCO/HP Brain Gain Initiative.

Detailed analysis:

The main operational activity of GILDA is to manage its infrastructure which provides users with several central services such as Workload Management System, VOMS server, User Interfaces, Information Indexes, Data and Metadata Catalogs, and a series of Computing Elements and Storage Elements for a total of more than 280 CPU cores and about 10 TB of storage. Another key component of GILDA is its Certification Authority, which is able to release X.509 certificates both to hosts/services and users. Host certificates are valid for one year, while user certificates have a default 15-day lifetime. During tutorials "generic-certificates" can also be created "on demand" and stored in a shared user area, ready to be accessed by trainers during the events. GILDA can offer free technical support to new pilot grid infrastructures, cloning itself locally and sharing the best practices and common policies in e-Infrastructure set-up and management.

Conclusions and Future Work:

In EGI, there will be no directly funded t-infrastructures and NGIs will support their own training by using their own national production resources. Both the present status and the next-future perspectives represented by the CUE project (if approved) will be presented.

Impact:

The main objective of this oral contribution is to present the activities carried out since the last EGEE User Forum in Europe and in the rest of the world (Asia, Africa and Latin America). In the last almost six years, GILDA has been used in more than 400 tutorials and induction courses all over the world and thousands of people have been trained on installing, using and exploiting gLite as well as other middleware. We report on the state-of-the-art of training in each continent and the impact which has been seen in several science domains. The relocable/clonable model which is GILDA's most important strength has favored the creation and fast deployment of local training infrastructures to support research institutions, universities and industrial communities.

Keywords:

Grid Computing, training, dissemination, user and application porting support

URL for further information:

https://gilda.ct.infn.it

Poster session / 111

A Grid-Enabled Problem Solving Environment for QTL Analysis in R

Authors: Carl Nettelblad[1]; Mahen Jayawardena[1]; Per-Olov Östberg[2]; Salman Zubair Toor[1]

Co-authors: Erik Elmorth [2]; Sverker Holmgren [1]

[1] *Department of Information Technology, Uppsala University, Box 337, SE-751 05, Uppsala, Sweden*

[2] *Department of Computing Science and HPC2N, Umeå University, SE-901 87, Umeå, Sweden*

Corresponding Author: salman.toor@it.uu.se

There exist a large number of applications in multidisciplinary research environments that require efficient, seamless, and easy-to-use infrastructures to address computationally intensive problems. In this work we present an approach to build Grid-enabled problem solving environments (PSE), that allow end-users to operate within familiar settings and provide transparent access to computational Grid resources.

Conclusions and Future Work:

The proof-of-concept implementation has illustrated that this type of system can be implemented with relative flexibility using pre-existing components. The PSE integrated into R can be further enhanced with existing code packages and tools, given proper interfaces and data translation services.

Detailed analysis:

Quantitative Trait Loci (QTL) analysis is identified as a suitable candidate for exploitation of large-scale computing infrastructures such as Grids. Our approach is illustrated in a proof-of-concept implementation of a PSE for multidimensional QTL analysis. The goal of such analysis is to determine a set of loci in the genome that are responsible for the genetic component of a quantitative trait, using phenotypic and genotypic data gathered from an experimental population. The analysis uses a powerful statistical model, and the computational requirements rapidly become very large when several loci are searched for simultaneously.

The prototype integrates a system built on the open source statistical software system R with the Grid Job Management Framework (GJMF) which provides middleware-transparent access to Grid functionality, allowing our PSE to not be limited to using a specific middleware, but rather offload Grid issues to existing arbitration layers.

Impact:

The QTL PSE is targeted towards end-users with limited experience of high-performance computing, and supports workflows where small tasks are performed locally on PSE hosts, while larger, more computationally intensive tasks are allocated to Grid resources. In this model, Grid computations are scheduled asynchronously, allowing client hosts to perform e.g. visualization and postprocessing of data in parallel with large-scale Grid computations. The proposed integration model provides the end-users with all the capabilities of the R environment coupled with transparent access to the computational power of grid environments. Use of the R-based environment allows end-users to operate in environments they are familiar with, and provides statistical processing capabilities. Use of the GJMF decouples the PSE from specific Grid middlewares and provides abstractive and reliable access to Grid resources through concurrent use of multiple Grid middlewares.

Keywords:

QTL, workflows, PSE, R, bioinformatics, middleware arbitration

URL for further information:

http://www.it.uu.se/research/project/ctrait, http://www.gird.se/gjmf

Demo Session 1, Welcome Drink / **112**

Impact of the synthetic seismogram archive for the CMT determination

Author: David Weissenbach[1]

[1] *CNRS / IPGP*

Corresponding Author: david.weissenbach@upmc.fr

The aim of this demo is to show the performance increase provided by an indexed synthetic seismograms archive. It is also an opportunity to identify partners wishing to distribute and consolidate this archive.

Detailed analysis:

The method used for CMT determination allows to keep intermediate results that can are used for results refinement but that also can be used for other earthquakes at the same location. This saves up to 80% CPU time (avg 600CPU hours) at the cost of less than 10GB per earthquake.
The use of a database interfaced with grid jobs (by use of the GRelC software) is the most accurate way to ensure that once recorded these synthetic seismograms will be used again if needed. The archive itself relies on grid storage facilities to allow full availability and consistent network usage.

Conclusions and Future Work:

The Grid already changed the utility of the application, allowing it to run effectively in a routine usage despite its heavy computation needs. The synthetic seismogram indexed archive not only improves the performance, but also the robustness and the flexibility of the application, and optimises its use of the Grid.

Impact:

The primary impact of the service is the tremendous decrease of time needed to complete a CMT determination once the synthetic seismograms are recorded. The added value of the database and the consolidated archive is to systematically reuse them. As these CMT are then useful for further studies, providing them to the community as early as possible is the main objective. As a side product, the database also provides statistics about the usage of the seismometer network by the application.

Keywords:

database earthquake seismogram archive

URL for further information:

http://www.geoscope.fr

Justification for delivering demo and/or technical requirements (for demos):

the demo gives the opportunity to show the significant impact provided by the database and archive running in the same time as the full computation as well as other tools (VTK correlation display)

Poster session / 113

Earth Science and Astrophysics Applications in Armenia: Present and Perspectives

Author: Hrachya Astsatryan[1]

Co-authors: Hovsep Keropyan [2]; Mikayel Gyurjyan [3]; Tigran Khotsanyan [2]; Vladimir Sahakyan [4]; Yuri Shoukourian [5]

[1] *IIAP NAS RA*

[2] *Junior Researcher*

[3] *Senior Researcher*

[4] *Director*

[5] *Scintific Adviser*

Corresponding Author: hrach@sci.am

The Institute for Informatics and Automation Problems of the National Academy of Sciences of the Republic of Armenia operates, supports and manages the national Grid Infrastructure and Academic Scientific Network of Armenia, which provides all core services to the users and consists of 7 Grid sites located in Yerevan and Ashtarak cities (424 cores). Armenia actively engaged in different international Grid (EU FP7 SEE GRID SCI, ISTC A-1606, ISTC A-1451) and connectivity (EU FP7 BSI) Projects, which make possible to deploy infrastructures and environments in the earth science and astrophysics.

Keywords:

ArmGrid, ArmNGI, gLite, WRF, Seismology

URL for further information:

http://www.grid.am

Detailed analysis:

The following applications from the earth science and astrophysics domains are already running on the national Grid infrastructure:
 The seismology platform consists of the seismic data (received from 30 seismic stations), seismic data server, programming tools and interfaces, and applications (ELF, SRA, etc.) developed within the SEE GRID SCI Project.
 The core of the numerical weather prediction system is the Weather Research and Forecasting model implemented and operationally used for the territory of Armenia by Armenian State Hydrometeorological and Monitoring Service. Initial condition data is taken from Meteo (downloaded from National Center for Environmental Prediction) database and the results of calculations are stored in the forecast database. The model serves as a basis for solving different problems (environmental, hydrological, etc).
 The Digitized First Byurakan Survey (digitized version of the Markarian Survey) has been ported on the National Grid infrastructure by using the standards of International Virtual Observatory Alliance. Each plate contains low-dispersion spectra of some 15,000-20,000 objects, and there are some 20,000,000 objects in the whole survey.

Conclusions and Future Work:

The implementation and usage of the suggested platforms are crucial for the local community in the earth science and astrophysics domains. The implementation of the mentioned infrastructure allows to integrate with the international and regional correspondent infrastructures, such as ORFEUS, NERIES, EnviroGrid, EGI SAFE, etc.. The work has been started to port the Community Multiscale Air Quality modeling system in order to improve the environmental management community's ability to evaluate the impact of air quality management practices for multiple pollutants at multiple scales.

Impact:

To study and analyze very large scale and complex problems in the earth science and astrophysics domains require to have data sets from different regional areas, experimental results from different research laboratories, some good data processing, analyzing software and above all high performance computational resources. The main impacts are:
 Facilitate access to the vast archives of continuous seismic waveforms collected in the Armenian regional datacenters and promote better standards in quality control of the data archives and procedures for data availability and exchange.
 The outputs of global numerical weather prediction models with coarse resolutions, which represent only broad features and patterns and are able to reproduce processes in the large scale. Weather forecasting would significantly benefit from information on the processes at spatial resolution much finer than the coarse resolution of global models (nested domain covers the South Caucasus region - spatial resolution 5km).
 Link astronomical archives and databases of the Armenian virtual observatory (VO) to the other national and international VOs, together with analysis tools and computational services.

Parallel Processing / 114

Tapas4Grid: a grid tool for performing parametric studies based on complex workflows

Authors: Antonio Gómez-Iglesias[1]; Francisco Castejón[1]; Marcus Hardt[2]; Michael Goetter[2]; Rainer Stotzka[2]

[1] *CIEMAT*

[2] *Karlsruhe Institute of Technology*

Corresponding Author: antonio.gomez@ciemat.es

One typical problem to be solved by fusion codes, which is suitable for grid and parallel computing, is the parameter scan. There are several codes whose main purpose is to get results for wide parameter range. So it makes sense to build a tool that allows one to perform this parameter scan in an automated way. We have build and used a new tool suitable for solving this kind of problems: TAPaS. This tool is, in principle, and thanks to its layer based design, suitable for both HPC s and grid computing, although it has only been used on the grid.

Detailed analysis:

TAPaS itself is able to manage a complex workflow that implies running the main code automatically with a wide variation of its input parameters and managing the results.

"TAPaS4grid" (Tool for Automatic Parameter Studies for Grid) provides a workflow for preparation and submission of extensive parameter studies for running on the grid (s.t.a.r.t.). As a case example, EMC3_EIRENE and BIT1 have been taken to test this tool, but it is possible to run other codes using this tool. Internally TAPaS4grid provides a workflow including all steps to run intensive parameter studies using any suitable code on the grid.

EMC3-EIRENE is a package of two coupled Monte Carlo codes, the 3D fluid edge plasma code EMC3 and the 3D neutral particle kinetic transport code EIRENE. An iterative procedure is applied to obtain a self-consistent solution for all plasma and neutral particle parameters.

BIT1 is an application to perform Particle-in-Cell (PIC) simulations. This application represents a powerful tool for plasma studies having a number of advantages like the fully kinetic description of high-dimensional plasma and the ability to incorporate complicated atomic and plasma-surface interactions.

Conclusions and Future Work:

We have developed a tool for automatic parameter studies that is being used to perform different tests with different applications in the nuclear fusion field. The application hides all the interaction with the grid environment, allowing the scientist to define the tests that he wants to run via a parameter's file and launch the tool. The obtained results show a set of relevant results for fusion community, and this tool opens a wide range of possibilities to perform more detailed tests.

Impact:

Even though TAPaS4grid has been initially developed to use the grid infrastructure, it could be ported to other infrastructures like HPC. The tool can be divided into two different layers:

- Parametric studies. This layer takes the parameter file and the input file of the application used for the studies, and generates a set of valid input files that will be used to run the application. This layer is standard and does not need to be changed.
- Infrastructure interaction. This layer is the responsible for the interaction with the grid. It creates the job, submits it, retrieves the results,... This layer should be ported to different infrastructures in order to use the tool.

Before starting TAPaS4grid it is needed to create a parameter file including the necessary input data for every job. In this file there can be added an unlimited number of different variables. During the workflow they will be replaced by the values specified in the parameter file.

With this parameter scan process we show how we can perform a deep study of the configuration parameters of a fusion device, modelling the behaviour of this device in many different ways.

Keywords:

Parametric, Grid, Automatic, START

URL for further information:

https://wiki.eu-euforia.eu/index.php/TAPaS4grid

Security / **115**

Sanctorum: a solution for safe management of digital certificates and passwords in a farm

Author: Stefano Dal Pra[1]

Co-author: Marco Verlato [1]

[1] *INFN*

Corresponding Author: marco.verlato@pd.infn.it

Sanctorum is a python tool created to aid sitemanagers to safely manage digital certificates and passwords for hosts in a farm. Motivation for the tool: when releasing a new host certificate, the Italian Certification Authority recommends the site to maintain two backup copies of the private key in a safe place not network reachable. This makes it quite difficult to both respect the CA rule and to efficiently manage the site, especially large ones. The proposed solution makes it possible to manage digital certificates in a comfortable and error free manner while still respecting CA recommendations.

Detailed analysis:

The reliability of common Authentication methods depends on how confidential a piece of sensible information can be kept. Host's Private keys are valid until they are inaccessible to all but the owner. Sanctorum aims to provide a comfortable way to manage host passwords, certificates and private keys while ensuring an adequate security level and preventing risk of human error. The basic idea to achieve this is to keep key/cert pairs in a cyphered database using raid1 disks in a machine with a disabled network. The machine should only be accessed through console or avocent like solutions. When a certificate or other data has to be transferred to/from a host, the network card gets enabled, a firewall rule is added to let the transfer happen, then the network card is deactivated and the fw rule removed. Apart from adding new hosts to the db, there is no frequent need to log on the sanctorum's host: a tool is provided to create new requests for expiring certificates and mail it to an operator; a second one is able to download a newly generated certificate from the CA website and stores it in the db after checking that it matches its key, then it copies the cert/key couple on the target host.

Conclusions and Future Work:

The Sanctorum tool has been designed and developed to allow for the easy management of digital certificates and host passwords in large scale computing grid farms. It complies with the security IGTF approved CA policies and CPS. It is currently adopted by the Italian Tier-1 facility and by several Tier-2 centres of the IGI infrastructure.

Impact:

Maintaining host certificates for a large number of hosts is a time consuming activity. If not strictly planned and performed it might let sporadic errors happen. These in turn may lead into possibly severe unscheduled down of a service. Common practice at site level is to "hide" all private keys into some local filesystem network reachable together with some script used to perform common management tasks. Should an attacker gain access to that hidden place all hosts identities should be considered compromised, all certificates should be revoked and requested again from scratch to the CA. This is why keeping keys not network reachable makes sense. Sanctorum offers a complete set of tools for certificate and password management (check, deploy, get, update etc.) and keeps private keys out of the operator hands. This can only be copied into its host or matched against its certificate. Any operation is strictly interactive. This avoids the need to build shell loops and gain or save knowledge on many hosts at once. Apart from the setup of new hosts, renewal operations are almost fully automated (the

only human intervention is to sign a request mail and bounce it to local RA).

Keywords:

X509, PKI, CA, Certificates

URL for further information:

http://www.italiangrid.org/grid_operations/site/documentation

Data Management / 116

Improved Task Scheduling for Distributed Data Analysis Systems

Authors: Daniel van der Ster[1]; Jakub MOSCICKI[1]; Massimo Lamanna[1]

[1] *CERN*

Corresponding Author: jakub.moscicki@cern.ch

In this work, we present an evaluation of a late-binding scheduler for HEP distributed analysis, where data is distributed globally with multiple replicas in a non-uniform way. We employ the late-binding technique in order to optimize job placements thereby minimizing the per-job time-to-completion. We evaluate different scheduling and prioritization strategies, and evaluate the approach using a prototype system implemented with the Ganga end-user tool and the DIANE scheduling framework applied to ATLAS distributed analysis.

Detailed analysis:

In the context of ATLAS distributed analysis, the presently used workload management systems do not optimally place jobs. At worst jobs are pre-assigned to a single execution site at submit time, and at best jobs are assigned to run on a limited number of closely-located sites (so-called “Atlas Clouds”). This preassignment of jobs to sites leads to two suboptimal behaviours:

1. While a job is waiting in a queue at a site, the data and resource availability can change and therefore the job’s placement at that site becomes less and less ideal.
2. The placement of an entire task (all jobs in the task) at a single or just a few sites can lead to other sites sitting idle or being utilized by lower priority jobs.

With the recent startup of the LHC, the urgency of achieving scientific results in a timely manner is on the critical path. Enabling quasi-interactive analysis on the grid is therefore essential.

Conclusions and Future Work:

Making effective use of resources for distributed user analysis is an important challenge with the potential to improve the user experience substantially. In the HEP community alone, physicists numbering in the thousands use the grid facilities to run jobs numbering in the hundreds of thousands daily.

In the current economic model of the Grid based on public funding and SLAs,
resource utilization of a site is a primary metric of its success. Therefore it is a responsibility of the user communities to apply job distribution strategies which reward sites according to their quality of service

Impact:

We analyze the impact of job scheduling strategies on handling of job output and
subsequent implications for overall data-management strategy for ATLAS collaboration.

To achieve a quasi-interactive distributed analysis system, we
consider a few metrics:

1. Time-to-completion. This is a measure of both overall performance
   for all users at all sites, and for individual users. The variance of
   this measure is used to evaluate the stability of the schedules.
   Timely delivery of partial results is also desirable to permit users
   to take corrective actions as soon as possible.
2. The correlation of job priority with time-to-completion. This is a measure of how well the priority values
   are respected by the workload management system.
3. Fairness to sites is an important property of the scheduling system as sites are rewarded for running jobs
   successfully. Therefore, sites should receive a share of the global jobs proportional to their quality, where quality relates to their
   efficiency and the popularity of their resident data.

Keywords:

data management, scheduling

URL for further information:

http://cern.ch/diane

Demo Session 2 / 117

Using gLibrary with mobile devices: a proof of concept with Apple iPhone

Authors: Antonio Calanducci[1]; Costantino Pistagna[2]

[1] *INFN, Sezione di Catania-Universita & INFN, Catania-Unknown*

[2] *Università degli Studi di Catania*

Corresponding Author: antonio.calanducci@cern.ch

gLibrary mobile is a native application for the Apple touchscreen devices that allows to access digital libraries, deployed over grid infrastructures, from mobile devices such as the iPhone and the iPod touch. gLibrary mobile is our attempt to offer a mobile client to interact with digital libraries created with the INFN grid digital library solution. It allows to easily browse libraries contents, inspect library items metadata and finally download and view in the device high resolution screen the closest replica of the selected object.

Detailed analysis:

gLibrary is the INFN digital library solution to create digital libraries on gLite grid: it allows to create, organize, populate, browse, search and access libraries of digital objects saved and replicated on storage

systems of grid infrastructures with a certificate-enabled intuitive and easy-to-use web 2.0 interface from any computer browser. The goal of gLibrary mobile is to offer access to digital libraries created with gLibray from anywhere. The iPhone and iPod touch devices, with their high-resolution touchscreen, multimedia capabilities and intuitive interface are the perfect devices to provide a mobile front-end to gLibrary. As the desktop counterpart, authentication and authorization on metadata and stored files are handled through X.509 certificate, loaded on the devices. The iPhone UIKit components fits perfectly to implement a cascading filter browsing system on digital objects metadata. Once the user has found the digital object he was looking for, with a simple tap on the screen he will start the download of the closer replica, selecting the location from a storage map and retrieving the current user position with the built-in GPS.

Conclusions and Future Work:

We plan to offer an upload features that will add the possibility to push digital contents generated from the iphone (like pictures, voice recordings, videos) to grid storage elements including the editing of the associated metadata.

Impact:

Being an advanced multimedia device, the iphone/ipod touch is suitable to access, on the road, different kinds of multimedia types such as videos, audio files and images. For example, the media player on the device is able to handle streams of movies and/or music files from a DPM HTTP/HTTPs enable storage elements. Moreover, those Apple devices are able to handle natively a lot of office formats such as .doc, .xls, .ppt, .pdf, .pages, .key. and .numbers. All these features provide the mobile platform to download and view suck kind of documents with a few taps. During the demonstration we will show how to search and stream movies trailers, music files as well as view some images and PDFs. In particular, we will demonstrate the access to the ancient manuscript repository of Federico De Roberto. Thanks to the built-in GPS, all downloads can happen from the physical closest location to the user.

Keywords:

iphone, digital libraries ,data grid

URL for further information:

https://glibrary.ct.infn.it

Justification for delivering demo and/or technical requirements (for demos):

demo + oral presentation

User Support and Services / **118**

Lessons learnt from the EGEE Application Porting Support activity

Author: Gergely Sipos[1]

[1] *MTA SZTAKI*

Corresponding Author: sipos@sztaki.hu

In its third phase, the EGEE project flagged application porting as a critical service for the successful adoption of grid solutions by new users and new communities and includes “Application Porting Support” among its generic user support activities. The goal of the EGEE Application Porting Support group is to aid developers in effectively porting Virtual Organizations’ (VO) applications to the Grid. The presentation gives an overview of the key activities of this support service, its evolution during the EGEE project, and the lessons learnt by the provision of this service.

Detailed analysis:

The EGEE-III application porting support group consists of grid middleware and high level tool experts from 6 countries from 3 continents. Group members worked together with scientists closely and became

well informed about the needs of grid newcomers, and the typical problems these people are faced with. The presentation will give an overview of the typical problems that an EGEE user has to deal while he/she interacts with the infrastructure, and will represent the results that the Application Porting Group achieved to resolve such issues. The support group was highly active in reporting and resolving infrastructure problems; further developing generic application porting tools; preparing application examples and success stories. All these efforts aimed at minimizing the overhead to become an active user of the EGEE grid. The presentation will give an overview of similar user activities of ARC and UNICORE communities and will outline how these services can be unified in the EGI era.

Conclusions and Future Work:

Members of the EGEE Application Porting Support Group, and similar representatives of the ARC and UNICORE communities, are aiming at a joint project to unite the most important user support services under a single, coherent service activity. The project proposal is still under evaluation but its result will be known by the time of the User Forum. The presentation provides a perfect opportunity to inform the European Grid user community about the project status, and outline the methodologies of similar services to be provided under the EGI umbrella.

Impact:

Efforts of the Application Porting Support Group resulted in the porting of several grid applications from various fields of science, thus significantly contributed to the visibility and scientific impact of the EGEE grid and of grid computing in general. These applications are all registered in the EGEE Application Database, and detailed technical explanation is provided on the Webpage of the Application Porting Support Group.

At the end of EGEE-III, at the beginning of EGI it is important for the wider European grid provider community to hear about the activities of the EGEE Application Porting Support Group, and to be aware of the possible ways of expanding such service towards the complete EMI middleware platform.

Keywords:

Application, porting, gridification, user support service

URL for further information:

www.lpds.sztaki.hu/gasuc

Earth Science / 119

Computational Requirement of Meteorological and Crisis Applications

Author: Juraj Bartok[1]

Co-authors: Andrej Lucny [1]; Ladislav Hluchy [2]; Martin Gazak [1]; Viet Tran [2]

[1] *Microstep-MIS*

[2] *Institute of Informatics, Slovakia*

Corresponding Author: viet.ui@savba.sk

We present several applications from the domain of meteorology and crisis management that we developed and/or plan to develop. In particular, we present the IMS Model Suite - a complex software system designed to address the needs of accurate forecast of weather and hazardous weather phenomena, environmental pollution assessment, prediction of consequences of nuclear accident and radiological emergency. We discuss the computational requirements and our experiences on how to meet them by grid computing.

Detailed analysis:

A pollution assessment and prediction requiresthe running of 3D meteorological model (4 nests with resolution from 50 km to 1.8 km, 38 vertical levels) as well as the running of the dispersion model performing the simulation of the release transport and deposition of the pollutant with respect to the numeric weather prediction data, released material description, topography, land use description and user defined simulation scenario. Several post-processing options can be selected according to the particular situation.
Another application is the forecasting of fog as one of the meteorological phenomena hazardous to aviation and road traffic. It requires complicated physical model and high resolution meteorological modeling. An installed fog modeling system requires a 4 time nested parallelized 3D meteorological model to be run four times daily. The 3D model outputs and a multitude of local measurements are utilized by the SPMD-parallelized 1D fog model run every hour.
The fog forecast model is a subject of parameter calibration before its real deployment. For each parameter, it requires re-running of the hundreds of historical situations and comparison with the observed data.

Conclusions and Future Work:

We found grid computing useful for our applications. We are satisfied with this technology and our experience encourages us to extend its use.

Within an ongoing project (DMM) we plan to include the processing of satellite images which extends our requirement on computation very rapidly. We believe that thanks to grid computing we are able to handle the job almost in real time.

Impact:

The architecture and inherent heterogeneity of both examples, and their computational complexity and their interfaces to other systems and services, make them well suited for decomposition into a set of web and grid services.
Such decomposition has been performed within several projects in which we have participated in cooperation with the academic sphere, namely int.eu.grid (dispersion model deployed as a pilot application), SEMCO-WS (semantic composition of web and grid services) DMM (development of a significant meteorological phenomena prediction system based on the data mining), VEGA 2009-2011 and EGEE III.
We present useful and practical applications of technologies of high performance computing. The use of grid technology provides access to much higher computation power not only for modeling and simulation, but also for the model parameterization and validation. This results in the model parameters optimization and more accurate simulation outputs. Since the simulations are used for the aviation, road traffic and crisis management, even a small improvement in the accuracy of predictions may result in a significant improvement of safety as well as a cost reduction.

Keywords:

High performance computing, meteorological model, hazardous phenomena, fog, environmental modeling

URL for further information:

http://www.microstep-mis.com, http://www.i2g.eu/, http://semco-ws.ui.sav.sk/

Regional Activities, International Projects and Collaborations / **120**

The South African National Grid Initiative : From prototype to production

Authors: Valeria Ardizzone[1]; bruce becker[2]

Co-authors: Albert Gazendam [2]; Albert van Eck [3]; Andrew Lewis [4]; Christiaan Kuun [2]; Francois Wolmarans [5]; Kriel Hannes [6]; Lambropoulos Stravros [5]; Riaan Oberholzer [6]; Sean Murray [7]; Tarirai Chani [2]; Timothy Carr [4]

[1] *Istituto Nazionale di Fisica Nucleare, Sezione di Catania*

[2] *Meraka Institute*

[3] *University of the Free State*

[4] *University of Cape Town*

[5] *University of Johannesburg*

[6] *North West University*

[7] *iThemba Laboratory for Accelerator-Based Sciences*

Corresponding Author: bbecker@csir.co.za

South Africa is has made significant investments in e-Infrastructure recently. These investments are partly in the form of centralised, centrally-funded initiatives such as the Centre for High-Performance Computing (CHPC) and South African Research Network (SANReN). However, a significant investment has also been made in parallel by universities and national laboratories of computing resources and manpower, dedicated to research support. The South African National Grid is a national federation of these efforts with the goal of deploying a production-quality regional e-Science infrastructure.

Detailed analysis:

During the year of 2009, a shift of activities in South Africa has been made in two directions. The first is the move from prototype development of national grid infrastructure to a more stable, production-level platform. The second is a shift from inward dissemination of the project towards an outward dissimnation effort of the project in an attempt to integrate operations with regional and global efforts. Several new sites have been formed and special service activities undertaken to ensure the national coherence of the platform. These include the development of a "clone" of the GILDA services, providing a virtualised pre-production and testing environment for new users and sites. Furthermore, an EUGridPMA-accredited Certificate Authority is under development for the Sub-Saharan Region. A core team of experts responsible for site operations has been formed and undergoes regular development in joint EGEE/SANReN/SAGrid training sessions.

Conclusions and Future Work:

SAGrid is now moving to production status, along with the SANReN network which will be fully operational by March 2010. With the basis now solidly set for research support in a federal infrastructure, the next step is to greatly increase participation and usage. The goal is to provide a tool for collaboration which is transparently accessible and supported, independant of where the user may be located in Sub-Saharan Africa, and which is fully interoperable with EGEE/EGI.

Impact:

This formalisation of the manpower dedicated to SAGrid has enhanced South Africa's participation to large collaborations, including that of the EPIKH project, which ran its first school for application porting in South Africa this year, which was run entirely on SAGrid services. Several memoranda of understanding with existing projects with FP7 funding have been signed in order to accelerate usage of the infrastructure and introduce South African researchers to the possibilities of collaboration with these projects. South African research efforts with great impact in the physics, geomatics and bioinformatics domains are already seeing great benefit. However, the greatest impact lies in the collaboration agreement with the HP/UNESCO project which could see the seeds of the first regional Afrian grid initiative.

Keywords:

SAGrid, gLite, EGEE, EGI, e-Infrastructure

URL for further information:

http://www.sagrid.ac.za

User Support and Services / 121

Training in EGEE - creation of a sustainable user service

Authors: Gergely Sipos[1]; Robin McConnell[2]

[1] *MTA SZTAKI*

[2] *NeSC*

EGEE-III Networking Activity 3 (NA3) has the responsibility of coordinating and delivering user training and induction. The activity aims to train a wide variety of users, both those within the EGEE consortium and those from outside the project who make use of the EGEE Grid infrastructure. The presentation will present the evolution of training activities in the EGEE series of projects, provide a summary of its current services and highlight activities that make gLite-related training in Europe a sustainable service

URL for further information:

http://www.egee.nesc.ac.uk

Detailed analysis:

The provision of training is an important segment in the EGEE virtuous cycle, the project's strategy for extending EGEE's base of users and communities.
This is realised with the use of an ever expanding portfolio of training material and resources designed for an ever-widening range of grid applications.
The activity provides central services to support the training teams. A web portal serves to assist advertising of events, and provide a gateway to the online facilities useful to trainers and users. A central digital library is also provided, allowing material created for and during EGEE-III training events to be made available to the community in a searchable way. The training activity provides a training infrastructure (GILDA), allowing realistic hands-on training.

Conclusions and Future Work:

With the end of the EGEE project approaching, the training activity is looking to leave a legacy of training services which can be used by the future training groups. This includes a database of accredited trainers,
It is apparent that face-to-face training is not sufficiently scalable for large numbers of new users; therefore developing e-Learning practices is becoming increasingly important.

Impact:

To date within the EGEE-III project, our trainers have delivered more than 150 training events attended by nearly 2000 participants. The federal management approach adopted seeks to satisfy local requirements such as language, timeliness and content. Local trainers have to be flexible to take advantage of potential training opportunities arising with short lead times.
The EGEE digital library continues to collect course material and documentation to add to the wealth of material already present, with nearly 9000 items currently available. The library also provides a home to recently developed eLearning modules to enable self-paced learning and provide content for the online International Winter School.
The GILDA infrastructure developed within EGEE is available round the clock for tutorial sessions and individual users, with the number of contributing sites in creased in EGEE-III.
Feedback obtained from course participants has consistently improved throughout the life span of EGEE indicating the increased quality of the training events.

Keywords:

Training user induction e-Learning digital library

Poster session / 122

Data Mining and Integration of Environmental Applications

Author: Ladislav Hluchy[1]

Co-authors: Martin Seleng [1]; Ondrej Habala [1]; Viet Tran [1]

[1] *Institute of Informatics, Slovakia*

Corresponding Author: viet.ui@savba.sk

In this paper we presents the data mining and integration of environmental applications in EU IST project ADMIRE. It briefly presents the project ADMIRE and data mining of spatio-temporal data in general. The application, originally targeting flood simulation and prediction is now being extended into the broader context of environmental studies. We describe several interesting scenarios, in which data mining and integration of distributed environmental data can improve our knowledge of the relations between various hydro-meteorological variables.

Detailed analysis:

The project ADMIRE aims to deliver a consistent and easy-to-use technology for extracting information and knowledge. Its main target is to provide advanced data mining and integration techniques for a distributed environment. In this paper, we will focus on one of its pilot applications with the target domain as environmental risk management. Several scenarios have been proposed including short-term weather forecasting using radar images, complex hydrological scenarios with waterworks, measured data from water stations and meteorological data from models. Historical data for mining are supplied mainly by the Slovak Hydrometeorological Institute and the Slovak Water Enterprise.
The main characteristics of data sets describing phenomena from environment applications are spatial and temporal dimensions. Integration of spatio-temporal data from different sources is a challenging task due to those dimensions. Different spatio-temporal data sets contain data at different resolutions and frequencies. This heterogeneity is the principal challenge of geo-spatial and temporal data sets integration – the integrated data set should hold homogeneous data of the same resolution and frequency.

Conclusions and Future Work:

In this paper we have demonstrated the data integration and mining platform developed in the ADMIRE project for environmental risk management. We have also shown complex meteorological and hydrological scenarios with different data sources in different formats and a simple way to make workflows in DISPEL for data integration and mining. The graphical editor for workflows is in progress and should allow experts in environmental applications to make data integration easily. This work is partially supported also by VEGA project 2009-2011 and APVV project DMM.

Impact:

The whole data integration and mining processing of scenarios are described in DISPEL, the language proposed in ADMIRE for data integration and mining. Data are stored in different forms (files with different formats, databases) and accessed via OGSA-DAI. A gateway will be responsible for processing the DISPEL file, sending request to the corresponding OGSA-DAI servers and collecting results. An easy to use graphical interface is being developed for the ADMIRE platform. The users can select existing components (activities), connect them together and make workflows of data integration and data mining with a visual editor. The development of environmental applications in the ADMIRE project is tightly collaborated with the VEGA project 2009-2011 and APVV project DMM.

Keywords:

data integration, data mining, environmental modeling, OGSA-DAI, modeling language

URL for further information:

http://www.admire-project.eu/

Poster session / 123

Development plans of the EGEE Operations Portal and dashboard towards and EGI/NGI infrastructure

Authors: Cyril L'Orphelin[1]; Helene Cordier[1]

[1] *CNRS/IN2P3*

Corresponding Author: helene.cordier@in2p3.fr

The operations portal of EGEE is undergoing a tremendous back-end evolution to wrap-up the 5 years experience gained on daily grid operational needs about procedures and tools developed for and made available to the different EGEE communities.This evolution is crucial to meet the changes to cope with the EGI structural needs and with the NGIs requirements for the regionalisation of the tools where applicable.

Detailed analysis:

The first phase of the restructuration of the Operations Portal has been completed with the release of the restructured regional operations dashboard: the gridops regional dashboard. It consists in a regional package including the Lavoisier module and a php web portal which allows high flexibility in the customisation. Also, support of an open source relational database management system fits the requirements of most regions who wish to deal with a standalone implementation of the tool.
Based on this first phase, the back-end restructuration enables a consistent evolution for the EGEE Operations Portal features that are considered as key for the EGI/NGIs infrastructure towards customization and regionalisation when relevant.

Conclusions and Future Work:

Future plans of the portal developments for the post-EGEE era include how we collaborate with all major existing partners such as GOCDB. Such collaboration has already been set-up to provide a seamless access to the users regarding the sites' and the VOs' information. The next phase of the Operations Portal restructuration work will also deal with the migration and re-engineering of the other existing features of the operations portal recognized as key for EGEE and EGI operations.

Impact:

The restructuration enables us to go on providing an up-to date, flexible and scalable integration platform that is now easier to maintain. Also, it is easier to implement the dedicated changes in order to cope with the needs of the new EGI structure and the new needs of the operational entities of the NGIs community in May 2010.
The operations model in EGEE has been relying on a sustainable and largely decentralized model for over 6 months now. Indeed, the regionalization implies a model that must allow for procedures and tools customization by the regions and at the same time must allow for coherence with the central layer, in order to provide consistent overview, reports and metrics at the EGI level.
Another key evolution for the challenge in the regionalisation relies on the scalability and the flexibility required regarding information source types and information handling. The normalised way to add a new source of data is to use Lavoisier, a standard-based information handling of heterogeneous sources, and the Web Interface will help the community to improve their own tool to cope with the evolution towards the structure of EGI/NGIs.

Keywords:

EGEE operations portal; gridops operations dashboard

URL for further information:

http://cic.gridops.org/

Workflow Management / 124

Services for advanced workflow programming on gLite with WS-PGRADE portal

Author: Gergely Sipos[1]

[1] *MTA SZTAKI*

Corresponding Author: sipos@sztaki.hu

After the great success of the P-GRADE Portal MTA SZTAKI has created the second generation P-GRADE portal, called WS-PGRADE. WS-PGRADE is a Web based environment for both developers of grid based dataflow applications, and end users who need transparent access to grid based application services. The presentation gives an overview of the application developer functionalities of the WS-PGRADE environment, details the various application patterns that can be realized in the system and executed on cluster based and PC based platforms.

Conclusions and Future Work:

WS-PGRADE has been developed as part of the EU CancerGrid project where it was used to port and run three applications on joint gLite-BOINC-database platform: Descriptor Calculation, Model Building and Property Prediction. The first generic and public release of gUSE has been opened for public in November 2009. MTA SZTAKI already collaborates with several research teams who need support for advance dataflow applications on top of gLite middleware.

Detailed analysis:

WS-PGRADE extends the workflow programming concept of P-GRADE Portal in many respects:
1. WS-PGRADE enables the development of nested workflows, recursive workflows, allow programmers to add control mechanism to the workflow graphs.
2. With the extended parameter study support data generator and data collector nodes can be placed anywhere in a workflow graph.
3. Components of a WS-PGRADE workflow can be Web services, jobs running on cluster grids (e.g. gLite, Globus, local cluster), tasks running on PC grids (e.g. BOINC), database queried performed on a JDBC compatible data repository.
4. WS-PGRADE supports the concept of abstract workflows, workflow templates and running workflow instances. These entities simplify the interaction among workflow developers, the collaboration of workflow developers with end users.
On top of that, the WS-PGRADE system is not a monolithic service any longer, it is a set of interoperable Web Services. Besides better scalability and fault tolerance, this approach also brings modularity and flexibility to WS-PGRADE programmers, who can from now replace different parts of the system with custom implementations.

Impact:

Since WS-PGRADE can support both advanced application developers who understand grid technology as well as end-users who do not necessarily know grid systems, WS-PGRADE can attract large number of new users for EGEE. The built-in application repository concept enables the collaboration of user communities to jointly develop new applications. Application developers can place grid-enabled applications into the repository and end-users can download and use those applications by simply defining the necessary parameters. The applications can be single jobs, parameter sweep jobs, single workflows, parameterized workflows, workflows including parameterized nodes, data generators and data collectors. These special node types can be used anywhere in the workflow providing a very flexible way of constructing highly sophisticated and very large workflows. WS-PGRADE workflows can be executed in various resources: local clusters, Web Services, LCG-2, gLite, GT2, GT4 and BOINC grids, and hence users can exploit many resources from different grids.

Keywords:

Grid programming, workflow, environment, dataflow, SOA

URL for further information:

www.wspgrade.hu

DCI programming / 125

Ganga/DIANE and SAGA for solving complex scientific problems

Authors: Jakub MOSCICKI[1]; Massimo Lamanna[1]; Ole Weidner[2]; Shantenu JHA[2]

[1] *CERN*

[2] *LSU*

Corresponding Authors: massimo.lamanna@cern.ch, jakub.moscicki@cern.ch

An increasingly important requirement for solving complex scientific problems is the simultaneous use of qualitatively different resources, from simple batch nodes to supercomputers. We present a project where Ganga and DIANE user-level middleware tools are used in connection with SAGA, a high-level Grid API, to achieve an improved quality of service when using resources across multiple computing infrastructures, such as EGEE, NDGF, OSG and TeraGrid. We discuss these advances in the context of LQCD simulations in theoretical physics and other applications in the current portfolio.

Detailed analysis:

This project make a complementary use of three software components, Ganga, DIANE and SAGA to support complex scientific applications. We extend the WMS/DIANE pilot-job approach to proivde a uniform policy-based pilot-job submission mechanism across these multiple Grids, and thus provide production level interoperabiltiy for a range of applications.

Ganga is a computational task-management tool , which allows for the specification, submission, bookkeeping and post-processing of computational tasks on a wide set of distributed resources. Ganga is primarily used by Atlas and LHCb experiments at LHC at CERN, and also by a number of smaller user-communities in the context of the EGEE project.

DIANE is a late-binding scheduler and optimizer of the task processing, taking advantage of the pilot-job paradigm. DIANE has been used successfully on a large-scale to provide an improved Quality of Service experience for end-users and to improve the reliability and efficiency of task processing.

Conclusions and Future Work:

The systems which we describe will become increasingly important for ROSCOE and INSPIRE applications.

Impact:

SAGA is a high-level API that is currently an OGF technical specification; it is targetted to become a standard in 2010. Several production grids worldwide currently support SAGA or are deploying it currently (including TeraGrid and UK's NGS). SAGA provides the ability to develop applications and tools that can work on and interoperate across different distributed systems.

SAGA is currently being deployed on several production grid
infrastructure – including the US-TeraGrid, UK-NGS, Japanese-NAREGI, and thus applications can seamlessly be extended to interoperate across these different infrastructure.

This project will ultimately take advantage of concurrent effort in developing SAGA as a uniform access layer to other European middleware such as gLite, ARC and UNICORE.

Keywords:

application support tools, pilot jobs, SAGA, Ganga, DIANE

URL for further information:

cern.ch/diane

Poster session / 126

Grid-CSIC project: porting astrophysic applications to Grid

Author: JOSE RAMON RODON ORTIZ[1]

Co-authors: ALICIA DESIREE BENITEZ YAÑEZ [1]; JOSE RUEDAS SANCHEZ [1]; SUSANA SANCHEZ EXPOSITO [1]

[1] *IAA. Instituto de Astrofísica de Andalucía*

Corresponding Author: rodon@iaa.es

Two important problems related to astrophysics applications are high computational cost and limited storage. The Grid-CSIC Project promotes the use of Grid infrastructure on the research institutes of Spain. The Institute of Astrophysics of Andalusia (IAA-CSIC) joined this project in order to provide scientific application support in the astrophysics area. Several applications have been implemented by the infrastructure work team.

Detailed analysis:

We show the architecture of the Grid node, the additional installed software, and how the work team has implemented the migrated applications. The astrophysical applications which have been ported to Grid infraestructure are varied, i.e: extragalactic astronomy applications, stellar physics applications, radioastronomy applications, galactic structure applications or solar system applications.

Conclusions and Future Work:

There are many applications in the astrophysics area, and we intend to continue migrating applications. to the grid infrastructure

Impact:

The work has resulted in a decrease of the elapsed computational time of the applications thanks to the distribution over the grid infrastrucure.

Keywords:

Grid, Astrophysic, porting

URL for further information:

http://www.grid.csic.es/

Parallel Processing / 127

Towards a robust and user-friendly MPI functionality on the EGEE Grid

Author: Jeroen Engelberts[1]

[1] *SARA Reken- en Netwerkdiensten*

Corresponding Author: jeroene@sara.nl

An MPI Working Group has been setup with the goal to investigate why so few Grid sites support MPI and why so few people are using it. The Working Group has come up with a recommendation document in which the current issues with MPI on the Grid are analyzed, the reasons for non-usage are investigated and the plans for future expensions are summarized.

Detailed analysis:

The installation of MPI on the Grid clusters is not default, nor straight forward. Therefore, many site administrators do not take the time to install MPI on their clusters. One of the reasons is that the MPI implementation depends on the type of interconnect between the nodes on the sites. Furthermore, there are several flavors of MPI, or message passing in general, which are not exchangable.

The users claim that MPI is hard to use, because of several reasons. Some indicate that they lack (central) support. Some say that they can't specify the configuration they require in their JDL script. Some say that their specific MPI flavour is not available. Others say that tracing MPI jobs on the Grid is extremely difficult.

Conclusions and Future Work:

From many e-mails, the survey done by the MPI Working Group and the presentations of the users on the EGEE'09 conference, it became clear that the focus should not only be on the fixing of the main issues with the MPI implementation (currently being tackled by the MPI Task Force), but on the development of new features for the user community (the new task of the MPI Working Group).

Impact:

After the MPI Session on EGEE'09, initiated by the MPI Working Group, a Task Force was set in place with the mission to solve the short term issues indicated by both the site administrator and user communities.

The current role of the MPI Working Group has therefore been adapted to focus on the longer term extension and improvement of the MPI implementation of the Grid. In this presentation ideas for the extension of the MPI implementation and in JDL are presented.

Keywords:

MPI Working Group

URL for further information:

http://grid.ie/mpi/wiki/WorkingGroup

Computational Chemistry / 128

MCTDH Quantum Dynamics in EGEE: Advances Made and Improvements needed

Author: Fermin Huarte[1]

[1] *University of Barcelona*

Our communication provides a detailed description of the present performance of grid-Fluss and grid/MCTDH, tools for the direct quantum-mechanical calculation of kinetic coefficients, implemented on the section of the production computing Grid of EGEE accessible to the COMPCHEM virtual organization. The performance and reliability of the method is illustrated by presenting the results of two computational campaigns: the N + N2 reaction and the CH4 + Ni(111) dissociative sticking process.

Impact:

The flux correlation method employed in our work can make quantum calculations suitable for implementation on computing grid and distributed computing. This possibility is exploited by propagating separately the several quantum state basis functions using the grid-MCTDH tool. In this way, an increase of the quantum dynamics basis to be propagated is transformed into a computational effort associated to the number of concurrent processes that need to be distributed.
The present state of the grid-Fluss and grid-MCTDH tools will be presented, focussing on two realistic applications. Special attention will be paid to the advantages provided by the
distributed nature of the grid-tools and the possible strategies for improvement.

Keywords:

quantum dynamics

URL for further information:

http://www.ub.edu/gdrq

Detailed analysis:

The direct calculation of chemical reactions thermal rates coefficients k(T) has increasingly attracted theoretical and computational work. The computation of reliable thermal rate
values can be of great use in several realistic multiscale simulations of complex Systems.
Whenever an elementary chemical reaction occurs through a reaction barrier and no long-living complex is formed, a rigorous method of calculating k(T) in a direct way by means of flux correlation functions. The method is based on a dynamics simulation confined in the region around the saddle point and has the advantage of decreasing significantly the numerical effort with respect to a full scattering simulation. The MCTDH scheme offers an
additional numerical advantage to the approach by expanding the time-dependent multidimensional wavefunction onto a basis of time-dependent functions.

Conclusions and Future Work:

Grid computing infrastructures offer new possibilities for Quantum Dynamics calculations, provided that the appropiate workflows are applied. The EGEE production section accessible to the COMPCHEM virtual organization has been employed to calculate thermal rate coefficients for the N + N2 reaction and and sticking probabilities of CH4 + Ni(111) .

Infrastructure Tools and Services / **129**

Visual GRID Activity Monitoring in the Dashboard Framework

Authors: Benjamin Gaidioz[1]; Elisa Lanciotti[2]; Jacobo Tarragón[1]; Julia Andreeva[1]; Lukasz Kokoszkiewicz[1]; Marco Devesas Campos[1]; Pablo Saiz[1]; Ricardo Rocha[1]; Sergey Mitsyn[3]

[1] *CERN*

[2] *PIC*

[3] *JINR*

Corresponding Author: marco.devesas.campos@cern.ch

It is sometimes hard to conceptualise the GRID and visualise its activity. But having a graphical representation of its workings and movements, we can more easily surpass these problems; coupling this with a system that provides accurate measurements of such activities, we obtain a powerful monitoring system that is visually appealing, intuitive to use and of great practical worth.

Our solution is supported on two, interconnected systems: the Google Earth Dashboard monitor provides the high-level, graphical view of the Grid while more detailed information is displayed by Siteview.

Detailed analysis:

The Google Earth application displays a highly detailed 3D model of planet Earth, on top of which various informations can be overlaid; the Google Earth Dashboard Monitor is a service that provides such information in near real-time about the GRID. It collects information about the GRID activities –job processing and data transfers– and transforms it to a format that Google Earth can understand and display. In particular, Sites are represented in their place on the 3D globe and one can see file transfers and job submission going through the skies between the intervening sites. This gives a very visual sense of the level of activity on the GRID.

For more detailed information about the status of the Sites, the Google Earth monitor takes us to Siteview. This service provides for every site, and for every VO on it, details of its current operational status. It is based on GridMap. Thus, one can with a simple glance assess the overall status of the Site and the activities running there. It also displays information hierarchically, enabling one to find information about relevant sub-activities, and, for more specific information, it takes one to the lower-level monitoring tools.

Conclusions and Future Work:

These higher-level monitoring tools, that visually show the GRID's activity, are invaluable not only in educating and presenting what the GRID is and how it achieves its goals but also in the daily monitoring activities of the GRID.

As such, work is underway that will extend these tools in terms of the quality of the visualisation and in the type and number of VOs supported. One particular point that we are particularly excited about is to provide real-time visualisation of the GRID activities. We would also like to support VOs other than the CERN experiments.

Impact:

The success of the Google Earth monitor is easily visible: it is currently on display at CERN's IT department and at the Atlas visitor center and it was one of the stars of the show at CERN's ITU booth. The reason is simple: it's visually attractive, educational and informative.

The Siteview service is currently used by Site administrators to check, from the VOs perspective, the status of their Sites and the on-going activities in them.

Keywords:

Monitoring Tools, GRID Visualization, Graphical User Interfaces

URL for further information:

http://dashb-earth.cern.ch

User Support and Services / 130

Latest achievements of the MTA SZTAKI Grid Application Support Centre

Author: Gergely Sipos[1]

[1] *MTA SZTAKI*

Corresponding Author: sipos@sztaki.hu

The Grid Application Porting Support Centre (GASUC) of MTA SZTAKI has been established in 2007 to support potential users of grids getting their applications to distributed infrastructures. The support team is the leader of the EGEE Application Porting Support Group since 2008 and has been ported about 15 applications since then to various production VOs of the EGEE grid. This presentation will give an overview of the latest applications and tools that the support centre has worked with, developed and operated.

Detailed analysis:

The GASUC team is currently working with or completed 6 applications, these will be introduced during the talk:

- OMNeT++ is a C++-based discrete event simulation package primarily targeted at simulating computer networks and other distributed systems. OMNET++ has been realized on gLite as a workflow application embedded into a scientific gateway.
- The main goal of the NMMC3D application is the quantitative study of the structure and surface manifestation of mantle plumes on the surface of the Earth. The Grid based NMMC3D has been ported to gLite with P-GRADE Portal as a parameter study workflow application.
- CFinder is a free software for finding and visualizing overlapping dense groups of nodes in networks, based on the Clique Percolation Method. The GASUC team is working together with the CFinder developers to enable the simulation on gLite and to offer it inside a science gateway for potential end users.

- Biochemists from Hungary and the Netherlands are collaborating with GASUC to make them capable of using already grid-enabled applications of the EGEE life science community, and to enable their own research code on gLite grids.

Conclusions and Future Work:

During EGEE-III GASUC will continue to provide application porting support activity to any user of the EGEE grid. As member of the EGEE Application Porting Support Group, GASUC is heavily involved in the transition to the EGI structure. With other members of the EGEE project, and with user support teams of the ARC and UNICORE communities GASUC are aiming at a joint project to unite the most important user support services under a single, coherent service activity. The project proposal is still under evaluation but its result will be known by the time of the User Forum.

Impact:

Efforts of the MTA SZTAKI Grid Application Support Centre resulted in the porting ofseveral grid applications from various fields of science, thus significantly contributing to the visibility and scientific impact of the EGEE grid and of grid computing in general. These applications are all registered in the EGEE Application Database, and detailed technical explanation is provided on the Webpage of the Application Porting Support Group. The GASUC team was also very active on the further development of porting tools, primarily on generic science gateway frameworks. The presentation will highlight these achievements from the above described applications' point of view.

Keywords:

Application porting, support, service, gridification, science gateway

URL for further information:

www.lpds.sztaki.hu/gasuc

Novel Architectures and Technologies / 131

Grids and Clouds Interoperation: Development of e-Science Applications Data Manager on Grid Application Platform

Author: WeiLong Ueng[1]

Co-authors: Cheng-Hsin Hsu [1]; Eric Yen [1]; Shi-Chung Chiu [1]; Simon Lin [1]

[1] *Academia Sinica Grid Computing*

The Grid Application Platform(GAP) is a middleware to reduce development efforts of e-Science implementation. The GAP development is stimulated by the systematic framework in which applications can easily well-formulated common components to build up new services and taking advantage of Grid without worrying about new technologies. Cloud technology of data management is of essential importance by its performance and resilience. HDFS is integrated into GAP for the value of a scalable, fault-tolerant distributed file system for an efficient data access services in any stage of an e-Science Apps.

Conclusions and Future Work:

We have successfully adopted and integrated HDFS on GAP to provide more options of heterogeneous data management mechanisms for gLite. This work also demonstrated lots of viable alternatives to Grid Storage Element, especially in terms of scalability, reliability, and manageability. Cloud technology enhances the capability of parallel processing and also versatile data management approaches for Grid. GAP could be a bridge between Cloud and Grid infrastructures and more computing frameworks from Cloud will be integrated in the future.

Detailed analysis:

Grid Application Platform is a lightweight framework for problem solving applications on the Grid. Design of the GAP system adopts a layered architecture to make the system easy to extend and to make lower implementation transparent to upper layers. The whole system of GAP, from bottom up, consists of Core Framework, Application Framework, and Presentation Framework. The best strategy for both Cloud and Grid is to deploy Cloud components onto a global gLite infrastructure to enable Cloud federation and to impose new versatile Cloud data management to Grid. Interoperation is one of the major achievements of EGEE while more efficient and flexible storage services are still in demand. HDFS is good at features such as default replication and RAID-free to retain performance and reliability by virtue of good failure management. With the layered design of GAP, it's straightforward to provide services a tthe levels of platform, software or application over gLite.

Impact:

Grid computing emphasizes the ability to share data across administrative sites using common protocols and possibly between completely different underlying storage systems. GAP provides higher-level Java APIs, which maps the problem domain to the programming model very easily. GAP also abstracts the grid middleware with a unified interface and could be extended with new storage technologies. The greatest value of gLite-based worldwide grid (WWG) is the global e-Science infrastructure. Deployment of Cloud computing components upon WWG is thus a reasonable evolution for both Cloud and Grid. Furthermore, with the advantages of Cloud computing technology, Grid is able to support a diversity of storage types for different application purposes. Single namespace and interface for different data management systems is enabled. Technology for large-scale hypertext data like Bigtable could be integrated. Moreover, distributed computing model such as Mapreduce could be supported by the Grid Application Platform. Applications of GAP Data Manager on earthquake data center and drug discovery have been exemplified and evaluated for further development.

Keywords:

Grids, Clouds, Grid Application platform, Hadoop File System, Data Management

URL for further information:

http://gap.grid.sinica.edu.tw

Medical Imaging / 132

Grid-enabled Medical Image Analysis: From Jobs to Workflow-based Service

Authors: Carsten Byrman[1]; Matthan Caan[1]; Silvia Olabarriaga[1]

Co-authors: Angela Luyf [1]; Antoine van Kampen [1]; Frans Vos [1]; Jan Just Keijser [2]; Lucas van Vliet [3]; Tristan Glatard [4]

[1] *Academic Medical Center of the University of Amsterdam*

[2] *NIKHEF and BiG Grid*

[3] *Technical University Delft, NL*

[4] *CNRS Creatis Lyon, FR*

Corresponding Author: s.d.olabarriaga@amc.uva.nl

Medical image analysis requires large computational effort. The construction of "atlases" specific for each study of a specific brain disease is a good illustration: typically all images need to be aligned (registered) to each other. In this work we explore the Dutch Grid (part of EGEE) to implement a service that automates the construction of such atlases for Diffusion Tensor Imaging (DTI), facilitating usage by neuroimaging researchers. Initial experiments show that the service is robust and valuable for medical imaging research.

Detailed analysis:

Initially we adopted conventional gLite command line utilities to run medical image analysis jobs on the grid, however these were considered too difficult by the researchers. Later we adopted the e-BioInfra platform developed in the VL-e Medical project, which is based on the MOTEUR workflow management system and the Virtual Resource Browser (VBrowser) front-end. Although this platform provided a high-level abstraction to describe and run large experiments on the grid, current limitations on the Scufl language force the separation of analysis into various workflows, imposing off-line actions and hampering usage by non-expert researchers and methodology dissemination. The developed service encapsulates all steps (workflows) needed to perform the complete image analysis, and also automates experiment monitoring and recovery. A web service interface was chosen for its implementation to enable usage in both from programmatic and interactive interfaces. A robot-certificate is available for users who do not own a grid certificate, facilitating usage from a web interface.

Conclusions and Future Work:

The system is used in two on-going studies (schizophrenia and neuroscience) involving typically a total of 30K jobs, 2500 hours of CPU, and 30 GB of data. Other services are planned in the near future, as well as training activities to the PhD students at the AMC. We also see possibilities to improve the system for example by incorporating into workflow management the general functionality that is currently implemented by the service. Moreover, the security model needs to be revised to address privacy and usability requirements in medical applications.

Impact:

In this project we started running individual jobs with command-line tools and evolved into a workflow-based service that automatically handles grid computation and experiment recovery. This service can be used directly by end-users that are not experienced with this methodology for DTI atlas construction, neither have knowledge nor access to such computing resources. It facilitates the dissemination of a sophisticated data analysis process into the medical imaging community. After the initial results, we already observe an increasing interest for grid-enabled applications in our closest user community, which starts to consider more ambitious experiments (for example large parameter sweeps) that would not be performed in their usual infrastructure. Additionally, the knowledge and tools developed for this specific service can be now reused to more rapidly port other (life science) applications and services to the Dutch Grid. For example, large multi-hospital studies to brain diseases could also benefit immensely from distributed management and processing of the data by multiple researchers.

Keywords:

Medical Imaging, Neuroimaging, e-Science, Grid workflows, grid applications, MOTEUR, VBrowser

URL for further information:

http://www.bioinformaticslaboratory.nl

Justification for delivering demo and/or technical requirements (for demos):

no demo

Demo Session 2 / 133

LHCb Data Production System based on the DIRAC project

Authors: Adrian Casajus Ramo[1]; Alexey Zhelezov[2]; Andrei TSAREGORODTSEV[3]; Andrew Cameron Smith[4]; Matvey Sapunov[3]; Ricardo Graciani Diaz[1]; Stuart Paterson[4]; Zoltan Mathe[5]

[1] *University of Barcelona*

[2] *University of Heidelberg*

[3] *CNRS-IN2P3-CPPM, MARSEILLE*

[4] *CERN*

[5] *University College Dublin*

Corresponding Author: atsareg@in2p3.fr

LHCb is one of the four experiments at the LHC collider are CERN. It has started recently to record the data coming from the real proton-proton collisions. The data processing chain is managed using the tools provided by the DIRAC project. All the operations starting from data transfers from the exprimental area up to the final user analysis distributed in all the LHCb Tier-1 centers are automated. MC data production is also managed within the same framework. The production controls as well as visualization of the monitoring information are performed through the DIRAC Web Portal.

Detailed analysis:

The LHCb experiment is using DIRAC middleware for all the activities related to distributed data processing. DIRAC provides a layer between the grid resources and services an the LHCb production system in order to increase the overall efficiency as seen by the LHCb users.
The WorkloadManagement System of DIRAC based on the Pilot Job paradigm is used to manage the payloads coming from both users and production managers. On top of it a higher level Production Management System is built. The Production Management system automates the tasks of data processing job submission according to predefined scenarios. The production job scheduling is data driven and allows to distribute the tasks to the Tier-1 centers according to the LHCb Computing Model.
The Data Management System of LHCb is also built in the DIRAC framework and allows for automated massive data transfers definition and scheduling. It uses FTS WLCG service for the transfers and provides means for failure recovery.
The DIRAC Web Portal provides secure access to all the DIRAC services. It allows to monitor the ongoing activities and behavior of various subsystems.

Conclusions and Future Work:

The LHCb Data Production System includes both Production and Data Management services and tools all built in the same DIRAC framework. The System is successfully used in the recent processing of the real LHC data. The experience shows the necessity of better support of the production managers activity. The future development will be devoted to better interactivity with the system through the Web portal, more convenient aggregation of the information necessary to solve particular problems, more focued notification system to allow early spotting of the problems.

Impact:

The Production Management System of LHCb was heavily used in the recent simulation data production and real data processing runs. It has shown good scalability properties managing up to 30K of simultaneously running jobs. Since both user and production jobs are passing through the same central Task Queue, it allows to manage efficienltly and accurately the relative priorities of user and production jobs including different policies for different user groups and various production activities.
With the start of the LHC operation and production of data from real proton-proton collisions the data flow becomes constant and necessitates versatile monitoring and management tools to be used by the production managers and members of the LHCb computing shifts. It is important that shifters are notified promptly of any abnormal situations and have tools to quickly identify problems and react accordingly. This is achieved by generating alarms and notifications to the shifters.
In the demo we intend to provide a view of the LHCb Production System with the DIRAC Web Portal. Operations and monitoring tools for the whole LHCb data processing chain will be presented.

Keywords:

distributed computing, grid, LHC, DIRAC, data management

URL for further information:

http://dirac.cern.ch

Justification for delivering demo and/or technical requirements (for demos):

This contribution presents a real interactive system best illustrated with an interactive demonstration. The demo session needs a computer connected to the network with a large monitor.

Security / **135**

Improving Grid User's Privacy with gLite Pseudonymity Service

Author: Henri Mikkonen[1]

Co-authors: John White [1]; Joni Hahkala [1]

[1] *Helsinki Institute of Physics HIP*

Corresponding Author: henri.mikkonen@cern.ch

The Grid computing model provides Grid users a way to use resources that are not usually owned by their parent organizations. The use of Grid resources entails a balance between the resource owner's need to oversee and account for the resource usage and the user's privacy requirements. From the user's point of view, complete anonymity is desirable, but not possible due to requirements like traceability without user intervention. The solution we propose is the concept of a lesser degree of anonymity, pseudonymity: the use of traceable pseudonyms as user identifiers.

Conclusions and Future Work:

We have presented a service implementation that enables pseudonymous Grid access, ie. the hiding of true identity of the user from the Grid resources. The service has been running in our test VO with our test online CA, and its users have been able to submit pseudonymous Grid jobs to our test cluster. However, wider deployments would allow us to gain practical experience of the system and identify problem areas. We also see Grid portals as potential pieces in the overall architecture of the system: the benefits and drawbacks need to be explored.

Detailed analysis:

The security of many of today's Grid middlewares, including gLite, is based on PKI certificates. In order to interoperate seamlessly with them, the pseudonymous identities must also be modeled as standard X.509 certificates, but their subject DNs needs to be anonymized and lifetime more limited. Users obtain pseudonymous identities from the Pseudonymity Service by using their existing Grid identity for authentication and authorization. The service registers the newly issued pseudonyms to the VO's attribute authority as aliases to the users' existing certificates. This enables the users to request VO attribute assertions to be used in conjunction with the pseudonymous identity. The service uses the gLite SLCS implementation and its XML-based request-response protocol as a basis on both the server- and client-side. The service runs on a Java servlet container and uses a relational database for storing the identity mappings. It is currently compatible with online CA software supporting CMC or CMP protocols and VOMS-admin version 2.5. The Java-based command-line client tool is used for communication with the service.

Impact:

After obtaining the pseudonymous certificate from the service, the user can initialize his pseudonymous Grid identity with the standard VOMS client. In the pseudonymous Grid access, the most important challenge is confidentiality: the true identity must remain hidden from the resources. In order to prevent correlation attacks, the lifetime of the pseudonymous identities should be short: ideally they are used for only one action in the Grid. As all the pseudonymous certificates must have a unique subject DN, the user authentication and authorization at the Grid resources must be based solely on the VO attributes. In fact, the bigger the group of users sharing the set of attributes is, the better individual user's identity is buried. On the other hand, big group sizes may also reduce privacy as the whole group can access the same data in the Grid, even though the real user identities behind the pseudonyms would be different.

Keywords:

Authentication, Authorization, Grid Security, Pseudonymity

URL for further information:

http://tek.hip.fi/Projects/IDM/GridPseudonymity

Demo Session 1, Welcome Drink / 136

Grid-enabled virtual screening service against the epidemiological disease

Author: HsinYen Chen[1]

Co-authors: Eric Yen [1]; Simon Lin [1]

[1] *Academia Sinica Grid Computing*

GAP Virtual Screening Service (GVSS), a large-scale in-silico drug virtual screening service, provides a system to speed-up the searching process among all conformations of a compound. Moreover, GVSS is a generic drug discovery framework over gLite by which users can upload their compounds and targets to do grid docking, compare the significance, and verify results by in vitro experiment afterwards. GVSS has proved to provide intensive computing power and effective data management over gLite infrastructure on neglected and emerging diseases, such as Avian Influenza and Dengue Fever.

Detailed analysis:

In the study of Biomedicines, the structure based molecular docking simulation is a common method for predicting potential interacting complexes of small molecules in protein binding sites. However, Massive molecular docking required intensive computing power and effective data management. A GRID computing framework was established for AutoDock 3.0.5 and evaluated for its quality of large-scale molecular docking process. GRID is an ideal environment, which can provide large-scale and on-demand resources, including computing and storage resources.

GVSS is a user-friendly graphical user interface enhanced desktop application for using this Grid-enabled virtual screening service. Through the GUI, the end-users can easily take advantage of GRID computing resources for large-scale virtual screening. Furthermore, they can even upload their own target and ligands, and do the same docking process, visualization and analysis with this drug docking scientific gateway, of course including the advanced refinement docking simulations. The end-users can finally have a real GRID-enabled desktop utility for the virtual screening service for their daily research.

Conclusions and Future Work:

To enhance the in silico pipeline processing for the application, we use another more accurate molecular simulation package APBS (Adaptive Poission-Boltzman Solver) in conjunction with the Autodock results. Furthermore we expect to foster the biomedical grid activities and promote the e-Science collaboration between partners in Asia and Europe.

Impact:

Inspired by the successful experiences on Avian Flu Data Challenges, ASGC developed the GVSS application package that incorporates the EGEE gLite middleware DIANE2 and AMGA. Therefore, ASGC coordinated the Dengue Fever Data Challenge via EUAsiaGrid VO in June 2009. The objective was to utilize the grid enabled high throughput screening for structure–based computational methods to identify small molecule protease inhibitors. 300,000 compounds from the ZINC CDI compound library, a free database of commercially available compounds, were selected for virtual screening.

For this activity, we allocated 268 CPU-cores computing resources from EUAsiaGrid partners including Taiwan, Thailand, Vietnam, MIMOS Malaysia, UPM Malaysia and CESNET. For the phase I, a total of 46 GB of data from execution of the 300,000 jobs were generated.

Keywords:

Grid Applications, Drug Discovery, Virtual Screening, Avian Influenza, Dengue Fever, Epidemiological

URL for further information:

http://gap.grid.sinica.edu.tw

Demo Session 2 / 138

Molecular Dynamics and Docking Simulations Using Parameter Sweep Workflows

Authors: Gabor Terstyanszky[1]; Gergely Sipos[2]; Pamela Greenwell[1]; Peter Kacsuk[2]; Tamas Kiss[1]

[1] *University of Westminster*

[2] *SZTAKI*

Corresponding Author: sipos@sztaki.hu

Carbohydrate recognition is a phenomenon critical to a number of biological functions in humans. Computer programs which can provide insight into such biological recognition processes have significant potential to contribute to biomedical research if the results of the simulation can prove consistent with the outcome of conventional wet laboratory experiments. In order to validate these simulation tools and support their wider take-up by the bio-scientist research community high-level easy to use integrated environments are required to run massively parallel simulation workflows.

Detailed analysis:

Modelling the interaction between receptors and ligands includes several steps where numerous scripts and simulation programs are utilized and the data is fed from one component to another. Working through such a complex scenario manually is a tedious process. Moreover, some of the steps require massive computing resources and the analysis of several parallel scenarios. First the receptor and the ligand file are selected, checked, energy minimized and validated. Next the docking parameters and the target grid-space are defined, followed by the actual docking simulations. The simulations are run on a predefined number of parallel branches (typically around 10 parallel simulations are executed) each of which perform a genetic algorithm of a set number of evolutions (typically around 100). The best results of these simulations are then selected and energy minimization and molecular dynamics is performed, again in a number of parallel branches, on every selected docking result. The above scenario has been implemented using the WS P-GRADE portal and its workflow engine, and the execution of components has been mapped to resources of EGEE and the UK National Grid Service.

Conclusions and Future Work:

The implemented solution efficiently automates the process of executing and visualising complex molecular docking and molecular dynamics simulations in a grid environment, and providing a rich and easy to use interface for the biologist end-users. Future work includes the creation of more customized workflow scenarios using the already developed and also new workflow components as building blocks. While the first stage workflow was created with significant support from the technical team, it is envisaged that future development will be based more on specialist end-user involvement.

Impact:

In order to increase the uptake of current e-Infrastructures new user communities with no Grid related knowledge have to be targeted. These researchers are reluctant to use this infrastructure due to its complexity and the very steep learning curve required. The aim of the above described work was to illustrate how high level Grid development and execution environments can be used to produce easy to use solutions with no or very minimal software engineering. These high level environments allow a small number of specifically trained researchers with more affinity towards software tools to serve the needs of larger user communities. The created workflow allowed researchers at the University of Westminster UK to validate and fine tune their wet laboratory experiments and to provide useful feedback for the developers of the software tools regarding the appropriateness and usability of the products. Moreover, the developed environment can be efficiently used in teaching to allow students without specific grid related knowledge to utilise docking and molecular dynamics packages and to support the in-vitro experiments.

Keywords:

Grid portal, grid workflow, molecular docking, molecular dynamics, simulation

URL for further information:

https://engage.cpc.wmin.ac.uk

Justification for delivering demo and/or technical requirements (for demos):

The demonstration will show how the simulation workflows can be parameterized, customized, executed and visualized using the extended WS-P-GRADE portal. The demonstration requires internet connection

Demo Session 1, Welcome Drink / 140

Porting Scientific Application to the EDGeS Service Grid/Desktop Grid Platform

Authors: Daniel Farkas[1]; Gabor Szmetanko[1]; Peter Kacsuk[2]; Robert Lovas[2]; Tamas Kiss[1]

Co-authors: Attila Marosi [2]; Gabor Terstyanszky [1]; Gergely Sipos [2]

[1] *University of Westminster*

[2] *SZTAKI*

Corresponding Author: sipos@sztaki.hu

The EDGeS project connects the g-Lite based EGEE grid to several BOINC and XtremWeb based Desktop Grid system. The EDGeS infrastructure successfully extends EGEE with volunteer and institutional desktop resources to be utilized by master worker or parameter sweep applications. The project has successfully ported several EGEE applications to the combined platform, including the WISDOM meta-middleware for molecular docking simulations, the ISDEP plasma fusion application, or the VisIVO server tool for the visualization of astrophysical data.

Detailed analysis:

The aim of the EDGeS project is to support parameter sweep or master worker EGEE applications with a large number of Desktop Grid resources collected from individuals or from scientific institutions. As part of the project several EGEE applications has been made capable to utilize Desktop Grid resources and can now be executed on the EDGeS platform. The demonstration will show applications from several application domains, including bio-molecular simulations using the WISDOM middleware, plasma fusion simulations using ISDEP, astrophysical visualization using ViSIVO, and the EMMIL e-marketplace application. The applications are submitted from either the command-line g-Lite user interface or from the P-GRADE or WS-P-GRADE portals. The usage of Desktop Grid resources is completely transparent for the EGEE users. The execution utilizes public Desktop Grid resources of the SZTAKI Desktop Grid and the EDGeS@home DG that was specifically set up by the EDGeS project to support EGEE application, and also institutional desktop resources in the University of Westminster Local Desktop Grid.

Conclusions and Future Work:

The European EDGeS project created a bi-directional bridging system between EGEE and large public and institutional Desktop Grid infrastructures. The project has successfully ported several EGEE applications and supported EGEE user communities in utilizing this hybrid platform. The aim is to disseminate these results and to raise the awareness of further communities regarding the usability of these additional resources to run computationally intensive applications faster and more efficiently.

Impact:

The EDGeS project has successfully demonstrated that a large number of EGEE applications can efficiently use Desktop Grid resources to support computationally intensive parameter sweep tasks. The resources can be collected from volunteer individuals or can come from more secure institutional desktop PCs. The sustainability of such Desktop Grid systems is much simpler and initiated by the actual resource owner, compared to the sustainability of large computing clusters. Moreover, the Desktop Grid resources can speed up the parameter sweep Grid applications and support the better utilization of cluster based Grid resources for more specific tasks, for example to run MPI applications. The ported applications are used by large user communities of EGEE. These communities gained a substantial number of new resources and a more efficient application execution by utilizing EDGeS.

Keywords:

desktop grid, application porting, grid portal

URL for further information:

http://www.edges-grid.eu/

Justification for delivering demo and/or technical requirements (for demos):

The demonstration will show the execution of several EGEE applications on the EDGeS platform using the g-Lite command line interface and the P-GRADE Grid portal.

Demo Session 1, Welcome Drink / 141

Earthquake Disaster Mitigation on EUAsiaGrid

Authors: Eric Yen[1]; YiTao He[1]

Co-authors: Simon Lin [1]; WeiLong Ueng [1]

[1] *Academia Sinica Grid Computing*

Although earthquakes are not predictable at the moment, with the aid of accurate seismic wave propagation analysis, we could simulate the potential hazards at all distances from possible fault sources by understanding the seismic wave velocity structure of the earth and the rupture process during large earthquakes. With the integration ofa strong ground-motion sensor network, an earthquake data center and seismic wave propagation analysis over gLite e-Infrastructure, we could have much better knowledge on the impact and vulnerability of potential earthquake hazards.

Detailed analysis:

Regional integration of earthquake sesnor networks could aid in fast event reporting and accurate event data collection. Federation of earthquake data centesr entails consolidation and sharing of seismology knowledge. Capability building of seismic wave propagation analysis implies the predictability of potential hazard impacst. With the gLite infrastrucuture and the EUAsiaGrid collaboration framework, earth scientists from Taiwan, Vietnam, Philippine, Thailand etc. are working together to alleviate potential seismic threats by making use of Grid technologies. Internet-based integration of sensor networks will be adopted for better data management and integration purposes. Standardization of data contents, services and federation would be enforced by the virtual organization and EUAsiaGrid infrastructure. A cross continental e-infrastructure, based on EGEE and EUAsiaGrid, is established for seismic wave forward simulation and risk estimation. Both the computing challenge on seismic wave analysis among 5 European and Asian partners, and the data challenge for data center federation have been exercised and verified. An earthquake hazard map prototype will be implemented and evaluated .

Conclusions and Future Work:

What we have accomplished could become a reference model for a global infrastructure to accumulate and integrate relevant data sources systematically for all disaster mitigation phases. To ease the access to all the services based on users workflow and to retain the maximal flexibility, a Seismology Science Gateway integating data, computation, workflow, services and user communities would be implemented based on typical use cases. In the future, extension of the earthquake wave propagation to tsunami mitigation would be feasible once the user community support is in place.

Impact:

A pilot framework made of a gLite infrastructure supporting an earthquake sensor network, data center and also simulation of seismic wave propagation for earthquake disaster mitigation was implemented. The objective of this research is to evaluate e-Science anda Grid infrastructure for seismic wave propagation analysis, and to understand the possible impacts beforehand by quantitative seismic hazard assessment. Based on this collaboration, an accurate regional seismic wave propagation analysis model could be built by integrating more detailed topographical data from each participating countries. This is indeed the key to construct a useful and unique risk analysis for disaster mitigation. With the support of an e-Science Infrastructure, a substantial understanding of potential impacts at all distances from possible fault sources could be achieved relatively easy compared with quickly assembled ad-hoc efforts. Thus researchers are able to answer “what-if” questions by exploiting the possible impact of potential

seismic events in advance. Implementation of such an infrastructure would benefit mitigating earthquake hazards by improving risk assessment and risk management.

Keywords:

Earthquake Simulation, Seismology, Disaster Mitigation, Grid technologies

URL for further information:

http://www.twgrid.org

Regional Activities, International Projects and Collaborations / 143

Instruments and sensors supported by e-Infrastructure

Authors: Alexey Cheptsov[1]; Franco Davoli[2]; Marcin Plociennik[3]; Michael Schiffers[4]; Norbert Meyer[5]; Roberto Pugliese[6]; ioannis Liabotis[7]

[1] *HLRS*

[2] *CNIT*

[3] *PSNC*

[4] *LMU*

[5] *Poznan Supercomputing and Networking Center*

[6] *ELETRA*

[7] *GRNET*

Corresponding Author: meyer@man.poznan.pl

The introduction gives the major objectives and the idea of instrument and sensor remote steering, controlling and monitoring. The second part presents the user experiences and requirements following a proposition of the general architectural framework. The way of integrating instrumention with existing e-Infrastructure is especially important for the ESFRI projects, where most are using some equipment for data capturing and computing resources for further simulations.

Impact:

The architecture proposed by the DORII project allows to integrate various type of applications, which daily use instruments to collect data. The data are used for further analysis and simulations. The natural integration is to have a unified infrastructure including sensors, instruments, grid, visualisation and data repositories. This is permitted by the DORII framework. It is worth mentioning that nearly all ESFRI communities are using some instruments, and the goal of these communities is the better use of the existing infrastructure.

Keywords:

instrument, sensors, experimental science, workflow

URL for further information:

www.dorii.eu

Detailed analysis:

The presentation will explore the e-Infrastructure requirements of experimental scientists and will provide an overview of the current projects involved in the field. The major goal is to demonstrate a prototype system developed by the project DORII (www.dorii.eu). Although the results are addressed to scientists who are using expensive equipment, the authors plan to use the same approach to control small devices and sensors which are used in future Internet technologies. The most obvious advantage of this approach is that one can exploit the storage and processing capabilities that the classical data and computing ecosystem offers. By representing the instrument as a service and integrating it

with other services through well-understood protocols, it becomes straightforward to directly store experimental data to arbitrary locations world-wide, replicate them in multiple locations, and perform post-processing, that previously took days or months, in only a fraction of the time.

Conclusions and Future Work:

The DORII community is aiming in the near future to define some ESFRI related use cases, and to adopt the current architecture for their needs.

Poster session / 144

Grid preprocessing application for Environment monitoring in Moldova

Author: Veaceslav Sidorenco[1]

Co-authors: Alexandr Golubev [1]; Dorian Gorgan [2]; Peter Bogatencov [1]

[1] *RENAM*

[2] *Technical University of Cluj-Napoca*

Corresponding Author: svv@renam.md

The aim of the EnviMon application is to organize data acquisition and processing for Environment state monitoring. A nation-wide distributed set of sensors are polled from a central station placed at the main site of the State Hydrometeo Service of Moldova (SHMS). The application provides data collection, filtering, storage and processing in order to produce synthetic reports and input that can be used as structured tables for database or input for geoinformation systems maps. .

Detailed analysis:

The application is developing by a joint team of specialists from MD-Grid NGI (Moldova) in cooperation with specialists from Romania and Hungary participating in the EC funded SEEGRID-SCI project. The applications are using resources of the regional SEE-GRID infrastructure. The application can be used both autonomously and as part of more common GreenView grid oriented application, that also is developed in the framework of SEEGRID-SCI project. The aim of the GreenView application is a refinement of surface- and vegetation parameters in SEE region based on processing of satellite images. Construction, usage and comparison of diverse satellite datasets will be performed. High resolution satellite measurements can be used for numerous environmental studies (climate-related or air pollution modeling). Using the sophisticated environmental data the change of the vegetation distribution in the Carpathian Basin and its climate-related causes will be investigated.

Conclusions and Future Work:

The EnviMon application will be used for National and regional environment monitoring by State Hydeometeo services of Moldova and Romania.

Impact:

Data collected from ground stations and decrypted by EnviMon Client is transferred to the GRID clusters servers' infrastructure. Based on collected data from different regions the Grid servers generates reports and calculates different parameters such as:
• Vegetation index(NDVI);
• leaf- area index (LAI);
• Index of green and dry biomass;
• Some other parameters
These data are compared with satellite data using a common GreenView application.

Keywords:

Regional Grid infrastructure; environment monitoring; satellite images; GreenView application

URL for further information:

www.renam.md

Earth Science / **145**

ES cluster towards ES VRC

Authors: Horst Schwichtenberg[1]; Monique Petitdidier[2]

[1] *SCAI Fraunhofer Institute*

[2] *CNRS*

Corresponding Author: horst.schwichtenberg@scai.fraunhofer.de

The ES community has actively participated in the EGEE projects as Strategic Discipline Clusters. In the framework of the DCI infrastructure and its related services deployed by the European Grid Initiative (EGI), the ES VRC will gather all ES partners to bring their contribution, i.e. resource and service providers, developers, support teams as well as end-users. This network of support and expertise is constituted within the ES VRC by a kernel of partners with Grid expertise in order to break down the entry barrier and overcome some of the technical difficulties that otherwise exist.

Detailed analysis:

The worldwide Earth science community, with its mosaic of disciplines and players (academia, industry, national surveys, international organizations, and so forth), provides a scientific basis for addressing societal issues. These capabilities require that the Earth science community utilize, both in real and remote time, massive amounts of data, which are usually distributed among many different organizations and data centers. It explains the interest of this community for Grid technology, the variety of applications ported as well as the tools developed. Besides the participation in EGEE through the ES cluster, other projects involving ES disciplines were or have been carried out as related projects to EGEE such as CYCLOPS, SEEGrid, EELA2. Projects outside of the context of EGEE are e.g., in the framework of WGISS/CEOS.
Partners from these projects have shown a strong scientific and technical interest to work together in the ES VRC. The transition from a central support-based model to a model driven by each VRC is ongoing and will lead to sustainable Grid infrastructure and VRCs. In order to achieve this transition, the ES VRC joins the A&A and Fusion VRCs in a common EU proposal

Conclusions and Future Work:

The transition has started with EGEE-III The sharing of experience in the framework of dedicated workshops is important to set up more concretely collaborations and specific objectives. New institutions especially from east Europe are included to the ES VRC. Today's ES VO is being opened to partners with less or no hardware resource. A knowledge database will be extracted from the ES-Cluster and continued after EGEE-III and the independent Website of the ES cluster will be extended to a Science Gateway. The ES VRC will organize and provide old and new services within the community.

Impact:

The ES VRC consists of 9 Virtual Organizations with partners in 23 European or associated countries and links to EELA2, EUAsia and African partners. This structure will permit to extend the Grid knowledge within the ES VRC and increase the collaboration among the communities, by sharing expertise. The transition period is focused on building a self-reliant and sustainable VRC. The goal is to increase the number of scientists that make use of the infrastructure through focused training and dissemination. To expand and support Grid uptake within ESFRI other related community scientific projects are important too. In this new organizational model, ES VRC will support its user community for domain specific applications, techniques, data sets and other non conventional Grid resources (e.g. instruments and network sensors), and will safeguard its expertise into a knowledge repository. ES VRC has to fulfill the requirements addressed by its user community; this could trigger the development of new components or interfaces for specific domains, and the re-design and implementation of tools and services related to the middleware evolution.

Keywords:

Earthscience VRC EGI Prace

URL for further information:

euearthsciencegrid.org

Novel Architectures and Technologies / 147

Managing Healthcare and Medical Information Utilizing Cloud Computing

Author: Charalampos Doukas[1]

Co-authors: Aristotle Chatziioannou [2]; Ilias Maglogiannis [3]

[1] *University of the Aegean*

[2] *National Hellenic Research Foundation*

[3] *University of Central Greece*

Corresponding Author: doukas@aegean.gr

Cloud Computing provides functionality for managing information data in a distributed, ubiquitous and pervasive manner supporting several platforms, systems and applications. This work presents the implementation of a mobile system that enables electronic healthcare data storage, update and retrieval using Cloud Computing. The mobile application is developed using Google's Android operating system and provides management of patient health records and medical images (supporting DICOM format and JPEG2000 coding).

Detailed analysis:

Cloud Computing provides the facility to access shared resources and common infrastructure in a ubiquitous and pervasive manner, offering services on demand over the network to perform operations that meet changing needs in electronic healthcare application. The location of physical resources and devices being accessed are typically not known to the end user. It also provides facilities for users to develop, deploy and manage their applications 'on the cloud', which entails virtualization of resources that maintains and manages itself.
Pervasive healthcare systems focus towards achieving two specific goals: the availability of e-health applications and medical information anywhere and anytime and the invisibility of computing. Applications and interfaces that are able to automatically process data provided by medical devices and sensors, exchange knowledge and make intelligent decisions in a given context are strongly desirable. Natural user interactions with such applications are based on autonomy, avoiding the need for the user to control every action, and adaptivity, so that they are contextualized and personalized, delivering the right information and decision any time.

Conclusions and Future Work:

The sharing of medical information resources (electronic health data and corresponding processing applications) is a key factor playing an important role towards the successful adoption of pervasive healthcare systems. Moreover, due to the mobility of the patients and the medical personnel, healthcare networks are increasingly equipped with capabilities to share healthcare-related information among the various actors of electronic health. In this context the concept of Cloud Computing will attract the interest of scientists and developers working in the field of biomedical informatics.

Impact:

In this context a pervasive healthcare information management system for mobile devices utilizing Cloud Computing and Android Operating System has been developed. The prevalent functionality of the application is to provide medical experts and patients with a mobile user interface for managing healthcare information. The latter interprets into storing, querying and retrieving medical images (e.g., CT scans, MRIs, US etc.), patient health records and patient-related medical data (e.g., biosignals). The data may reside at a distributed Cloud Storage facility, initially uploaded/stored by medical personnel through a Hospital Information System (HIS). In order to be interoperable with a variety of Cloud

Computing infrastructures, the communication and data exchange has to be performed through non-proprietary, open and interoperable communication standards. Utilizing Web Services connectivity and Android OS supports the following functionality: Seamless connection to Cloud Computing storage utilizing Web Services, Patient Health Record Management, DICOM image viewing support, JPEG2000 viewing support, Image annotation support and proper user authentication and data encryption.

Keywords:

Cloud Computing, pervasive healthcare management, mobile interfaces

URL for further information:

not available yet

DCI programming / **148**

Providing uniform and standardised access to European Grid middleware with SAGA

Authors: Andre Merzky[1]; Ole Christian Weidner[1]; Shantenu Jha[1]; Steve Fisher[2]

[1] *Louisiana State University*

[2] *Science and Technology Facilities Council*

Corresponding Authors: sjha@cct.lsu.edu, ole.christian.weidner@cern.ch

The Simple API for Grid Applications (SAGA), has demonstrated its ability to provide a simplified and uniform multi-language access layer to heterogeneous grid middleware stacks. SAGA already provides support for Unicore and gLite, and extensions to ARC via OGSA-BES as well as generic DRMAA support are already being developed or planned. First SAGA application integration tests within EGEE have been successful.

Detailed analysis:

LSU's C++/Python implementation of OGF's SAGA standard (GFD-R-P.90) has been used successfully as a unified Grid access layer for scientific applications running on the TeraGrid (Globus, Condor, LSF, PBS), UK-NGS (GridSAM) and cloud infrastructures (EC2, Nimbus, Eucalyptus). Our current focus is on the development of extensions (adaptors) that enable SAGA to interface to three of the major European grid middleware projects: gLite, Unicore and ARC. Unicore access through GridSAM and parts of the gLite stack are already available. Adaptors for gLite have been developed as implementations of the SAGA Service Discovery extension (GFD-R-P.144). Support for further gLite components as well as a generic OGSA-BES extension that will allow native access to both Unicore and ARC are currently under development. To demonstrate viability of this effort, SAGA has been integrated with Ganga and Diane and successfully executed Lattice-QCD jobs simultaneously on EGEE and TeraGrid infrastructures.

Conclusions and Future Work:

We believe that our current and future work will align well with future European efforts in a sense that it provides the right technology with the right level of abstraction that is required to leverage the development of truly distributed and portable applications in an expanding heterogeneous environment. In order to support a seamless integration, our future work will focus on additional bindings to gLite (WMS & CREAM) as well as native OGSA-BES and DRMAA bindings.

Impact:

Although this work is still under development, its long-term impact on the European middleware can be inferred from our on-going work that currently fosters interoperability between the US TeraGrid, UK-NGS and NAREGI, as well as from first EGEE application integration tests. We have shown that using SAGA as a high-level middleware access layer not only makes application development and tooling easier (and faster), but also naturally results in portability and interoperability on a local, national and international level. We have demonstrated how a simple SAGA plugin enables Ganga/Diane to not only execute jobs on local CERN infrastructure, but also on the TeraGrid - an infrastructure that previously

wasn't accessible through Ganga/Diane. Similar advances are being made with the UNICORE-based DEISA infrastructure. Extending further (e.g. to NAREGI and ARC), requires developing the appropriate SAGA adaptors and does not require modifying the application.

Keywords:

SAGA OGF standard API interoperability

URL for further information:

http://saga.cct.lsu.edu/

Regional Activities, International Projects and Collaborations / **149**

Building Scientific Workflows for the Fisheries and Aquaculture Management Community based on Virtual Research Environments

Author: Pasquale Pagano[1]

Co-authors: Leonardo Candela [1]; Pedro Andrade [2]

[1] *CNR-ISTI*

[2] *CERN*

Corresponding Author: pasquale.pagano@isti.cnr.it

The management of resources related to the Fisheries and Aquaculture (FARM) domain is complex and involves different scientific activities. The people involved in the domain are distributed worldwide: scientists in the field, regional statistics departments, national governing bodies, etc. Distributed information must be produced, analyzed, processed, shared, and preserved by all multiple actors through a number of scientific workflows. The D4Science-II project provides a number of Virtual Research Environments (VREs) to address the collaboration needs of this community.

Detailed analysis:

VREs are integrated environments providing seamless access to resources as well as facilities for communication, collaboration and interaction among scientists and researchers. D4Science supports the building of VREs by dynamically aggregating the needed constituents through a controlled on-demand hiring. The Food and Agriculture Organization (FAO) of the United Nations and the WorldFish Center are currently exploiting the D4Science production e-Infrastructure through a number of VREs. The AquaMaps VRE provides a map-generation environment where species distribution maps are produced by using a choice of algorithms and datasets. The ICIS VRE is a repository of statistical datasets that can be 'sliced and diced' to generate table-objects and/or graphs. The FCPPS VRE provides report templates that can be used to create country profile reports. By bringing these environments closer and enabling seamless exchange of information across VREs, complex scientific workflows can be established. For example a reporting environment (FCPPS) can present maps from large data collections (AquaMaps) and create objects such as graphs from statistical information (ICIS).

Conclusions and Future Work:

The services FAO and WorldFish are provided with an excellent example of a synergic use of D4Science VREs and EGEE technologies. Scientists have access to innovative research environments, well integrated each other, facilitating data sharing and the execution of scientific workflows.
The expansion of such integrated environment to external communities, not belonging to the same Virtual Organizations, is currently the major objective of the D4Science infrastructure.

Impact:

D4Science proposes a different approach consisting in the definition of Virtual Research Environments replying to several requirements for confidentiality, sharing, and collaboration of scientific workflows. Objects produced or retrieved in one environment are easily consumable in another in the context of

a trusted sharing environment. These environments support the re-utilization of information to be shared, simplify data loading (load once, use everywhere), and simplify cross VRE-data-analysis and presentation.
On another hand, VREs control the sharing of resources across the members of a VO allowing a subset of VO users to consume a subset of the VO resources easily and for a limited timeframe. VREs are created on demand without additional costs on Site and VO managers.
Furthermore, a number of the tasks to be executed under these environments require access to a large computational capacity such as the one offered by the EGEE production infrastructure. The support from EGEE production sites to the D4Science VO brings such benefit and enhances the computational power of the FARM VREs.

Keywords:

Virtual Research Environments, Scientific Workflows, Virtual Laboratories

URL for further information:

http://www.d4science.eu

Workflow Management / **150**

Workflows Description, Enactment and Monitoring through SAGA

Author: Ashiq Anjum[1]

Co-authors: Andre Merzky [2]; Irfan Habib [3]; Richard McClatchey [1]; Shantenu Jha [2]

[1] *UWE Bristol, UK*

[2] *Louisiana State University, USA*

[3] *UWE Bristol,UK*

Corresponding Author: ashiq.anjum@cern.ch

Most existing workflow systems are tightly integrated with middleware, which limit wide scale adoption and may limit efficient execution of workflows. Whereas the use of SAGA enables the use of multiple heterogeneous resources, there exists an initial but somewhat limited support for workflows using the SAGA specification. Thus to reconcile the advantages of the SAGA paradigm with broad usage of workflows as a way of composing applications, this project investigates approaches that extend the SAGA-workflow package (Digedag)to support workflows such as neuroimaging analysis in neugrid.

Detailed analysis:

Digedag will be extended to support the efficient execution of broad range of applications such as nueroimaging in neugrid, in an extensible,flexible and robust manner. Currently digedag does not have support to coordinate with a EGEE workload management system to schedule a series/sequence of jobs as per the requirements of a workflow. This limits the effective management, monitoring and executions of workflows. The proposed system will implement the following in the Digedag workflow engine:

1. Users may write workflows using the SAGA workflow package which can be enacted and scheduled using appropriate middleware adapters.
2. The submitted workflows can be managed and monitored as a single entity.
3. SAGA workflow package will take care of all the dependencies and execution will be subject to user preferences and dependencies
4. The sequence of jobs and order will remain intact even if a workflow is distributed across sites for execution.

Although we will initially establish and demonstrate these benefits in the context of neugrid based neuroimaging, the advantages will extend to a range of DAG-based and other workflow applications, thus impacting a number of communities.

Conclusions and Future Work:

Using neuroimaging as an exemplar, SAGA will provide support for workflow enactment. Users will translate workflow descriptions into a SAGA API, which can then be enacted using the appropriate execution engine and scheduled, distributed and executed in multi-site environments.

Impact:

The proposed solution will reduce job submission and scheduling latencies. The user will not have to break the workflow into jobs and then submit and schedule them individually. The user will write the workflows in a workflow authoring environment which will be translated into appropriate SAGA structures. Users will also be able to monitor and manage hundreds of jobs in a workflow as a single entity. This will help the users to keep track of execution, failures and outcome of dependencies. Scheduling will also become effective as a higher level enactor, Digedag, will coordinate with a Grid wide scheduler to plan and schedule the jobs.

As a consequence of the proposed project, user workflows can be developed programmatically, and can be enacted to use any available middleware without restricting the specification to a particular execution engine. This flexibility and extensibility along with robustness is critical to enable applications to utilize infrastructure at scale. Users will be able to get features such as “write once and run anywhere”.

Keywords:

SAGA, Workflow Description, Monitoring and Enactment

URL for further information:

http://forge.gridforum.org/sf/projects/saga-rg

Novel Architectures and Technologies / **151**

EDGeS infrastructure for the EGEE user community

Authors: Peter Kacsuk[1]; Robert Lovas[1]

[1] *SZTAKI*

Corresponding Author: rlovas@sztaki.hu

There is a strong collaboration between EGEE and EDGeS in order to extend the EGEE infrastructure with volunteer and institutional desktop grids (DG) and to support EGEE users to migrate their application to the EDGeS infrastructure. The talk explains to EGEE users how this integrated infrastructure works, what the benefits are and how their applications can be ported and run on this infrastructure. The talk also explains how individual VOs can extend their VO resources with connected DGs. Experiences in related projects (IberCivis, EELA) using EDGeS technology will be presented.

Detailed analysis:

EDGeS has created a bridging technology by which service grids (SG) and desktop grids (DG) can be interconnected to both in SG->DG and DG->SG direction. The first option enables that EGEE users can submit large parameter sweep applications from EGEE VOs into connected DG systems. In this way any EGEE VO can be easily extended with additional local DGs organized from the existing desktops of the participating institutes. The other direction enables for the DGs to send work units to the EDGeS VO that was established by EDGeS in order to support connected DG systems. In this way university level DG systems can gain additional resources from EGEE VOs. The talk will explain all these possibilities and how the EGEE user community can take advantage of this new heterogeneous grid technology that can interconnect SGs and DGs. Particularly, the application porting method will be explained by

which existing EGEE applications can be ported to the EDGeS infrastructure. Already more than 20 applications have been ported to EDGeS from many different application areas (physics, bioinformatics, engineering, image processing, etc.)

Conclusions and Future Work:

EDGeS contributes to the significant extension of available EGEE resources by extending gLite grids with volunteer and institutional DG systems. EDGeS created a production infrastructure that enables the submission of large parameter sweep jobs from gLite VOs into connected desktop grids. Future work will include the extension of ARC and Unicore grids with DG systems as well as support for data-intensive applications.

Impact:

EDGeS made a significant impact not only in Europe but world-wide. First of all the 3G Bridge solution of EDGeS by which service grids and desktop grids can be interconnected has successfully been adopted by the EELA-2 project in order to interconnect gLite and OurGrid. Furthermore, the EDGeS VO supports several large desktop grid projects like IberCivis, SETI@home, AlmereGrid, etc. Several EGEE user communities (e.g. fusion, WISDOM, etc.) have successfully ported or adapted their applications to EDGeS. The EDGeS technology was successfully been used by several companies (Atos Origin and Correlation Systems). National grids have adopted the EDGeS technology (Kazakh National Grid) or they have been considering its adoption (Russian, Taiwan and Ukrainian Grid).

Keywords:

service grids, desktop grids, grid interoperability, application porting

URL for further information:

http://www.edges-grid.eu/

Fusion / **152**

Overview of the grid fusion applications and the possible workflows among them

Author: Francisco Castejón[1]

Co-authors: Andrés Bustos [1]; Antonio Gómez-Iglesis [1]; Antonio Juan Rubio [1]; Daniel López-Bruna [1]; Esther Montes [1]; Igor Semenov [2]; José Luis Velasco [1]; José Luis Vázquez-Poletti [3]; Manuel Aurelio Rodríguez [1]; Marcin Plociennik [4]; Marcus Hardt [5]; Miguel Cárdenas [6]; Nikolai Marusov [2]; Rafael Mayo [1]; Rainer Stotzka [5]; Álvaro Cappa [1]

[1] *CIEMAT*

[2] *Kurchatov*

[3] *UPM*

[4] *PSNC*

[5] *KIT*

[6] *CIEAMT*

Corresponding Author: francisco.castejon@ciemat.es

Grid computing developments have increased substantially in fusion research.
A great variety of applications have been ported to the grid in the frame of EGEE, EELA and EUFORIA projects. These applications belong to different plasma physics domains and different strategies have been used to port them, depending on the nature of the applications. These codes are ready to establish workflows among them in order to produce new relevant physics results that could not be achieved running the applications separately

Impact:

The grid-ported applications have been chosen to play a role in the different fields of fusion research in order that the grid produces relevant results for different fusion subcommunities. Moreover, different tools have used to port the applications in order to take advantage of the developments that have been performed by the grid scientists.
The onset of these applications on the grid architecture allows the establishment of complex workflows between both grid and HPCs applications, allowing the researchers to join different physical models.
Several new relevant scientific results have been obtained using the grid, which opens the window to new research activities. All these developments have had a demonstration effect that shows that grid computing can be useful in a large variety of research activities. But, finally, the main impact of this work is that an increasing number of fusion scientists is using the grid for their customary research.

Keywords:

Fusion, Workflows, EUFORIA, EGEE, EELA

URL for further information:

http://www-fusion.ciemat.es

Detailed analysis:

Plasma Physics research is composed of a large variety of problems. Many branches of Physics are at work in fusion research and, therefore, a large variety of applications is used by fusion scientist, many of them suitable for running on distributed architectures like the grid.
In this work, a review of the fusion applications running in the grid is presented, taking into account their different structures and their different applications on Plasma Physics. We have ported Monte Carlos codes, parameter scan applications, and parallel MPI-based ones. To accomplish this task, several tools and pieces of middleware have been used, including both standard tools and ad-hoc developments. Some instances of the standard tools are Kepler, gLite, Gridway, MPI, and DRMAA. While examples of developed applications are START, TAPAS4GRID and VASHRA-T.
Several possible complex workflows among these applications are under development, showing that the grid is flexible enough to perform such workflows. Kepler is the standard workflow engine used among the fusion community for establishing workflows among different architectures.

Conclusions and Future Work:

A large variety of applications ported to the grid in the frame of different projects has been presented in this work, showing the use of different techniques for the grid-porting. The applications have been chosen to show that the grid can play a role in different research fields, opening the window to new developments and new research activities. Several possible workflows have been also shown, as a preliminary result of one of the main future activities in this direction, since complex workflows among fusion applications will allow the scientist to relate different physics models.

Poster session / 153

Configuration and optimization for a gLite grid computing farm.

Authors: Giacinto Donvito[1]; Giorgio Pietro Maggi[2]; Vincenzo Spinoso[1]

[1] *INFN-Bari*

[2] *INFN and Politecnico di Bari*

Corresponding Author: giacinto.donvito@cern.ch

We will illustrate the optimization work performed on a mid-sized gLite grid-computing farm, testing various solutions, to improve the services offered to both grid and local users. Particular attention was devoted to the storage infrastructure and its configuration trying to improve the overall farm performance, reliability and to minimize the manpower required for its management. Improvements were also pursued by optimizing the configuration of the batch system. Experiences from both LHC users and non LHC ones will be reported.

Detailed analysis:

Details on the tests performed and the results achieved while attempting to improve the performance of a typical tier2 farm while serving both LHC and non LHC VO's. Important improvements on the storage performance have been achieved through a patient optimization of several layers starting from the hardware to the storage manager software. In particular the performance and the functionalities obtained using two different storage manager solutions, dCache and StoRM/Lustre, will be presented.
Requirements coming from the different users groups' force also an optimization of batch system, Torque/MAUI, in order to allow the final user to exploit the farm resources with grid, local, interactive and MPI jobs. New users also require a clear framework in order to simplify the porting of new applications to grid. Examples of applications ported (coming from different community: theoretical physics, astro-physics, bioinformatics, etc) and the performance achieved will be provided.
We will report also about the performance and the scaling of MPICH based applications, and on experience about using Worker Nodes of the batch farm in order to support interactive sessions of the end users.

Conclusions and Future Work:

Through this infrastructure the users could start achieving know how on how to submit their application to a batch farm. This is often the first step needed to help users to migrate their application to a distributed grid infrastructure such as EGI production grid, as soon as the workload for a given application increases and could not be executed within a single batch farm.

Impact:

We will show how this configuration could improve the overall usage of the available CPUs in a typical batch farm.

In this work we will give also some short information about the scientific results obtained.

The work spent into optimizing the farm configuration and performance will improve the user satisfaction and increase the number of researchers that could solve their problems by means of using such computational infrastructure and will allow them to solve problem that could not be solved with typical desktop computers.

Keywords:

gLite, batch farm, LHC, MPI, storage, Lustre, dCache

URL for further information:

http://webcms.ba.infn.it/cms-software/index.html/index.php/Main/GridBariFarm

Poster session / 154

Maintaining Corpus Download Facilities with Changing Technology

Authors: Daan Broeder[1]; Mischa Sallé[2]; Tobias van Valkenhoef[1]

[1] *Max-Planck Institute for Psycholinguistics*

[2] *NIKHEF*

Corresponding Authors: dieter.vanuytvanck@mpi.nl, daan.broeder@mpi.nl

Introduction of federated login using Shibboleth for a linguistic data archive created a problem by making existing local tools for downloading data-sets stop working. To address this problem the use certificate based authentication in combination with a SLCS service was setup. The application domain demanded that special attention was paid to shielding the user from the complexity of working with certificates and making the SLCS handshake as easy as possible. Parallel access to the archive using both Shibboleth and client certificates required careful configuration.

Detailed analysis:

At the Max-Planck Institute for Psycholinguistics (MPI) [1] a large set of linguistic corpora from MPI researchers and also external projects is archived and made on-line accessible. Already some years ago, also within the linguistic community, we saw the need to integrate the existing archiving infrastructures, since this would be the basis for a viable e-Science infrastructure for the linguistic domain. A small EU project of four archiving institutions "DAM-LR" [2] was created that aimed at integrating the archives at different levels, including AAI. A requirement was formulated for federated login for all the archives' users such that for instance SSO for distributed collections would become possible. To realize this, the DAM-LR identity federation was created and the archives Shibbolized their web servers. A consequence of this was that local tools that had been used to copy data-sets from the archives to local storage stopped to function since Shibboleth only addresses access by web browsers. This was an unsatisfactory situation especially since further big EU integration projects in the linguistic domain like CLARIN [3] all plan to use federated login and Shibboleth.

Conclusions and Future Work:

Currently we have a working application, and the developed software can also be used to enable other tools to access resources in a similar way. Applying this in a EU wide context like the CLARIN project, imposes the question of the status of the SLCS service. Should it be organized on the basis of national NRENs, on a EU wide basis or as part of a virtual organization platform? This is of course a question that may be of interest for similar integration projects.

Impact:

Together with the BiG Grid project [4] and SURFnet [5], a project was setup to test if the use of SLCS obtained certificates [6] could be a solution for this problem. In this context, SURFnet set up a SLCS service, accessible by the members of the SURFfed identity federation. Secondly the MPI's repository apache server was configured with mod_ssl and mod_rewrite to allow client certificate-based authentication in parallel with Shibboleth based authentication. Thirdly, the "IMDI-Browser" a local tool that was originally used to download data sets from archives running the LAT archiving software was modified to perform a handshake with the SLCS to obtain the certificate and use it to download the items of a data set.

[1] MPI archive, http://mpi.nl/resources/data
[2] DAM-LR, http://www.mpi.nl/DAM-LR
[3] CLARIN, http://www.clarin.eu
[4] BiG Grid, http://www.biggrid.nl
[5] SURFnet, http://www.surfnet.nl/
[6] SLCS, SWITCH SLCS service, http://www.switch.ch/grid/slcs/

Keywords:

Authentication, Shibboleth, SLCS, archiving, linguistics

URL for further information:

http://www.nikhef.nl/pub/projects/grid/gridwiki/index.php/User:Msalle

Demo Session 1, Welcome Drink / 155

Cancer Sentinel project: a grid network for distributed medical data management on Cancer.

Author: Paul De Vlieger[1]

Co-authors: David MANSET [2]; Jean-Yves BOIRE [3]; Jérôme REVILLARD [2]; Lydia MAIGNE [1]; Vincent BRETON [1]

[1] *LPC CNRS/IN2P3*

[2] *maat-Gknowledge*

[3] *ERIM*

Corresponding Author: vlieger@clermont.in2p3.fr

The Sentinel project involves several medical actors related to cancer: screening structures, medical laboratories and both regional and national public health authorities.
The project builds upon grid technologies to create a federation of medical data sources related to cancer. The main purpose of the project is to enable secured medical data exchanges between cancer screening structures and cancer analysis (pathology) laboratories. The architecture and tools used to deploy such a network are also relevant to distributed medical images diagnosis, global health and epidemiology.

Detailed analysis:

The French national programs for early cancer diagnosis (breast, colorectal, cervical) is carried by associations which are in charge of inviting a targeted population to be screened. In case of positive result, a detailed medical report is created and registered in a local database.
In principle, these reports should be transferred to the screening structure. However, due to data ownership issues, there are no electronic exchanges. The grid technology is particularly well fitted to address this issue as a grid can federate data sources and provide a secured framework where patient data are stored in the laboratory and made available in a secured way to authenticated external users.
The grid security mechanisms allow fine-grained management of rights: a data holder can allow a customer to access and query his databases. Of course, users must beforehand be registered and trusted by the virtual organization. Therefore, providers keep the complete control of their data as nothing is massively extracted from the medical structure. Moreover, the data sources are always up to date, offering new opportunities for (near) real-time data analysis.

Conclusions and Future Work:

This project aims at proving the feasibility and the reliability of a grid-based surveillance network for cancer screening using grid technology developed within the EGEE and AuverGrid projects. The network infrastructure will be enlarged in a near future to host a large panel of laboratories, medical structures and public health institutes in order to perform epidemiological statistics on cancer.
The deployment of the grid network started in June 2009 and the first prototype was released in December 2009. The extension to medical images (mammography) is scheduled for 2010.

Impact:

The development started in mid-2009 in collaboration with Maat-Gknowledge and a prototype was released in December. The grid network is currently deployed with a dedicated Virtual Organization and its own specific grid services. the first objective was to offer an access to electronic pathology reports for cancer screening associations.
In a second step, the national public health could access to the medical data in order to produce epidemiological statistics on cancer incidence in Auvergne and potentially nationwide if the grid is extended beyond Auvergne.
The grid security framework has been modified to accept health-professional smartcard certificates to authenticate users on the grid.
Confidentiality and patient protection are critical to fulfill legal requirements on data privacy. The central issue consists in correctly mapping a patient identity through the network, avoiding false identification and offering a good quality of data linkage. For this purpose, a new distributed way to identify a patient using pseudonymisation techniques and data mining mechanisms is currently being tested over the network.

Keywords:

Surveillance network, cancer, database federation, epidemiology, data linkage, security, smartcard

URL for further information:

www.e-sentinelle.org

Justification for delivering demo and/or technical requirements (for demos):

The grid infrastructure is now operational between three main actors of cancer screening in the French Auvergne region. No special requirements: (Internet access)

Bioinformatics / 156

Job submission tool: bioinformatic application porting on the gLite grid Infrastructure.

Authors: Giacinto Donvito[1]; Giorgio Pietro Maggi[2]; Guido Cuscela[1]

[1] *INFN-Bari*

[2] *INFN and Politecnico di Bari*

Corresponding Author: giacinto.donvito@cern.ch

The Job Submitting Tool provides a solution for the submission of a large number of jobs to the grid in an unattended way. Indeed the tool is able to manage the grid submission, bookkeeping and resubmission of failed jobs.
It also allows the monitor in real time of the status of each job using the same framework. In this work we will introduce same key new features and application that we have added to this tool. In the work several already executed challenges will be reported together with a logical description

Detailed analysis:

In this work we will present the status of development of the Job Submission Tool both in terms of new functionalities and new applications ported within this framework.
The main new feature added are: automatic distribution of input data using standard gLite Data Management Tools, the possibility to exploit the dependency between tasks belonging to the same run, the possibility to create on the fly the needed task from both existing portals or application running over the grid (without any human intervention).
The Job Submission Tool is an easy and highly customizable framework that speeds up the process of porting new application to the grid glite infrastructure.
The applications submitted using this framework could be executed both using robot certificate or user certificates.
In the work we will present also the portal built on top of the framework in order to allow final users to submit large collection of jobs by filling up a small number of web form.
We will also describe the auto-adaptive submission algorithm that could guarantee to distribute the load of the calculation over all the farm available within the grid infrastructure.

Conclusions and Future Work:

This framework for submitting control and managing the jobs over the glite grid infrastructure could sensibly reduce the man power needed to run large challenges. With the new features it allows the final user to deal easily with large input files anda huge collection of output file in a unattended way. It was mainly tested with bioinformatics applications but it surely could work with other applications coming from different science, thanks to it flexibility and to the possibility to be customized to meet the end users needs both for serial and for workflow driven applications.

Impact:

Using this tool the end user could easy create a collection made up a large number of jobs, where the complete run could take months to be executed also in a large grid infrastructure like EGEE .
In this case one needs to solve problems like: detecting application failures, resubmit the failed jobs, collect the output back. Moreover in some bioinformatics applications the input files are fairly large, and this could easily become the main bottleneck if those files are available on one or a few sites. In this case it could be difficult to know where the files should be replicated as it may depend on where the CPU are available. In those cases it could be interesting to have a procedure to move the data where the CPU are free at a given time, and for the framework do this automatically without human intervention, in order to run on several farms, using data locally available.
Moreover as the number of the jobs needed to finish the whole run could easily be several thousands it is useful to be sure that the submission procedure does not overload a single or few sites and that the jobs are spread among the largest number of available sites.

Keywords:

Job Submission, Workflow, gLite, bioinformatics, Data Management

URL for further information:

http://webcms.ba.infn.it/cms-software/index.html/index.php/Main/JobSubmissionTool

Astronomy and Astrophysics / 157

CTACG - Cherenkov Telescope Array Computing Grid

Author: Giovanni Lamanna[1]

Co-authors: Cecile Barbier [1]; Thierry Leflour [1]

[1] *LAPP*

Corresponding Author: lamanna@lapp.in2p3.fr

Gamma-ray astronomy is one of the most active topics in astroparticle physics involving both particle physicists and astrophysicists leading to the design of new types of research infrastructures. The Cherenkov Telescope Array - CTA - is a proposed new project for ground based gamma-ray astronomy. This communication aims at providing a report of the most relevant activity carried out and in progress within the CTA consortium in EGEE and a view of perspectives and implications of ICT-based infrastructures such as the future EGI for the needs of a project like CTA.

Conclusions and Future Work:

CTA is proposed as a next-generation open observatory for very high energy gamma rays. The high data rate together with the large computing power requirements for data analysis demand dedicated resources, thus EGEE-Grid infrastructures for distributed data storage, data analysis and data access are considered the most efficient solution for the CTA e-infrastructure. The CTACG project aims to deploy end-to-end services in support of a wide scientific community for open data access for CTA.

Detailed analysis:

The CTA Computing Grid (CTACG) project aims at optimizing the application of Grid technology for the CTA simulation, data processing and storage, offline analysis and the Virtual Observatory interface through a dedicated global CTA EGEE Virtual Organization. Positive experiences have already been achieved through applications of Grid technologies in the context of services for distributed computing resources exploitation for Monte Carlo studies. The main issues inherent to the observatory work flow, which could benefit of Grid applications and which concern the CTACG project are: Monte Carlo simulations; Data flow, data transfer and storage; Data reduction, data analysis and open access.
This communication will report mainly about recent experiences in scheduling, executing and managing Grid jobs for massive Monte Carlo simulations as well as improving the distributed analysis of the simulated data. The development of a web-browser application to provide a CTA-Monte Carlo-Dashboard running a CTA dedicated MySQL database and the implementation of applications oriented to all different levels of user requirements are the main first achievements of the CTACG project.

Impact:

Intensive Monte Carlo simulations are an essential tool in order to optimize the best configuration and performance of the CTA installation. Even exploring only a few possible configurations, simulated CPU time consumption for that purpose is on the order of 10 CPU years per configuration. Despite many tricks that improve the efficiency of these simulations the needed memory storage cannot be reduced to less than 200-300 TB. A large number of degrees of freedom to explore as well as different levels of specific needs in the analysis process demand an efficient worldwide access to simulated output files and computing resources for MC data analysis to all members of the CTA consortium. The need to fulfill such requirements has motivated the Grid-approach. Furthermore a dedicated database and a Monte Carlo dashboard are two important services implemented to strength the adoption of Grid solutions for a part of the scientific community never involved before in an EGEE virtual organization. This communication would like to tell about the successful impact of the Grid application in CTA, one future research infrastructure in the ESFRI roadmap.

Keywords:

Gamma-rays; CTA; Monte Carlo; Dashboarc

URL for further information:

http://lappwiki01.in2p3.fr/CTA-FR/doku.php?id=cta-computing-grid-public

Bioinformatics / 158

Grid-based International Network for Flu Observation

Authors: Ana Lucia DA COSTA[1]; Tung DOAN TRUNG[2]

Co-authors: Aurélien BERNARD [3]; Quang Nguyen Hong [2]; Thanh-hoa LE [4]; Vincent BRETON [3]; Yannick LEGRE [5]

[1] *HealthGrid*

[2] *IFI, Vietnam*

[3] *CNRS-IN2P3, France*

[4] *Institute of Biotechnology, Vietnam*

[5] *HeallthGrid, France*

Corresponding Author: ana.dacosta@healthgrid.org

Since the H1N1 outbreak lately, there has been a worldwide effort to isolate and sequence flu virus genomes. Specimens with a positive result are sequenced and deposited in influenza databases. The present EUAsiaGrid application, called g-INFO (Grid-based International Network for Flu Observation), shows the integration of existing data sources towards a global surveillance network for molecular epidemiology, based on Service Oriented Architecture and Grid technologies. Its relevance is being tested through the current H1N1 outbreak.

Impact:

Results are made available to the research community in the corresponding website: http://g-info.healthgrid.org/, providing a real identity card of the concerned virus strains. Thanks to the molecular specificities highlighted (site for protease cleavage, glycosylation sites, epitopes and binding site), experts have in their possession promptly elements allowing them to take the most appropriate decisions relating to the transmission and the geographical expansion of the epidemic.

Keywords:

flu, surveillance network, grid, Service Oriented Architecture, epidemiology

URL for further information:

http://g-info.healthgrid.org/

Detailed analysis:

The current prototype is using the NCBI database (National Center for Biotechnology Information). Everyday the NCBI-FTP server is updated with new sequences of H1N1 segments, with 7 files: sequences of nucleotide, protein and coding region and corresponding metadata. A grid database (AMGA) is populated with such data through an automatic synchronisation. The pipeline starts with a sequence preparation in correct format, then a multiple alignment using Muscle followed by a curation with G-blocks to identify conserved blocks. From this step a phylogenetic analysis is performed to obtain a branching diagram. Based on virologists requirements, the selected sequences from this diagram are subjected to further analysis in order to identify key features related to pathogenicity such as the site for protease cleavage, the glycosylation sites, the epitopes or the binding site.

Conclusions and Future Work:

Future developments will involve additional influenza databases within the network. By being constantly attentive to the virologists and epidemiologists requirements, the data processing can be adapted accordingly. The final goal is to have the grid-based surveillance network ready to impact the next pandemics.

Astronomy and Astrophysics / 159

The Grid for Astronomy and Astrophysics in Italy and Europe

Authors: Claudio Vuerli[1]; Giuliano Taffoni[2]

Co-authors: Fabio Pasian [1]; Giuseppe Longo [3]; Ugo Becciani [4]

[1] *INAF - SI*

[2] *INAF - IASFBO*

[3] *Università Federico II Naples*

[4] *INAF - OACT*

The Astronomical and Astrophysical (A&A) community is extremely active concerning the use of the EGEE Grid infrastructure in Italy as well as in Europe. We briefly provide a summary of current AA Grid activities and we present selected compute and data intensive applications, which have been recently ported to Grid. Those applications are related either to large international key projects (as Planck ESA satellite) or to small-scale regional scientific activities. In particular, the AA community efforts toward the set up of the Italian NGI will be also presented.

Detailed analysis:

Italian Astronomers are facing different kind of problems that involves data reduction and analysis, modeling of physical observations, theoretical simulations and comparison of theoretical and observed data. The EGEE production Grid provides the computational frame for compute-intensive tasks, such as the transformation of raw instrument data into calibrated and catalogued data, or to produce theoretical simulations. There are some important examples of the use of Grid computing by Italian Astronomers, that we mention here: the pre-launch numerical simulations of ESA Planck Satellite, the simulations used to populate the Bag of Stellar Tracks and Isochrones astrophysical catalogue, and the numerical simulations for galaxy formation.
AA research activities are also focused on the need to integrate AA databases in the Grid and to create proper science gateways to bridge the EGEE e-Infrastructure with the European Virtual Observatory (EuroVO) “data-grid”.
Italian Astronomers also participated in a number of initiatives related to the use and development of e-Infrastructure for science and research (for EGEE, EuroVO, Grid.IT, DRACO, Cometa, TriGrid), giving them the possibility to develop a well-established and successful Grid community.

Conclusions and Future Work:

The A&A is a mature community concerning the knowledge and use of the Grid. This expertise is documented by means of the applications that are using the infrastructure, the participation in Grid related projects and the large number of Grid related publications. In the future we aim at increasing the number of researchers making use of the Grid infrastructure, and to support them. To reach this goal, in Italy we will actively participate in the organization of the Italian NGI and, more generally, at European level, in coordinating the national A&A Communities with their respective NGIs. We will also collaborate with the EuroVO projects.

Impact:

During the last few years, the A&A community has grown both in terms of astronomical research groups and related applications. A&A activities involve more than 1000 researchers in Europe distributed in several VOs (ASTRO, MAGIC, PLANCK, DCA and others) and they have an impact on the research of a number of Astronomers in EU countries (for instance, in the case of Planck simulations, Grid activities affect more that 1000 astronomers involved in Planck consortium). Many Astronomers are involved in the EuroVO projects. EuroVO is an example of an operational data and service grid. A crucial research activity is the interoperability between EuroVO and EGEE grid infrastructure. This produces a ”cyber-infrastructure” that will support Astronomers in any aspect of their research activity, from data discovery and query to computation, from data storage to sharing resources and files. The use of HPC resources in the Grid infrastructure is crucial to fulfil the needs of theoretical astronomers that use numerical simulations for their research activity. In Italy this problem has been successfully faced in the framework of the Cometa and TriGrid projects.

Keywords:

Astronomy, Astrophysics, Italy, Europe, NGI, Grid

URL for further information:

wwwas.oats.inaf.it

Scientific Gateways and Portals / 160

P-GRADE grid portal family

Author: Peter Kacsuk[1]

Co-authors: Krisztian Karoczkai [1]; Robert Lovas [1]; Tamas Kiss [2]; Zoltan Farkas [1]

[1] *SZTAKI*

[2] *Univ. of Westminster*

Corresponding Author: rlovas@sztaki.hu

Science gateways are important tools to provide user friendly access to various grid systems for various user communities. The most popular way of creating science gateways is the establishment of grid portals through which users can access grid facilities without any grid installation. In many cases science gateway portals completely hide the underlying grid infrastructure and some of them provide high level programming concepts like workflow programming. The P-GRADE portal family belongs to this class of science gateways. It has three main members: P-GRADE, NGS P-GRADE and WS-PGRADE.

Detailed analysis:

P-GRADE Grid Portal is an open source, generic-purpose, workflow-oriented, graphical Grid front-end. It supports workflows composed of sequential jobs, parallel jobs and application services. P-GRADE can be used to develop, execute and monitor workflow applications on Grid systems built with Globus, EGEE (LCG or gLite) and ARC middleware technologies. Workflows and workflow based parameter studies defined in P-GRADE Grid Portal are portable between Grid platforms without learning new systems or re-engineering program code.
The NGS P-GRADE Portal significantly extends the portal with a legacy code repository and submission engine (GEMLCA), with support towards various Grid based file and database management solutions (SRB and OGSA-DAI support at the level of individual portlets and at workflow level), and enables the embedding of third party workflows (e.g. Taverna, Triana or Kepler workflows) into native P-GRADE workflows.
WS-PGRADE (newest member) has numerous new features (scalable SOA architecture, seamless access to various types of Grids and resources) and increased capabilities in many areas.
The talk explains the highlights of these portals and their relationship.

Conclusions and Future Work:

P-GRADE portal has been developed by the P-GRADE portal Developer Alliance. This is an open organization and any institution can join it. The aim of the alliance to further develop the portal family according to the needs and feedbacks of the user community. Recent feedbacks from the user community requires the porting of P-GRADE from GridSphere to Liferay portal framework and hence this will be one of the near future works. The portal will also be extended with Shibboleth support as requested by the UK NGS and Swiss Grid.

Impact:

In the recent years more and more communities select members of the P-GRADE portal family as their science gateway. P-GRADE portal became particularly popular due to its robustness and easy-to-use interface. It has been established for national grids (Grid Ireland, Swiss Grid, Belgian Grid, Turkish Grid, Kazakh Grid, Hungarian Grid, Croatian Grid, Armenian Grid, Malaysian Grid, etc.), for regional grids (SEE-GRID, VOCE, Baltic Grid, White Rose Grid) and for application specific grids (Chemistry Grid, Math Grid, EGRID - Economics VO, etc.). The NGS P-GRADE portal has been serving the UK NGS for several years. WS-PGRADE portal has been used by the EU FP6 CancerGrid project and the UK ProSim projects.
P-GRADE is a GPL licenced OSS that can be accessed and downloaded from sourceforge.
P-GRADE portal has been extended with a new layer that enables the rapid creation of application

specific science gateways. Based on this feature the following community specific portals have been set up: Rendering portal for Blender community, OmNet++ portal (for network simulation community), Traffic simulation portal, e-market place portal, seizmology portal.

Keywords:

science gateway, grid portal, workflow, open source

URL for further information:

http://portal.p-grade.hu/, http://sourceforge.net/projects/pgportal/, http://www.guse.hu/

Earth Science / **161**

Geospatial and Grid infrastructures interoperability in enviroGRIDS

Author: Dorian Gorgan[1]

Co-authors: Anthony LEHMAN [2]; Denisa RODILA [1]; Gregory GIULIANI [2]; Karel CHARVAT [3]; Nicolas RAY [2]; Victor BACU [1]

[1] *Technical University of Cluj-Napoca*

[2] *University of Geneva*

[3] *Czech Centre for Science and Society*

Corresponding Author: dorian.gorgan@cs.utcluj.ro

EnviroGRIDS (Black Sea Catchment Observation and Assessment System supporting Sustainable Development) is a 4-years FP7 Project aiming to address the subjects of ecologically unsustainable development. The geospatial technologies offer very specialized functionality for Earth Science oriented applications as well as the Grid oriented technology that is able to support distributed and parallel processing. One challenge of the enviroGRIDS project is the interoperability between geospatial and Grid infrastructures by providing the basic and the extended features of the both technologies.

Detailed analysis:

According with the Service Oriented Architecture concepts and the requirements of interoperability between geospatial and Grid infrastructures each of the main functionalities is visible fromthe EnviroGRIDS Portal and consequently, by the end user applications such as Decision Maker/Citizen oriented Applications. The Geospatial Oriented Level in the enviroGRIDS architecture supports the integration of the compatibility between Grid (i.e. gLite) and the OGC Web Services concerning mainly with: data management (e.g. metadata, catalogues, and data access), resource management (e.g. storage elements, computing elements, services, etc), job execution (e.g. process synchronization, fault recovering, workflow management, etc), security (e.g. certificates, single sign-on mechanism, etc), and other issues arisen from conceptual, technical and technological particularities.
Interoperability between geospatial and Grid infrastructures provides features such as the specific geospatial complex functionality and the high power computation and security of the Grid, high spatial model resolution and geographical area covering, flexible combination and interoperability of the geographical models.

Conclusions and Future Work:

The research carried out in enviroGRIDS is closely connected to achievements of the standards and organizations such as OGC, G-OWS, gLite, INSPIRE, GEOSS, EGEE, etc.

Impact:

The single access Web portals and Earth Science oriented applications could combine the features provided by the interoperability between the geospatial and Grid infrastructures. The main features consist

of complex geospatial functionality and the Grid oriented high power distributed computation and secured resource management, which support the development and the simulation of large geographical and high resolution spatial models.

Keywords:

Geospatial infrastructure, Grid infrastructure, interoperability, Web and Grid services, enviroGRIDS

URL for further information:

http://www.envirogrids.net/

Justification for delivering demo and/or technical requirements (for demos):

it is appropriate for an oral presentation and debate.

Novel Architectures and Technologies / 162

Parameter sweep job submission to Eucalyptus clouds

Authors: Attila Marosi[1]; Peter Kacsuk[1]

Co-author: Robert Lovas [1]

[1] *SZTAKI*

Corresponding Author: rlovas@sztaki.hu

Large parameter sweep applications require large number of resources. Unfortunately the average number of processors in EGEE VOs is between 500-800 processors that is far not enough for large parameter sweep applications. The situation can be improved if jobs of such parameter sweep applications can be distributed to available cloud resources, too. The interconnection of P-GRADE portal with clouds enable this execution scenario. The talk explains how the interconnection of P-GRADE with Eucalyptus has technically been solved and how this new system can be used by the end-user communities.

Conclusions and Future Work:

Future work will include the designation of a smart scheduler than can optimize the number of jobs to be sent to the clouds and to the grid according to the available grid and cloud resources. Furthermore, the model will be extended for commercial models including the consideration of payment for cloud resources.

Impact:

The mixed use of gLite-based academic grids and Eucalyptus-based clouds is enabled for large parameter sweep simulations by interconnecting P-GRADE with both gLite grids and Eucalyptus clouds. As a result any existing VO of EGEE can be extended with cloud resources in a seamless way.

Keywords:

clouds, parameter sweep applications, portal, scheduler

URL for further information:

http://www.lpds.sztaki.hu/

Detailed analysis:

The 3G Bridge is an extensible generic grid-to-grid bridge developed in the EDGeS project. Its modular plugin
architecture allows extending it to connect arbitrary grid middleware. The Eucalyptus/ Amazon EC2 plugin for the 3G-Bridge allocates resources from the cloud to a pool accessible via a standard job manager (e.g. Condor) for job
submission. The resources are allocated into a single pool (cluster) where the size can be dynamically

adjusted up to a pre-configured limit (for security reasons). Theoretically this allows to grow the cluster to an arbitrary size, but since cloud resources are charged by hour this is undesirable. The goal is to find the "sweet spot" the balance between price and performance. Not utilized resources need also to be removed from the pool after a period of inactivity, taking into account that adding a resource to the pool has a high overhead (cost and time required for network transfer and deployment). Recently P-GRADE portal has been connected to 3G Bridge. As a result of this new development jobs of a parameter sweep applications can be seamlessly distributed between EGEE resources, desktop grid resources and cloud resources.

Novel Architectures and Technologies / 163

A Business-Driven Cloudburst Scheduler for Bag-of-Task Applications

Author: Francisco Brasileiro[1]

Co-authors: David Candeia [1]; Raquel Lopes [1]; Ricardo Araújo [1]

[1] *UFCG*

Corresponding Author: fubica@dsc.ufcg.edu.br

Opportunistic peer-to-peer (P2P) grid computing infrastructures have been proven to be cheap and effective platforms to run Bag-of-Tasks (BoT) applications. They operate on a best-effort model, which is understandable given their negligible cost. However, for some applications that could benefit from these infrastructures, having some guarantees on the time that takes to complete is an important requirement. In this work we investigate the possibility of combining a P2P grid with a cloud computing provider to provide some quality of service (QoS) guarantees.

Detailed analysis:

We consider the problem of scheduling a BoT application in a hybrid computing infrastructure whose processing elements are comprised of in-house machines, virtual machines acquired from a cloud computing provider, and remote virtual machines made available by a best-effort P2P grid. The applications that run in this hybrid infrastructure are characterised by a utility function. The utility yielded by the completion of an application depends on the time taken to execute it. We take a business-driven approach to manage this infrastructure, aiming to maximise its profit, that is, the utility produced as a result of the applications that are run minus the cost of the computing resources that are used to run them. Each of these resources has a different cost basis. Whenever possible, applications are run using computing power just from the in-house capacity and from the P2P grid. Any extra capacity required to improve the profitability of the infrastructure is purchased from the cloud computing spot market on an on-demand basis. We propose a heuristic that is able to schedule BoT applications in this hybrid infrastructure.

Conclusions and Future Work:

We are currently testing the cloudburst scheduler with the certification infrastructure of the EELA-2 project and an Eucalyptus-based cloud computing provider maintained by ourselves. Our next step is to identify potential users among the EELA-2 user community and the grid community at large.

Impact:

The core of the scheduler is the module that decides when to buy new resources from the cloud provider. When an application is submitted for execution, the scheduler has no information about the QoS of the P2P grid, thus, it starts by trying to obtain as many machines as the number of tasks it has to run. At the end of each subsequent turn it computes the throughput of the P2P grid and estimates if the throughput needs to be enhanced by acquiring extra resources at the cloud provider. This scheduler was implemented as a broker of the OurGrid middleware and is instrumented to acquire extra resources from a cloud computer provider when needed. We currently have implementations for both Amazon EC2 and Eucalyptus. Results of measurement experiments show that the proposed scheduler is able to substantially increase the profitability of the infrastructure, in comparison with a traditional broker

that does not consider the hybrid infrastructure. We also compared the cloudburst heuristic with an optimum scheduler and notice that the profitability of the former is close to the profitability achieved by the optimum scheduler.

Keywords:

cloud computing; opportunistic grid; peer-to-peer grid; bag-of-task applications

URL for further information:

http://redmine.lsd.ufcg.edu.br/projects/show/ourgrid

Infrastructure Tools and Services / **164**

Workflow repository, user specific monitor, and vulnerability analyzer in SEE-GRID

Author: Robert Lovas[1]

Co-authors: Akos Balasko [1]; Miklos Kozlovszky [1]; Sandor Acs [1]; Zoltan Balaton [1]

[1] *MTA SZTAKI*

Corresponding Author: rlovas@sztaki.hu

This talk introduces some high-level services, which extend the gLite-based SEE-GRID infrastructure in order to ease several tasks of end-users, application developers, and grid operators. The Common Workflow Repository Extension (CWRE), and the User/application Specific Grid Infrastructure Monitoring Extension (USGIME) of P-GRADE portal can provide efficient tools for user and application developer communities. With the Grid Site Software Vulnerability Analyzer (GSSVA) the grid operators can inspect the vulnerability/security level of the grid infrastructure with minimized intrusion.

Detailed analysis:

The talk focuses on three important aspects of Grid exploitation and maintenance/operation.
(1) The CWRE provides support for users to share application workflows, and enhances the information dissemination and collaborative development in this way. The solution is based on the popular D-SPACE digital repository technology. The repository is accessible from the new Upload/Download portlets of P-GRADE portal, and via the D-SPACE native graphical interface from any web browser.
(2) Handling remote grid files is exceptionally critical and acts as source of frequent troubles and vast amount of user complaints. The new PAKITI based USGIME is able to monitor and check various scenarios of remote file handling systematically, and manage these user/application specific tests by the end-user.
(3) The primary goal of GSSVA is to provide accurate online status information about the vulnerability/security level of the grid infrastructure. The service provides a centralized system, which is monitoring the grid infrastructure from security point of view using traditional grid protocols in order to work on every site without modification in configuration/installing any new software.

Conclusions and Future Work:

The presented services can contribute to easier, more secure, and less error prone execution of complex applications on the Grid. The future plans includes (among others) the development of enhanced interfaces for visualization of historical information (GSSVA), improvements towards more WEB2 functionalities (CWRE), and enhancements based on the new users feedbacks (USGIME).

Impact:

The presented solutions can be easily adopted not only in SEE-GRID infrastructure but in other gLite-based production grids and VOs, and can be a solid base of porting to other Grid systems. These are ensured by several ways; all the presented services have been released under open source license (GPL), the P-GRADE portal related achievements are available on the sourceforge.net, and all these tools are already in production.

(1) CWRE repository can be exploited as a bridge between more than 15 different P-GRADE portal installations worldwide, fostering the creation of new application developer communities from the individual developers, and provides more visibility of research achievements.
(2) USGIME can assist the users to understand better the reasons for common critical failures and enable the execution of application specific tests systemically.
(3) GSSVA addresses vulnerability issues of Grids, which can efficiently help administrators increase the security level of the site and leaving less chance for various attacks.

Keywords:

SEE-GRID, glite, worklow, repository, vulnerability, monitoring

URL for further information:

http://wiki.egee-see.org/index.php/JRA1_Commonalities

High Energy Physics / 165

Distributed Computing and Data Analysis in the CMS Experiment: the Spanish and Portuguese case

Author: Jose Flix Molina[1]

Co-authors: Ana Yaiza Rodríguez [2]; Antonio Delgado Peris [1]; Arnau Bria [3]; Esther Acción [1]; Fernando López [3]; Francisco Javier Rodríguez Calonge [1]; Francisco Martínez [1]; Francisco Matorras [2]; Gerard Bernabeu [1]; Gonzalo Merino [1]; Gonçalo Borges [4]; Ibán Cabrillo [2]; Isidro Gonzalez [5]; Joao Paulo Martins [4]; Jorge Gomes [4]; José María Hernandez [1]; Juan José Rodríguez Vazquez [1]; Luis Pedro Menezes [4]; Manuel Delfino [3]; Marc Caubet [1]; Mario David [4]; Miguel Oliveira [6]; Nicanor Colino [1]; Nuno Almeida [4]; Pablo Orviz [2]; Rafael Marco [2]; Vanessa Acin [3]; Álvaro López [2]

[1] *Cent. Invest. Energ. Medioamb. Tec. (CIEMAT)*

[2] *Instituto de Física de Cantabria, IFCA (U. Cantabria - CSIC), Santander, Spain*

[3] *Institut de Física d'Altes Energies, IFAE, Bellaterra, Spain*

[4] *Laboratório de Instrumentaçâo e Física Experimental de Partículas (LIP), Lisboa, Portugal*

[5] *Universidad de Oviedo, Oviedo, Spain*

[6] *Laboratório de Instrumentaçâo e Física Experimental de Partículas (LIP), Coimbra, Portugal*

Corresponding Author: jose.flix.molina@cern.ch

By the end of 2009, the CMS Experiment at the Large Hadron Collider (LHC) has already started data taking with proton collisions at 450 GeV and 1.2 TeV per beam. CMS has invested a few years to build a robust Distributed Computing system to meet the expected performances in terms of data transfer/archiving, calibration, reconstruction, and analysis. Here, we will focus on the readiness of the Spanish and Portuguese computing centers (PIC [Tier-1], CIEMAT, IFCA, LIP-Lisbon, and NCG-INGRID-PT [Tier-2s]) to first LHC collisions and the measured performance in the Iberian region by April 2010.

Detailed analysis:

Being prepared for LHC data taking does not come for free: in recent years, CMS has conducted a series of Computing, Software, and Analysis challenges to demonstrate the functionality, scalability and usability of the relevant components and infrastructure. These challenges have been designed to validate the CMS distributed computing model and to run operations in quasi-real data taking conditions or even beyond expected levels. Additionally, CMS extensively and routinely test all relevant aspects of a Grid site, such as the ability to efficiently use their network to transfer data, the functionality of all the site services relevant for CMS and the capability to sustain the various CMS computing workflows (Monte Carlo simulation, event reprocessing and skimming, data analysis) at the required scale. Some figure of merits: around 32 PBs of data transferred in 2009 among distributed >60 computing centres worldwide; millions of jobs run, so far, for CMS Monte Carlo simulation, event reprocessing and skimming, data analysis...

Conclusions and Future Work:

After all the preparation work, we can proudly state (and believe) that the CMS Distributed Computing is prepared for data taking. However, the real challenge is coming now, with hundreds of users analyzing the data and central CMS Data Operations and Facilities Operations teams coping with workflows of real data coming from the detector in a quasi-continuous mode. For sure, difficulties will emerge during 2010, resulting in a very exciting first year of LHC data taking.

Impact:

In February 2010, after LHC-protection and beam commissioning, first collisions at 3.5 TeV per beam are expected. CMS is confident to cope with all the computing load resulting from this long 2010 run, with hundreds of physicists using the infrastructure to easily access and start analysing the long awaited LHC data. For the first time, the CMS Distributed Infrastructure will cope with data, resulting from LHC proton collisions at the detector. The Spanish and Portuguese communities have prepared to receive and analyze the wave of real data. This contribution focus on the description of the CMS distributed system, the preparation for data taking, and the overall result of using the infrastructure, particularly focussing on first data taking experiences on the Spanish and Portuguese computing centres.

Keywords:

CMS Distributed Computing Grid Analysis

URL for further information:

https://twiki.cern.ch/twiki/bin/view/CMS/WorkBookComputingModel

Demo Session 1, Welcome Drink / **166**

neuGRID - A Grid-Brained Infrastructure to Understand and Defeat Brain Diseases - 3rd Neuroscientific Data Challenge

Authors: Alberto REDOLFI[1]; Baptiste GRENIER[2]; David MANSET[3]; Giovanni FRISONI[1]; Jerome REVILLARD[3]; Yannick LEGRE[2]

Co-authors: Alex ZIJDENBOS [4]; Richard MCCLATCHEY [5]

[1] *I.R.C.C.S Fatebenefratelli*

[2] *HealthGrid*

[3] *maatG*

[4] *Prodema Informatics*

[5] *University of the West of England*

Launched early 2008 by the EC Research Infrastructure Unit, the neuGRID project has established a distributed e-Infrastructure interconnecting major clinical research centres in Europe, supplying neuroscientists with the most advanced ICT to defeat Alzheimer's Disease (AD) and neurodegenerative pathologies in general. Based on EGEE gLite, neuGRID has developed a harmonized and powerful environment to design, test and assess new disease markers.

Detailed analysis:

The infrastructure now offers access to popular neuroimaging and data mining toolkits, which can be further composed and executed within the grid.

Capitalizing upon its former successful data challenge, presented at EGEE'09 in Barcelona, neuGRID will this time demonstrate the largest ever grid-based analysis in the field. This concluding challenge will aim at analysing both the US-ADNI and Australian-ADNI datasets (largest AD imaging datasets in the world) using a complex combination of the 3 mostly utilised cortical thickness extraction pipelines (i.e. markers of AD) with the ultimate objective of statistically comparing pipelines' outputs.

Conclusions and Future Work:

Capitalizing uppon its former prototype infrastructure, successfully demonstrated at the EGEE'09 Conference in Barcelona, neuGRID has delivered a third generation production quality environment.
It is now looking to extend its portfolio of research tools to become a recognised Scientific Gateway for neuroscientists within the European Research Area eco-system.

Impact:

neuGRID adresses the needs of a large community of scientists. In particular, it aims to become the reference infrastructure of the European Alzheimer's Disease Consortium (EADC) representing more than 35 clinical research centres throughout Europe, and to deliver a European public facility for neuroscientists in general.
Recently, neuGRID has kicked-off an international cooperation, so called outGRID, bringing together the best known neuroimaging centres in the world, i.e. the Laboratory of NeuroImaging from the University of California in Los Angeles in the U.S.A, developping the LONI Pipeline software (pipeline authoring interface widely used in the community) and the Montreal Neurological Institute from the McGilll University in Canada.
By pulling this expertise together, neuGRID intends to integrate a large portfolio of international leading edge neuroscientific tools which will benefit to the entire community.
To address this challenge, NeuGRID leverages on EGEE/gLite as its solid middleware foundation coupled with GEANT as its broadband high capacity network.

Keywords:

grid, e-Infrastructure, neuroscience, Scientific Gateway, GEANT, EGI

URL for further information:

http://www.neugrid.eu

Demo Session 2 / **167**

Using Spare Space in Worker Nodes and Desktop Disks to Build a gLite Storage Element

Author: Francisco Brasileiro[1]

Co-authors: Alexandro SOARES [1]; Jonhnny Wesley SILVA [1]; Thiago Emmanuel PEREIRA [1]

[1] *UFCG*

Corresponding Author: fubica@dsc.ufcg.edu.br

We present how idle disk space available at desktops and worker nodes in computing elements can be used to implement the file system back-end of a gLite storage element. We developed the BeeFS to federate the distributed disks. Like some special-purpose file systems (eg. GFS), it uses a hybrid architecture that follows a client-server approach for serving metadata and manage file replicas, and a peer-to-peer one for serving data. This allows BeeFS to aggregate the spare space of disks to build a single logical volume on top of which a general purpose POSIX-compliant file system is implemented.

Detailed analysis:

A BeeFS installation consists of a single queen-bee server that handles naming, metadata and replica management operations and a number of honeycomb servers that store the actual files. The queen-bee and the honeycombs provide service to many honeybee clients. The queen-bee server is deployed in a dedicated machine. It is responsible for providing a global file namespace with location-transparent access for files, access control, resource discovery and placement coordination services. On the other hand, it is not involved in data storage at all. Honeybee clients contact it in order to obtain the location of the honeycomb servers that store the files. After that, they fetch/send data directly from/to the appropriate honeycomb server. The role of the honeycomb servers is to collaboratively store files, providing basic read and write primitives. Honeycomb servers are conceived to be deployed over a set of desktop machines or a nodes in a cluster interconnected by a LAN. This hybrid architecture mixes

aspects of client-server and peer-to-peer systems in a fashion that simplifies the design and facilitates the administration of the system.

Conclusions and Future Work:

Our implementation has been checked against the version of Pawel Jakub Dawidek's POSIX file system test suite maintained by Tuxera and has successfully executed all the 3,061 tests that comprise the suite, giving us confidence that it is, indeed, fully POSIX-compliant. We are currently deploying a storage element in the EELA-2 infrastructure backed up by a BeeFS system that was able to harness more than 1.5Tbyte of spare disk in the desktops of our lab. Our future work includes the execution of MapReduce-like applications exploring the distributed implementation of this storage element.

Impact:

We have developed an implementation of BeeFS that runs on Linux machines. BeeFS exposes the POSIX API for file system service; this is especially important for reasons of applications compatibility, allowing a standard gLite SE to use it. Programming a POSIX file system on Linux, usually requires coding at the VFS (Virtual File System) level. Instead, we have implemented BeeFS at the user level using the Java programming language. The coupling between the user level application and the Linux kernel file system modules was done via FUSE. In order to measure the file system performance in a wide range of typical operations, we ran the well-known Andrew benchmark. This benchmark emulates a software development workload. In average, BeeFS outperforms NFS execution time in 74% for write operations and 30% for read operations in the best case. In the worst case, BeeFS results in a 56% improvement in write operations and 20% for read operations when compared with NFS. A storage element that uses such a distributed file system as the storage back-end is particularly suited to executed MapReduce applications in an efficient way, provided that appropriate scheduling mechanisms are in place.

Keywords:

distributed file system; hybrid file system; POSIX; MapReduce applications

URL for further information:

http://redmine.lsd.ufcg.edu.br/projects/show/ddg

Justification for delivering demo and/or technical requirements (for demos):

Running a demo showing how to efficiently execute MapReduce applications taking advantage of a storage element that uses the BeeFS as its back-end file system should be of interest to many EGEE users.

Bioinformatics / 168

A Protein Tertiary Structure Prediction Service in the ProGenGrid System

Author: Maria Mirto[1]

Co-authors: Alessandra Ferramosca [2]; Giovanni Aloisio [1]; Vincenzo Zara [2]

[1] *CMCC, University of Salento, Lecce*

[2] *Di.S.Te.B.A, University of Salento, Lecce*

Corresponding Author: maria.mirto@unile.it

Proteins are central to all biological processes: a very important problem in proteomics is the prediction of the three-dimensional (3D) structure of proteins from their amino acid sequence, because this information can be useful for determining the protein function, given by a specific spatial conformation that the protein assumes when it reaches the active state. We have integrated a routinely expert dependent strategy, based on the homology modeling procedure, in an automatic tool that may facilitate the generation of carrier models at low resolution, exploiting the EGEE infrastructure.

Detailed analysis:

We have defined an “in silico” workflow for the tertiary structure prediction of carrier proteins based on the homology modeling procedure, starting from the amino acid sequence. As case study, we modelled the structure of the DIcarboxylate Carrier (DIC) of Yeast, using the ADP/ATP carrier protein of Bos Taurus heart mitochondria. The protein tertiary structure prediction process involves several data and software tools in the following phases: a) template/s search with the PSI-BLAST tool; b) secondary structure computing: the topological model is calculated by using the ENSEMBLE predictor; c) multi alignment between the target sequence and the template with ClustalW and Yap tools; d) models building with Modeller; and e) models check with CE and Procheck tools. These tools are all sequential and in particular Modeller is computationally expensive if a great number of models is requested as output. We run the alignment tools on Globus machines of the SPACI grid infrastructure and other tools on gLite, by using our tool of submission and monitoring of jobs. It is a meta scheduler that has several plug-ins for scheduling and monitoring of jobs on different grids.

Conclusions and Future Work:

This paper presented a service for tertiary structure prediction implemented in the ProGenGrid system. The validation of the system has been performed by predicting the structure of the DIC protein. ProGenGrid has been used both for designing the workflow prediction graph and for scheduling it onto a Computational Grid, based on gLite and Globus grid resources. Moreover a virtual reality environment, based on the X3D and Ajax3d technologies, has been built in order for interacting with the protein. We plan to extend the service for the prediction to other kind of proteins.

Impact:

We have integrated a set of software and data involved in the tertiary structure prediction process of several carrier proteins, producing good models with RMSD equals to 1.2-1.3 A. Several tools have been ported on the EGEE infrastructure and integrated in the ProGenGrid system, a Grid PSE. A main contribution of this work is thus a Grid-based tertiary structure prediction service based on the previous cited grid-enabled bioinformatics tools. The developed service offers the following features: 1) prediction of the protein three-dimensional structures; 2) optimization and 3) visualization of the structures, taking into account several formats such as PDB and X3D.
With respect to other tools such as Swiss-Model and Geno3d, that are suited for general purpose prediction of the models, our system has been developed and optimized for transmembrane proteins, using specific tools such as ENSEMBLE and YAP for generating the secondary structure prediction.
Tests were made by using the ProGenGrid Grid Portal, that allows the launching and monitoring of the simulation in a Grid environment based on gLite and Globus middleware, by using Biomed VO and SPACI facilities for testing the service.

Keywords:

Tertiary Structure Prediction, Grid Problem Solving Environment, Grid Portal, Workflow

URL for further information:

https://sara.unisalento.it/cgi-bin/bioinfo/enter

Poster session / 169

Earth Science applications of the ES Virtual Research Community

Author: Monique Petitdidier[1]

Co-author: Horst Schwichtenberg [2]

[1] *IPSL/Latmos*

[2] *FhG/SCAI*

Earth Science is an all-embracing term for sciences related to the planet earth covering a large user community from academy, industry and organizations. To provide a scientific basis for addressing societal issues is more and more computing resources are neededboth for real and remote time applications. In the framework of the European Grid Initiative, EGI, the ES VRC gathers the partners

of 9 Virtual Organizations located in 23 European or associated countries and is linked to EELA2, EUAsia and African partners. Overview of applications and tools are presented.

Detailed analysis:

The applications already ported on EGEE are from different ES disciplines like atmospheric chemistry, biodiversity, climatology, hydrology, meteorology, pollution, seismology domains. Atmospheric chemists and meteorologists use these applications to model ozone or pollution, or predict regional weather patterns. Solid earth physicists within this community use applications to study earthquakes. A number of hydrology applications focus on predicting floods. Several climate applications deal with access and distribution of climate model outputs for further studies like model inter-comparisons, downscaling or impact of the climate change on agriculture, water resources.... The forecasting and monitoring of natural risks such as flood and fire have driven the use of OGC (Open Geospatial Consortium) components implemented on top of Grid middleware. The number of applications and tools developed is very impressive. However the interface between the ES environment and Grid middleware is not simple for many applications and developments. There persist significant gaps due to complex computing protocols in ES that will be addressed by the ES VRC.

Conclusions and Future Work:

Demonstration of the interest of Grid technology in ES domain is done through this large panel of applications and tools. Then the next step is to expand and support Grid uptake within ESFRI, as well as other related community scientific projects and propose them collaboration to use the Grid. In order not to discourage the new user the ES Grid community needs to continue providing user-friendly tools and helping the porting of applications. For projects a gateway customerized to the needs and requirements of the concerned community will be proposed.

Impact:

Success stories are always a good way to attract new applications and projects. Forecasting or monitoring of risk and weather are applications running regularly on a medium or long term basis following the scientific and technological evolution. The use of the Grid in the other applications lasts when results are obtained and published in international journal and/or in PhD dissertations. The ES VRC is expecting to bring new applications and projects. The increase of new applications will permit to create collaboration among people working in the same field and then to develop more common tools, and/or among people having applications based on the same requirements concerning input or output data, or numerical approach. An important point is that the use of Grid technology is addressed in several infrastructure EU projects related to the ES community and in the Climate field.

Keywords:

Earth Science Applications

URL for further information:

www.EUEarthScienceGrid.org

Demo Session 2 / 170

Applications using Remote Instrumentation in Grid environment - DORII project

Authors: Andrea Del Linz[1]; Emilio Diaz[2]; Katarzyna Bylec[3]; Marcin Plociennik[3]; Milan Prica[1]; Ángel David Gutiérrez Barceló[2]

[1] *Sincrotrone Trieste S.C.p.A.*

[2] *University of Cantabria*

[3] *PSNC*

Corresponding Author: marcinp@man.poznan.pl

We present here the final results obtained in the context of the EU FP7 project Deployment of Remote Instrumentation Infrastrucutre project. This is one of the EGEE related projects. We present here the oceanographic and coastal observation and modeling using an imaging use case for demonstrating the usage of the remote instrumentation infrastrucure. The remote instrumentation techniques and services has been integrated with grid - glite middleware, giving new possibilities to the scientists. The provided solutions are generic and so could be re-used for different applications.

Detailed analysis:

The DORII infrastructure provides several services, ranging from general purpose to more specific ones, to carry out daily activities in a specific community context, to communicate with remote collaborators and to control and manage remote resources and equipment integrated with the traditional Grid resources. The proposed use case - oceanographic and coastal observation and modeling using imaging - has several cameras along the north coast of Spain. These cameras provide important information about the coastal state. The DORII infrastructure allows users to easily access the images provided by these cameras and apply different kinds of image processing algorithms to them. Image retrieval allows users to access and process huge sets of data in different manners, such as initial and final date, image type, the specific beach and camera. The same infrastructure and services have been applied in DORII also for another applications in the filed of Network-Centric Seismic Simulations, Earthquake early warning system, Experimental applications like: SAXS beamline, SYRMEP beamline, environmental applications .

Conclusions and Future Work:

We present here a very important extension of the gLite middleware towards accessing remote instruments. Remote Instrumentation is one of the emerging technologies that, combined with the grid gives a lot of new opportunites, and DORII is the major project in FP7 that takes care of such services.The proposed demonstration having good visual aspects could attract the visitors. For the future work DORII is in last stage of the project so we are planning to support the applications on the production level.

Impact:

Remote Instrumentation is an emerging technology that combined with grid capabilities gives a new perspective. It has become particlarly important with respect to ESFRI projects, where remote access to very unique and expensive installations will be essential for many scientists. Many of the ESFRI projects have already raised their interest in the technologies that have been developed by DORII. Integration of instrumentation with other elements of the e-infrastructure is a key point to enable not only better research projects based on acquisition of experimental data, but also to promote an adequate framework for inter-disciplinary initiatives. The access to remote instruments was also mentioned as one of the subcall subject in the last infrastructure call. Such an dissemination possibility - during the user forum - could result in a higher number of researchers accessing the resources. It is worth to mention that the necessity of supporting remote instrumentation was reflected in the last e-IRG (e-Infrastructure Reflection Group, http://www.e-irg.org) 2009 White Paper with a special chapter devoted to “Remote Instrumentation”.

Keywords:

grid, remote, instrumentation, dorii,

URL for further information:

http://www.dorii.eu

Justification for delivering demo and/or technical requirements (for demos):

This demo will show broad range of the RI infrastrucutre usage. DEMO will show integrated whole gateway collaborative platfrom including scientific workflows. It has nice visualisation capabilities

Poster session / 171

Workflow support for multinode tests of Grid components

Author: Robert Lovas[1]

[1] *MTA SZTAKI*

Corresponding Author: rlovas@sztaki.hu

Desktop Grid solutions —such as Internet-based distributed or volunteer computing infrastructures— usually collect non-reliable and vulnerable resources from the donors (desktop PC owners) for some selected, grand challenge projects. In this paper we discuss one of the key issues, the software build, test, and validation procedures for such (and similar) heterogeneous environments, based on the workflow based description of multinode tests and ETICS-2 services.

Conclusions and Future Work:

According to our first experiences, the application development methods of EDGeS, together with ETICS-2 build and workflow-based test facilities, can provide an efficient solution for the above described efforts. For the future work, we plan to apply these new methods to more applications, and to disseminate the results.

Impact:

The presented solution tries to offer best practices for other Grid initiatives and projects, and increase the reliability and portability of Grid applications among various flavours of platforms.

Keywords:

testing, deployment, workflow, ETICS

Detailed analysis:

The aim of the validation phase is to assure that the application causes no harm to the computers of Desktop Grid donors and also that it conforms to the generic aims of the target Desktop Grid system. This stage is inevitable in order to deploy the application on a Desktop Grid platform, where individuals or institutions offer their volunteer resources for the computation.

In the framework of the EDGeS project MTA SZTAKI has started utilizing some high level, workflow-based solutions for these purposes developed in the ETICS-2 project.

The main goal of the high level workflow tool is to ease the multi-node test design, when the required services are deployed on different machines, and where the data exchanges between different phases of complex multi-node deployment on different nodes are crucial.

URL for further information:

http://etics-project.eu

Bioinformatics / 172

The Nordic BioGrid project – Bioinformatics for the grid

Authors: Ann-Charlotte Berglund Sonnhammer[1]; Bengt Persson[2]; Erik Sonnhammer[3]; Inge Jonassen[4]; Joel Hedlund[5]; Josva Kleist[6]; Kimmo Mattila[7]; Michael Grønager[6]; Olli Tourunen[8]; Steffen Möller[9]

[1] *Linnaeus Center for Bioinformatics (LCB), Uppsala University, S-751 05 Uppsala, Sweden*

[2] *IFM Bioinformatics, Linköping University, S-581 83 Linköping, Sweden; Dept of Cell and Molecular Biology, Karolinska Institutet, S-171 77 Stockholm, Sweden; National Supercomputer Centre (NSC), S-581 83 Linköping, Sweden*

[3] *Stockholm Bioinformatics Center (SBC), Stockholm University, S-106 91 Stockholm, Sweden*

[4] *CBU, Bergen Centre for Computational Science, N-5020 Bergen, Norway*

[5] *IFM Bioinformatics, Linköping University, S-581 83 Linköping, Sweden*

[6] *Nordic Data Grid Facility, Kastruplundgade 22, DK-2770 Kastrup, Denmark*

[7] *CSC – IT Center for Science Ltd, P.O. Box 405, FI-02101 Espoo, Finland*

[8] *Nordic Data Grid Facility*

[9] *Institut für Neuro- und Bioinformatik, University of Lübeck, D-23538 Lübeck, Germany*

Corresponding Author: bpn@ifm.liu.se

Life sciences have undergone an immense transformation during the recent years, where advances in genomics, proteomics and other high-throughput techniques produce floods of raw data that need to be stored, analysed and interpreted in various ways. Bioinformatics is crucial by providing tools to efficiently utilize these gold mines of data in order to better understand the roles of proteins and genes and to spark ideas for new experiments.

Conclusions and Future Work:

The BioGrid has already contributed to provide computational power for analysis of the medium-chain dehydrogenase/reductase (MDR) superfamily. The size and complexity of this superfamily has recently been shown to far surpass the means of subclassification that have traditionally been employed for this task. Instead, more computationally demanding methods must be employed, such as profile Hidden Markov Models, implemented in the HMMer package.

Impact:

Regarding databases, the frequently used databases UniProtKB and UniRef have been made available on the distributed and cached storage system within the Nordic grid. A system for database updating has been deployed in a virtual machine hosted by NDGF. The database PairsDB updates have been run on BioGrid & M-grid resources. Further applications are in the pipeline to be gridified including molecule dynamics and phylogeny calculations.

Keywords:

bioinformatics, grid, hidden Markov models, large-scale analyses, distributed storage

URL for further information:

http://wiki.ndgf.org/display/ndgfwiki/BioGrid

Detailed analysis:

BioGrid is an effort to establish a Nordic grid infrastructure for bioinformatics, supported by NDGF (Nordic DataGrid Facility). BioGrid aims both to gridify computationally heavy tasks and to coordinate bioinformatics infrastructure efforts in order to use the Nordic resources more efficiently. Hitherto, the widely used bioinformatics software packages BLAST and HMMer have been gridified. Furthermore, the multiple sequence alignment programs ClustalW, MAFFT and MUSCLE have been made available on the grid.

Novel Architectures and Technologies / 173

Pros and Cons of cloud adoption in the Scientific Data Infrastructures - the D4S-II case

Author: Andrea Manieri[1]

[1] *Engineering Ingegneria Informatica S.p.A.*

Corresponding Author: manieri@eng.it

The D4S project is going to provide as major product the gCube middleware. It is a grid-enabled service oriented middleware enabling the creation and operation of Virtual Research Environments, to serve the management and exploitation of scientific data. As part of the D4S-II project a study on the impact of cloud technology (mainly Virtualisation) and cloud capabilities (on-demand availability and scalability) on the gCube software will be performed. The study will also analyse the impact of cloud technology from a commercial point of view (market positioning, possible exploitation strategies).

Detailed analysis:

During the session will be presented the results of the study eventually showing hands-on results that will be developed as experimentation during the study period.
The proposed study will be divided into three phases:
1) Market analysis and positioning - aiming at clarifying the position of gCube with respect to the market of cloud services. During this stage the team will identify competitors, commercial potentials, suppliers in the value chain costs and potential revenues.
2) gCube installation and performance analysis – the system will be installed at ENG premises (a virtualised infrastructure) in order to test its installability and measure its performance.
3) Analysis of gCube architecture with respect to cloud technology - aiming at identifying ways to exploit existing cloud technology to boost performance or processes, and/or analyse how on-demand availability and scalability are provided by gCube. In more detail two opportunities will be taken into consideration: the uses of an IaaS (e.g AWS or blade server) instead of EGEE resources and the use of a PaaS (e.g. Google Apps or MS Azure) to implement gCore functionalities.

Conclusions and Future Work:

During this session the participants will be informed about the results of a study comparing the gCube solution, developed within the framework of DILIGENT and D4S projects, and the cloud technologies and models.
The results of the study will be exploited at the level of the D4S-II project in order to improve the present solution and provide an up to date interoperability solution for developing Virtual Research Ecosystems. The study will also contribute to the discussion at the level of e-Infrastructures on the potential impact of cloud technology on e-Science.

Impact:

The study will provide useful information to the different project activities, in particular to:
- potential customers to understand the type of service provided, to clarify SLA and any related cost for the potential customers in adopting the solution, and to structure and to infrom about the sustainability strategy.
- Site manager to know about the the installation and other operational procedures, how to reduce the need of dedicated servers and how to measures the cost of providing the D4S-II services
- gCube developers to understand on how to improve the gCube software and related services, identify any bottleneck, understanding performance of the entire system and how to intervene to improve them.
- All participants to learn about the cloud technologies, how to exploit them to further services offered to the communities.

Keywords:

cloud, business models, technology impact, gcube, virtual research environment

URL for further information:

http://grid.eng.it

Technical Plenary / **174**

The computational challenges in Life Sciences

Author: Modesto Orozco[1]

[1] *Director Joint Research Program on Computational Biology. Institute of Reserach in Biomedicine and Barcelona Supercomputing Center.*

Corresponding Author: modesto.orozco@irbbarcelona.org

Life sciences is one of few areas of research where the complexity of the problems grows faster than our ability to deal with them from a computational point of view. Massive genomic projects have yield detail sequence information on all the species of human interest and we are now deriving metagenomic information of complex ecosystems, including for example the human digestive system. In a near future we might face the pangenomic scenario, where sequence information of all the human

beings might be accessible. In parallel, to the genomic revolution, system biology projects are putting all this sequence data in the proper cellular context, building entire networks aimed to explain the chemical complexity of life. The third axis structural genomic projects are increasing the structural knowledge on biological machines opening then the possibility to modulate on their actions and by extension on the entire cellular functioning.

During my talk I will try to draw the exciting, but scaring scenario of the research in biology, trying to outline the main topics were help from computers is needed. I will provide examples of how current computer platforms can help to derive biologically-relevant information and I will emphasize the complexity of the problem from the computational point of view.

Technical Plenary / **175**

Addressing Complexity in Emerging Cyber-Ecosystems - Exploring the Role of Autonomics in E-Science

Author: Manish Parashar[1]

[1] *Rutgers University*

Corresponding Author: parashar@rutgers.edu

Significant strategic investments are quickly realizing a pervasive computational cyberinfrastructure that integrates large-scale computing, high-speed networks, massive data archives, instruments, observatories, experiments, and embedded sensors and actuators, and are catalyzing new thinking, paradigms and practices in computational science and engineering – those that are collaborative and information/data-driven. However the ability of scientists to realize this potential is being severely hampered by the complexity of the applications and infrastructure, which together present unprecedented development, configuration and management challenges. Autonomic computing has the potential to fundamentally address these challenges. The goal of autonomic computing is to design and engineer systems and applications that are capable of managing themselves, adapting their resources and operations to workloads and execution context, and anticipating needs, all with minimal involvement of users. In this talk I will explore the role of autonomics in computational science and engineering, both in managing systems and applications as well as in enabling new application formulations. I will then describe specific research efforts aimed at enabling autonomic scientific and engineering applications that can address the challenges of (and benefit from) pervasive cyber-ecosystems.

Technical Plenary / **176**

Advanced computing for fusion simulation and modelling

Author: Pär Strand[1]

[1] *Chalmers University of Technology*

ITER is the next generation of fusion devices and is intended to demonstrate the scientific and technical feasibility of fusion as a sustainable energy source for the future. To exploit the full potential of the device and to guarantee optimal operation for the device, a high degree of physics modelling and simulation is needed even in the current construction phase of the ITER project. The detailed modelling tools that are needed for an adequate description of the underlying physics cover both a wide range of timescales and spatial orderings and are in general very demanding from a computational point of view. Current modelling activities rely on local or national computational resources and an improved access to computing infrastructures will be instrumental in advancing a pan-European modelling activity for ITER to a very competitive status in relation to the ITER partners.

The EUFORIA project is developing a comprehensive framework and infrastructure for core and edge transport and turbulence simulation, linking grid and High Performance Computing (HPC), to the fusion modelling community.

EUFORIA is enhancing the modelling capabilities for ITER sized plasmas through the adaptation, optimization and integration of a set of critical applications for edge and core transport modelling targeting different computing paradigms as needed (serial and parallel grid computing and HPC). Deployment of both a grid service and a High Performance Computing service have proven essential to the project. A novel aspect is the dynamic coupling and integration of codes and applications running on a set of heterogeneous platforms into a single coupled framework through a workflow engine, a mechanism needed to provide the necessary level of integration between the physics applications. This strongly enhances the integrated modelling capabilities of fusion plasmas and will at the same time provide new computing infrastructures and tools to the fusion community in general.

Current status, lessons learned and future prospects are reviewed in the project.

Technical Plenary / **177**

LHC: an Example of a Global Scientific Community

With the advent of large scale scientific instruments, the need for global collaborations in science has arisen. High Energy Physics and LHC in particular has long experience in collaboration on this scale and can serve as a model for other communities which now have to face similar challenges. This talk will give an overview of the need for LHC and the collaborative efforts that have built the accelerator, the experiments and the computing. The computing has benefitted from the e-science infrastructures built in Europe and elsewhere while pushing those infrastructures to unprecedented levels of scale and performance. With the arrival of first data in the LHC experiments, the investments in the computing are paying off in being able to very rapidly analyse the data. The talk will discuss the role of High Energy Physics and CERN in the European and worldwide scientific landscape and how this might develop in the future.

Opening Plenary / **178**

The Swedish e-Science Research Centre

Author: Juni Palmgren[1]

[1] *Stockholm University and Karolinska Institutet*

Corresponding Author: juni@math.su.se

As a result of a strategic research initiative by the Swedish government in the area of e-Science, a consortium of four universities have been able to substantially increase their research activities in the area. The Swedish e-science Research Centre (SeRC) is formed by the universities in Stockholm and Linköping – KTH, Linköping University (LiU), Stockholm University (SU) and Karolinska Institutet (KI) – around the two largest high-performance computing (HPC) centres in Sweden: PDC at KTH and NSC at LiU. Research at SeRC is focused on the collaboration between tool makers and tool users, and brings together a core of nationally leading IT research teams with expertise in e-Science method development and leading scientists in selected application areas.

SeRC will provide a platform for increased collaboration between applied and method-oriented groups, with the aim to constitute a visionary e-Science node with a national scope and strong international ties. Work will evolve along three lines: (i) Formation of e-Science Communities that connect application groups with relevant core e-Science groups and computer experts at PDC and

NSC, (ii) Research in core e-Science methods such as distributed resources, database technology, numerical analysis, visualization and interaction, mathematical modelling and parallel algorithms, focusing on problems critical for several e-Science communities, (iii) Closer collaboration between PDC and NSC, and a substantial increase in advanced support staff, which will turn the two centres into comprehensive e-Science enablers.

The general structure and scope of SeRC will be presented including examples of collaborative work in some of its communities.

EGI Session - Board: None / **179**

The EGI-InSPIRE Project: Community Interactions

Corresponding Author: steven.newhouse@cern.ch

The EGI-InSPIRE project has two activities focussed on interactions with the community: External Relations and User Community Coordination. This presentation will provide an overview of the dissemination and policy development tasks that form the basis of the external relations activity in EGI-InSPIRE, and the user community activity that encompasses a support team based at EGI.eu and those located within the community.

EGI Session - Board: None / **180**

EGI-InSPIRE: Operations

The production grid infrastructure in EGI-INSPIRE is built from resources provided by the participating partners. These are brought together into a secure, integrated, reliable infrastructure by EGI.eu through the federated operations of independent infrastructures. This presentation will provide an overview of the operational infrastructure and the development of the operational tools used to manage it.

EGI Session - Board: None / **181**

EGI-InSPIRE:Technology

The EGI-InSPIRE project will rely on external software providers to develop the technology that meets the needs of EGI's user and operations community. The EGI.eu staff will work to collect and prioritise requirements across the community from the virtual research communities using the infrastructure and from the operations staff deploying the technology. It is expected that the technology will evolve rapidly during the project, and the technology team will define the functional and performance requirements that need to be achieved by the software providers for their technology to be deployed into production.

Closing Plenary / **182**

Cloud Computing at Microsoft and Venus-C

Detailed analysis:

Cloud computing is an emerging new approach to distributed computing. Microsoft has developed a new Cloud computing platform named Azure. The talk will describe the overall Cloud Computing strategy and current plans for Microsoft Research collaboration in Europe within the EU project Venus-C, submitted to the recent FP7 call for proposal.

Opening Plenary / **183**

Welcome from EGEE Project Director

Opening Plenary / **184**

Welcome from Local Organisation

Opening Plenary / **185**

Presentation of Uppsala University and history of the building

Opening Plenary / **186**

e-Infrastructures as enablers of the ”fifth freedom”

Even if the scientific community is often not aware of it, the free movement of knowledge (also known as the “fifth freedom”) is anchored in the Lisbon Treaty. Its Article 179 reads “The Union shall have the objective of strengthening its scientific and technological bases by achieving a European research area in which researchers, scientific knowledge and technology circulate freely […]”.

The talk will reflect on the role of e-Infrastructures in enabling the “fifth freedom” and devise some of the challenges ahead.

Closing Plenary / **187**

Best demo and poster awards

Closing Plenary / **188**

Outlook and farewell

Astronomy and Astrophysics / **189**

Final discussion

EEF meeting (Closed) / **190**

Introduction

Corresponding Author: robert.jones@cern.ch

EEF meeting (Closed) / **191**

Review of EEF draft report

EEF meeting (Closed) / **192**

Next steps

EEF meeting (Closed) / **193**

Preparations for the ESFRI session tomorrow (Tuesday)

EEF meeting (Closed) / **194**

AOB

Computational Chemistry / **195**

Molecular and material science innovative applications on the grid

Corresponding Author: cozzini@democritos.it

ESFRI session / **196**

Introduction

Corresponding Author: robert.jones@cern.ch

ESFRI session / **197**

SIOS

Corresponding Author: geha@forskningsradet.no

ESFRI session / **198**

CESSDA

Corresponding Author: jwshep@essex.ac.uk

ESFRI session / **199**

Analysis of ESFRI projects e-infrastructure requirements by European e-Infrastructure Forum

Corresponding Author: robert.jones@cern.ch

ESFRI session / **200**

Discussion and next steps

EGI Session / **201**

NGIs and VOs Readiness Survey

Corresponding Author: gabriel.zaquine@cern.ch

EGI-InSPIRE / **202**

EGI Collaboration Board

Corresponding Author: steven.newhouse@cern.ch