EGI USER FORUM 2011

Book of Abstracts

held in conjunction with the

EMI TECHNICAL CONFERENCE

EUROPEAN MIDDLEWARE INITIATIVE

VILNIAUS UNIVERSITETAS · 1579 · UNIVERSITAS VILNENSIS

EGI USER FORUM 2011

held in conjunction with the
EMI TECHNICAL CONFERENCE

11–14 APRIL 2011, VILNIUS, LITHUANIA

BOOK OF ABSTRACTS

First published in 2011 in the Netherlands by EGI.eu
Science Park 105
1098 XG Amsterdam
The Netherlands
www.egi.eu

ISBN 978 90 816927 1 7

EGI-InSPIRE ("European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe") is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years.

Editors: Sara Coelho, Neasan O'Neill
Design: Viviane Li
Forum poster: Ori Ginale

FORUM HOSTS

EUROPEAN MIDDLEWARE INITIATIVE

VILNIAUS UNIVERSITETAS · 1579 · UNIVERSITAS VILNENSIS

Critical IT infrastructure

PROGRAMME COMMITTEE

Alberto Aimar CERN
Sergio Andreozzi EGI.eu
Steve Brewer EGI.eu
Daniele Cesini INFN
Diana Cresti INFN
Michel Drescher EGI.eu
Florida Estrella CERN
Tiziana Ferrari EGI.eu
Laurence Field CERN
Patrick Fuhrmann DESY
Catherine Gater EGI.eu
Francesco Giacomini INFN
Emidio Giorgio INFN
Maria Girone CERN
Balázs Kónya LUND
Yannick Legre HealthGrid
Rob van der Meer EGI.eu
Alberto Di Meglio CERN
Johan Montagnat CNRS
Steven Newhouse EGI.eu
Morris Riedel JUELICH
Horst Schwichtenberg FRAUNHOFER
Jamie Shiers CERN
Gergely Sipos EGI.eu
Giuliano Taffoni INAF
Claudio Vuerli INAF
John Walsh TCD
Cecile Germain CNRS
Ioannis Liabotis GRNET

Contents

VIRTUALISATION AND CLOUD COMPUTING 78

USER ENVIRONMENT 92

Dear Delegate,

On behalf of the European Grid Infrastructure community, specifically the EGI-InSPIRE and European Middleware Initiative (EMI) projects, and the Vilnius University we would like to welcome you to the EGI User Forum and EMI Technical Conference in Vilnius, Lithuania.

April 2011 is the last month in the first year in both the EGI-InSPIRE and EMI projects – and what a year it has been! April will see the release of EMI 1 – the first public release of an integrated software offering from the gLite, ARC, UNICORE and dCache consortia – after which planning starts on the next release in a year's time.

For EGI-InSPIRE, the first year of its four-year project plan has seen the transition to a sustainable European Grid Infrastructure continue, with the establishment of EGI.eu in Amsterdam and the transition from the regional operational model established within EGEE to the national model within EGI.

The combined EGI User Forum and the EMI Technical Conference, provides a perfect opportunity to catch up with a range of activities taking place within the community:

- The EMI sessions provide an overview of some of the technologies being released in EMI 1 and in the developments planned for file transfer and access, the software quality assurance procedures that are employed in the project, and the plans for standards adoption.
- The User Support Services sessions give an overview of the services provided by EGI.eu and its partners, and the domains specific support services provided by various communities to help their users to make the best use of the infrastructure.
- Tailored User Environments are frequently provided to a community to simplify their use of the infrastructure. A variety of scientific gateways and portal technologies are described, in addition to applications, tools and programming libraries that can simplify access to distributed resources. The use of these environments across a range of scientific disciplines is also reported upon.
- The Virtualisation and Cloud Computing sessions provide details of the technologies and user experiences coming from European and national activities in this area.
- The Technologies for Distributed Computing sessions highlight some of the challenges in providing federated access to e-Infrastructures to achieve interoperability.
- Some of the high-performance Data Management solutions coming from the WLCG community are reviewed.

In addition there is a rich selection of demonstrations being shown each day in the exhibition area, training sessions for resource administrators and end-users on a variety of different technologies, and workshops from European projects aligned with the issues facing EGI.

Our thanks to the local organisers (LitGrid, BAIP and the Vilnius University) and the members of the Programme Committee for their suggestions and support during the preparation of this meeting.
We hope you have an enjoyable week.

Overview

Steven Newhouse
Project Director, EGI-InSPIRE

Alberto Di Meglio
Project Director, EMI

User Support Services

This track presents the User Support services provided by EGI.eu and the wider EGI-community. EGI provides user support through the GGUS ticketing system which is aimed at capturing issues relating to the operational e-Infrastructure. The Requirements Tracking system offers a mechanism for all members of the community to submit details of their future needs for the e-Infrastructure. In addition to capturing new requirements, the User Community Support team at EGI offers a range of services aimed at meeting the broader needs of research communities.

A key focus in EGI is to provide communities of like-minded researchers to work together as Virtual Research Communities in order to better communicate their common needs and also to benefit from common solutions. Furthermore, the EGI support team collaborates closely with the various NGIs and other partners to assist them in their support of researchers in their own countries. This is done by making the EGI shared services available, by sharing the knowledge of user needs captured from the Virtual Research Communities, as well as by learning directly from the NGIs what their own direct needs are.

Of course, many support services are highly specialised and relate specifically to particular projects and communities. A number of the presentations within this track describe such initiatives. The specialist subject areas covered include: neurodegenerative diagnostics enhancement, digital cultural heritage, gamma-ray astronomy, heliophysics and structural biology. Other technical areas covered include various accounting and virtual organisation monitoring services, gLite's Logging and Bookkeeping service and the HammerCloud automated testing service. The GÉANT Project's dedicated task to Liaison and Support is also described.

Steve Brewer Chief Community Officer, EGI.eu

Benefits and vision for the VRC community model

OVERVIEW

Scientific research is no longer being conducted within a single research group in a single institution in a single country but in Virtual Research Communities (VRCs) that span national borders encompassing many different organisations but with a need to share ICT resources. EGI provides the e-Infrastructure that enables the researchers within a VRC to collaborate, communicate, share resources, access remote computers or equipment and produce results as effectively as if they, and the resources they require, were physically co-located. Becoming an identified VRC within EGI offers several benefits for researchers, scientists and the developers of distributed scientific applications. The talk overviews the EGI VRC model, the benefits of establishing EGI VRCs for scientific communities and the landscape of already established VRCs.

DESCRIPTION

EGI is a partnership between National Grid Initiatives (NGIs), EIROs and a coordinating body – named EGI.eu – to operate a sustainable, pan-European grid infrastructure for international scientific communities. EGI.eu is seen as the glue enabling coherence between the NGIs for the benefit of their users and members of Virtual Research Communities or VRCs. VRC membership in EGI offers the following key benefits for researchers, scientists and the developers of distributed scientific applications:

1. VRCs can access computing, data storage and other types of resource made available by EGI stakeholders through open source middleware software solutions. VRC members can store, process and index large datasets and can interact with partners using the secured services of the EGI infrastructure.
2. The user support units of the NGIs and EGI.eu help VRC members during the routine usage of the systems and provide assistance to access and utilise the largest multi-disciplinary grid in the world.
3. VRCs will have the ability to establish their own Virtual Organisations (VOs) as collections of hardware, software and human resources configured in order to share capacities, to collaborate with partners and to run data intensive simulations. The VOs can benefit from the resources provided by NGIs and other VRCs.
4. NGIs provide trainers and technology specialists for VRCs to support them during the integration and adaptation of their legacy applications and datasets to the EGI infrastructure. The combination of VRC's own training resources together with EGI's infrastructure-related modules can provide comprehensive packages for VRC members in an efficient and timely manner.
5. VRCs can influence the evolution of EGI's services through representation in the User Community Board and the User Services Advisory Group. Based on requirements collected from its members, VRCs can advise EGI on its planning and operational priorities.

IMPACT

The VRC is defined as an organisational grouping that brings together transient Virtual Organisations within a persistent and sustainable structure. A VRC must be a self-organising group that collects and represents the interests of a focussed collection of researchers across a clear and well-defined field. Named contacts are agreed upon by the VRC to perform specific roles and these then form the communication channel between the VRC and EGI. EGI welcomes any scientific community to submit a proposal for the accreditation of an EGI VRC. VRC proposal must demonstrate that it represents a community of researchers that has an established existence outside of the VRC, i.e. that it has structure (such as an ESFRI project, EIROFORUM laboratory, national research structure. professional organisation or affiliation, etc.) and that this body represents this particular community. This community must also show that it has an established governance model and open mechanisms for new participants to enter (or leave) the organisation, and that all members of this organisation will have access to all the services offered by the VRC, i.e. beyond those who are just part of the proposal. These conditions enable EGI.eu to recognise the VRC as being the 'voice' of a particular community of users within the infrastructure. EGI can provide help and advice on suggested best practices for such organisational models if needed.

CONCLUSION

The presentation gives an overview of the benefits of EGI VRCs, the VRC accreditation process and the profile of already established VRCs. VRCs benefit from the services provided by members of the EGI collaboration, as well as from the services provided by other VRCs. VRCs can influence the evolution of EGI software and services through various mechanisms, primarily through the User Community Board. The EGI VRC model contributes to the sustainability of European Distributed Computing Infrastructures and to the sustainability of European scientific communities to a large extent.

Become an EGI VRC:
www.egi.eu/user-support/becoming_a_community

Steve Brewer steve.brewer@egi.eu EGI.eu

EGI VO Services activities

OVERVIEW

The EGI-InSPIRE project aims at supporting both Heavy User Communities (HUC) and new emerging communities. While HUCs are directly supported through specific tasks in the EGI SA3 activity, new emerging communities lack this organisation and resources and take-off can fail if assistance is not provided. The VO Services subtask, developed in the context of the EGI-InSPIRE TNA3.4 Technical Services activity, aims at providing documentation and services to these emerging communities, allowing them to start up their own Virtual Organisations (VO), Virtual Research Communities (VRC) and boost their integration and usage of the production infrastructure.

DESCRIPTION

The VO Services offer new emerging communities the documentation to guide VO managers through the various steps of the VO life-cycle process; support VO managers in deploying, using and adapting technical services which may simplify their user community's access to the infrastructure and promote collaboration within the VO; and collect feedback from VOs and VRCs regarding the used services, mapping them into requirements for improvement.

Documentation is presently offered through a wiki-page that gathers in a central and unique point information regarding protocols and procedures relevant for VO management activities, such as the VO registrations process in EGI, VO Management Frequently Asked Questions wiki and cost/effort evaluations (from the operations point of view) of several services and tools that were found to be useful for new emerging communities.

Several monitoring frameworks for the VOs are already available under the EGI framework, and information is compiled in the Wiki pages. Specific effort has been put on testing a Nagios monitoring instance which could be used by one or more VOs in monitoring their infrastructure.

Finally, an official communication channel has been established via a support unit in GGUS where VO Managers can ask for clarifications on using some special tool, handling some special operation or just expressing a requirement that they would like to see fulfilled.

Future work will continue with the investigation of the possibility to extend the use of dashboards for new emerging communities. Presently, the available dashboards used in the HUC framework are based on commercial software databases, which prevents the possibility to extend its use to a broader range of users. Other alternatives are being explored.

IMPACT

One of the main achievements during the first activity period was the production of an inventory of basic services and monitoring tools needed by every VO, and emphasising characteristics considered important like purpose of the service; hardware requirements; known issues; services capacities and costs from the perspectives of the major players involved: service provider, VO operations and User Support. The impact of such documents intends to raise the awareness of VO Administrators regarding the range of services needed to support standard VO requirements, and understand their associated costs. VOs should understand that providing services to their VO members requires an effort that should be at least partially assumed by the VO. This information is an important tool to assist VO Managers in deciding and planning the range of services needed, and whether these services can be operated by the VO or provided through EGI catch-all instances.

Another major achievement was the official establishment of a communication tool (via a dedicated GGUS support unit) where VO managers can ask for help and clarification. The impact is that dedicated staff are ready to provide guidance through the VO life cycle; assist VO managers through the process of deploying technical services; and receive VO managers' requests (in terms of services/tools).

Finally, the VO management documentation was revised, and a Frequently Asked Questions wiki-page was created to properly address new comers' VO issues (https://wiki.egi.eu/wiki/VO_Services).

All these achievements are contributing so that VO operations become more transparent from all points of view: VO management, operation and usage. The future actions are also pointing in that direction with the testing of VO tools for monitoring.

CONCLUSION

The technical integration of any research community as a Virtual Organisation in a production grid environment is a demanding task from the operations point of view. In EGI, that activity has further constraints due to the decentralised political organisation, where NGIs play a leading role in the project governance model. VO Services activity intends to address many of the difficulties which new emerging VO managers face in the operation and technical instantiation of their VO. During this first period of the EGI project, this was achieved through the evaluation of services and tools which were found to be useful for VOs and VRCs. Simultaneously, some steps of the VO Management process have been revised, and improvement requirements have been communicated to the developers of the services and tools used during that process. The enhancement of VO activities management is expected to continue through future activities as well as the testing of new VO monitoring framework offered from Nagios tool.

https://wiki.egi.eu/wiki/VO_Services

Gonçalo Borges goncalo@lip.pt LIP • **Jorge Gomes** jorge@lip.pt LIP • **Ignacio Blanquer** iblanque@dsic.upv.es UPV

Applications Database

OVERVIEW

The EGI Applications Database (AppDB) stores information about tailor-made computing tools for scientists to use, and about the programmers and scientists who have developed them. The applications and tools filed in AppDB are finished products, ready to be used on the European Grid Infrastructure. Its main mission is to provide all the necessary information about the applications running on EGI, thus enabling people to search for applications matching a pattern (such as a scientific domain) and also to contact the corresponding authors for guidance on application usage or further development. Storing pre-made applications means that scientists don't have to spend research time developing their own software. Hence, by storing pre-made applications and tools, AppDB aims to avoid duplication of effort across the EGI user community, and to inspire scientists less familiar with programming into using the European Grid Infrastructure.

DESCRIPTION

Activities since the beginning of the project have been primarily focused on migrating existing data from the EGEE era. Support for all EGI-endorsed middleware was added and the concept of tools was introduced, by porting the entries of the existing RESPECT programme into the database. The storing of personal profiles for each application developer was introduced, aiming at simplifying the search for experts. In order to provide the user with a consistent and responsive end result, such architectural design patterns and techniques as MVC and AJAX have been employed. Finally, another important feature of AppDB, around which much of the development effort was concentrated, was providing write-enabled access to registered users, by linking AppDB to the EGI SSO system, thus giving them control and responsibility over the data.

For the next version of the AppDB, there are several enhancements planned related to the information provided, such as adding an 'experience' section to people's profiles, and keywords and a tagging mechanism to applications. Another important planned feature is integration with GOCDB/CIC through their web APIs, in order to provide detailed NGI and VO information, which will require architectural changes and careful planning. Integration with the EGI Community Software Repository is also a feature planned which will enable developers to register and submit software releases through the AppDB interface.

Apart from enhancements, another key aspect planned for the next release is the further development of an XML-based web API. This requirement has stemmed from the need for NGIs and VRCs to be able to provide their own localised, custom interfaces to the service, which could, for example, display entries relevant to their county or discipline only, or in their own native language. Upon successful deployment of such a web API, the possibility of developing an alternate REST API on a dedicated FQDN will be considered, depending on feedback and impact.

IMPACT

At the time of this writing, the AppDB counts 272 application and tool entries, and 505 people profiles. Since the public release of the latest major version in mid-November 2010, there has been interest from various countries such as Switzerland, Ireland, and Norway for the testing and addition of entries. The site's traffic has been monitored through the use of Google Analytics, and during this one-month period, there have been 182 visits and 1151 page views in total, from 25 countries. Additionally, the GISELA project has declared its intention to share 52 of its applications with the EGI AppDB, an intention which may bring further collaboration between initiatives in the future.

It should be noted that the AppDB is unique as a service in the EGI infrastructure, and through its web API, which is under development, new localised instances are expected to spring out. This is in turn expected to help promote scientific work more efficiently by targeting specific niches, while the AppDB remains the master instance.

CONCLUSION

The AppDB subtask has used the first six months of the project to migrate and transform existing data from the EGEE era, to initially provide a new read-only portal to the community, with subsequent write-enabled authenticated access. For the next six months, there exist plans for additions to the data and object model, additions to the portal's functionality, mainly in order to integrate with other services such as the GOCDB. The CIC portal, and localised versions of the portal itself, and integration with the EGI Community Software Repository.

Further requirements have been laid out for later development, such as integration with monitoring and information services of the various supported middlewares, in order to provide information about application availability and resources in sites across NGIs. However, due to the lack of a uniform manner for information retrieval, such requirements have been put on hold, until there arises a way to overcome the technical difficulties.

http://appdb.egi.eu

William Vassilis Karageorgos wvkarageorgos@iasa.gr IASA

NGI-based user support in the European Grid Infrastructure collaboration

OVERVIEW

The stated goal of the European Grid Infrastructure is to facilitate the work of established and emerging multi-national user communities. The growing user demands have provided, and will continue to provide, the necessary push for development and extension of the grid infrastructure. Therefore, the active support for these communities is a primary concern for the members of the EGI collaboration. The presentation introduces the goals of EGI user support activities, the structure and linkage of user support teams and gives an overview of the services and tools that the different stakeholders provide for multi-national user communities.

DESCRIPTION

The primary stakeholders of EGI user support services are the User Support Teams of National Grid Initiatives/Infrastructures (NGIs) and EIROs. While the primary responsibility of these NGI User Support Teams is to serve users and user communities inside their countries, through their involvement in international collaborations – under the EGI umbrella – they share and expand their services for the benefit of large, multi-national user communities. The User Community Support Team (UCST) of the Amsterdam-based EGI.eu organisation provides coordination for this activity. The UCST ensures that the national efforts fit together at the European level, they satisfy the demands of multi-national scientific collaborations and that the evolution of support tools and mechanisms is aligned with the long term goals of communities.

The user support services are implemented as a mixture of human services, software services and feedback mechanisms. While users are mostly served by human services (e.g. by trainers who provide courses, by consultants and technicians who help porting of applications) software services allow remote interactions with experts (e.g. a helpdesk system) or facilitate the provision of human oriented services (e.g. a training event registry).

NGIs are heterogeneous in terms of the size and composition of their user communities, the size and expertise of their user support teams. Consequently, NGIs are not equally active and strong in all the previously described user support fields. The EGI.eu UCST has defined and runs mechanisms to understand the strengths and weaknesses of each of the NGI USTs and to map their offerings and capabilities to the requirements of multi-national user communities. These mechanisms, namely the Wiki pages, the knowledge bases, the email lists, the Web pages, surveys, web forms, Request Tracker system and various user-facing technical services will be reviewed in the presentation.

IMPACT

Efficient support for users is the primary interest of all the EGI stakeholders. The operation of efficient support activities is a non-trivial task, given the highly distributed nature and heterogeneity of infrastructure components, user support teams and user communities. EGI has managed to define and setup processes and tools to deal with this complexity. The presentation provides an excellent opportunity for all the EGI stakeholders to learn about user support mechanisms and to see how their partners and the various NGIs are involved in this process.

CONCLUSION

The presentation gives an overview of user support processes and activities from NGIs' perspective: how the NGIs are involved in the process and how they can improve their commitment even further. The key focus of the integrated activities is to support sustainable international research groups through the Virtual Research Community (VRC) model. An EGI VRC represents a community of researchers with an established presence in a scientific or other research field and with the shared goal to develop themselves into self-sustainable grid communities. The long term sustainability of EGI depends on the success of organising existing EGI users and emerging scientific communities into VRCs.

www.egi.eu/user-support
https://wiki.egi.eu/wiki/TNA3.3_NGI_User_Support_Teams

Gergely Sipos gergely.sipos@egi.eu EGI.eu

CTACG – The Cherenkov Telescope Array Computing Grid

OVERVIEW

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 project is driven by the CTA consortium, comprising 132 institutes in 25 countries. The CTA Computing Grid (CTACG) uses grid technology to perform heavy Monte Carlo simulations and to investigate the potential of grid computing for the future data reduction and analysis pipeline. This talk presents the tools developed in this framework, the performance achieved and the lessons learned.

DESCRIPTION

CTA is currently in the preparatory phase and no dedicated computing infrastructure is available. Using grid infrastructure the computing centres of participating institutes can provide computing resources for CTA. Currently, the CTA virtual organisation (vo.cta.in2p3.fr) comprises 8 computing centres with several 1,000 CPUs and about 500TB of storage. The participating computing centres are Tier 1 and Tier 2 centres and provide very heterogeneous computing resources.

The CTA computing grid is currently used for massive Monte Carlo simulations. For job submission and monitoring we use EasiJob (EASy Integrated JOB submission), a tool developed within the MUST framework (Mesocentre de calcul et de stockage ouvert sur la grille EGEE/LCG), the computing infrastructure from LAPP (Laboratoire d'Annecy-le-Vieux de Physique des Particules). This tool includes a web interface to define and configure a grid production.

The configuration and submission is currently being done manually by the production manager. In the future we plan to set up an automated data pipeline where each newly produced simulation file will be automatically processed. This automated data processing will be based on a data management system. We are investigating the potential of AMI (ATLAS Metadata Interface) for the needs of CTA. The outcome of this investigation will be presented in the talk.

IMPACT

In 2010, 60,000 files containing 109 simtelarray showers have been produced, outnumbering by far any production on dedicated computer clusters.

The current production jobs are very demanding in memory size (up to 4GB RAM) and disk space (5-10GB scratch). Not all available computing centres meet these requirements. Therefore only a subset of all sites is currently used. Future plans for upgrading computer centres will take into account the CTA memory requirements. The necessary scratch disk space and disk operations can be reduced by running processes parallel within one job and piping the output directly. It will be necessary to require several cores on multi-core machines on job submission.

The grid provides the necessary infra-structure to help participating institutes to easily supply their computing power to the CTA community. By relying on official grid standards the work load on the technical staff at the computing sites is reduced and the development of the grid can be easily followed. In this way CTA can concentrate nearly all its efforts on scientific study.

CONCLUSION

The combined power of computing centres at different CTA institutes, grouped in the grid infrastructure, allows massive Monte Carlo simulations. Our aim is to keep the development for CTA at a minimum by staying grid middleware compliant and relying on tools already used in the grid community. Current simulation production is much more efficient than anything being done at single computing centres and allows the estimation of the future CTA computing requirements. The grid infrastructure can be recommended for large scale computing in CTA.

http://lappwiki01.in2p3.fr/CTA-FR/doku.php?id=cta-computing-grid-public

Nukri Komin komin@lapp.in2p3.fr LAPP, CNRS/IN2P3 • **Giovanni Lamanna** lamanna@lapp.in2p3.fr LAPP, CNRS/IN2P3 • **Sabine Elles** elles@lapp.in2p3.fr LAPP, CNRS/IN2P3 • **Cecile Barbier** barbier@lapp.in2p3.fr LAPP, CNRS/IN2P3

DC-NET – Digital Cultural Heritage Network

OVERVIEW

The cultural heritage domain (libraries, museums, archives and other cultural bodies) has seen quite a revolution in the last twenty years. Increasing amounts of cultural heritage data are being converted to digital form and made available via the World Wide Web to researchers, students and the interested public. This represents an important opportunity for the improvement of the quality of life of the European citizen, for the delivery of cultural education, for the promotion of cultural tourism, for the development of the creative industry, for increasing the value which society enjoys from its common patrimony.

However, there has been no revolutionary technological leap forward since the digitisation and online publication began. Bandwidth, storage and search power are all major constraints.

The establishment of a dialogue and a factual relationship between the cultural sector, the research sector and the research infrastructures in Europe is the key factor to overcome these limits.

DESCRIPTION

Digital Cultural heritage has several exciting and important challenges which need to be addressed in order to increase the value that society as a whole derives from its heritage. These include:

- Multilingual and multi-national collaboration. Combining national cultural heritage research initiatives to European-wide and international scale;
- Access to everything. Enabling online access to more than a small proportion of holdings;
- Delivering the best. Providing access to the highest quality multimedia materials online, including audio-visual, 3D, text and images;
- Searching beyond text. Enabling users to use images, sounds, shapes and other aspects of items as search criteria;
- From active to passive. Allowing interaction and manipulation of digital items;
- Maintaining control and rights management. Cultural heritage material belongs to a huge number of autonomous entities, each with its own intellectual property rights;
- Long-term preservation of digital materials. Both digitised and born digital.

Each of these challenges can be met by exploiting the results derived from R&D projects for ICT applied to the cultural heritage, including national, regional and EC initiatives and by redeploying these results to e-Infrastructures.

In the framework of the DC-NET project, eight governments and research agencies are working to reach an agreement on a set of common priorities for digital cultural heritage research and are discussing with e-Infrastructures to identify how best to carry out this research using the National Research and Education Networks and the other grid providers.

However, the scope of DC-NET goes beyond the current ERA-NET project. The community of the digital cultural heritage is naturally larger than the eight Ministries participating to DC-NET. The network enlargement is on-going and many cultural institutions, governmental agencies and Ministries in Europe and beyond are in the process of formalising their participation to DC-NET.

IMPACT

At political level, there are two main bodies that should be mentioned: the Member States Experts Group for Digital Libraries established by the European Commission (MSEG); and the Joint Programming Initiative for Cultural heritage (JPI-CH) established by the Member States. The DC-NET community is fully involved in MSEG, where many experts of DC-NET participate as national representatives. With regard to the JPI-CH, DC-NET shares with NET-HERITAGE the operative foundation of the Initiative, being the former target to the research for the digital heritage and the latter to the research for the tangible heritage.

At project level, in parallel with DC-NET, the community is active at international level with the INDICATE project targeted to the Mediterranean region and in supporting Europe with many projects, among which are the two Best Practice Networks ATHENA and LINKED HERITAGE.

The next step in the vision of DC-NET is the deployment of a persistent and robust data infrastructure for digital cultural heritage, built on the facilities available from the e-Infrastructures sector. This data infrastructure should be able to deliver services and facilities that enable researchers to treat collections of data from many heterogeneous data sources as a continuum, overcoming linguistic, institutional, national and sectorial boundaries. It should also provide services and facilities for collection owners, so that they can store, protect and present their collections online without any loss of identity, ownership or intellectual property. Finally, it should address the most common causes of fragmentation in the digital cultural heritage sphere, including variations in policies, languages, metadata models and vocabulary. The design and implementation of such a data infrastructure dedicated to the digital cultural heritage is central to the vision of DC-NET and is at the core of its Joint Activities Plan.

CONCLUSION

DC-NET coordinates interaction and liaison across borders and also across the interface between cultural heritage and e-Infrastructure. Across borders, the Cultural Ministries of the EU work together to identify and prioritise the new actions which should be undertaken in order to bring digital cultural heritage to a new level of development. Across sector boundaries, the cultural heritage experts and the e-Infrastructure experts match the capabilities of the e-Infrastructure and the ambitions of the digital heritage community. The European Ministries of Culture, owners and managers of the national programmes for the research in the field of the digital cultural heritage interact with their respective national executive bodies responsible for the e-Infrastructure development to develop an action plan that will strengthen their cooperation at national and cross-border level, in pursuit of a common objective: a European research infrastructure for digital cultural heritage.

www.dc-net.org

antonella fresa fresa@promoter.it ICCU – Italian Ministry for Cultural Heritage and Activities

FP7-EU project 'DECIDE' (Diagnostic Enhancement of Confidence by an International Distributed Environment): a national prototype of e-services for the assisted diagnosis of neurodegenerative functional studies

OVERVIEW

Aim of the Diagnostic Enhancement of Confidence by an International Distributed Environment (DECIDE) project is to design, implement, and validate a grid-based e-Infrastructure building upon neuGRID and relying on the pan-European backbone GÉANT and the NRENs. Over this e-Infrastructure, a service will be provided for the computer-aided extraction of diagnostic markers for Alzheimer's disease and schizophrenia from medical images. DECIDE will offer access to large distributed reference databases (850 and 2,200 datasets from normal subjects and from neurological subjects, respectively), high computation and storage resources (more than 1,000 CPU core processors and 70 TB of storage) and intensive image processing tools. DECIDE will be open to the implementation of future algorithms based on other neuroimaging methods; in a longer-term perspective, the infrastructure could be extended to other research and diagnostic algorithms relating to other diseases of the brain and other organ systems.

DESCRIPTION

The project objectives are:

- Provide the neuroscientific and medical community with a dedicated grid-based e-Infrastructure building upon the FP7 e-Infrastructure neuGRID (www.neuGRID.eu) and relying on the pan-European backbone GÉANT and the NRENs; different competences and communities (neurological research, medical community, grid, networking) will be gathered, to implement the e-Infrastructure in an innovative way;
- Deploy on this e-Infrastructure a secure and user-friendly service for the early diagnosis and research on dementia and other brain diseases, and exploiting large distributed reference databases of multimodal neuroimages;
- Validate the e-Infrastructure and service through application to real patient cases (Alzheimer's disease, neurodegenerative dementias) and validate the DECIDE model on cutting-edge diagnostic conditions (e.g. schizophrenia);
- Propose a long-term vision for the sustainability of the e-Infrastructure and its extension to new communities; ethical issues related to the management and exploitation of sensitive patient clinical data will be specifically addressed as well as the business model for the wider exploitation of the service;
- Disseminate the results to promote the adoption of the DECIDE e-Infrastructure and service by the clinical community at large.

The planned tools will extract (i) diagnostic markers of established value in the medical community and might be used immediately in clinical settings, and (ii) markers currently in an earlier stage of validation, that will be brought forward by DECIDE.

The DECIDE applications will be implemented into a grid middleware that will allow:

1. authorisation and secure access to largely distributed databases for reference images,
2. computationally intensive processing,
3. image processing on patient images residing locally, compliant with the strict data-sharing hospital policies.

IMPACT

The potential impact for clinical use and research of the proposed e-Infrastructure will be on a large scale by enabling clinicians from hospitals with no access to sophisticated computational algorithms, resources, and large sets of reference images to carry out analyses remotely and efficiently. This will be reached by the use of a centralised user friendly web-grid service. The service will be accessed by the medical community through a science gateway portal.

The project will define with national and European regulatory agencies protocols and rules for the qualification of experts enabled to use the service for clinical diagnostic purposes.

The DECIDE consortium strongly believes that a European, and even wider approach is needed to tackle the problem of Alzheimer's disease and neurodegeneration in general. Furthermore the rapidly ageing society, that Europe and other developed economies has to cope with, make the societal emergency of the phenomenon more dramatic and resource-consuming. In this perspective, the project consortium and pilot infrastructure should be regarded as an extensible, proof-of-concept platform, whose usage should be extended to other communities in Europe.

CONCLUSION

The implementation of the DECIDE infrastructure and service should be regarded not only as a step towards streamlining and enhancing confidence in early diagnosis of neurodegenerative pathologies, but as a concept that can be successfully extended to other pathologies and communities, not only in the field of neurology (e.g. brain cancer), but to other organs as well (e.g. cardiology or traumatology). An international approach will be especially beneficial in view of such extension in scope, as it will help bring together other groups who are working on the subject in different fields, sharing with them approaches and results, and making the best of the state of the art in this field.

www.eu-decide.eu

Federica Tanlongo federica.tanlongo@garr.it GARR • **Laura Leone** laura.leone@garr.it GARR

GÉANT and EGI: enabling seamless distributed computing infrastructures through high-speed, dedicated networking

OVERVIEW

Over the last 15 years, science and research have become increasingly based on cross-border collaboration between researchers across the world. Sharing experiences, resources and facilities is, in many situations, the only way to ensure sustainability. In addition, science, research and education are making intensive use of high-capacity computing to model complex systems, to process and analyse experimental results, and to deliver them in real time in different locations around the globe. Computing grids for e-Science and Research and Education (R&E) Networks have emerged to respond to the requirements of the most demanding scientific disciplines to share and combine the power of distributed computation and the reliability of dedicated networks. The work presented aims at describing this synergy, discussing how advanced R&E networking provides the reliable connectivity which is enabling a scalable, sustainable distributed grid computing infrastructure in Europe and across the world.

DESCRIPTION

Distributed computing over large geographical areas and dedicated Research and Education (R&E) Networks have emerged to respond to the requirements of the most demanding scientific disciplines by combining the power of distributed computation and the reliability of dedicated networks.

GÉANT is the pan-European data network dedicated to the research and education community. Together with Europe's National Research and Education Networks (NRENs), GÉANT connects 40 million users at over 8,000 institutions across 40 countries. The resulting network provides an ideal infrastructure for any scientific and research application with demanding network needs in terms of sustained bandwidth, low latency (or one-way delay) and low jitter (latency or one-way delay variation). Grid computing is probably one of the most interesting and complex of these applications. The reliability of such a distributed computing facility is in fact, directly related to the reliability of the connections among computing elements, storage elements, resource brokers and grid user interfaces. The footprint of the GÉANT and the NREN networks, their coverage across Europe and the international links to networks in North America, Latin America, Africa and the Middle East, the Southern Caucasus, Central Asia and the Asia-Pacific region all enable a seamless, reliable, competitive infrastructure for the highly distributed grid computation EGI/NGIs aims to provide.

Finally, the GÉANT Project (GN3) has a task dedicated to Liaison and Support which focuses on providing assistance and support to projects, initiatives and organisations such as EGI, to facilitate and encourage a productive use of Research and Education Networks.

The presentation proposed will describe the importance of the crucial relationship between the network and the computing/application layers and how the GÉANT project is working in collaboration with EGI to constantly improve the research experience of scientists.

IMPACT

The presentation is mainly addressed at EGI/NGIs' grid computing developers, users and at site administrators, network administrators and PERT (Performance Enhancement Response Team) engineers.

Dealing with how the network layer is integrated with the distributed computing facility, the work submitted aims at considering the network as a stable, solid, yet dynamic and customisable resource to support grid-enabled research communities.

Finally, it will highlight how network-related activities such as monitoring and performance measurements can improve the perception of the grid computing resource from the perspective of the end users (researchers and scientists). Checking the status of crucial links between distributed computing and storage resources and regularly assessing the transmission cost of big data files across the network can effectively improve the overall computing operations within the EGI community.

CONCLUSION

Computing grids and Research and Education (R&E) Networks are increasingly growing to respond to the evolving requirements of demanding scientific and education disciplines. Their synergy is achieved by combining the power of geographically distributed computation and the reliability of advance networks dedicated to the research community.

The work presented aims at describing both the immediate and far-reaching effects of the interaction, discussing how advanced R&E networking provides the reliable connectivity which enables a scalable, sustainable distributed grid computing infrastructure in Europe and other parts of the world. Finally, the submitted presentation will describe the role of the GÉANT Project's dedicated task to Liaison and Support. This task is devoted to providing assistance and support to projects, initiatives and organisations such as EGI in order to facilitate and encourage productive use of the network. Some examples of support given to grid projects will be provided

Domenico Vicinanza domenico.vicinanza@dante.net DANTE

Logging and Bookkeeping – what can it do for you?

OVERVIEW

gLite's Logging and Bookkeeping service (L&B) has been with us for almost ten years as a robust monitoring service which gathers, aggregates and archives information on infrastructure behaviour from the perspective of users' tasks, and provides a single (per VO) endpoint for accessing that information. Numerous use cases have emerged over the years, some well known and widely applied in gLite-based grids, some not so widely known but possibly important for the user community.

The current situation gives us an opportunity to stop and take time to review L&B's potential for various usage scenarios and to plan its continued development according to real-world needs. We would like to involve current and potential L&B users (end users, groups, VO managers, monitoring tool developers/designers) in this process. To help this happen we give an outline of current L&B features and use cases (sometimes not widely known but ready to use or prototyped) and possible directions of further development.

DESCRIPTION

The most common user of L&B are the glite-wms-job-status and glite-wms-job-logging-information commands, querying the status of jobs handled by gLite WMS. It is less known that monitoring applications such as RTM, Experiment Dashboard and the grid observatory heavily rely on L&B, too.

Starting with its recent release 2.1, L&B is able to track native (i.e. non-WMS) CREAM jobs. Experimental support for PBS jobs and Condor is also available. L&B can also work with data transfers (e.g. sandbox transfers), and the job status information is enhanced with the state of associated data movements. We believe that a great opportunity to provide L&B users with a more accurate and descriptive view of their jobs lies in tracking the dependencies between computational and data transfer jobs, either directly through job state L&B queries or indirectly through monitoring tools providing VO- or site-specific views.

L&B can store application-specific information (job annotations, metrics, or status) as user tags. The tags can be queried, and used to build application-specific dashboards, as we have demonstrated several times. We also provide support for tracking generic user workflows currently used in medical image processing (a subject of a standalone contribution).

We can also demonstrate successful use of L&B to track different entities – CA revocation lists. Individual CRLs are registered in L&B in place of jobs. Whenever the CRL is updated, it is reflected in a state update of the corresponding 'job' in L&B, which triggers delivery of L&B notifications to the sites subscribed for receiving CRL updates. The payload of the notification carries the actual CRL update then.

IMPACT

The main purpose of this talk is to trigger discussion with the user community. It gives an overview of existing usage patterns involving the L&B service, emphasising less well-known use cases and applications often well outside the scope of computing job tracking. We invite users and user communities to provide their input. They are welcome to approach the L&B team with their priorities and requests to support other interesting use cases.

As a starting point for the discussion we present several topics L&B wishes to address in the near future. Feedback on priorities or possible extensions of the list is more than welcome.

- interesting entities to be tracked, besides gLite WMS jobs (e.g. ARC and UNICORE CE jobs, data transfers, SRM operations…),
- the ability to track dependencies among such entities (e.g. a computational job is blocked by the wait for transfer of its input),
- the desired level of complexity of queries (all jobs of a particular user, the user's jobs within a given time interval, failed/successful jobs on a given CE, up to full SQL/XQuery/etc., power on the job data),
- output formats to be supported (Glue-conformant WS interface, simple key=value text format, human-readable HTML, etc),
- history of the data to be kept (a day, a week, a month, etc),
- types of aggregate information, e.g. average queue traversal time, job failure rate etc., level of aggregation (per user, VO, grid service instance).

CONCLUSION

The grid environment has matured in recent years and essential functionality of the services is, more or less, available. The current development activities, besides pursuing standardisation and harmonisation of parallel solutions, aim at providing advanced functionality, which can be leveraged by the emerging grid user communities.

In this talk we present, besides a summary of less-known usage patterns of the service, a view on possible directions of further development of the L&B service as we ourselves can foresee its advanced use. The talk is intended to trigger discussion with the user community, which will result in a more specific and eventually extended work plan, better aligned with the expected user needs.

http://egee.cesnet.cz/cs/JRA1/LB

Zdenek Sustr sustr4@cesnet.cz CESNET • **Jiri Sitera** sitera@civ.zcu.cz CESNET • **Ales Krenek** ljocha@civ.zcu.cz CESNET • **Frantisek Dvorak** valtri@civ.zcu.cz CESNET • **Ludek Matyska** ludek@civ.zcu.cz CESNET • **Daniel Kouril** kouril@ics.muni.cz CESNET • **Jiri Filipovic** fila@ics.muni.cz CESNET

Supporting Distributed Computing Infrastructures

OVERVIEW

By establishing EGI.eu to create and maintain a pan-European Grid Infrastructure with the National Grid Initiatives as building blocks, Europe has taken a major step towards providing a long-term sustainable production quality grid infrastructure for multi-disciplinary use by international, national and regional user communities. With the aid of the EGI-InSPIRE project, EGI.eu coordinates between user communities, resource providers and technology providers within Europe and with partners from outside Europe.

As all these stakeholders are structured in their own organisations or projects, it is necessary to enable formalised communication between them. Here, EGI.eu plays the integrating role.

In the area of user support the integration platform is the Global Grid User Support (GGUS). GGUS is used as the central helpdesk for EGI, as well as for linking support infrastructures of other projects into the overall EGI set-up.

DESCRIPTION

Developed during the EGEE project, GGUS has also become the central integration platform for user support for EGI. Project-wide support is handled through the system and it acts as a central information hub for user, operations and technology issues. Integrating existing tools rather than replacing them, GGUS enables support teams from all areas to communicate with each other.

All NGIs need to interface their helpdesk systems to GGUS. To provide an easy way for small NGIs to do this, the xGUS helpdesk template has been developed. Instances of this template, while being hosted and maintained by the GGUS team, can be customised and updated by NGI staff. xGUS comes with the standard helpdesk features and the interfaces to GGUS already included. xGUS could also be of interest for user communities of the EGI infrastructure wanting to structure their internal user support workflows and to link them to the EGI support.

Several workflows that are needed between EGI and the technology providers are also realised through GGUS. Support requests by users are analysed within EGI by the Deployed Middleware Support Unit (DMSU) and are transferred by the DMSU to the respective technology providers' support teams, in case their problems are due to a bug in the middleware.

Releases will be announced by the technology providers by submitting a GGUS ticket of a special category. EGI then uses the same ticket to communicate its feedback to the technology provider. A similar workflow has been defined for middleware bugs found in production. The bug will be reported in a GGUS ticket to the technology provider, who will then use the ticket to communicate the date and number of the release fixing the bug in the same ticket. An additional benefit of handling these diverse workflows in the same tool is the possibility of creating statistics and reports on the metrics of the support process.

IMPACT

EGI is a major part of a larger environment of Distributed Computing Infrastructures (DCIs) from within and outside Europe. These infrastructures rely on the same or similar technologies and tools. The workflows defined in EGI for the communication between user communities, technology providers and resource providers can therefore be extended and adapted to be utilised other DCIs. This would have the benefit of providing the technology providers and user communities with one channel through which they receive or submit requests from or to the various DCIs.

Such a support infrastructure servicing various DCIs would help to harmonise the landscape for scientific computing on the European scale and could also ease the integration of extra-European DCIs. A major benefit would be that the user communities would be provided with one contact point regardless of which DCIs they are using.

Currently, various projects exist that aim at linking multinational grid infrastructures from various regions around the globe to the European Grid Infrastructure. The relationship between these initiatives and EGI, once established and formalised, needs to be aided by a technical infrastructure linking their regional support for user, operations and technology issues to EGI. GGUS could also play the role of the integration platform for these initiatives, by providing tools to set up the regional support, as well as by linking it to EGI. As a prototype a helpdesk for the EUMEDGrid project has been implemented using the xGUS helpdesk template.

CONCLUSION

GGUS has shown during EGEE and now in EGI that it has the capabilities of acting as this integration layer between user, resource providers and technology providers. Extending this support infrastructure to cover other DCIs as well would help to harmonise the European distributed computing landscape, would ease its multi- and interdisciplinary use and its extension beyond Europe.

This presentation will give an overview of the support infrastructure currently in place and plans how to further improve, extend and consolidate it in the near future.

Torsten Antoni torsten.antoni@kit.edu Karlsrühe Institute of Technology

HammerCloud: an automated service for stress and functional testing of grid sites

OVERVIEW

Many research communities rely on EGI resources to process vast quantities of data rapidly. Performance-critical activities such as this give motivation for tools which aid in the design and configuration of a grid site to ensure that its capabilities meet or exceed the requirements of the foreseen user applications.

HammerCloud (HC) is an automated testing service with two main goals. First, it is a stress-testing tool which can deliver large numbers of real jobs to objectively aid in the commissioning of new sites and to evaluate changes to site configurations. Second, it is an end-to-end functional testing tool which periodically sends short jobs to continually validate the site and related grid services.

In addition to delivering the test jobs to grid sites, HC provides a user-friendly web interface to test results. The service presents plots of job efficiency and many performance metrics to enable quick comparisons to past configurations and between other sites.

DESCRIPTION

HammerCloud was first developed for the ATLAS experiment at CERN. Prior to the turn on of the Large Hadron Collider (LHC), members of this research community were often participating in coordinated efforts to stress test the grid sites by simultaneously submitting their physics analysis jobs. Collating the results of these multi-user tests was time-consuming, yet the site performance evaluations these tests enabled were critical during the commissioning phase of many grid sites. In the end, these manual tests motivated the development of an automated service to carry out the stress tests and present the results via a web interface.

The HC service is composed of a back-end which submits and monitors the test jobs, and a user front-end which allows users to schedule on-demand tests and to watch the progress of running tests or review completed tests. Jobs are submitted and monitored using GANGA; this tool's grid Programming Interface provides an efficient framework to develop a grid service which needs the flexibility to submit arbitrary applications to any grid back-end.

The web interface is developed using Django and is designed to provide common web views in a core HC library while allowing Virtual Organisations (VO) to customise their web views in VO-specific plugins. Example metrics provided by HC include job success rates, timings of the various steps of a grid job (e.g. preparing input files, execution, and storing output files), and I/O metrics including storage latency and throughput values.

The current users of HC include three LHC experiments. ATLAS is the heaviest user, having used the service since its inception in late 2008, while HC plugins for CMS and LHCb were developed in 2010. The transition of HC from a single- to multiple-VO service required the generalisation of core components and the development of a plugin architecture. As a result of this work, HC can flexibly accept further additions of plugins for new communities.

IMPACT

The HammerCloud service empowers site administrators to undertake detailed studies of their site's capabilities without requiring any VO-specific knowledge or permissions. With only three clicks, HC users can schedule a test and shortly thereafter performance metrics are made available.

The experiences of the ATLAS experiment in using HC demonstrate the potential of such a tool. Since late 2008, ATLAS has invested more than 200,000 CPU-days processing HC test jobs globally. The primary focus of the ATLAS stress-testing has been on optimising the data access method at the sites. In particular, HC was used to compare strategies such as copying input files to a local disk (e.g. using lcg-cp, dccp, or xrdcp) against reading files directly from the site storage element using the local access protocol (e.g. rfio, dcap, xrootd). During large-scale global stress tests such as STEP09, HC was used to simulate the resource requirements of hundreds of real users by delivering a constant stream of up to 15,000 concurrent jobs throughout a two week period. Tests like this have led to I/O optimisations in the ATLAS software resulting in improvements to the overall job throughput on the grid.

For functional testing, HC delivers around ten types of test jobs for ATLAS. These tests validate not only the basic functionality of the grid sites, but are also used to test remote database access, to validate release-candidates of the grid middleware, and to compare data-access methods. Further, a subset of the ATLAS HC functional tests are deemed as critical – consecutive failures of these jobs result in HC taking action to blacklist the site from receiving user jobs. While the site is blacklisted, HC continues to send tests; when the jobs succeed again the site is informed and they can reset their site online at their convenience.

CONCLUSION

In summary, HammerCloud is an automated testing service which can be used for on-demand stress testing and continuous functional testing. Such a tool has been demonstrated to be useful during grid site commission and to evaluate changes to site configurations.

The LHC VOs have a strong interest in the service. ATLAS continues to use the service for on-going stress tests and site validation tests, while CMS and LHCb are currently incorporating HC into their grid operations procedures. CMS has thirteen HC users who can schedule on-demand stress tests, and a migration to HC from a CMS-specific tool for functional testing is planned. The LHCb instance of HC has been used to validate the rollout of new storage software at RAL, and further integration to the grid operations is in progress.

The flexible architecture of HC will allow future development of new plugins for research communities who have interest in the service.

http://hammercloud.cern.ch

Daniel van der Ster daniel.colin.vanderster@cern.ch CERN

Catering for different user profiles in the HELIO project

OVERVIEW

HELIO, an FP7 funded project, offers a set of tools to the heliophysics community to support scientific research. HELIO uses services specific to HELIO and also external services for storage and computation; the various services of which HELIO is composed are orchestrated through the means of workflows.

HELIO's architecture is complicated by a variety of factors mainly driven by the large variety of users it intends to support: while the normal users should be able to easily select pre-defined workflows, more advanced users should be able to define their own procedures either as workflows or as scripts and programmes. Also, scientists and the services back-ends have very different approaches towards security and privacy, all of which have to be catered for.

To meet all these requirements HELIO is built on a SOA, multi-layered architecture based on the following principles: service duality, workflow agnosticism, need-to-know policy, flexible deployment and policy compliance.

DESCRIPTION

The architecture of HELIO is composed by components arranged in multiple layers that span the conceptual space from the resources to the user. The services, their interfaces and the back-ends (that only some services rely upon) are of different kinds:

1. services that allow searches in external data sources;
2. services that allow an easy access to diverse data sources: they offer a layer of abstraction to the existing data sources;
3. services that offer functionalities such as computation and storage;
4. services that offer functionalities that are of no direct interest to the user (Authentication, Logging, etc...) but that are important to the system;
5. services that allow discovery and monitoring of HELIO services.

To grant access to these services in compliance of the HELIO principles, HELIO offers an ‘access layer’, an abstraction used to describe the different ways in which the services can be accessed.

There are three possible implementations of the access layer:

1. The HELIO user interface, which allows the user to invoke the various HELIO services independently or to connect them in patterns defined on the fly. The user interface also supports an instance of a workflow engine used to execute pre-defined workflows,
2. A desktop instance of a workflow engine used by the scientist to orchestrate the HELIO services in a user-defined fashion, and,
3. Direct access that allows the user to write code to directly use the HELIO services.

All the different implementations of the access layer connect to the services through the HELIO API that includes the stubs of the various services and offers an optimal entry point to the services. In particular, the grid services are designed to cope with the complexity stemming from the negotiation between users and back-end policies.

Security is based on authentication tokens with a low and high level of security to cater for all the permutations of user profiles and implementations of the access layer

IMPACT

The impact of HELIO covers two main aspects. The scientific aspect covers the impact that HELIO will have on the heliophysics community. To allow a profitable feedback with the users, special events named CDAW (Complex Data Analysis Workshop) are scheduled after every major release of HELIO. The first such workshop was held in Orly in November 2010 and showed a good interest of the scientific community in the features offered by HELIO.

The other aspect is technological: some of the problems faced by HELIO are of a wider interest as the need of negotiation between the needs of the different user profiles and the policies of the back-ends is common among infrastructures for the scientific community.

CONCLUSION

HELIO is now well into its second year of development, most of the query services have been developed and connected to the HELIO Query Interface with success, enabling services that provide security. Processing and storage facilities are now being prototyped and will be connected to HELIO in the next release in January 2011.

www.helio-vo.eu

Gabriele Pierantoni pierang@cs.tcd.ie Trinity College Dublin

WeNMR: the tale of virtual research community in NMR and structural biology

OVERVIEW

Structural biology and life sciences in general, and NMR in particular, have always been associated with advanced computing. The current challenges in the post-genomic era call for virtual research platforms that provide the worldwide research community with both user-friendly tools, platforms for data analysis and exchange, and an underlying e-Infrastructure. WeNMR, a three-year European Commission co-funded project started in November 2010, groups different research teams into a worldwide virtual research community. It builds on the established eNMR e-Infrastructure and its steadily growing virtual organisation, which is currently the second largest VO in the area of life sciences. WeNMR provides an e-Infrastructure platform and Science Gateway for structural biology. It involves researchers from around the world and will build bridges to other areas of structural biology.

DESCRIPTION

The main objective of WeNMR is to establish an e-Infrastructure-based global virtual research community for structural biology in the life sciences. To this end six objectives are defined:

1. to operate, maintain and further develop a user-friendly science gateway for the NMR and SAXS communities,
2. to establish a virtual research platform to serve as a digital knowledge repository, data exchange medium, and forum for (interaction with) the user community,
3. to provide support to software developers, users and other e-Infrastructure projects in an e-science knowledge and training centre,
4. to foster the adoption and use of e-Infrastructure on a global scale by supporting a wide range of flanking disciplines within the life sciences,
5. to operate and consolidate the eNMR grid infrastructure in line with NGIs and the EGI, and to extend it to interoperate with other worldwide grid initiatives,
6. to develop a model to ensure sustainability of the project.

The computational grid infrastructure is central to the services offered by WeNMR; its operation, maintenance and consolidation are therefore core service activities. The infrastructure is being expanded to worldwide grid initiatives, which is an important factor to both increase the available computational resources and the user community. WeNMR also operates a Virtual Research Community portal that will address the needs of the various user communities, from general public to experienced users and software developers. For each community, a number of services, information channels and communication tools is being deployed. WeNMR also aims at expanding the extensive human collaboration network successfully implemented within the e-NMR project geographically, i.e. involving users and other stakeholders from many new countries, in size and scientifically by involving new scientific communities. The inclusion of the SAXS community is a first step in that direction.

IMPACT

WeNMR will strengthen the European ties with National Grid Initiatives, the new EGI and PRACE initiatives toward an effective sharing of the offered service. Collaborations will be established with the Asian, South-African, South- and North-American grid initiatives, in close contact with existing EC projects, to extend and open the WeNMR Life-Science Gateway at a global level. WeNMR will further establish and promote best practices in life sciences and offers (virtual) training services to worldwide researchers.

CONCLUSION

WeNMR has established itself as a lively and growing Virtual Research Community in the life sciences and structural biology area. It already represents one of the largest if not the largest Virtual Research Organisation (VO), accounting for about 20% of the CPU on the grid within the life science area. Its global character is well illustrated by the international nature of its user community covering more than 30 countries, including South Africa and Latin America. To date, more than 10% of the users come from outside Europe. WeNMR has thus developed into a successful Virtual Research Community in structural biology for researchers from around the world.

www.wenmr.eu

Alexandre Bonvin a.m.j.j.bonvin@uu.nl WeNMR – Utrecht University • **Chris Spronk** chris@spronknmr.eu WeNMR – SpronkNMR • **Marco Verlato** marco.verlato@pd.infn.it WeNMR – INFN • **Antonio Rosato** rosato@cerm.unifi.it WeNMR – CERM University of Florence • **Harald Schwalbe** schwalbe@nmr.uni-frankfurt.de WeNMR – Goethe-University Frankfurt • **Ernest Laue** e.d.laue@bioc.cam.ac.uk WeNMR – Cambridge University • **Geerten Vuister** g.vuister@science.ru.nl WeNMR – Radboud University Nijmegen • **Gert Vriend** vriend@cmbi.ru.nl WeNMR – Radboud University Nijmegen • **Dmitri Svergun** svergun@embl-hamburg.de WeNMR – EMBL Hamburg

A suite for distributed execution of general purpose applications on the grid for small- and mid-size VOs

OVERVIEW

Many research activities from several fields, such as high energy, nuclear and atomic physics, biology, medicine, geophysics and environmental science, rely on data simulation and analysis tools, which are high CPU-consuming and/or data-driven. Sequential computation may require months or years of CPU time, so a loose-parallel distributed execution is expected to give benefits to these applications. In fact, the large number of storage and computation resources offered by a grid environment allows to consistently reduce the amount of computation time by splitting the whole task in several parts and executing each part on a single node.

High Energy Physics (HEP) experiments pioneered work on this but their results are hardly available to small and mid-size organisations that may have similar computational requirements, mostly due to the large amount of technical expertise needed. The goal of our work is to provide a suite allowing an easy, quick and highly customisable access to the grid.

DESCRIPTION

This work started from the need of the SuperB project (a HEP experiment) to simulate detector systems.

The effort in developing a distributed simulation production system capable of exploiting multi-flavour grid resources resulted in a general purpose design based on minimal and standard sets of grid services, capable to fit the requirements of many different Virtual Organisations (VO).

The system design includes a central EGI service site providing the web interface to job and metadata management tools. A database system is used to store the environment-related metadata; it works as a back-end for all the suite sub-services and communicates via RESTful protocol with running jobs on remote sites. The whole suite is based on standard grid services such as the Workload Management System as job brokering service and grid flavours interoperability element, the Virtual Organisation Membership Service as authentication and accounting system, the Large Hadron Collider Computing Grid File Catalog and Utils, GANGA as job management system and the Storage Resource Manager as the protocol for data access.

A web-based user-interface has been developed for database interactions and job preparation; it also provides basic monitoring functionalities. The web interface initialises the bookkeeping database and provides a submission interface allowing an automatic submission to all the available sites or a fine grain selection of job submission parameters.

The submission workflow includes the offline sites set up (job input file transfers to Storage Elements and VO software installation), and the bulk job submission via web tool suite. Jobs communicate status information to the central database; at job completion the output files are transferred to the central site data repository, metadata and log files are stored and location information registered into the LFC service.

IMPACT

The prototype suite we have developed has been successfully used in summer 2010 during a large Monte Carlo simulation production of SuperB.

The centralised access site has been the CNAF computing centre, which has provided the web-interface, the submission point, the bookkeeping database and hosted the data repository. Fifteen remote sites in Italy, France, UK, USA and Canada, which deploy three different grid middleware, have also been involved.

More than 11 billion simulated events have been produced. Over an effective period of four weeks, approximately 180,000 jobs were completed with an approximate 8% failure rate, mainly due to executable errors (0.5%), site misconfigurations (2%), proxy expiration (4.0%), and temporary overloading of the machine used to receive the data transfers from the remote sites (2%). The peak rate reached 7,000 simultaneous jobs with an average of 3,500. The total wall clock time spent by the simulation executables is about 195 years.

The distributed infrastructure and the developed suite have been fundamental in achieving the SuperB production cycles goals.

The online and offline monitor included with the web-interface keeps the metadata information stored in the bookkeeping database available for querying and further processing and analysis.

Our suite can be seen as a light-weight general-experiment framework which focuses on basic functionalities, designed specifically for organisations that cannot afford the use of the more specialised HEP frameworks but that still require an easy-to-use interface to the grid. Customisation of the web-interface, the bookkeeping database, job executable and site requirements are considered the key points to achieve the goal as well as small installation and configuration footprint.

CONCLUSION

The prototype suite has proven to be reliable and efficient although it still requires a careful initial configuration and needs a deeper abstraction from the application specifics.

Further development is needed to provide such an abstraction layer both in the web-interface and in the bookkeeping database. For instance, a user-defined database schema should be made available by providing an admin web-interface that allows the environment configuration, table modifications and setting of display and logic preferences. Moreover a better configuration structure is needed in order to allow site, job script, data management and executable customisation. At a higher level, the minimal set of grid services and protocols must be redefined in terms of standard compliance, lastingness and easiness, in order to expand the user base. An authentication management layer will be added in order to provide a direct access to the grid resources from the suite.

www.fe.infn.it/gridprod

Luca Tomassetti luca.tomassetti@fe.infn.it University of Ferrara and INFN • **Armando Fella** armando.fella@pi.infn.it INFN Sezione di Pisa • **Eleonora Luppi** eleonora.luppi@fe.infn.it University of Ferrara and INFN

Worker node software management: the VO perspective

OVERVIEW

Grid jobs often run software that is domain-specific, VO-specific and sometimes even user-specific. They depend on a software stack beyond what is generally available by default on grid nodes. During the Dutch VL-e project an additional distribution of software packages was created, the Proof of Concept (PoC), which is also distributed to the WNs on the BiG Grid infrastructure (Dutch NGI). However this was still very generic and it had very long update cycles and needed to be done by site administrators. For scientific experiments more specific stacks are required with shorter cycle times and with less (or no) participation of site administrators.

The toolset that we have developed offers centralised package management complete with dependency resolution, without the need for administrator's privileges. It uses the VO software area, which makes packages available on all the worker nodes.

DESCRIPTION

The grid is a (potentially) heterogeneous environment. Therefore simply distributing binaries is not always sufficient. We have therefore chosen to use a model where binaries are compiled remotely per site. This assumes that all resources in one site are homogeneous.

This very issue comes down to dependency management. Our solution allows a VO manager to maintain a stack of software libraries independent of what is provided by the WN software, so it can be updated asynchronously.

Also, multiple versions of libraries can be maintained side-by-side without interference through the use of the Environment Modules package.

With the help of the Environment Modules package [http://modules.sourceforge.net] we manage all the paths and environment variables that are needed so that the software can be correctly found and executed by the job.

Further tooling is developed to update, monitor and test installations at multiple sites, so the users can check the validity of all systems in a single view and decide which sites can run their jobs. All these tools run as normal grid jobs that are submitted to the site with the permissions of the VO manager.

IMPACT

Maintaining dependencies of software is a time-consuming task, but it is a very generic problem at the basis. By pooling from a large available set called pkgsrc [http://www.netbsd.org/docs/software/packages.html] and building the tooling for installation and maintenance on grid sites, the problem is reduced to defining packaging instructions and dependency resolution for the domain-specific parts. Any VO that needs to maintain a local software stack can adopt this method.

By transferring management of software packages from the infrastructure providers to the VO managers, the responsibility for the maintenance of the stack now lies with those who have the proper domain knowledge.

The toolset that we have developed offers centralised package management complete with dependency resolution, without the need for administrator's privileges. It uses the VO software area, which makes packages available on all the worker nodes.

Additionally, these methods are not limited to VO managers. Individual grid users can make use of the same tools to fit their grid jobs with the right environment, although they will have to do this on a job-by-job basis as they lack the privilege to make such environments persist between jobs.

CONCLUSION

We have presented a solution that allows VO managers to maintain a stack of software applications and libraries independent of what is provided by the WN software.

Although we have a working prototype, there are still some improvements that we want to make to address the following issues.

Many VOs in the EGI project might have similar software wishes, but because of communication barriers those similarities might be unknown. We therefore have some thoughts about a more collaborative distribution model that enables reuse of porting effort.

Our current implementation is only tested within the infrastructure of one NGI. We are therefore looking forward to testing our software with a VO that uses resources from multiple countries.

Mark Santcroos m.a.santcroos@amc.uva.nl Academic Medical Center Amsterdam • **Dennis van Dok** dennisvd@nikhef.nl Nikhef

A new integrated monitoring tool for grid site administrator

OVERVIEW

In this work we will show the development and the work carried on in order to build a monitoring tool that gives an aggregate view of all the users' activities on a given grid site.

The tool is able to show the job submitted by each user together with information about the file accessed on the storage system. Also in a farm with posix-like parallel file-system, the tool is able to track down both SRM standard operation and the local 'posix' file access.

We will put a particular highlight on how this monitoring system works in a mixed environment like a farm used both via grid and with local job submission. Moreover, it could easily work with different types of computing elements and batch systems, as it is highly modular and customisable.

This monitoring system will help the sys-admin to have a complete and detailed view of what is happening with the computing centre.

DESCRIPTION

The monitoring system was built starting from different agents and monitoring services. There is a central database system that takes care of storing, aggregating and presenting the information gathered from each monitored node.

In particular each computing element has its own agent in order to send the information about the jobs; indeed, this agent provides information about: user DN, FQAN, grid-jobid, local-jobid, queue, local user, VO.

Also the StoRM and the gridftp servers provide information about the file accessed both from the farm itself and/or from remote sites. Also in this case the monitoring agents provides: DN, FQAN, name and path of the file, VO. For each file accessed locally through lustre file-system, the local user that access the file, the node from which the file is accessed, the pid of the process accessing the file, the name of the process accessing the file.

Thanks to sensors installed in all the nodes of the farm it is possible for the site admin to know each accessed file over a Lustre/GPFS parallel file-system. The monitoring agent is as lightweight as possible in order to run it every one or few minutes.

By design the monitoring system allows the sys-admin to change, add, or switch off each plug-in used to find out the data.

The database schema is built in order to keep track of the dependency between several observed values: for example, it is easy to match between the job running on a given node, and the files accessed by the user on that machine. In order to gather as much information as possible and in order to be easily adapted to new batch system, we use pre- and post-exec scripts.

We have already developed sensors for several services: LCG-CE, CREAM-CE, StoRM, Gridftp servers, Xrootd servers, Torque/Maui, Lustre.

The web interface that allows the site admin to look at the status of the farm exploits newly available web technologies in order to give to the final user an advanced experience.

IMPACT

All information gathered in the central database could be aggregated and presented by users, host, processes etc., in order to have a much clearer view of what is happening on the computing nodes. Using this monitoring facility indeed makes it far easier to track down misuse of the computing facilities from a given user. The database was designed to speed up the procedure of building web pages. This was obtained by means of an accurate data movement between on-line and near-line tables and a deep study on how to speed up the query needed to build the interfaces. Using this monitoring tool it is also possible to keep under control the disk space used by each VO or user. This monitoring infrastructure was already used with success at the INFN-Bari site in order to have statistics on CMS dataset usage: this is particularly useful in order to take decisions about dataset deletion. It is also easy to monitoring the usage of the Wide Area Network bandwidth as it is possible to have a clear view of the users that are transferring files into (or out of) the farm, and it is also possible know which files are transferred. Looking at the job submission, using the web interface, it is possible to find out information on how many jobs are submitted from a given DN, or from a given VOMS group, which executables are running on the farm, and which nodes they are using.

It was already evident, while using it at a medium-large site made by 1000 CPU and 700TB of storage, that the system is able to handle a huge amount of data without decreasing the responsiveness of the web interface. The monitoring system also integrates several status views for each service that should be up and running in order to have normal operational conditions on the farm. This monitoring system is capable (differently from other systems already developed) to give a complete and aggregate view of the status of the farm together with the capability to have historical information on each observed metric.

CONCLUSION

In every grid computing farm, there are several services and each of those has different logging or monitoring system. It is important to have a unique point in which the site admin could aggregate and analyse all the information in order to have a clear view of what is going on.

This is more important in those computing farms that are shared among several VOs, in order to guarantee that activities of a VO are not blocked or affected by other users.

Using this tool the site administrators in the INFN-BARI farm were able to better fulfil the user requirements in terms of performance and reliability and also detect proactively problems and failures that could stop or affect the users' activities.

Giacinto Donvito giacinto.donvito@ba.infn.it INFN

Accounting and SLA support in KnowledgeGrid Malaysia by P-GRADE portal

OVERVIEW

KnowledgeGrid Malaysia is based on gLite. Since KnowledgeGrid Malaysia is used by several universities and user communities it was important to solve the problem of accounting and SLA support for the users. The single entry point for scientists is based on P-GRADE portal and hence the accounting and SLA support is built into the portal as new portlets. The talk will explain how KnowledgeGrid Malaysia and P-GRADE portal was extended with accounting and SLA support.

DESCRIPTION

The accounting portlet of KnowledgeGrid Malaysia portal presents accounting information from the connected grid sites on a per job and per user basis. The main feature of the accounting portlet is the traceability of consumed CPU time. In the accounting portal we can define two user roles: generic user and system administrator user. The generic user is able to check CPU time consumed by his/her jobs. The system administrator is able to check CPU time consumed by any user belonging to the grid.

The solution is based on the APEL accounting system, an optional module used in some EGI VOs. APEL serves as a back-end service for the newly developed account portlet. APEL is capable of collecting resource consumption statistics from EGEE Computing Elements and publishing this information into a central database. APEL can be installed as an optional component for EGEE grids.

The SLA framework can offer users of P-GRADE portal a guarantee that his/her application would start on the KnowledgeGrid Malaysia before a given deadline. The features of this SLA solution are the following:

The application start-up deadline is generated by the portal based on dynamic values that originate from the system state instead of being defined by the portal administrator. Several parameters are taken from the environment dynamically, such as:

1. Current number of users and application instances that run simultaneously within the portal and could delay the start-up time of the current application.
2. The number of resources that are available in the system and can be assumed that will operate correctly in the near future.
3. The length of the applications that are currently running in the grid.

The talk will explain the architecture of the accounting and SLA support system as well as their user interface within the P-GRADE portal.

IMPACT

KnowledgeGrid Malaysia is used by several universities and user communities and hence the balanced use of the grid among these communities required the measurement of the use of the grid sites as well as the provision of a guaranteed start-up time for the users.

The solution is generic enough that it could easily be adopted for any gLite-based VOs of the EGI community.

CONCLUSION

This talk explains the specifications and implementation of the applied accounting system and the complex SLA solution offered by P-GRADE portal for KnowledgeGrid Malaysia. This complex SLA allows the portal to calculate deadlines for workflow start-up before the submission of the workflow. Based on the calculated deadline, the user can choose if he or she wants to run the workflow within the given deadline or not. If yes, the jobs of the workflow will be sent to high-availability Tier 1 sites. If not, the jobs will be excluded from Tier 1 sites. The complex SLA uses dynamic values for calculating the deadline formula.

Peter Kacsuk kacsuk@sztaki.hu MTA SZTAKI • **Zoltan Farkas** zfarkas@sztaki.hu MTA SZTAKI • **Mohd Sidek Salleh** sidek.salleh@mimos.my MIMOS • **Nazarudin Wijee** naza.wijee@mimos.my MIMOS

UNICORE accounting service for Belarusian Grid Infrastructure

OVERVIEW

UNICORE middleware is widely used in the Belarusian Grid infrastructure, but it does not have any accounting and usage data collecting functionality. At the beginning of 2009, the development of a UNICORE accounting service was launched to monitor the employment of national grid resources. The service provides various accounting information from grid sites through a web-interface and helps users and administrators to know about every job of every user on every site. The presentation is mainly focused on implementation and technical details of the new service, but also describes some benefits of using it in the Belarusian grid.

DESCRIPTION

The service was developed with several requirements in mind:

- Ability to record the actual time of computing resources usage in normalised units based on sites' performance indexes
- Ability to identify commercial software usage
- Ability to interact with different categories of users and provide relevant information
- Securing all connections with user and host certificates

Basically, the accounting service consists of three main components:

1. Application server;
2. Database server and
3. Accounting agents.

The application server is based on Apache httpd server with mod_ssl and mod_php modules included. This solution makes it possible to run the accounting service under Linux or Windows systems. Application server is the main component with only passive behaviour. It does not make any network connections. It only responds to two kinds of requests:

1. user's requests and
2. accounting agents' requests.

They both communicate with the application server using the HTTPS protocol.

MySQL 5.x is used as database server for accounting application. Data tables use InnoDB storage engine. Performance tuning was done, but no extra patches were applied. To avoid data loss, binary logs and everyday backups to external storage were set up. For best performance, a database server was set up on a dedicated physical server.

An agent is used to extract accounting information from different sources, match it with XUUDB user record and upload results to the billing server, which verifies the format of the sent data. It is possible to develop a custom agent's version for any source, by bringing information into the recognised state. Each record of sent data must contain the following information:

1. task id;
2. certificate id;
3. task termination time and
4. used resources.

Data is uploaded via HTTPS request. Each agent had its own certificate signed with root UNICORE CA. An agent's certificate should be saved on the accounting server. Otherwise, the service would not recognise the agent and the connection would be rejected.

IMPACT

UNICORE accounting service is the solution of accounting problem in UNICORE grids. The service has been working in Belarusian grid infrastructure since 2009 and has already provided much useful information to users and admins.

With the help of the service, our users are able to monitor their own usage statistics and prediction in tables in charts, to generate reports, bills, and then print them. Admins can do much more. They are able to see the whole grid statistics and make decisions based on it. There are several scenarios for example:

1. One of the sites is under heavy load. Admins will find that many users need an application installed only on that site and then will make a decision to install the application somewhere else too.
2. One of the users generates a heavy load and makes other users' jobs have to wait for their turn. Admins can make a decision to set up priorities or even provide dedicated resources for active users.

Collected statistics help the developers of the UNICORE resource broker (another project of United Institute of Informatics Problems) to tune their service to make load balancing in the Belarusian grid more accurate.

CONCLUSION

The current version of accounting service due to technical requirements had been implemented as a third party web application written in PHP, so it cannot be one of the UNICORE parts. The possible future of this project is to implement the billing system as UNICORE official add-on using obtained experience with the full developing guidelines compliance.

Anyway, the service is an essential part of the Belarusian grid infrastructure, it provides flexibility and will be especially important if it will be decided to provide a paid access to grid resources.

http://account.grid.by

Andrew Lukoshko andrew.lukoshko@gmail.com United Institute of Informatics Problems of NASB

EMI

EMI represents the challenge taken up by the four major European middleware providers (ARC, dCache, gLite and UNICORE) to work on the long-term sustainability of the middleware as an essential tool for European research infrastructures like EGI. Sustainability can only come if important aspects of simplification, standardisation and quality are addressed in a collaborative way.

The first of three years of EMI has brought important achievements in the software engineering processes, the standardisation of the middleware and its functionality. The EMI sessions offer an extensive range of presentations about the activities of the past twelve months. Improvements in the way the middleware is deployed and configured have been introduced. Standard guidelines from mainstream operating systems have been applied and better compatibility with off-the-shelf package repositories like EPEL has been introduced. Standardisation activities have brought consolidation in a number of important areas, like the new EMI Execution Service providing for the first time common interfaces to the Compute Element systems; like the pervasive use of standard SSL technology; or advances in the definition and implementation of specifications like GLUE 2.0, Storage Accounting Records, WebDAV and NSF 4.1. There are technological improvements, like the extension and integration of MPI support across ARC, gLite and UNICORE and the introduction of solutions to long-standing problems in data catalogues synchronisation. At the core of these initial, but important results is the EMI software engineering process, based on standard ITIL methodologies and continuously monitoring the quality of software and procedures, collecting and analysing metrics, running tests and producing reports.

It has been an important first year for EMI and it has laid the foundations for even more important achievements along the hard road of sustainable, reliable and easily accessible infrastructures.

Alberto Di Meglio Project Director, EMI

What's new in EMI 1

OVERVIEW

The European Middleware Initiative (EMI) is a close collaboration of four major middleware providers, ARC, dCache, gLite and UNICORE.

The project aims to deliver a unified middleware distribution which will harmonise the various middleware products and focus on standards, interoperability, ease of use and reliability, while simultaneously improving functionality and responding to the needs of the user communities.

The first release EMI 1: Kebnekaise is due for release at the end of April. This session introduces some of the new key features of the middleware, and explains how the new version differs from the previous separate middleware releases.

DESCRIPTION

EMI 1: Kebnekaise introduces a number of changes which impact not only the functionality, but also the mechanisms for obtaining and installing the software, support channels, etc.

EMI 1 focuses on simplification and adoption of established best practices. It introduces significant improvements in interoperability between the various middleware products, for example the support for standard authorisation patterns which eases interoperability of grid credentials.

The process of obtaining and installing the middleware has also been simplified, with a single repository for all of the middleware packages, and EPEL-compliant component installation.

This session outlines the mechanisms for obtaining and installing the new distribution. It also introduces the new support mechanisms which the EMI project has put in place. The session will serve as an introduction to the hands-on EMI training workshops which will take place throughout the week and will cover issues relevant to all the EMI products.

IMPACT

EMI 1 represents a significant step forward in creating a unified middleware distribution with enhanced standards-compliance, interoperability, ease of use and reliability of the EMI products.

CONCLUSION

This session highlights some of the main changes introduced by EMI 1.

Kathryn Cassidy kathryn.cassidy@cs.tcd.ie TCD

EMI-ES: a common interface to ARC, gLite and UNICORE computing elements

OVERVIEW

The EMI Execution Service (EMI-ES) represents an agreement between ARC, gLite and UNICORE for a common web service interface to computing elements. This service deals with the submission and management of single-job activities on a computing element.

It covers important use cases from both high throughput and high performance computing.

DESCRIPTION

The EMI-ES specification has been developed by a team consisting of leading developers from the three grid middleware systems ARC, gLite and UNICORE.

The specification covers roughly the same space as the existing OGSA-BES and JSDL specifications, i.e. creation and management of single activities on a computing element.

However, there are evolutions and differences as well:

- re-designed activity description XML schema, covering data staging, execution, runtime environments, parallel applications and a simple yet powerful resource description (slots, memory, etc.). It can be seen as a merger of the JSDL 1.0 schema and some critical and often-used extensions
- support for bulk operations (e.g. bulk job submission)
- flexible data staging model, synthesising the capabilities of the involved middlewares
- simple state model
- consistently uses GLUE 2 as the information model for resource and activity descriptions

The technical realisation of the interface will be done using standard SOAP web services.

IMPACT

The EMI-ES uses the experience from the existing JSDL, OGSA-BES and some related specifications, combining them into a single coherent interface.

The EMI-ES will allow transparent job submission to the different middlewares, based on a common web service interface.

Further standardisation activities with the relevant OGF groups (PGI, JSDL, BES) are foreseen, in which this agreement between the European players is expected to have a strong impact. Conversely, we foresee future modifications of the specification, both necessitated by external requirements, and through the experience gained from implementing the EMI-ES specification, and doing interoperability testing.

CONCLUSION

The execution service agreement represents a significant milestone for the EMI project, which will enable transparent job submission to the different middlewares that make up EMI. Initial implementations are expected mid-2011.

https://twiki.cern.ch/twiki/bin/view/EMI/EmiExecutionService

Bernd Schuller b.schuller@fz-juelich.de Juelich Supercomputing Centre • **Balazs Konya** balazs.konya@hep.lu.se University of Lund • **Aleksandr Konstantinov** aleksandr.konstantinov@fys.uio.no Universitetet i Oslo • **Massimo Sgaravatto** massimo.sgaravatto@pd.infn.it INFN Padova • **Luigi Zangrando** luigi.zangrando@pd.infn.it INFN Padova

A uniform MPI user experience with MPI-Start in EMI

OVERVIEW

Execution of parallel applications on a grid environment is a key task which requires the cooperation of several middleware tools and services. Current middleware stacks provide basic access to such applications, which may be one of the reasons for its limited use. EMI faces the challenge of providing extended and seamless access to the execution of MPI jobs across the middleware it supports. This is being done building on an existing project, MPI-Start, originally implemented in the context of the int.eu.grid and which has been used in the gLite/EGEE grid over the last years. In this work we present the effort EMI is putting on MPI-Start, that is targeted on the provision of a uniform abstraction for the execution of parallel jobs and the integration of the tool with all the middleware stacks.

DESCRIPTION

Middleware support for MPI applications is usually limited to the possibility of allocating a set of nodes. The user still needs to deal with low level details that make the task non-trivial. Furthermore, the heterogeneity of resources available in grid infrastructures aggravates the complexity that users must face to run their applications.

MPI-Start provides users with an abstraction layer that simplifies the execution of MPI and other types of parallel applications on grid resources. By using a modular and pluggable architecture it manages for the user the details of:

- Local Resource Management Systems
- File Distribution
- Application Compilation
- Application Execution

Although MPI-Start is designed to be middleware agnostic, its original target was the gLite environment. In this work we have integrated it into ARC and UNICORE middlewares using their specific mechanisms, Runtime Environments of the ARC CE and Execution Environments of UNICORE.

The latest developments of MPI-Start have also introduced a new architecture for extensions that allows sites and users to modify the MPI-Start behaviour.

Moreover, MPI-Start now supports the basic mapping of logical processes to physical resources by allowing the user to specify how many processes they want to execute on each of the available hosts.

IMPACT

User communities have frequently stated that the support for MPI applications is too coarse and unwieldy for their needs. The processes of job submission and management should be simple, more transparent and should support generic parallel job types.

The adaptation and integration MPI-Start in the middleware stacks of EMI was performed by using the particular mechanisms of each middleware – Runtime Environments of ARC, Execution Environments of UNICORE and user wrapper scripts in gLite. Therefore users get a uniform experience for running their jobs with MPI-Start, while at the same time they maintain the interface for submitting the jobs of their middleware.

In most cases the automatic environment detection mechanisms of MPI-Start provide support for the most common MPI implementations and resource managers without any administrator intervention, thus allowing sites with little or no previous MPI experience to support these applications without major efforts.

The possibility of defining how user processes are mapped to physical resources and the modular architecture of MPI-Start, allows the execution of new kinds of parallel jobs. OpenMP jobs are now supported using the common interface of MPI-Start.

CONCLUSION

The unique interface of MPI-Start simplifies the task of starting the jobs by handling transparently the low level details of the different Local Resource Management Systems, execution frameworks and file distribution methods. The most common batch systems and MPI implementations are supported by MPI-Start and its modular architecture allows the easy extension of the tool to support new implementations.

With the integration of MPI-Start into the three computing middleware stacks of EMI, users have a unified experience for running their parallel jobs.

The availability of multi-core architectures opens the possibilities of new types of parallel jobs. MPI-Start's ability to define the way the user logical processes introduces the possibility of new types of applications in the tool, such as hybrid MPI/OpenMP applications.

Enol Fernandez del Castillo enolfc@ifca.unican.es CSIC

EMI Data, the unified European Data Management Middleware

OVERVIEW

At the time of the Vilnius event, one third of the EMI project life will have elapsed. Although EMI is built upon a rather complex structure, our customer infrastructures can expect that all the administrative gears are finally interlocked, that realistic plans exist for future software development, maintenance and deployment and that significant work has been done to achieve the goals planned for the first 12 months.

This presentation will focus on the work done within the data management part of EMI (aka EMI-Data). This includes, but is not limited to, the improvement of existing components, the integration of middlewares, the work on cross platform projects to enable interoperability and to provide standard interfaces to make EMI data components compatible with corresponding industry products.

DESCRIPTION

Besides improving individual components, EMI-Data has identified about a dozen cross middleware activities aiming to guarantee interoperability and to make EMI data components compatible with common standards. This includes work on the SRM specification as well as on its security mechanism. In terms of data access standards NFS 4.1 and WebDav will be offered by EMI storage elements. In addition to the work on interoperability of components and middlewares, flaws in the design of the currently deployed systems will be fixed as there is the problem of the inconsistency between the name spaces of file catalogues and storage elements. Pre-EMI activities like the migrating from GLUE 1.3 to 2.0, as well as the integration of the Argus authorisation framework, will be finalised within EMI-Data. Other pressing topics, like the specification and application of a common accounting record for data and a common monitoring framework will be tackled. Whenever possible, EMI-Data will use existing standards or will feedback specifications to standardisation bodies, like the Open Grid Forum. EMI-Data takes great care that modifications in specifications will be communicated and agreed on by grid frameworks outside of EMI, e.g. the US Open Science Grid.

IMPACT

EMI-Data is aiming for a universal data management system suitable for use in large and highly versatile infrastructures. It combines the expertise of the teams, working on well-established products, which have already been in production for many years. Under the financial and organisational umbrella of EMI, those teams now have the chance to build a complete system, covering the needs of infrastructures like EGI and PRACE and to attract new customers for the EMI distribution.

CONCLUSION

This presentation will report on the work of EMI-Data and on activities planned for the rest of the project. It will focus on the implications of EMI-Data work for EGI and its user communities. This contribution will also include details on those EMI-Data activities which do not have their own presentation.

https://twiki.cern.ch/twiki/bin/view/EMI/EmiJra1T3Data

Patrick Fuhrmann patrick.fuhrmann@desy.de DESY

Supporting SRM over SSL

OVERVIEW

SRM v2.2 is a standard protocol that allows end-users not only to manage their stored data, but also to negotiate transfer protocols between data exchanging parties, independent of which software is deployed. Storage software that supports SRM v2.2 is widely deployed within WLCG: there are currently over 400 SRM v2.2 endpoints registered.

The SRM protocol uses Globus' Grid Security Infrastructure (GSI) protocol as both a secure transport for sending messages and mechanism to allow delegation. Delegation is the process where an SRM server obtains a short-lived credential from a user, allowing the server to act on behalf of that end-user. Such delegation is required for some SRM operations.

GSI has several drawbacks: it isn't a standard, there is only one provider of GSI software libraries, and it is incompatible with SSL, the current standard secure transport. This leads to duplication of code, as some grid software uses SSL while others use GSI.

DESCRIPTION

Within the EMI project, work has started to first understand how SRM may be used over a standard SSL transport and then to extend the existing SRM clients and servers to achieve this. Since SSL does not provide a delegation mechanism, this work also involves establishing an alternative delegation mechanism for SRM. This work is conducted in collaboration with the SRM software providers outside EMI.

IMPACT

Allowing clients and servers to use the standard SSL protocol brings several benefits. It allows storage software providers to choose which security implementation to use, potentially bringing performance gains. It also allows better tools for diagnosing problems and consolidation of software towards standards: one of the aims of the EMI project.

CONCLUSION

Using the SRM protocol over standard SSL brings several advantages. In this paper, we discuss these advantages, present the progress in supporting SRM over SSL and detail the future plans.

Paul Millar paul.millar@desy.de DESY • **Alex Sim** asim@lbl.gov Lawrence Berkeley National Laboratory • **Jean-Philippe Baud** baud@mail.cern.ch CERN • **Ricardo Rocha** ricardo.rocha@cern.ch CERN • **Giuseppe Lo Presti** giuseppe.lopresti@cern.ch CERN • **Patrick Fuhrmann** patrick.fuhrmann@desy.de DESY

The ARC data client in EMI 1

OVERVIEW

The libarcdata2 library and the corresponding arc* data clients are successors of the widely used ng* data clients and provide a uniform way to move data from source to sink, with support for a number of data transfer protocols.

DESCRIPTION

The libarcdata2 library has a modular structure to support various data transfer protocols. The core libarcdata2 library builds on standard libraries and does not introduce any additional external dependencies. The plug-ins (DMC's) for specific data access protocols can however have various external dependencies. This separation of external dependencies from the core library minimises the prerequisites for an ARC installation, while allowing the support for additional access protocols requiring special dependencies to be installed by those who need it. Most of these components and the clients are also available on several platforms (Linux, Windows, Mac, Solaris), and libarcdata2 also has API's for Python and Java.

IMPACT

The libarcdata2 library and the arc* data clients replace the ng* data clients currently in production and will be available in the EMI 1 distribution.

CONCLUSION

The ARC data clients can be used as a general-purpose data client for all grid-related needs. The architecture, status and plans for the ARC data clients will be presented and usage examples will be shown.

Jon Kerr Nilsen j.k.nilsen@fys.uio.no UIO

Parallel NFS: EMI data in competition with high-end industry storage solutions

OVERVIEW

Providing storage solutions for large international e-science infrastructures, like EGI, is an interesting challenge, as such infrastructures are supposed to support a large variety of diverse science communities. Most of those communities have well established but unfortunately different patterns with respect to data management and data access. The only way to tackle this problem, in a generic way, is to provide industry standard mechanisms in these areas. Besides other solutions, the data management part of EMI is approaching the idea of standards-in-storage by offering access to its storage elements through the NFS 4.1 (pNFS) protocol currently being adapted by leading industry storage providers. (e.g. IBM, Panasas, NetAPP etc.). This approach makes EMI data solutions competitive with expensive industry solutions and as a side effect, EMI and the corresponding infrastructures are no longer in charge of maintaining the data access software on their data client entities and with such, resources are freed.

DESCRIPTION

For the last years, the Centre for Information Technology Integration (CITI), as part of the University of Michigan has been working on a specification of a successor of the overwhelmingly successful Network File Protocol (NFS), focusing on today's needs in terms of speed, reliability, security and modern distribution pattern of data sources. The effort was supported by the major storage solution vendors, including Sun Microsystems, IBM, NetAPP, Pansas, Mircosoft and dCache.org, as an unfunded partner. This work resulted in a specification which is now being implemented by all participants of the effort.

Parallel NFS (pNFS) is a part of the NFS v4.1 standard that allows clients to access storage devices directly and in parallel. The pNFS architecture eliminates the scalability and performance issues associated with NFS servers in deployments today. This is achieved by the separation of data and metadata, and by moving the metadata server out of the data path.

The design of pNFS perfectly matches the design of the storage elements provided by EMI. dCache, being part of the CITI group for some years, already offers a production version of pNFS with its current releases. DPM provides a beta version of pNFS available for early testers. A production release is expected before EMI-2. StoRM will provide pNFS as soon as IBM makes pNFS available for GPFS which is expected in 2011.

Besides the intense development work for DPM and dCache in this field, dCache.org has been putting significant efforts into testing of pNFS with available Linux clients. A small but realistic Tier2 has been built including CPU and storage resource at DESY. pNFS client server interactions have been tested with plain I/O, ATLAS and CMS analysis software and with specific setups provided by the ROOT team. DPM will catch-up in terms of testing by making the DPM/pNFS implementation available to DPM-sites in the UK for community feedback.

IMPACT

In current deployments of EMI data storage elements, specific proprietary software has to be installed and maintained on data client nodes. Moreover for RFIO, dCap and others, libraries have to be linked with the applications, which makes this approach useless for communities where the source code of the application is not available. For other products, data clients might even have to be purchased. (e.g.GPFS).

With the common availability of pNFS in EMI storage elements, client software will be provided by the OS vendors, which will free resources in EMI and very likely in the infrastructure support (EGI) as well. Furthermore, with pNFS, the repository of EMI storage elements can be mounted into the file system as any other local storage resource. Any unmodified application will have access to data using plain POSIX I/O, an essential requirement for non-HEP communities.

Last but not least, pNFS enables EMI and with that EGI, to provide storage solutions which can easily compete with expensive industry products and can be used as a drop-in replacement.

CONCLUSION

The just-in-time availability of pNFS in industry (IBM, NetApp, Oracle) and Open Source products (Linux kernel) allow EMI to provide professional storage elements, which can easily compete with expensive high-end storage solutions. The discomfort of currently deployed systems, in terms of product dependent protocols and client software which needs to be installed and maintained, will be overcome. Significant testing of pNFS in EMI data products as well as in industry products is in progress and more results will be presented at the time of the Vilnius meeting.

www.dcache.org/manuals/2010/20101102-hepix-patrick-nfs41.pdf

Patrick Fuhrmann patrick.fuhrmann@desy.de DESY/dCache.org

gLite File Transfer Service in the EMI data area

OVERVIEW

We propose a new gLite File Transfer Service (FTS) that implements the latest requirements of experiments and serves the needs of smaller (Tier 2 and Tier 3) sites as well.

DESCRIPTION

We summarise the limitations of the current FTS architecture which pre-supposes a strict tier model which the experiment workflows have now largely moved beyond. To address these problems, to enable user initiated workflows and to promote efficient resource usage by a robust and easily maintainable file transfer service, we propose redesigning FTS based on standard building blocks, in three steps.

Step 1: remove the channel model, taking into account the published load on an individual SE and the network, while supporting other database back-ends

Step 2: introduce a pull model of job processing based on messaging. Main advantage: a load balanced service where the client does not need knowledge of which FTS to contact. The current push model will be maintained for backward compatibility.

Step 3: Merge FTS and LCG_Util functionalities. Besides the advantages in resource handling and unified UI, FTS users get access to the catalogue interface as well.

IMPACT

With the new FTS and LCG_Util, EMI will have a scalable set of data transferring tools with simple configuration, lower administration burden and better usage of storage and network resources. FTS will serve the needs of T2-T3 sites as well. In addition, ordinary users (not only VO admins) will also be able to schedule file transfer jobs.

CONCLUSION

We have proposed redesigning FTS and (in the long run) merging its functionalities with LCG_Util using standard building blocks. We believe that the new FTS would extend the usability of our data transfer services, lower administration and maintenance costs and serve a much broader set of use cases according to the latest requirements of the experiments.

Zsolt Molnar zolt.molnar@cern.ch CERN

Storage catalogue consistency

OVERVIEW

This talk describes the status of the work performed in EMI to find practical solutions to the problem of maintaining consistency in the grid data catalogues. This activity aims to provide grid Storage Elements and Central Catalogues the means to interact and keep their metadata synchronised.

.

DESCRIPTION

The EMI Catalogue Synchronisation Working Group has produced a prototype system built on standard messaging tools that successfully integrated with LFC and DPM. This was used as a proof of concept to show that the objective can be achieved and extended to the other Storage Element implementations in EMI. Now the activity of the working group is headed towards developing a more comprehensive specification of the messages that are needed to fulfil the objective, and towards starting to address the aspects related to security and deployment.

The initial prototype has generated new ideas that may yield a higher degree of integration with the whole distributed storage infrastructure.

IMPACT

Within the EMI project, finding a way for Storage Elements and catalogues to interact in real time in order to keep their metadata synchronised will give a higher degree of robustness to the data infrastructure. The deployment of such a common synchronisation subsystem will help the overall data management systems in giving a better service to the users.

CONCLUSION

The EMI data group has proposed a draft specification and a prototype system built on standard building blocks with the goal of making Storage Elements interact for metadata synchronisation purposes. The ideas behind these steps will be presented, together with some possible enhancements and directions.

Fabrizio Furano furano@cern.ch CERN

Software quality assurance in EMI

OVERVIEW

In EMI, the definition and monitoring of software quality assurance is the responsibility of the SA2.2 team. This talk will present the procedures defined to carry out this activity and how to monitor their application within the EMI software lifecycle.

DESCRIPTION

SA2.2 is the task responsible for defining and monitoring the software quality assurance process within the EMI project. One of the documents produced by SA2.2 is the Software Quality Assurance Plan (SQAP) which specifies the tasks that are needed to ensure quality, the responsibilities for quality monitoring, the required software component documentation and the procedures that will be followed to manage the software process. In practice the SQAP specifies the manner in which the EMI project is going to achieve its quality goals that are defined in terms of quality factors like usability, maintainability or portability. In order to better describe the software process, a set of guidelines have been also written giving specific details to product teams in various stages of the software lifecycle: integration, configuration, packaging, testing and certification, change management and release management. The process of defining the guidelines has involved experts from the middleware distributions and the other activities within the project. SA2.2 is responsible for managing changes in these guidelines and announcing them to the project members to make sure all product teams are up to date with the documentation governing the software lifecycle. SA2.2 also works very closely with the quality control tasks in development, maintenance and support activities. SA2.2 defines the release criteria for the EMI software components that will be checked by the quality control tasks. SA2.2 produces regular periodic reports evaluation the application of the existing procedures and the availability of proper documentation for each software component. This contribution will present the main aspects of the EMI SQAP and will focus on those guidelines related to development, change management and release of software, providing information useful for software developers and testers.

IMPACT

This presentation describes how adopting standard software engineering and quality assurance processes within academic research environments can become a realistic and beneficial activity. The work performed in EMI is not only targeted at improving the quality and timeliness of software delivered to the infrastructures, but also at establishing a more standard and efficient software production lifecycle as the base to make the development and maintenance of software more sustainable.

.

CONCLUSION

This presentation is a good opportunity to understand how software quality assurance has been implemented within EMI, understanding the procedures, roles and responsibilities involved in the software lifecycle and the monitoring activities carried out to control the quality goals.

Maria Pradillo maria.alandes.pradillo@cern.ch CERN

Software quality assurance tools in EMI

OVERVIEW

One of the important goals of the EMI project is the unification and standardisation of the software engineering and quality assurance process. This can be made possible only by providing a single and unified tool chain throughout the project which will be used as a standard environment to enforce common procedures and constraints. This tool chain needs then to become a stable, trustable and reliable foundation for the production of high quality software throughout the project lifetime.

DESCRIPTION

In EMI, the selection, integration and maintenance of quality assurance tools, starting with the identification of the tools initially used by each middleware distribution and the identification of new tools which may be required by other activities is the responsibility of the SA2.4 team. The main functions of this team are:

1. to support all the required tools and service installations for the duration of the project
2. to plan the introduction, change and removal of services as smoothly as possible providing a stable and reliable infrastructure
3. to integrate different tool chains used by each distribution in a single tool infrastructure able to provide at least the same functionalities
4. to enable continuous integration and testing processes by selecting and maintaining tools and resources for building and testing software either within the project or in collaboration with external resource providers

IMPACT

The EMI project started as a consolidation activity of the four European middleware distributions, namely ARC, dCache, gLite and UNICORE. Each of these distributions started several years ago from specific needs dictated by their own scientific communities. Moreover the distributions developed along the years in different and independent ways aiming at solving different scientific problems and adopting different strategies on software architecture and design. As a result, their culture, processes and tools appear today diverse and in some aspects, contrastive. The unification of this heterogeneity to provide a single and unified quality assurance infrastructure is the first challenge the quality assurance team must face.

CONCLUSION

This presentation shows the status of this activity by providing a description of the tools selected, their integration and their use in the project. Starting from the build system, going through ticketing systems, quality assurance tools, testing infrastructure and dashboards, ending with the production package repositories, every aspect of the software engineering process and quality assurance is illustrated from the tools point of view.

Lorenzo Dini lorenzo.dini@cern.ch CERN

Definition, reporting and analysis of EMI software metrics

OVERVIEW

Within EMI, the definition, continual collection, reporting and analysis of software quality metrics are handled by the SA2.3 team. The team provides information to the project's management bodies on the status of the software as an instrument to taking corrective actions. The success criteria of this task are the regular production of reports and their use to identify areas for improvement, as shown in this contribution.

DESCRIPTION

Metrics in EMI are to be generated internally per month and visibly per quarter. The metrics are broken into two main groupings: process metrics and software metrics. The process metrics relate to each middleware bug-tracker, the GGUS bug-tracker and each software release. The software metrics relate to building, testing and static code analysis. Exceptions to these two groups appear when a metric is confined to a software component, usually normalised according to the number of bug-tracking defects open against that component.

This work is seemingly easy at first glance, but is complicated by the diverse number of software languages in the EMI middleware. In addition, each of the four middleware involved in EMI uses its own bug-tracking system in very different ways. However, metrics must be calculated in exactly the same way across multiple product teams, product components and bug-trackers to compare like with like. For this reason, an XML schema was generated to map each bug-tracking system into one unique format which can be understood by the metrics collection mechanism. Additionally, many procedures have been implemented to ensure that software metrics are produced coherently across all products.

Practically speaking, the visualisation of specific statistics in metric reports is currently used to present a very precise overview of specific parts of the process. For example, with careful construction of suitable metrics in EMI, it is not only possible to say how many components are building (as in more build systems), but it is also possible to suggest the parts of the project that are creating bottlenecks in the release process. For instance, a normalised histogram presenting the number of successful/failing software components in a product team, singling out the external dependencies and the dependencies provided by other product teams components, will help define the sources of any delay.

IMPACT

Attempting to collect metrics changes the landscape of the working environment of a project. The inability to generate certain required metrics has resulted in project procedural changes, which then allowed the metric calculation to progress. The knock-on effect therefore, is that metrics are automatically causing the cohesion of the four middlewares into one distinct middleware. The metrics will help pinpoint problems on a weekly basis as well as a quarterly basis, by generating reports for the EMI Executive Management Team (EMT). The definition of carefully chosen metrics such as: success builds per product team, average time to fix a top priority bug, number of untouched open bugs older than two weeks old, delays in release, back-log in bug handling per quarter, bug severity density, FindBugs plots per product team are all helping to produce a more heavily tested, stable middleware.

CONCLUSION

Metrics are a vital part of the quality assurance process for large scale projects such as EMI. Results will be presented for priority, severity, density and delays for the bug-trackers of each middleware. The comparison of results per product team, products, components inside products gives an overall view of how the project is progressing, whilst highlighting the areas/products that require fixing or streamlining. Box and whisker charts showing medians and quartile ranges as well as histograms showing means resulted in plots that were more informative and easier to analyse. The uniformity required to generate metrics across many products and product teams gave rise to changes in procedure for each middleware. In summary, metrics help the middlewares to highlight problems in specific products and enforce bug-tracking uniformity and correct usage.

Eammon Kenny eamonn.kenny@cs.tcd.ie TCD • **Gianni Pucciani** gianni.pucciani@cern.ch CERN

EMI Integration and large scale testing infrastructure

OVERVIEW

Releasing software with production level quality, i.e. satisfying end user expectation in terms of functionality and performance, is the final goal of every software collaborative project. Software continuous and effective testing is a key step in the software production process to match quality targets. The presentation will describe infrastructural resources provided by the EMI quality assurance work package for the project continuous integration and large scale acceptance testing. The two testing infrastructures follow different implementation approaches reflecting the different testing goals. Integration testing testbed must provide a snapshot of release products to allow for testing correct mutual interaction of all related products in the release. Large scale testbed must reproduce end-user production environment for specific test use case, therefore requiring an on-demand model where end users are involved in specific testing scenarios.

DESCRIPTION

The set up of integration and large scale acceptance testing infrastructures for EMI projects was a three step process:

1. testing objectives and requirements collection;
2. implementation model definition;
3. actual set up and revision in collaboration with testbed users.

Integration testing is the part of testing and certification process of a software product where the product functionalities and expected behaviour are tested against other related grid service components. In EMI's decentralised software development model, testing and certification are charged to different product teams, each responsible for one or more software components. Taking place after functionalities, testing of products in insulation has been successfully carried out. Integration testing then represents the first centralised point of contact among different products. Therefore the solution focused on permanently deploying instances of both production and release candidate versions of all products components for every EMI release. Hence a continuously evolving snapshot of all released and upcoming versions of all EMI products is provided to product team developers. For flexible and dynamic creation of testbed subset views, configurable central information system instances publishing resources in the testbed were set up.

A different model was adopted to implement Large Scale testbed infrastructure, necessary to perform acceptance, interoperability and scalability tests. Here the main objective is to stress the service under test reproducing conditions as similar as possible to real production environment in terms of geographical distribution or different deployment scenarios or scale. Therefore a demand (requests for specific testbeds defining use case and target scale) and supply (collaborative effort from a community of partners available to participate to specific tests campaigns) approach was adopted allowing to set up testbeds fitting the specific test use case needs.

IMPACT

The EMI project is a close collaboration of major European grid middleware providers (ARC, gLite, UNICORE, dCache) whose mission includes the harmonisation and evolution of the middleware, ensuring it responds effectively to the requirements of the scientific communities relying on it. Therefore the impact of setting up a centralised EMI testing infrastructure in the project's first year is two-fold:

1. providing the testing and certification infrastructure, a central service for the project, and
2. somehow having a living laboratory where actual integration and interoperability problems can be practically experienced, identified and approached.

Considering the first, and the more visible of the two mentioned outcomes, the infrastructural and operational resources put in place consist of:

1. hardware resources amounting to roughly 100 server instances deployed across five geographically distributed sites including monitoring, resource publishing and testing security utilities;
2. support handling, activity coordination infrastructure and documentation;
3. model definition and collaborative effort coordination channels set up for the large scale testing infrastructure

Notice that also providing a testing infrastructure for both intra- and cross-middleware integration testing has posed new implementation challenges. On the other hand, the large-scale infrastructure model, though still under validation about effectiveness and sustainability, brings both a flexible approach, adaptive to variable testing conditions required and a useful communication feedback channel with enduser communities. Finally a less visible but equally important outcome of the work was the contribution to identification and analysis of practical issues concerning the merging and harmonisation process among different middlewares converging into EMI. As an example consider cross-middleware resource discovery and publication or common authentication/authorisation framework for testbed accessibility.

CONCLUSION

The main outcome of the work was the model definition and implementation of EMI's centralised testing infrastructure together with the operational resources necessary for its daily usage and continuous update. Also an 'on demand' model for implementing large scale acceptance testbeds has been defined to match specific testing use cases with end-user community resource providers.

Among the outcomes of the work it is useful to mention the identification of key open issues to be addressed for assuring the testbed homogeneous usability by all middlewares converging into EMI. Future work will then focus on tracking and integrating in the testbed the solutions to these issues, among which the most relevant with respect to direct effort contribution is the adaptation of the testbed infrastructure implementation to the planned automation of testing process.

https://twiki.cern.ch/twiki/bin/view/EMI/TestBed

Danilo Dongiovanni danilo.dongiovanni@cnaf.infn.it INFN

The gLite Data Management continuous integration and testing process

OVERVIEW

The continuous execution of builds and tests on newly developed software is a key activity for any software development team. To properly implement such procedures it is fundamental to automatise all tasks related to software building, deployment, and testing.

The gLite Data Management (DM) product team has put in place a continuous integration and testing procedure to properly support the development and release of the product team software products: DPM/LFC, FTS, and GFAL/lcg_util.

DESCRIPTION

A continuous integration and testing procedure must integrate the execution of source code builds, the deployment of packages, and the execution of different types of tests under a common continuous execution plan. The execution plan defined for the DM product team defines the following tasks: building the source code from SVN (taking code from three repositories: LCGDM, LCGUTIL, FTS); creating platform specific YUM repositories containing all generated packages (including meta packages which represent the services provided by the product team); verification of the correct deployment and configuration of all services; execution of a number of functional, regression, and stress tests; and finally the creation of reports for the product team members. This execution plan is executed n-times on all supported platforms (currently slc4_ia32_gcc346, slc4_x86_64_gcc346, sl5_ia32_gcc412, and sl5_x86_64_gcc412).

The automatic execution of these tasks is managed by a new tool, called SAKET (Swiss Army Knife for ETICS Test), which not only orchestrates the entire process but provides reports to the product team members. These reports are sent by email and archived in a web accessible location for historical reference. SAKET also relies on other existing tools:

- ETICS: to run builds and tests on remote infrastructure nodes;
- YUM: to make available the packages generated by the build;
- YAIM: to deploy the available packages and to configure the services;
- YAIMGEN: to execute the pre-defined tests and to generate reports.

Besides these tools, the continuous integration and testing procedure also relies on a well-defined controlled environment to ensure the reproducibility of the build and test executions. This environment is composed by the ETICS infrastructure nodes where the builds and tests actually run and the EMI certification testbed which provides other external services, such as BDII.

IMPACT

Currently, the DM product team runs every night this procedure on the latest code produced by developers. As such, all changes committed by developers during the day are immediately built and tested every night. Furthermore reports are automatically produced and sent to developers informing them on the impact of their changes. Potential problems are discovered in advance and the code quality increases.

The same process can also be executed during the release of new software versions. In this case the advantage is not on the continuous verification of the code but rather on the easiness to verify if the new release passes all existing tests. This directly contributes to a decrease in the time needed to complete the release process.

CONCLUSION

A procedure to continuously verify the integrity and stability of the products being developed in the gLite DM PT has been put in place and is being exploited on a daily basis. This procedure relies on and integrates a number of existing tools making more efficient and easier the development and release of new gLite DM releases.

Pedro Andrade pedro.andrade@cern.ch CERN • **Oliver Keeble** oliver.keeble@cern.ch CERN • **Alejandro Ayllon** alejandro.alvarez.ayllon@cern.ch CERN

ARC tools for revision and nightly functional tests

OVERVIEW

Identifying errors in the early stage of software development ensures a lower cost development process and higher quality final software release. Automated testing procedures in EMI performed during ARC middleware development are done by two independent tool sets. First is the ETICS build system and second is the ARC automated build and functionality test system.

DESCRIPTION

Automated builds enable us to track errors which could break the integrity of the final code. For each code change committed to ARC subversion the tool checks status of following build steps: Autogen, Configure, Make, Make Check, Make Dist, and Build. The check is done for several platforms (including Mac OS X) and subtrees of ARC sources codes.

IMPACT

The functionality tests are periodically running carefully chosen test scenarios on ARC “nightlies” or on specific code revisions. The goal of these test scenarios is to check a basic functionality of ARC clients and services. During these tests we additionally check integration of clients and servers. The test scenarios, build and functionality test results are stored in a database so they can easily be visualised and exported in any form using simple PHP or python scripts. The request for a new test scenario is done through a simple web form.

CONCLUSION

We present automatic tools which are used to detect build and functional errors in the early stages of the software development process.

http://arc-emi.grid.upjs.sk/revisionTests.php

Jozef Cernak jcernak@upjs.sk UPJS Kosice • **Marek Kocan** kocan.marek@gmail.com UPJS Kosice

EMI and dCache: the vanguard of standards adoption

OVERVIEW

dCache is a hugely successful storage software that is currently deployed in eight out of 11 Tier-1 centres in WLCG and providing over 40% of the total storage capacity available to the LHC experiments. With dCache as a founding member, the EMI project has great experience in storage.

The High Energy Physics (HEP) community has a tradition of building their own protocols for LAN and WAN access to data. This is because the existing protocols have proven unable to sustain the required throughput.

Unfortunately, these non-standard protocols introduce a major barrier to new communities adopting grid computing. These communities may have custom software that would need to be made ‘grid-aware’; the end-users may be using commercial off-the-shelf analysis software that cannot be modified.

Recent developments in the industry mean that a large storage system, such as dCache, must adopt industry standards to remain competitive with solutions from HPSS, NetApp, EMC^2, Panasas, Dell-DDN, etc., without sacrificing throughput.

DESCRIPTION

The EMI project includes a programme where the storage elements are adopting industry standards. Within this programme dCache has taken a trailblazer role: demonstrating the feasibility of adopting these industry standard protocols and demonstrating their functionality in production environments.

WebDAV is an industry standard protocol that allows end-users to read and write data. dCache has been adopting WebDAV, allowing end-users to read and write data over the WAN. Within dCache, access via WebDAV is no different from other protocols: end-users experience a coherent namespace and are subject to the same set of access controls.

For local access dCache has implemented the NFS v4.1 protocol with pNFS, one of the first storage solutions to do so. Through the use of realistic tests, dCache has demonstrated deploying NFS v4.1 for the high-throughput analysis work typical of HEP communities. There are sites in the process of rolling out NFS v4.1 support.

IMPACT

dCache's adoption of WebDAV means that end-users can access their data on their choice of platform, with their choice of access tool. There are WebDAV clients for all major computing platforms, which are often supplied as part of the operating system.

Several sites are deploying dCache with WebDAV specifically to allow end-users outside the HEP community to have access to large data storage; for example, dCache and WebDAV forms the basis of the Swedish National Infrastructure for Computing (SNIC) national storage facility: SweStore. SweStore aims to support non-HEP communities such as climate modelling, bioinformatics and bioimaging.

With the availability of native support for NFS v4.1 in the various operating systems, a site-administrator can simply mount the large data stored in dCache on local analysis machines. Programmes running on those machines may access data like any other file, without modification. This permits dCache to support analysis work-flows that use software that cannot be modified to support HEP-specific protocols, which permits analysis facilities and compute farms to support a much wider range of analysis activity.

Adopting NFS v4.1 also brings advantages to site administrators. Adoption of standard protocols allows a mix-and-match approach, where storage from different vendors may be combined, based on their relative benefits, to produce a coherent storage solution. The standard also provides better tools for monitoring.

CONCLUSION

By adopting standard protocols, such as WebDAV and NFS v4.1, dCache is allowing end-users to be agnostic to the nature of the storage. This lowers the barrier for new communities that need access to large storage, allowing them to use grid resources without modifying their software. This work is already bringing benefits to end-users.

Paul Millar paul.millar@desy.de DESY • **Patrick Fuhrmann** patrick.fuhrmann@desy.de DESY • **Tigran Mkrtchyan** tigran.mkrtchyan@desy.de DESY • **Gerd Behrmann** behrmann@ndgf.org NDGF • **Dmitry Litvintsev** litvinse@fnal.gov FermiLab • **Tanja Baranova** tatjana.baranova@desy.de DESY • **Christian Bernardt** christian.bernardt@desy.de DESY • **Karsten Schwank** karsten.schwank@desy.de DESY • **Antje Petersen** antje.petersen@desy.de DESY • **Per Lundqvist** perl@nsc.liu.se National Supercomputer Centre in Sweden

Proposing a Storage Accounting Record standard

OVERVIEW

To enable a shared storage infrastructure, it must be possible to account for and report on the resources consumed by persons and groups in a common format. While usage record formats for computing jobs have been defined by OGF (GFD98) and used for several years, there is no corresponding definition for storage accounting records. The EMI data group, with data experts from ARC, dCache, gLite and UNICORE, has seen the need and defined StAR – a common storage accounting record to be proposed to OGF.

DESCRIPTION

The StAR definition describes a format, which can account for resources consumed on a storage system. The format allows for granularity on a storage system level (e.g. separating disk and tape consumption) as well as reporting on both a per-user and a per-group basis. The different properties enable a variety of aggregation scenarios.

Due to a general approach the StAR format is not restricted to usage in a distributed (grid) environment but can be applied in any storage environment.

IMPACT

Storage accounting records, as defined in the final StAR specification, will be implemented in ARC, dCache, gLite and UNICORE, thus providing means to monitor and account for distributed storage in a standardised fashion.

CONCLUSION

The EMI data group has proposed a standard for storage accounting records. The ideas behind the StAR definition and the envisioned implementation and usage will be presented.

Jon Kerr Nilsen j.k.nilsen@fys.uio.no UIO

Implementation of GLUE 2.0 support in the EMI Data Area

OVERVIEW

We report on the implementation of support for the GLUE 2.0 information model in the EMI Data Area.

DESCRIPTION

The EMI data area team is committed to the implementation of GLUE 2.0 support across the relevant services and clients. This has involved reaching agreement on a common EMI interpretation of the schema to ensure a consistent implementation across components.
We describe the various tasks required to realise the GLUE 2.0 support and describe the status of the work. In addition, detail is given on validation, testing and rollout as well as plans to further exploit the potential of the new schema.

IMPACT

The GLUE 2.0 schema was designed to allow a more accurate and natural description of the resources available on grid infrastructures. By supporting this, the EMI data team is promoting the optimal use of grid resources.

CONCLUSION

EMI has committed to the implementation of GLUE 2.0 support across the relevant services and clients.

Oliver Keeble oliver.keeble@cern.ch CERN

Virtualisation and Cloud Computing

Cloud computing has emerged in the last years as the new paradigm for the provision of on-demand computing resources. More generally, virtualisation is growing, in terms of use cases and interest from both user communities and service providers.

Many user communities have very specific configuration requirements for their software. Using virtualisation the applications runtime environment can be ported through different sites. Virtual networks favour communities that have particular privacy requirements, by configuring Virtual Private Networks as an additional layer of security. Virtual resources can be accessed through the grid interface, as well as the cloud paradigm; the latter should focus on different use cases, for instance interactive applications as well as accessing SaaS (Software as a Service) resources.

Infrastructure Providers can run their services over the cloud to improve availability, scalability, and maintainability of their grid and non-grid services. Moreover, physical resources costs can be optimised by deploying services on private clouds. Data Centers can expose their resources as an 'Infrastructure as a Service' (IaaS) cloud, as well as deliver 'Software as a Service' (SaaS). Generally they can achieve a better exploitation of their resources and can offer a diverse set of services, attracting in this way different user communities. For a grid site deployed on a private cloud it is also easier to include resources from public (commercial) clouds.

The enormous potential of these technologies has attracted the interest of e-science communities, and now many projects (e.g. StratusLab, Venus-C, ECEE) are working on the cloud itself and on the integration between cloud and grid.

Peter Solagna Operations Officer, EGI.eu

Operating grid services on the StratusLab cloud

OVERVIEW

Cloud computing has emerged in recent years as the new paradigm for the provision of on-demand distributed computing resources. The inherent capabilities of the cloud architecture for optimised resource management and consolidation appear very appealing for grid resource providers. StratusLab is a European-funded project that started in 2010 with the purpose to investigate possibilities, implications and optimal solutions for the combination of cloud and grid technologies. The project is integrating its own cloud distribution based on the OpenNebula cloud management middleware. The first versions of the distribution have already been released and a reference cloud service is being offered to the public. This public cloud currently hosts a minimal but complete grid site. In this paper we present the first experiences gained by operating this grid site on the StratusLab cloud and the impact on users and their applications.

DESCRIPTION

The StratusLab project started in June 2010 with the purpose to investigate the impact of the emerging cloud computing paradigm in the provision of grid computing services. The primary motivation was the belief that these two technologies are complementary and their combination can offer new capabilities and optimise resource utilisation. Where cloud offers flexibility in resource management the grid provides high-level services that enable collaboration and resource sharing among dispersed scientific teams. StratusLab focuses on the Infrastructure-as-a-Service (IaaS) cloud paradigm, which implies the usage of virtualisation technologies for the provision of computing resources.

The project is integrating a cloud distribution, specifically designed with the purpose to host grid services. During the design phase the specific requirements and/or restrictions of grid services are taken into account in order to provide optimised cloud environments for deploying virtualised grid sites. These requirements are both technical and operational. In the heart of the StratusLab cloud distribution lies the popular OpenNebula (www.opennebula.org) cloud management toolkit which is being extended with additional capabilities, either developed within the project or integrated from the ecosystem of existing add-ons. The first version of the StratusLab distribution was released in October 2010. Incremental or bug-fixing versions are released every six weeks. The distribution is used by the project itself to set up and provide a reference cloud service.

Currently two capabilities are available to the public: a cloud computing IaaS service giving the ability to users to instantiate and manage VMs, and an appliance repository where the VM images are stored. This reference cloud service is used also internally by the project as testbed for deploying grid sites and in order to investigate potential implications of their operation over the cloud.

IMPACT

Cloud technologies can significantly alter the way grid services are currently deployed and operated. By exploiting the inherent capabilities of IaaS clouds for on-demand provision of elastic computing services, grid resource providers can optimise the utilisation of their physical computing resources.

The primary user of StratusLab is the grid site administrator, yet the adoption of the cloud paradigm is expected to indirectly have a positive impact to the VO managers and end-users. The project also prepares and makes available through the appliance repository a set of VM images, with grid middleware pre-installed, for the basic grid components (CE, SE, WN). The usage of VMs accelerates the instantiation of a grid site and makes it easier to try out new features or validate middleware updates.

Moreover, for VOs and end-users the ability to use VMs with pre-installed applications and scientific software, provides additional flexibility for deploying grid services customised for specific application domains. Also, end-users are expected to be impacted positively by the expected enhanced stability and availability of grid sites running on virtualised environments.

CONCLUSION

StratusLab is developing a grid-optimised cloud distribution based on OpenNebula. This distribution is used to offer a reference public cloud service which is used as a testbed for deploying virtualised grid sites. The operation of these grid sites is helping us evaluate the applicability of cloud technologies, their impact on end-user applications and to gather requirements for grid middleware. As the distribution evolves new capabilities will be added enhancing the current cloud solution. In parallel the public cloud service is attracting real-life applications which will help us benchmark the technologies and validate the applicability of StratusLab cloud distribution for the operation of production-level grid sites.

www.stratuslab.eu

Vangelis Floros efloros@grnet.gr GRNET

Data management at VENUS-C

OVERVIEW

VENUS-C is an FP7 project aimed at bringing together existing DCIs and commercial cloud providers to create a sustainable ecosystem for e-science applications. One of the key aspects of this ecosystem is data management.

More and more academic groups store and work on their data outside the university domain. While most cloud storage systems offer convenient access APIs, they are rarely fully compatible with each other. This makes migration between back-ends, or integration of data located in various places, more complicated. A common answer to this vendor lock-in problem lies in adopting common standards. Although there are no commonly adopted standards for cloud storage, yet, the potentially most promising direction at the moment is SNIA's Cloud Data Management Interface (CDMI), which defines the semantics of handling data objects and the required metadata.

In this talk we present our approach to managing data in a heterogeneous environment: a prototype implementation of CDMI that offers a transparent integration layer on top of both cloud systems and local storage infrastructures.

DESCRIPTION

VENUS-C data management solution consists of two main parts:

1. CDMI-compliant (http://cdmi.sniacloud.com) server that is used for bridging together a number of back-ends, both from the public cloud computing offerings and local infrastructures. On the technical level it works by separating metadata storage (datastore) and data abstractions (blobs, message queues). Current list of supported back-ends includes for blobs: POSIX, AWS S3, MS Azure Blob; for message queues – AMQP, AWS SQS, MS Azure Queue; and for datastore: CouchDB, AWS SimpleDB and MS Azure Table.
2. Client-side SDKs in several languages that allow interaction with the CDMI-server.

CDMI-compliance means supporting of standard operations on blob data items and message queues (table storage support is planned in the next version of the standard). Apart from the core CRUD functionality for those objects, more advanced features, specific for the VENUS-C offering, are foreseen that take into account aggregated meta-information, for example, optimisation of storage costs and transparent encryption of data when storing outside the domain of trust.

Implementation is based on Python and uses Twisted networking framework. We also rely on the services provided by Twisted for implementing access controls (Twisted Guard) and logging. For connecting to the components' back-ends, we either rely on the existing SDKs, for example, boto for Amazon Web Services and py-amqplib for talking to AMQP-compliant queue engines, or implement the missing link ourselves.

IMPACT

With our solution we are aiming at providing a way for applications to benefit from sustainable public cloud storage offerings while avoiding the interface-level vendor lock-in. Moreover, integrated storage management creates a marketplace for possible CDMI service deployments, as back-ends could be potentially brokered from other domains (both commercial and academic).

CONCLUSION

During the advent of cloud computing there was a lot of discussion about the traditional DCI model being extended or, in some cases, replaced by commercial offerings. Our work builds on the evolving standards in the field of data management in clouds and aims at providing an integration solution for both DCIs and cloud providers.

www.venus-c.eu

Ilja Livenson ilja.livenson@gmail.com KTH

Mantychore – dynamic networks for e-Infrastructure

OVERVIEW

Mantychore aims to provide dynamic networking as a service for e-Infrastructure users.

The e-science community has expressed a need for flexible and application-driven networks. Traditional methods for configuring networks for research users are impractical when the scale is international and involves several networks.

Mantychore takes an 'Infrastructure as a Service' (IaaS) approach to networking to enable National Research and Education Networks (NRENs) and other e-Infrastructure providers to deploy a configurable, automatable network that allows virtual research communities to control routers, switches, optical devices, and IP networks as necessary to meet the needs of their applications.

Grid and cloud computing will provide some of the first real uses for this technology. Mantychore's facilities should make it easier to share computing and storage resources between grid sites or to create a distributed computing cloud spanning several institutions.

DESCRIPTION

The Mantychore FP7 project aims to allow the NRENs to provide a complete, flexible network service that allows research communities to create an IP network under their control, where they can configure:

Layer 1, optical links. Users will be able to get permissions over optical devices like optical switches, and configure some important properties of its cards and ports. Mantychore will integrate the Argia framework which provides complete control of optical resources.

Layer 2, ethernet and MPLS. Users will be able to get permission over Ethernet and MPLS (Layer 2.5) switches, and configure different services. In this aspect, Mantychore will integrate the ETHER project and its capabilities for the management of Ethernet and MPLS resources.

Layer 3, IP. Mantychore FP7 suite includes a set of features for:

- Configuration and creation of virtual networks.
- Configuration of physical interfaces.
- Support of routing protocols, both internal (RIP, OSPF) and external (BGP).
- Support of QoS and firewall services.
- Creation, modification and deletion of resources (interfaces, routers) both physical and logical.
- Support of IPv6. It allows the configuration of IPv6 in interfaces, routing protocols, networks.

Mantychore FP7 is the natural evolution of its parent projects: MANTICORE and MANTICORE II. MANTICORE implemented a proof of concept which proves the idea of IP network and router as a Service (IPNaaS), all this management in Layer 3 where MANTICORE works. MANTICORE II continued its steps to implement stable software with the feedback, expertise and know-how received. Also, MANTICORE II added new and improved capabilities to the software and ran a success pilot project where MANTICORE II partners got the chance to run some trials on their own network and equipment. The Mantychore FP7 project consolidates this trend.

IMPACT

Mantychore FP7 will carry out pre-operational deployments of the IP network service at two NRENS: HEAnet and NORDUnet. Part of the project effort will be dedicated to consolidate and enhance the community of providers (NRENs but also commercial) and users of the IP network service.

Initially three communities of researchers will benefit from this service: the Nordic Health Data Network, the British Advance High Quality Media Services and Grid-Ireland, the Irish NGI.

The Grid-Ireland Operations Centre (http://grid.ie) in Trinity College Dublin runs Ireland's national computing grid infrastructure which connects grid nodes at many higher-education institutions around the country and is part of the European grid infrastructure.

Grid and cloud computing use-cases will be explored by the Grid-Ireland Operations Centre. As Grid-Ireland has resources distributed across Ireland and connected via HEAnet it is an ideal testbed for Mantychore.

The Grid-Ireland team will explore novel approaches to sharing and connecting computing and storage resources. This will include connecting grid worker nodes at multiple physical sites into a single virtual grid site over a dynamically configured virtual private network. Further investigation will lead to creating a distributed computing cloud between institutions, with an intended use as a platform for agent-based computing. This will raise issues relating to adding and removing sites, and relating to dynamic provision of extra links for data migration for virtual machine migration and backups.

In general, by working with research communities in higher-education and research institutions Mantychore will gather real-world requirements for connecting dynamic and user-controlled networks to traditionally managed networks.

CONCLUSION

In addition to the central development and deployment of dynamic networking as a service for e-Infrastructure, Mantychore will innovate the business model used for services based on IaaS, establishing a marketplace where all infrastructure providers can announce their available resources and all customers can automatically negotiate SLAs getting the best resources combination for their needs. Furthermore, Mantychore services will be used to contribute to the research performed in the GreenStar Network project to enable carbon-neutral infrastructures.

The Mantychore FP7 project is committed to incorporating as many viewpoints and potential uses as possible in order to reach a more complete and valuable software and expertise pool. For this reason, the project resources are open and available for any interested individual or community to join.

www.mantychore.eu

David O'Callaghan david.ocallaghan@cs.tcd.ie Trinity College Dublin • **Peter Lavin** lavinp@cs.tcd.ie Trinity College Dublin • **Sergi Figuerola** sergi.figuerola@i2cat.net i2cat • **Pau Minoves** pau.minoves@i2cat.net i2cat • **Victor Reijs** victor.reijs@heanet.ie HEAnet • **Dave Wilson** dave.wilson@heanet.ie HEAnet • **Per Nihlen** per@nordu.net NORDUnet

Federated cloud computing environment for malaria fighting

OVERVIEW

The research group of Dr. Hugo Gutierrez de Teran, Galician Public Foundation for Medical Xenómica (FPGMX) and member of Scientists Against Malaria, needed to perform, using different computational methods, virtual screening of a chemical library of 350,000 chemical compounds to study a selected therapeutic target to further develop, based on the findings, new drugs against the parasite Plasmodium, which causes the disease.

A cloud computing collaboration with FCSCL, using the framework of the cooperation agreement signed months ago between the presidents of the regions of Galicia and Castilla-León was used as the fastest way to solve this challenge.

DESCRIPTION

The aim is to generate a list of potential candidate molecules to pass Bioassay in parallel using two independent methods of ligand-receptor coupling or docking. In the field of molecular modelling, this is a method that predicts the preferred orientation of a molecule that will be linked to another to form a stable complex. Knowledge of the preferred orientation can in turn be used to predict the strength of the association or the binding affinity between two molecules.

Prior to this study, researchers from the group of Dr. Gutierrez de Terán generated a collection of possible molecular species for each of the chemicals, considering all the possible isomers (compounds with the same atomic composition but different structural formula), tautomers (two isomers that differ only in the position of a functional group) and protonation states (addition of a proton to an atom or molecule) which led to a collection of 1,238,000 processed molecules. This type of calculation requires a large amount of computational resources, which are only possible using large data centres.

CESGA technicians used the cloud developed in the framework of the NUBA project, extending it to use resources from FCSCL. The management of both sites was performed using computer software developed in Spain, OpenNebula, which manages the nodes in both locations, deploying virtual machines in either one or the other depending on the criteria of optimising computing resources.

IMPACT

This is one of the first international experiences in which two supercomputing centres create a set using HPC Cloud computing and storage equipment distributed between the two institutions. For the connection, the centres used the network of academic research and RedIris, which provided the necessary bandwidth.

CONCLUSION

The collaboration between both super-computing centres in the implementation of this cloud, in addition to being essential support in the development of international research groups, is a step in the research of technologies for cloud computing services. The success of the experience opens the door to new collaborations between national and international research groups, which could deliver very demanding computer applications for research projects almost immediately.

Carlos Fernandez carlosf@cesga.es FCTSG

Searching Eclipse behind the cloud

OVERVIEW

Emerging cloud infrastructures offer new ways to develop dynamic services. Management, provisioning and inter-operability are the main challenges among others. Within the Eclipse community, there is a clear trend from the desktop usage to a general runtime system. With the help of the Eclipse ecosystem (www.eclipse.org), a toolbox of solutions exists already today to develop new services. The Equinox runtime of Eclipse – based on the OSGi standard – can be seen as the basis for an interoperable cloud platform in the future.

An introduction to the various Eclipse projects will be given and distributed search applications will be set up on cloud resources (AWS or/and Eucalyptus) during the talk.

DESCRIPTION

Emerging cloud infrastructures offer new ways to develop dynamic services. Management, provisioning and interoperability are the main challenges among others. A short time-to-service is important to all kinds of applications on grid and cloud infrastructures. With the help of the Eclipse ecosystem (www.eclipse.org), a toolbox of solutions exists already today to develop new services. As the underlying technology is based on standards (JAVA and OSGi) the application is independent from the underlying infrastructure.

This talk will explain the overall picture of OSGi applications on the cloud. Furthermore it will demonstrate the set up of a search application in the cloud; including the management of the cloud, the dynamic provisioning of the application itself and the provisioning of an advanced AJAX-based user interface on the cloud. This will be performed with the help of the following Eclipse projects:

- g-Eclipse will be used to manage and configure the virtual cloud resources based on its general Cloud model. Access to Eucalyptus and AWS resources is supported today.
- p2 will be used to deploy the search application.
- SMILA (SeMantic Information Logistics Architecture) is an extensible framework for building search applications for data like office documents, emails, images, audio and video files, blogs, etc. One of the features of SMILA is the parallelisation of processes/workflows, so the natural deployment environment of SMILA is similar to the distributed environment of the cloud.
- RAP will be used to create a simple search user interface for the application based on JAVA script and AJAX technologies.

This talk demonstrates existing goodies from various Eclipse projects, which can help to build cloud applications independent from underlying infra-structures. It will show the potential power of Eclipse technology in the cloud.

IMPACT

The Eclipse community is one of the biggest open source communities worldwide with strong support from industrial players. In an open community process the important trends for the next year were defined and they contained amongst others:

- Cloud and browser-based delivery of applications is the future.
- Cloud and browser-based developer tools needed to support this future process.

The demonstration will show what is already possible today with existing open source technology in this area. An alternative approach for dynamic service delivery will also be shown.

CONCLUSION

The set up of distributed OSGi-based applications and the management with an integrated user frontend is possible today. The combination of existing Eclipse technology helps to build a full cloud service stack including IaaS, PaaS and SaaS layer.

www.eclipse.org/geclipse
www.eclipsecon.org/2009/sessions?id=336

Harald Kornmayer harald@kornmayer.eu DHBW Mannheim

GPU computing in EGI environment using a cloud approach

OVERVIEW

Recently GPU computing, namely the possibility to use the vector processors of graphics cards as computational general purpose units of High Performance Computing environments, has generated considerable interest in the scientific community. Some communities in the European Grid Infrastructure (EGI) are reshaping their applications to exploit this new programming paradigm. Each EGI community, called Virtual Organisation (VO), often requires specific environments, making it necessary for each grid site to enable an efficient system to fulfil VO's software requirements.

Cloud computing and more generally the opportunity to transparently use computational resources, together with the consolidation of virtualisation technologies, provides end-users with the required environment for their activities.

The present work aims to provide an on-demand GPU environment (GPU framework, Operating System and libraries) for each VO and to make it accessible via the production infrastructure using the cloud.

DESCRIPTION

In the present work, the possibility of enabling the convenient usage of GPU devices for VO users, exploiting the capabilities of EGI, and the emerging paradigm of cloud computing are explored.

A strategy to provide on-demand execution environments has been proposed through the joint usage of traditional and widespread gLite components and the popular standard EC2 web-service APIs.

An entire job flow that enables the Local Resource Management System (LRMS) to discriminate the GPU resources requests, through GLUE Schema parameters, has been defined to allocate, in a dynamic fashion, the required resources on a cloud-like infrastructure either public, private or hybrid.

To achieve this goal, part of the work has been devoted to the virtualisation of the physical GPU resources in order to make them available in an Infrastructure as a Service (IaaS) private cloud.

To this end a centralised mechanism, responsible to listen for events generated by the LRMS like job scheduling and termination, has been implemented to keep track of each request.

These events are then used to carry out the required actions as follows: once a job is received and identified as a GPU usage request, it is treated as an event that triggers the allocation of virtualised resources according to simple leasing rules. In a similar way the termination of jobs are notified to a daemon that releases the execution environment.

In order to develop and test the whole infrastructure, a fully working test bed has been built with the adoption of the Eucalyptus software system to implement a private cloud over the cluster.

We have also addressed the need to create Virtual Machine Images that match the requirements of the execution of GPU-dependent jobs, such as CUDA, OpenCL libraries and gLite middleware.

IMPACT

GPU computing is increasing in the EGI communities, starting from computational chemistry (COMPCHEM VO) to theoretical physics (THEOPHYS VO), as well as satisfying the needs of other communities.

The main purpose of the present work is to dynamically provide a ready-to-use GPU environment for the communities using EGI to share GPU resources, since a single GPU environment does not satisfy the different requirements of all communities (such as operating systems, compilers and scientific libraries). For this reason, the developed system provides dynamical environments with the aim to optimise GPU resources usage.

Contextually, the cloud computing opportunity allows taking into account the GPUs as a Service (IaaS). From a cloud point of view, the project carries out a feasibility study to understand how the next evolution of the grid computing to cloud computing, or better, how the switch from Batch Model to Service Model could be done.

The approach adopted in this system is not focused only on GPU computing and it can be easily extended to other special hardware devices and, in general, to other environments (as in the IaaS Model).

CONCLUSION

In this work we presented a system to provide to EGI users a specific on-demand GPU environment to transparently execute jobs on GPU devices.

The proposed system uses a cloud approach, based on EC2 compliant clouds, in order to control the specific GPU-enabled VO's environment from EGI middleware interfaces.

The system is currently in its testing phase at UNI-PERUGIA grid site and supports COMPCHEM VO. In this phase the LRMS is used to join the EGI infrastructure and the UNI-PERUGIA private cloud. We are planning to decouple the on-the-cloud allocation mechanism from the LRMS and place it at the computing element level using for example the CREAM Architecture capabilities.

This will allow a fine grained control over the Virtual Instances and the Accounting.

Alessandro Costantini alex@dyn.unipg.it UNIPG • **Osvaldo Gervasi** osvaldo@unipg.it UNIPG • **Flavio Vella** flavio.vella@dmi.unipg.it UNIPG • **Riccardo Cefala'** riccardo@cefala.net UNIPG

User Environments

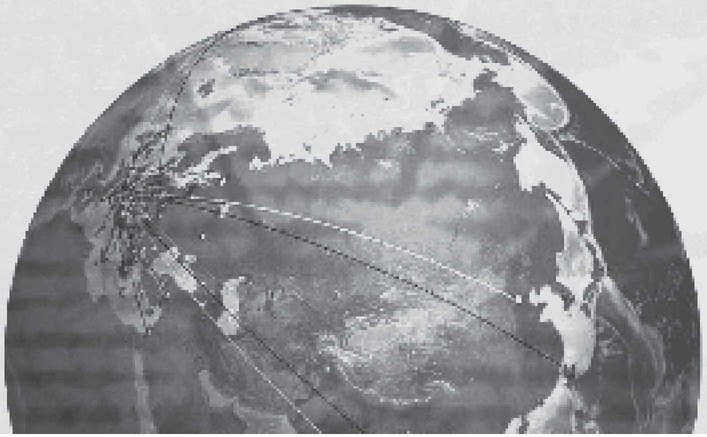

End-users of Distributed Computing Infrastructures typically expect that these systems can be smoothly and easily integrated into their routine research workflows. DCI-based ‘e-workplaces’ should allow users to customise their data, computational and software services, and make these available within easy reach. While highly customised user environments, built on top of e-Infrastructure services can certainly satisfy these needs, the development and maintenance of such solutions can rarely justify the investment. Highly customisable systems can attract more users, however these systems require involved configuration, maintenance and training from the parties.

Presentations within the User Environments track introduce various tools and projects concerned with how end-users interact with the European Grid Infrastructure. The number of user environments available to EGI communities is high. The track aims to bring together the developers, providers and users of these environments and enable them to share experiences, issues, solutions and trends with each other and with the EGI stakeholders.

The integration of National Grids into a single European Infrastructure triggered the need to extend existing user environments – whether these are implemented as portals or client-side applications – with support for multiple middleware technologies. Only those tools that are interfaced with a variety of workload manager systems can successfully satisfy the needs of multi-national communities.

The emergence of desktop PC and cloud-based computing within the EGI ecosystem has been recognised by the providers of user environments. To achieve the transparent integration of users’ application and data across different infrastructures, the latest generation of user tools can connect to cluster, desktop and cloud infrastructures.

Gergely Sipos Senior User Community Support Officer, EGI.eu

A science gateway for molecular simulations

OVERVIEW

Nowadays, scientists in multiple research areas like material science, structural biology, and drug design are supported by invaluable molecular simulation tools. These tools allow analysis of increasingly complex chemical structures on high-performance computing facilities. However, the lack of graphical user interfaces, the limitation of the usability of the tools and the complexity of infrastructures demands intuitive user interfaces.

The project MoSGrid (Molecular Simulation Grid) addresses these issues by combining an easy to use portal-based infrastructure with expert knowledge on the correct use of complex methods. The emerging portal will support users in all stages of the simulation process with easy access to data repositories storing information about molecular properties, and the possibility of creating, editing and invoking workflows. The project integrates the UNICORE 6 grid middleware and the cloud file system XtreemFS into the workflow-enabled grid portal WS-PGRADE.

DESCRIPTION

MoSGrid is an interdisciplinary project which aims to facilitate molecular simulation tools in grid infrastructures via a science gateway. This science gateway is based on the workflow-enabled grid portal WS-PGRADE which is the highly flexible graphical user interface (including a graphical workflow editor) of the grid User Support Environment (gUSE). gUSE is a virtualisation environment providing a set of services including a workflow engine, so-called submitters (Java-based applications), a workflow storage, an application repository, monitoring and logging. Submitters implement the grid service interface of gUSE and exist for classical service and desktop grids, clouds and clusters, unique web services. MoSGrid added a submitter for UNICORE 6 using the UCC (UNICORE Commandline Client) libraries. The submitter enables to invoke jobs of WS-PGRADE workflows on UNICORE 6 infrastructures and additionally the invocation of existing UNICORE workflows via WS-PGRADE.

Besides the extended workflow features of WS-PGRADE, MoSGrid has been developing intuitive portlets for predefined basic workflows for the most popular molecular simulation tools within the MoSGrid community. Furthermore, users will be supported by repositories which store information about initial structures for the simulation process, parameters of workflows like used molecular simulation tools and resulting structures. The repositories will allow external referencing of simulation results and will be available replicated in XtreemFS via the science gateway.

In general the authentication process for grid infrastructures relies on X.509 based certificates. In MoSGrid WS-PGRADE has been extended for the use of SAML (Security Assertion Markup Language). The user is enabled to create an assertion file from his X.509 certificate via an applet. The certificate remains on the user's computer, solely the assertion file is stored on the science gateway server and handed over to the UNICORE infrastructures.

IMPACT

The emerging science gateway offers a single point of entry to the whole process of analysing data for molecular simulations facilitating UNICORE 6 infrastructures. It will support multiple molecular simulation tools, workflows on molecular simulations and repositories which store the information on structures, workflows and simulation results. The supported tools are selected via a survey in the MoSGrid community including about 100 working groups in the fields of computational chemistry and bioinformatics. Workflows and results can be exchanged in the community via a global repository. Workflows can be created, re-used and changed by the community and additional knowledge will be obtained by cross-referencing different result data files.

Compared to similar existing activities like ProSim, G-FLUXO, the e-NMR portal and projects based on EnginFrame, the MoSGrid science gateway will offer a larger scale of tools and methods and will support the possibility to refer to molecules and results in a global repository under an open source licence.

EnginFrame is also a workflow-enabled grid portal with existing projects for molecular simulations but under a commercial licence model for non-academic users. ProSim and MoSGrid can benefit from each other. ProSim offers workflows and workflow templates for carbohydrate recognition in a portal on top of WS-PGRADE with access to several grid infrastructures. Users could re-use workflows of ProSim in the MoSGrid science gateway and the ProSim portal could offer carbohydrate recognition in UNICORE 6 infrastructures. G-FLUXO provides Gromacs workflows based on P-GRADE (the first generation of WS-PGRADE) and visualisation of simulation results. The workflows and the visualisation could be integrated into the MoSGrid science gateway. The e-NMR portal offers also multiple molecular simulation tools with access to a number of grid resources but does not offer a repository of results.

CONCLUSION

The emerging MoSGrid science gateway will support the users of molecular simulation tools in grid infrastructures. At the current stage of the prototype, the UNICORE 6 integration for WS-PGRADE workflows and the SAML extension is operational. Furthermore, there are archetypical workflows available for quantum chemical simulations with Gaussian and molecular dynamic simulations with Gromacs in single portlets.

Next steps include the definition of workflows for the MoSGrid community and the uptake of the Application Specific Module (ASM) for WS-PGRADE in the developed portlets. The latter will serve as templates for further portlets on additional tools. The portlet developers do not have to become acquainted any more to the UNICORE integration in detail and the users are enabled to monitor their different workflows on one page.

Furthermore, the repositories based on XtreemFS will be integrated into the science gateway. Hence, XtreemFS will be extended for authentication with SAML.

http://mosgrid.informatik.uni-tuebingen.de:8080

Sandra Gesing sandra.gesing@uni-tuebingen.de University of Tuebingen • **Peter Kacsuk** kacsuk@sztaki.hu MTA SZTAKI, Budapest • **Miklos Kozlovszky** m.kozlovszky@sztaki.hu MTA SZTAKI, Budapest

Proteomics portal: an e-science gateway for Swiss bioinformatics

OVERVIEW

Bioinformatics is a science that routinely makes use of large computing infrastructures including grids to understand and analyse the data and to prepare for the next wave of instruments producing even more data. Not only the data volume but also its complexity has increased. The combination of complex analysis tools and large-scale infrastructure makes it difficult for the scientists to analyse the data without expert help. There is a strong need to establish automated processes and workflows that can be executed easily without any knowledge of the complexity of the underlying infrastructure. Thanks to a joint project between SystemsX.ch, ETH Zurich and MTA SZTAKI, a new web-portal for easy to use automated proteomics analysis was built using the P-Grade Portal technology developed by MTA SZTAKI. The portal has been set up to be able to execute programmes on local clusters of ETH and on the national distributed grid, the Swiss Multi-Science Computing Grid based on the ARC middleware.

DESCRIPTION

The P-Grade Portal provides a web-based user interface where the users can develop/manage applications on various types of DCIs. The portal is able to submit jobs using gLite or ARC-based middlewares as well as on LSF or PBS-based clusters, a capability that was developed for this project. The flexibility and the workflow and parallelisation capabilities coming with the user-friendly interface were the reasons to choose P-Grade as the baseline technology for this specialised proteomics e-science gateway. To start with, three different types of workflows were developed based on commonly used proteomics tools and applications set up to be executed on the large clusters. Then corresponding portlets were developed providing easy-to-use web interfaces tailored to the needs of the end-users. These portlets hide the complexity of the workflows and the DCIs allowing the end-users to focus just on the parameters that are important to their own research. The three initial portlets are interfaces to the Transproteomic Pipeline, a workflow to perform label-free quantification and a quality control metric calculation. But also much more complex workflows have been implemented by more advanced users of the platform for dedicated research problems and several new portlets are in the process of being finalised.

There was also a clear need to make the usage of the portal's security components as easy as possible with no shortcuts taken in terms of security. This was achieved by the integration of the Swiss national AAI infrastructure based on Shibboleth2 into the portal and providing access to the grid automatically using the SLCS service to generate user certificates based on the user's AAI login to the portal. This way the users only have to log in once and the portal can then submit jobs to the grid with no additional steps necessary as in previous portal instances where users needed to generate a proxy certificate outside of the portal and upload it to a myproxy server.

IMPACT

Data analysis in the domain of proteomics has made a lot of advances in the recent years with a very large number of new emerging tools to identify and quantify proteins and peptides through mass spectrometry and liquid chromatography experiments. The problem is that many tools do not adhere to standard data formats, although such standards do exist.

Through the Swiss Proteomics Gateway the usage of these tools is automated and very much simplified for the end-user. Experts take care of data transformation into the usable formats, chaining the tools to get the most relevant information and inclusion of the most recent methods. With the gateway scientists will be able to share their analysis workflows and the best methods to look at the data. More importantly, they can automate the workflows for repeated analysis with changed parameters, which was a manual, slow and very error prone process in the past. Developers of algorithms can easily interface with the existing services and test new ideas on real data. Collaborations between lab scientists and algorithm developers are very much facilitated, nationally and internationally; the scientists can focus on their core scientific discoveries as opposed to spending time in the details of data transformations. Due to these advantages the whole Swiss research community interested in proteomics analysis can set up workflows and run applications on the local and national DCIs.

There is a strong synergy with the medical sciences, i.e. between the proteomics gateway community and the medical gateway, to be later developed by AMC (Amsterdam Medical Centre). A strong interest exists to extend the gateway to be used to analyse other types of data, i.e. imaging, microscopy or genomics datasets. The concepts are identical, but of course the data types and algorithms would need to be adapted. By the use of P-GRADE portal technology it is possible to link this effort with the efforts of AMC, giving this development a European dimension.

CONCLUSION

Due to the recent joint development between the ETH Zürich, SystemsX.ch and MTA SZTAKI, a new web-based portal is available to researchers in proteomics. The complexity of the inner component systems is hidden, therefore the end-users can focus on their science, only parameterising and executing common proteomics analysis workflows by the push of a button. This increases their productivity very efficiently.

https://www.imsbportal.ethz.ch

Akos Balasko balasko@sztaki.hu MTA SZTAKI

SOMA2 – gateway to grid-enabled molecular modelling workflows in www-browser

OVERVIEW

SOMA2 gateway is a molecular modelling workflow environment developed and deployed by CSC – IT Centre for Science Ltd. The SOMA2 environment is used with WWW-browser and it allows users to combine scientific applications into unique application workflows, which are automatically executed in the underlying computing system. SOMA2 offers a flexible framework for integrating and executing molecular modelling applications, and facilitates automated molecular data exchange. SOMA2 source code is distributed under the GPL open source licence.

For end-users, SOMA2 offers a secure, personalised and easy to use environment for utilising computing infrastructure. In SOMA2, scientific applications are presented and configured via web forms, which guide users to correctly configure a programme by supplying default values, thresholds, runtime help and validation. For experts, SOMA2 offers a framework to make virtually any molecular modelling application accessible to the end-users.

DESCRIPTION

SOMA2 enables communication and data exchange between applications by employing a common data exchange format, CML (Chemical Markup Language). The common data format is crucial for seamless integration of different applications within the SOMA2 system. The SOMA2 web application handles user's initial molecular data by converting the data to CML format and preparing it for the application workflow. With the web application, the user is able to set up the workflow, configure chosen applications, submit projects for computation, review the status of the project and access the results.

In SOMA2, scientific applications and their execution is described in pluggable capsules, which make use of interfaces to manage the internal data. A capsule consists of an XML description, used for example, to generate an application web form, and scripts and file templates to enable scientific application execution and processing of the programme output.

SOMA2 system includes a workflow manager programme which is responsible for internal data transfers and controlling execution of the capsules.

Recently we have added DCI-support in SOMA2. In SOMA2 web application and its utilities, this includes handling of users' X509 certificates, requesting available resources from the grid and generating proxy certificates upon project submission. We make use of grid middleware to submit the jobs. This work has been conducted as part of EGI-InSPIRE project's WP6-SA3. Currently, we have support for Nordugrid Arc middleware but other middleware could be used as well.

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 on the actual scientific task instead of dealing with technical issues requiring manual work.

With the new DCI-support features of SOMA2, we can really say that SOMA2 does not only integrate applications but also different computing infrastructures making complex computing environments reachable for all users. We think that technical complexity is still a major bottleneck, which prevents users from fully benefiting from the distributed resources. SOMA2 system can alleviate this problem a lot.

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 programme descriptions and basically no real programming skills are required to create a SOMA2 capsule. Flexibility, application oriented approach for reusable workflows and open source licence make SOMA2 system very unique in its domain.

CONCLUSION

SOMA2 source code is distributed openly for all interested parties. At CSC, SOMA2 is also available as a service for CSC's academic users providing access to 14 different molecular modelling applications, which are seamlessly integrated within the system. The system is fully integrated with the local computing infrastructure.

Within EGI-InSPIRE WP6-SA3, we plan to extend the current SOMA2 service to include DCI-enabled applications. This service will be first introduced to EGI. Later on, we plan to make the service available for other communities in EGI.

In addition to common enhancements in the web application, SOMA2 development plans include investigating other middleware to be used in SOMA2 DCI-integration and new applications to be integrated in SOMA2.

www.csc.fi/soma

Tapani Kinnunen takinnun@csc.fi CSC

Data-driven portal for hadron therapy

OVERVIEW

PARTNER grid prototype is a data-driven platform for recording and sharing medical data within the hadron therapy community. The idea is to provide secure cross-border access to heterogeneous repositories while leaving data where it was created and not copying it centrally (data warehousing) due to data privacy and ownership concerns. Our solution is to describe local data using metadata and use federated queries to provide users with a unified view over distributed resources as a virtual database.

DESCRIPTION

Using the Vine Toolkit [http://fury.man.poznan.pl/vinetoolkit], we connected Liferay portal [www.liferay.com] with VOMS [https://voms.ific.uv.es:8443/voms/vo.partner.eu] to enable strong grid security and to use a DPM Storage Element as a grid resource to store DICOM images. For domain specific data recording we developed forms using web-services, Java and Flex portlets. We work on a data virtualisation layer that coupled to a metadata registry [https://cancergrid.org/] enables semantic queries over different data repositories.

IMPACT

The portal is an easy to use tool for all the actors involved in the hadron therapy domain: medical doctors can securely import, transfer and view patient related information and researchers can construct complex queries. Moreover, patients could record their quality of life during and after the treatment. Specific roles and privileges are protecting data and actions of actors using policies enforced by VOMS credentials.

CONCLUSION

We demonstrate how we use grid services to securely integrate medical data from various sources. The combination of enterprise portal Liferay and Vine Toolkit provides an easy to use grid-enabled environment where additional services could be prototyped and tested. To improve the portal services, data provenance and data mining aspects of our platform we hope to establish close interactions with other communities and projects in EGI.

http://partner.web.cern.ch

Faustin Laurentiu Roman faustin.roman@cern.ch CERN

Job Submission Tool, web interface and WebDAV data management

OVERVIEW

The Job Submission Tool (JST) allows the exploitation of the grid computing power to many research communities with highly compute-intensive applications. The tool helps to sub-divide large applications into single independent tasks and execute them on the grid nodes in an optimised time. Furthermore with the implementation of an ad hoc graphic interface, users have the opportunity to use the grid without knowing all the technological details and without taking care of its complex authentication methods. So far, several bioinformatics applications have been successfully executed onto the grid using JST, allowing the users to reach important goals on their research.

DESCRIPTION

Several bioinformatics scientific questions often require months or even years to be solved using the computational power provided by a single cluster. In such cases it is important to exploit the grid in order to reduce the execution to a few days or hours. The collaboration between bioinformatics researchers and grid experts has been found to be very productive: the grid expert takes care of splitting the application into independent tasks, filling the central DB with the independent tasks, submitting the tasks to the grid and retrieving the output. In other words, the final user does not have to deal with all the grid technicalities which are taken by the grid expert. Several important challenges have been successfully executed with this approach.

We have tried to go one step behind, trying to enable final users to use the grid on their own. For this purpose we have developed a tool, JST, which uses a web graphic interface, written using PHP, Javascript and XSLT, where a user can authenticate using a username and password. Using the same graphic interface the user uploads the input files. Since it was found not feasible to upload big input files (up to 10GB) using the http protocol, a new mechanism based on the 'webdav' protocol has been implemented to upload input files on a file server in an efficient way and with an automatic registration of the files on the grid Storage Elements. The user can then customise the application (changing the configuration parameters) and specify the name of the output files. After the submission of the web form, JST executes all the steps required for the submission of the application to the grid (slicing of the problem, user authentication by means of a robot certificate, submission of the jobs to the grid, retrieving of the output files). At the end, JST sends the link of the output files to the user by e-mail.

IMPACT

Researchers are often not familiar with X509 certificates and grid technology: in recent years JST has proved the capability to solve the problem of submitting huge challenges on the grid with little grid competence and small human effort.

The first goal achieved by JST was to split applications born to be executed on single machines, in order to reach the same results executing single independent jobs in a parallelised way.

The JST team ported over the grid about 20 different bioinformatics applications, submitting around half a million jobs representing about 50 years of CPU.

After the release of the graphic interface more than 50 challenges have been submitted to grid through the web. A total of 359 days of CPU usage on the grid worker nodes has been reached for the completion of all the tasks. The submission with the GUI has produced 19,656 correctly executed jobs and 583 failed ones. Considering other similar tools that allow access to the grid, we have tried to completely mask the underlying technology to provide the end-users the chance to use a simple portal that takes care of the complex steps they would face if using the grid directly. Furthermore we have overcome the limit that a classic portal could give for input file dimensions, setting up the new webdav server for JST that offers a new way to manage very big files. We decided to use webdav as it is a standard and well supported protocol, indeed each operative system has its own client (this is true for Linux, MacOS and Windows).

CONCLUSION

The JST has shown a good reliability and efficiency and few improvements are needed. The development team is just working on the integration of the interface in other portals. Since JST has been developed inside the BioinfoGRID and LIBI project, so far all the work planned has been focused on the bioinformatics applications, but the generalisation of the tool allows us to open our technology to other communities.

Giacinto Donvito giacinto.donvito@ba.infn.it INFN

New developments of gUSE and WS-PGRADE to support e-science gateways

OVERVIEW

E-science Gateways are the primary solutions dedicated to support end-users without grid knowledge to do their research on distributed grid infrastructure. Due to the wide range of research fields, generic grid workflow management systems/portal services should be flexible enough to support the different requirements both at front-end and back-end. Previously the importance of science gateways was overlooked. This required a huge (time and skill) overhead in usage and development due to the required complete understanding of various distributed infrastructure technologies and solutions. As a result adoption of these e-Infrastructures has been below expectation until today. Recognising the problem many attempts have been made to create user/community-specific and user-friendly science gateways for the various e-science communities. In this oral presentation we introduce the newly developed features of gUSE a Grid User Support Environment and the WS-PGRADE portal solutions.

DESCRIPTION

The gUSE and the WS-PGRADE solutions are tailored for supporting grid application development and building up e-science gateways.

gUSE is an easily usable, highly flexible and scalable co-operative grid application development infrastructure, connecting developers and end-users of grid applications with computational resources of different technologies. It enables the design, integration and submission of sophisticated (layered and parameter sweep enabled) workflows spreading parallel among grids, web services, clusters, and eventual local resources of gUSE. It provides a collaborative, community-oriented grid application development environment where grid application developers and end-users can share workflow graphs, workflow templates and ready-to-run workflow applications via a workflow repository.

WS-PGRADE portal is a second generation, generic purpose e-science portal that introduces many advanced features both at the workflow and architecture level. In this presentation we are focusing on the following feature set of the WS-PGRADE/gUSE solution offered for the various types of user and NGI user communities:

- Scientific gateway services. We will show how the solution supports unified access to and seamless integration of the underlying networking, computing and data infrastructures and services for all the major DCI infrastructures (local and remote clusters, supercomputers, local and volunteer DGs, grids, public and private academic or commercial clouds).
- Developer and end-user services. We will detail how the solution provides support for the design, development and deployment of user-friendly interfaces which abstract service provision from the underlying infrastructure complexities and specific implementations through the use of various web technologies.
- Baseline technologies. We will enumerate the additional feature set, which facilitates community level collaboration among developers and end-users, such as support of various repositories.

IMPACT

As the available grid infrastructure achieves larger community adoption and higher performance, there are an increasing number of user communities with well described needs for an application-specific gateway. Recently various e-science communities tried to create user-specific and user-friendly science gateways in their own way. However, in many cases the attempt failed since the community underestimated the effort needed to create a stable, secure and robust gateway to be used simultaneously by many community member as a reliable 24/7 service. The scattered activities among communities also led to lot of redundancy and duplication of efforts. At the end in many cases they have realised that it is much more cost-effective and faster to build the gateway based on an existing generic-purpose DCI gateway. Even those communities that have already successfully deployed their first gateways have identified new requirements and would like to use the latest technologies to keep their research competitive. With our newly developed WS-PGRADE/gUSE feature set we have targeted the following emerging user community requirements: easy and reliable way to build new gateways, upgrade existing gateways to modern technologies, strong customisation capabilities, and seamless infrastructure compatibility.

As a result gUSE and WS-PGRADE is already popular for many communities since it helped to lower barriers to entry in e-science environments for researchers. It also increases the potential for e-Infrastructure usage by non-specialists, as demonstrated in the FP6 and FP7 projects (CancerGrid, SHIWA, HP-SEE, MoSGrid).

CONCLUSION

The developed new feature set of WS-PGRADE/gUSE can be categorised into scientific gateway services, developer and end-user services and baseline technologies. It provides support to build up, customise and integrate easy and reliable new gateways and to upgrade older existing gateway solutions to recent technologies.

It also further facilitates and speeds up the development of community owned application-specific gateways and supports application developers, end-users/e-scientists and e-Infrastructure operators in their work. New user communities are supported strongly during their e-science gateway development in a collaborative way. As a result, a wide range of other European user communities benefit already from the improved e-science services provided by the WS-PGRADE/gUSE solutions. In addition, the developed feature set also bridges into the commercial cloud application domain, allowing scientists and commercial users to use clouds and grids together.

www.guse.hu
www.wspgrade.hu

Dr. Miklos Kozlovszky m.kozlovszky@sztaki.hu MTA SZTAKI

Grid computing for biomechanical markers extraction

OVERVIEW

Visual evaluation of an object (in images) in biomechanical and orthopaedic processes is a usual method of contemporary medical diagnostics. In some cases visual diagnostics is relatively simple (for example, in the case of non-complicated bone fracture) and does not require complex medical methods and calculations. However, such cases are rare and also mere visual inspection is sometimes not sufficient. The diagnosis may require more complicated measurements, as well as extraction of biomechanical parameters of an object. However, this is a time-consuming process also prone to mistakes. This article presents computer vision algorithms, developed and implemented to improve diagnosis and increase the reliability of medical decisions. Algorithms, set in a grid environment, are able to accelerate and unify the diagnosis, leaving to a physician mainly a function of interpretation (assessment of medical condition), selection of an appropriate treatment method and evaluation of its efficiency.

DESCRIPTION

Different dysfunctions of an acetabulofemoral joint result in pain and/or motility restrictions. Such dysfunction may be congenital, caused by trauma or joint operation conditions (such as long-term joint strain exceeding joint stress limits), natural wear on joints or may be a result of other pathologies. Modern medicine can help all patients suffering from these types of joint dysfunctions; however, treatment resources are not sufficient. Establishment of pathological diagnosis in the pelvic area (such as dysplasia, arthrosis, dislocation or fracture) and selection of treatment methods relies on a number of methods based on calculations of biomechanical parameters of objects under analysis. These parameters can be conveyed in geometrical expressions: points, distances among objects, positions of objects to each other and angle formations. In general, the diagnostic process involves the following steps:

1. pain diagnostics,
2. anamnesis,
3. patient examination (temperature, visual changes, palpation, pace, stand, movement extent),
4. radiology test and extraction of biomechanical parameters,
5. additional diagnostic methods (computed tomography, radioisotope bone scan, ultrasound testing, magnetic resonance imaging, arthroscopy).

This article presents a number of computationally-based methods for extracting biomechanical parameters, intensively using a computer vision approach. Such parameters are used in primary diagnostics, certain calculations are made when observing the course of treatment, while other methods are employed only in exceptional cases when standard methods fail to provide enough information. Grid platform provides 'on demand' computational resources, as well as suitable services for such computing procedures.

IMPACT

The approach researched in this article influences medical procedures to:

1. specify more precisely a set of recognisable objects and calculate/separate biomechanical parameters,
2. examine the adaptability of the developed methods for diagnostics of other body areas which is based on the analysis of bone structures visible in x-ray images,
3. examine the adaptability of the developed methods for diagnostics of other areas which is based on the visual analysis of objects in 3D images.

CONCLUSION

While extracting biomechanical markers by various computational procedures a number of methods have to be applied when analysing radiographs. The most easily recognisable objects are femurs. When searching for clearly visible pelvic arches a set of additional thresholds for parameters due to peculiarities of the equation describing pelvic arches and the radiograph itself are done on the basis of data about detected femoral ridges. After defining the parameters, pelvic bones are successfully detected. Various sequences of calculations and parameters of detected objects involve more sophisticated diagnostic methods, involve drug selection and shorten diagnostic procedures. This also opens broader suggestions for further explorations. These computational procedures are obviously suitable to implement in cloud computing or virtualisation environments.

Algimantas Juozapavicius algimantas.juozapavicius@mif.vu.lt VU • **Ramunas Markauskas** ramunas.markauskas@mif.vu.lt Vilnius University

Socially impacting grid services from the chemistry Heavy User Community

OVERVIEW

The work of the community of computational chemists and molecular scientists is at the core of several innovations and technological advances of modern society. To this end we have joined our efforts (carried out within the COMPCHEM VO) with those of some members of the GAUSSIAN (CYFRONET) belonging to the same computational chemistry area to prepare the ground for joining other communities like MoSGrid (UCO) and ENEA-GRID (ENEA) in building a Heavy User Community (HUC). Our effort has focused on adopting (with appropriate adaptations) some basic grid tools, including those designed and implemented by other HUCs to support their users, by assembling suitable workflows for chemistry and molecular science applications, by managing related repositories and proposing appropriate mechanisms to foster the formation of collaborative research within the community.

DESCRIPTION

The work focused on the design and development of ab initio grid empowered molecular simulators of complex physical phenomena and technologies to the end of assembling study cases for building grid services of social relevance.

Among the software tools and computer applications considered for this purpose are:

- Middleware, workflows and frameworks: this is a set of tools like FARO, KEPLER and GriF that are being analysed and compared with other products already implemented in EGI. The first of them is targeted to foster the communication between the EGI grid and other platforms (like that of ENEA). The second of them is a workflow used for the assemblage of the grid empowered molecular simulators. The third one is a framework based on a JAVA Service Oriented Architecture (SOA) designed to facilitate the use of the grid by non-specialists by properly selecting the user interface, the computing element and the storage element.
- Electronic structure: this is a layer of application software made of highly successful quantum chemistry packages (like GAMESS, GAUSSIAN, MOLPRO, DALTON, MOLCAS, NWCHEM, TURBOMOLE, etc.), which are at present solid foundations of any determination of molecular structures and properties.
- Dynamics: this is a layer of molecular dynamics codes like either VENUS, DL_POLY, GRO-MACS, CPMD, CP2K, NAMD, etc. treating molecular systems as an ensemble of particles or RWAVEPR, MCTDH, FLUSS, ABC, etc. treating molecular systems as quantum waves and therefore requiring a larger size of node memory.
- Multiscale and statistical treatments: this is a set of services based on grid simulations of natural structures and processes based on multi-scale and multi-physics approaches starting from the nanoscale level and including when appropriate fluid dynamics, randomisation, virtual reality, data mining, etc.

IMPACT

The impact of our work has materialised into the offer to the members of our community of GEMS, a prototype simulator able to start from first principles and end up, in a complete a priori fashion, by evaluating the signal of crossed beam experiments and the values of reactive cross sections and rate coefficients. The tools used for that purpose (the KEPLER workflow, the data models of Q5cost and D5cost, the grid framework of GriF) drive the user through the massive calculations needed by the above mentioned packages to allow a realistic modelling based on a complex combination of building blocks in which ab initio treatments, dynamics, kinetics, fluid dynamics, statistical sampling can be combined together with statistical treatments and rendering techniques to build versatile user centric instruments.

Examples of the application of this approach have been worked out for software devoted to the construction of scenarios for the production of secondary pollutants, to the labelling of chemical products and to the assemblage of distributed repositories for learning objects. The development of grid implemented versions of this software that is being taken care of in our laboratory is aimed at making sustainable, within the chemistry community, the development of grid tools and applications and at producing as well high social and economic benefits by tackling areas of applications relevant to environmental, health and educational aspects.

The improvement and further development on the distributed platform of EGI of the molecular level components of this software is, in fact, a clear citizens-oriented target that impacts not only the work of chemists, materials technologists, engineers, biologists etc. working on the development of faster and more efficient modelling of complex systems starting from the microscopic level but also (and even more) the everyday activities of students and generic citizens.

CONCLUSION

The porting of key ab initio approaches to molecular simulations onto the grid, using tools developed by other communities and developing new tools ad hoc designed for the chemistry and molecular science, has shown to be highly effective not only for further development of grid technologies and related chemical applications but also for paving the way for developing higher education virtual campus technologies, governing emergencies, designing innovative processes and materials for the quality of life, energetic problems and environmental issues of high social impact.

Antonio Lagana lagana05@gmail.com University of Perugia

e-Social Science development in Taiwan: scalable social simulation using commodity computing

OVERVIEW

This talk describes the overview of the process of developing the agent-based modelling of simulation, the current development of model construction, how to utilise multiple CPU cores and an investigation of the scalability of the resulting code using grid/commodity computing.

DESCRIPTION

This project aims to build up social simulation models in several application domains, including models of demographics and migration in Taiwan, and the electorate voting model for political science. In this project, we describe experiences made in the development of an example model that utilises multiple CPU cores and investigate the scalability of the resulting code. The feasibility study and stability analysis of the model has been completed. We argue that commodity compute resources and commoditised simulation frameworks can now be used to simulate real-world populations and to use these simulations to investigate social phenomena such as migration.

IMPACT

The e-Social Science community in Taiwan has been built up within the EUAsiaGrid project. The population-scale social simulation model using commodity computing power can be tractable under current simulation frameworks. The agent-based simulation has been encouraged as a complementary social science research method for the social scientists in Taiwan.

CONCLUSION

A fruitful international collaboration of e-Social Science has been established between Europe and Asia. This project also shows a promising result for a feasibility study of a scalable social simulation.

Hsin-Yen Chen hychen@twgrid.org ASGC • **Jing-Ya You** jingya.you@twgrid.org ASGC

Access earth science data from the EGI infrastructure

OVERVIEW

Accessing data stored outside of EGI during the execution of a job has been in discussion since the grid was first used by Earth Science (ES). There is, for example, high interest in the community in using the OPeNDAP protocol and corresponding Hyrax Data Server. Due to the many different technologies, data centres, standards and pseudo standards, however, it seems that no general solution can be found. Even if Open Geospatial Consortium (OGC) compliant web services are adopted, interfaces are required.

DESCRIPTION

We will present the work of the earth science activity of EGI-InSPIRE SA3, which is devoted to earth science data access. The activity is investigating and following the development of specialised data access and management methods in the area of earth science.

Accessing data stored outside of EGI during the execution of a job has been in discussion since the grid was first used by ES. There is, for example, high interest in the community in using the OPeNDAP protocol and corresponding Hyrax Data Server. Hyrax offers many interesting features that go beyond high performance access to distributed datasets, such as an extensible component-based architecture, multiple data representations, static and dynamic THREDDS catalogues and more. Due to the many different technologies, data centres, standards and pseudo standards, however, it seems that no general solution can be found. Even if OGC compliant web services are adopted, interfaces are required.

The activity can be seen as the follow-up to the work that was carried out in EGEE III based on a memorandum of understanding with the Ground European Network for earth science Interoperations–Digital Repositories (GENESI-DR) project. The work of GENESI-DR is continued in the GENESI Digital Earth Community (GENESI-DEC) project that has started in 2010 and aims to extend and improve both the infrastructure and target audience of the projects network. The activity tries to act as a bridge between the project and the earth science community in EGI.

Problems, recent developments and different approaches to solution will be discussed in this session.

IMPACT

The access to data centres from the EGI infrastructure supports existing and new Earth Science application and will be used for new deployments.

CONCLUSION

We will discuss the available OGC services and will show that an OPeNDAP solution is upcoming.

Horst Schwichtenberg horst.schwichtenberg@scai.fraunhofer.de Fraunhofer SCAI

Grid computing for electrophysiology

OVERVIEW

In electrophysiology many diagnostic processes are based on digital signal detection and processing procedures. In epilepsy diagnosis and treatment the electroencephalography (EEG) signals are one of the main tools. However the inspection of EEG in a 'manual way' is very time consuming. Automatic extraction of important EEG features provides an entire new level for doctors, for disease diagnostics and exploration, it also saves doctors' time. EEG analysis using data mining methods gives powerful tools for further exploration. In this article, methods to extract some of these features – drowsiness score and centre-temporal spikes are analysed. For spike detection a new method based on morphological filters is used. Also a database design is proposed in order to allow easy EEG analysis and provide data accessibility for data mining algorithms to be developed in the future. Grid computing platform is an efficient way for implementation of algorithms needed.

DESCRIPTION

One of the main subjects of research is measurement of brain activity recorded as electroencephalograms (EEGs). A multichannel EEG signal from one person is usually a collection of about 12 to 128 signals, which are recorded with 200Hz rate or similar and may last from some minutes to some hours. This data requires large amounts of storage, sophisticated data manipulation, complex data transformations. EEG is a summation of electrical activities generated by cortical neurons and it's widely used in diagnosis of neurological disorders related to epilepsy. The seizure of epilepsy leads to transient disturbances of the EEG signal. Thus usually epileptic EEG data contains transient components and background activities. The epileptic transient activity usually appears as sharp spikes in signal which occur randomly with short duration of 20-70ms. One of the main tasks in EEG signal processing is spike detection, as their recognition is significant for clinical diagnosis of epileptic disorders. In hospitals a method of visual/manual detection of spikes in EEG recordings by an experienced neurologist is widely used. In the case of long EEG recording, this process is more time consuming. Thus automatic spike detection methods are needed. Such a detection method, based on mathematical morphology algorithms, was developed by the authors, EEG data being provided by Vilnius University Children's hospital. Input files are of European Data Format (EDF). The grid infrastructure is used to create a virtual repository for such data, making data accessible for authorised manipulations by hospitals, doctors, computing processes. The functionality of a virtual data repository includes 'raw' data storage (together with suitable indices), data anonymisation, data segmentation into attributed segments, data warehousing for complex algorithms, which are required by doctors for diagnostics, selection of drugs, shortening decision time.

IMPACT

The grid computing services designed and developed for EEG analysis, are implementing the operational process for electrophysiology. The segmented EEG parts, as initial objects of the database of EEG signals, together with suitable computational procedures, reflect the following attributes:

1. correlation of spike number in EEG and spike parameters with type of epilepsy and clinical data,
2. spikes and segments clustering,
3. locating damaged areas of the brain,
4. numerical estimation of amount of waves of different rhythms (wave types of EEG) and their correlation with the type of epilepsy, sleepiness, etc.

Thus grid resources offer a possibility to manipulate data, to process big amounts of data and to produce various analysis on it. Grid computing procedures are produced on EEG segments (from the database), on entire EEGs, primary and secondary attributes of them (such as KDS, KSS, spike index, slow eye movement recognition, distribution of spike appearance time, amplitude, duration, shape, others). Theoretical models of such data manipulation and computational modelling are based on attributed graphs representation, where nodes and edges have many parameters and pre-computed attributes. In the case of full EEG data (for one patient) the computational model to be performed requires a lot of CPU hours (simultaneously involving hundreds of cores). Grid technology allocates such resources dynamically, actually producing all such computing 'on demand'. Grid is also able to allocate enough resources for data mining analysis (data clusterisation, pattern recognition, etc.) of all data accumulated in the data repository.

CONCLUSION

The Lithuanian grid computing and technical resources, together with the combination of grid technology and HPC procedures, enabled us to use minimal amount of efforts and in a short range of time to create and develop the efficient model of electrophysiological processes, as well as to support doctors in their diagnostics and drug selection processes to save them time. Such a system is highly evaluated by medical researchers and doctors.

Algimantas Juozapavicius algimantas.juozapavicius@mif.vu.lt VU • **Gytis Bacevicius** b.gytis@gmail.com Vilnius University

Grid-empowered molecular simulators of crossed beam signals

OVERVIEW

In the last years, several theoretical and experimental efforts have been spent in the field of quantum reactive scattering and this has fostered the development of molecular simulators. Our group is working at extending GEMS (Grid Empowered Molecular Simulator), to simulate the experimental signal of Crossed Molecular Beams (CMB) experiments. In this communication we report on the implementation of a prototype version of the GEMS block generating the virtual intensity of the product beam of a CMB reactive scattering experiment out of the integration of a large number of quasi-classical trajectories for a three atom reaction.

DESCRIPTION

GEMS is made of different blocks each of which is devoted to a particular task. In particular, the skeleton of the simulator (the production of ab initio electronic energy values, the fit of these values to work out a potential energy surface and the block devoted to dynamics calculations) has already been designed and implemented whereas the block devoted to the assemblage of the virtual experimental observables is still in the test phase. For this reason, as already mentioned our efforts have been spent in implementing the last block of the simulator. As a matter of fact, we have designed and implemented on the grid a Fortran code (called gmtsigma.f) best fitting the measurements of the intensity of the CMB experiments taken as a function of the angle and the velocity of the scattered products (expressed in the laboratory frame) starting out of the related computed dynamical quantities in the Centre of Mass (CM) frame (angular distribution and translational energy distribution).

As a study case we have chosen the $Cl+H2$ system. To this end, an accurate ab initio Potential Energy Surface (PES) has been incorporated inside the quasi-classical trajectory programme (VENUS96.F). In order to simulate the observables of the experiment, we have implemented inside the VENUS programme the calculation of the fixed angle velocity distribution of the products in the laboratory frame out of the theoretical CM angular and translational energy distributions.

Production runs of the simulator were performed for the system proposed for the study case. To this end, the modified version of VENUS implemented on the grid was run for 400,000 trajectories at different values of the rotational quantum number j. Finally, the comparison between calculated and experimental observables has been carried out.

IMPACT

The main impact of this work is to be recognised in the possibility of simulating a chemical reaction from first principles and comparing calculated results directly with those of the related CMB experiment. This affects both the theoretical and experimental sides of dynamical studies. In fact, thanks to this approach the feasibility of a CMB experiment could be predicted a priori by simulating cross sections and possible product distributions. At the same time, on the theoretical side the validity of a proposed potential energy surface could be characterised to a very fine detail if an access to grid platforms can be obtained. Moreover, the development of the chemical simulator on a grid environment is also important from a computational science point of view because it prompts the production of specific innovative technological solutions.

CONCLUSION

The design and the implementation of the last segment of the GEMS simulator has been carried out. This means that the simulator is now ready to be used as a powerful instrument in quantum reaction dynamics. We are still in the test phase with respect to the observables block of the simulator in order to extend the workflow to four atom reactions.

Antonio Lagana lagana05@gmail.com University of Perugia

Earth science oriented data processing tools and applications for the Black Sea catchment basin

OVERVIEW

EnviroGRIDS (Black Sea Catchment Observation and Assessment System supporting Sustainable Development) [www.envirogrids.net] is a 4-year FP7 Project aiming to address the subjects of ecologically unsustainable development and inadequate resource management. The project develops a Spatial Data Infrastructure of the Black Sea catchment region. The geospatial technologies offer specialised functionalities 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 extended features of both technologies. The enviroGRIDS system resources are accessible to the large community of users through the BSC-OS Portal that provides Web applications for data management, hydrologic model calibration and execution, satellite image processing, report generation and visualisation, and a virtual training centre.

DESCRIPTION

The portal publishes through the Web applications the geospatial functionality provided through Web technologies and the high power computation supported by the grid technologies. There are five categories of users such as data providers, earth science specialists, decision makers, citizens, and system administrators.

The portal consists of a set of Web applications through which the users access the system resources such as spatial data, hydrologic models, environmental scenarios, data processing tools, visualisation facilities, environmental reports, and training materials.

The main user application categories provided by the portal are:

1. data management – provides the user with spatial data management and operations. The user may enter data and metadata, visualise, modify, update, and remove spatial data from the data repositories;
2. hydrologic model management – provides earth science specialists with hydrologic model configuration, scenario and model development, model calibration and scenario running. One of the water quality models that will be used is SWAT (Water Assessment Tool) [http://www.brc.tamus.edu/swat/index.html]. SWAT is a model designed to estimate impacts of land management practices on water quantity and quality in complex watersheds. The SWAT model requires specific information about weather, soil properties, topography, vegetation, and land management practices of the watershed;
3. Satellite data processing – the specialist may process satellite data and images in order to search for relevant information (e.g. land cover, vegetation, water, soil composition, etc.);
4. Data visualisation and report – the specialists visualise various spatial data in different formats and views and compose environmental reports for decision makers and citizens;
5. Decision maker and citizen application – provides the decision makers with the interactive and graphical tools to access the private environmental reports. The user may visualise data that make possible statistical analysis and predictions.

IMPACT

The regular users visualise the reports generated by the specialists as a result of executing different environmental scenarios. The input data for the reports are built up by the specialists by running hydrological models of the Black Sea catchment area and by processing related satellite data. All data sets required to build up the hydrological models, environmental scenarios, and spatial models are provided and entered into the system by the data providers.

EnviroGRIDS functionality gathers services provided by various technologies such as SWAT related modules, Collaborative Working Environment (CWE), Uniform Resource Management (URM), gProcess and ESIP Platforms, and experience on other research projects like GiSHEO, SEE-grid-SCI, and EGEE.

EnviroGRID system is developed on the gLite middleware available on EGI, SEE-grid and enviroGRIDS VO infrastructures.

CONCLUSION

The presentation mainly highlights the issues and the experiments on interoperability between geospatial and grid platforms in order to support high power computation requirement applications. The SWAT model covers the huge area of the Black Sea catchment region. It is a quite complex model that requires the high performance resources of the grid.

The experiments have revealed the potential of the grid related technologies. New user interaction techniques and platform interoperability are the main directions of the future research work.

enviroGRIDS Project: *www.envirogrids.net*

http://users.utcluj.ro/~gorgan

Dorian Gorgan dorian.gorgan@cs.utcluj.ro Technical University of Cluj-Napoca

The grid rocks

OVERVIEW

The development of web based technologies has seen a proliferation of e-science related projects, globally. One such project is the Paleobiology database (http://paleodb.org). This system "has been organised and operated by a multi-disciplinary, multi-institutional, international group of paleobiological researchers. Its purpose is to provide global, collection-based occurrence and taxonomic data for marine and terrestrial animals and plants of any geological age, as well as web-based software for statistical analysis of the data." (Source http://paleodb.org.)

We intend to work closely with the Paleodb.org as described below.

DESCRIPTION

As PaleoDB has grown, the primary methods for accessing the data have not kept pace with the user demands. In collaboration with the Earth Sciences department at the University of Glasgow, we have deployed a mirror of the database inside our existing grid cluster, and worked on ways to provide combined compute and data access to this via grid interfaces.

An additional Earth Science project is the deployment of modelling software within the cluster at Scotgrid, Glasgow to allow for the detailed analysis of prehistoric movements in tectonic plates including geo-seismic activity. This research will be conducted at both the regional and continental scales. A full analysis of the technical implementation and issues surrounding this type of research will be discussed including the benefits and drawbacks of batch computing in a grid framework on this type of research.

IMPACT

The primary benefits of these projects are that they will allow earth scientists at the University of Glasgow to investigate the potential of batch computing for this type of analysis. An additional benefit is that these project types lend themselves to the wider European Earth Science community through grid deployments over specific departmental projects.

CONCLUSION

In addition to facilitating use of the database for larger and more complex analysis, this work offers insights into how existing web based databases could be transitioned to the grid when they outgrow their existing infrastructure, and how the user experience may be re-examined as part of such a transition. The overall performance of the database within a grid environment rather being solely web-based will be discussed and limitations of both systems explored.

www.scotgrid.ac.uk/Projects.html

Mark Mitchell mark.mitchell@gla.ac.uk GridPP

e-science on earthquake hazard mitigation

OVERVIEW

In EUAsiaGrid, ASGC coordinates the construction of a pilot virtual research environment (VRE) on earthquake disaster mitigation, by streamlining the primary research process and components; seismic sensor network, seismic wave propagation simulations and earthquake data centres. Based on these, Strain Green's Tensor (SGT) methodology was also deployed to support hazard modelling more efficiently in Taiwan. All the VRE and e-Science applications are now supporting Asia-wide collaboration and customisation.

DESCRIPTION

Under the EUAsiaGrid framework, we have been implementing a series of forward wave propagation simulations and analysis applications. From the generation of accurate earthquake wave propagation, on-demand simulation of seismograms for any place on earth for any event, to the cost reduction of seismic wave propagation analysis by using the pre-calculated Strain Green's Tensor (SGT) function method. The work is not just porting designated computational models to gLite-based Grids, but also establishing research-oriented production services and long-term collaboration mechanisms among the partners. Given the SGT database (a 50-station 20 second case study of Taiwan running for 76,800 CPU hours has been done), hazard mapping and disaster potential analysis could be done much more efficiently without re-calculating the required seismograms. Development of the VRE on seismic hazard mitigation with the aforementioned applications will be described in this presentation.

IMPACT

With accurate regional seismic data, domain-specific services and the e-Infrastructure, we are able to greatly enhance earthquake hazard alleviation and conduct unique seismology research in Asia.

CONCLUSION

This work is a great example of the added value of e-Science and regional collaboration. Further federation of various regional resources and customisation is planned, which will be compatible with the European Grid Infrastructure.

Eric Yen eric@sinica.edu.tw Academia Sinica • **Hsin-Yen Chen** hychen@twgrid.org Academia Sinica

An overview of CMS workload management for data analysis

OVERVIEW

The Compact Muon Solenoid (CMS) is one of the four main experiments (VOs) at the LHC and relies heavily on grid computing. The CMS computing model defines how the data is to be distributed and accessed to enable physicists to efficiently run their analysis over the data. The CMS Remote Analysis Builder (CRAB) is the specific Workload Management System, that allows the end-user to transparently access data in a heterogeneous pool of resources distributed across several continents.

DESCRIPTION

CRAB adopts a client server architecture implemented in Python. The client, which provides the user with a batch-like command line application, has the responsibility of generating the user's proxy, packaging up the pre-compiled user code, and submitting the job package to the CRAB server. The intermediate Analysis Server has the responsibility of submitting the jobs to the chosen grid middleware, resubmitting jobs as needed, and caching the users' output. CRAB's design allows a transparent interaction with different grids and batch systems. CRAB has been in production and in routine use by end-users since Spring 2004. It has been extensively used between the initial definition of the CMS Computing Model in 2004 and the start of high energy collisions in 2010, attending numerous scaling tests and service challenges.

CMS is currently in the process of replacing its workload management system. This next generation tool is named WMCore, a Python based workflow library which is used for so-called task lifecycle management. The workflow library is the result of a convergence of three subprojects that respectively deal with scientific analysis, simulation and real time data aggregation from the experiment. CRAB will be part of this migration.

IMPACT

CRAB is currently starting the re-writing process with the aim to better cope with the foreseen growing amount of data and user community. This process is providing an extremely important opportunity for developers to share a set of functionality already implemented by other Workload Management Systems (WMS). These latter are the systems which have been developed by other LHC Experiments to perform grid submissions, such as Panda, Dirac, Alien and GANGA.

CONCLUSION

The CRAB Workload Management System is widely used for accessing CMS distributed data by end-users. During the first year of the data taking the system coped well with the physics needs. In order to reduce the operations load and improve the scalability, CRAB is going to be re-written. This process has also the aim to share common functionality with other existing WMS. The system is being developed in a very close collaboration with the developers of other Workload Management Systems as well as the Experiment Dashboard team.

Daniele Spiga daniele.spiga@cern.ch CERN

Getting a grip on the grid: a knowledge base to trace grid experiments

OVERVIEW

Information management is challenging in science due to the variety of data produced by the physical instruments and the amount of information generated daily by scientists. In addition to supporting the experiment execution, it is currently crucial for the scientific applications to trail their mechanisms of performing experiments, so that it is possible to trace back the resulting scientific data.

In this work, we describe an approach for building a knowledge base for the scientific experiments performed using the e-Infrastructure for bioscience (e-Bioinfra). The e-Bioinfra platform provides grid workflow management and monitoring services for biomedical researchers that use the Dutch grid. Our approach focuses on gathering meaningful information from these services and populating it into the knowledge base, within its proper context. For this, an agent-based software tool is designed and developed to retrieve, classify and transform existing data into meaningful information.

DESCRIPTION

A comprehensive knowledge base gathers relevant information to help scientists clarify their research questions and to validate operational tasks. However, building and populating such a knowledge base, with proper and detailed information resulting from different sources, is a challenge in itself. Although the information is usually accessible, e.g. in logs, it is not trivial to correlate pieces of data.

Manual data collection is an error-prone task and requires enormous manpower, due to the amount of logs registered by the processes. An automated solution is needed.

Our approach to building a knowledge base is a three-folded mechanism. First, the EbioCrawler is designed to gather automatically already existing logs that contain information generated by different application systems (e.g. MOTEUR and DIANE), including workflow descriptions and execution reports, system outputs, communication accounts and status reports.

Secondly, a provenance repository is built around the notion of graphs outlined by the OPM model. The repository is defined using a relational database schema that captures the concepts of the OPM model but extending it to support Events. Event-driven systems are common in scientific environments, but their provenance is not well captured by the OPM alone.
The provenance data collected by the EbioCrawler is stored into the repository using the Application Programme Interface (API) from the Provenance Layer Infra-structure for E-Science Resources (PLIER), allowing developers to build, store and share (by XML serialisation) graphs using the OPM model.

The third mechanism enables the analysis of the provenance data. Using the homogeneous view of the information in the repository based on the OPM ontology, the provenance can be now transformed, or serialised, into specific formats (RDF, XML, etc.) or other representations.

Examples of viewing mechanisms are a graphical interface to analyse the provenance graphs and a query interface to find events of interest.

IMPACT

The described solution delivers a set of tools that have a potential diverse audience. We foresee three main areas of impact: scientific, usability, operational.

Because of the wide and distributed scale of the resources, the 'grid' often becomes a black box to users, due to the fact that:

1. jobs may get scheduled arbitrarily on distributed resources,
2. log-files will be generated on the remote nodes selected by the scheduler,
3. the user may lose access to the remote nodes where those files reside, and
4. the logs may get cleaned without prior notification. In addition, with the increased volume of data, resources, methods and collaborators in the domain of e-science, it becomes more and more difficult to perform repeatable, scalable, and traceable experiments.

The distributed information is now collected in a structured way that can be queried for various goals like experiment results, operational statistics, tracing error, customised reporting, etc. By tracing the whole history of the resources up to the current state, it is possible to confirm or provide evidence of the scientific work done. Analysis and comparative mechanisms, scientific opinions or interpretations, and the results of various kinds of examinations may provide further ways to improve (or debug) the actual scientific experiment (workflow).

Lastly, we think that the insight that our solution provides in all the data and resources, will present an important added value to the operational aspects of running a science gateway. These benefits vary from data retention strategies, assistance for workflow debugging, and producing all kinds of usage statistics.

CONCLUSION

The knowledge base for e-Bioinfra was achieved by implementing OPM and DC data models into a relational database management system. The rich and complex data available within the e-Bioinfra application were challenging enough to test and validate our approach. These data relate to a few thousand scientific experiments, which have been performed using Moteur/Diane workflow system.

Provenance data for e-Bioinfra, including workflow descriptions and log files, is collected automatically from the various e-Bioinfra components and transformed into a knowledge base repository. This knowledge base can be accessed for analysis of the experiment provenance in various forms (e.g. GUI for provenance graphs)

Although establishing the knowledge base is fundamentally for documentation, specific tools need to be implemented in order to better explore the information in the knowledge base. Each tool could be tailored to the specific needs of the application domain and the type of users.

www.ebioscience.amc.nl

Ammar Benabdelkader ammar@pccuva.nl PCC UvA • **Mark Santcroos** m.a.santcroos@amc.uva.nl Academic Medical Center Amsterdam • **Victor Guevara Masis** vguevara@nikhef.nl Nikhef

Shared services and tools based on the Ganga job definition and management framework

OVERVIEW

Ganga is a user-targeted job management tool designed to provide a homogeneous environment for processing data on a variety of technology 'back-ends'.

We will present examples that illustrate how Ganga, originally developed by and for LHC experiments, has since been adopted by other, non-HEP, communities and how it can play a key role within a sustainable support model.

DESCRIPTION

Initially developed within the high-energy physics (HEP) domain, Ganga has since been adopted by a wide variety of other user communities as their default analysis and task-management system. The modular nature of Ganga means that, if desired, disparate communities are able to develop their own suite of tools that remain independent of both the core code and those of other communities.

This presentation will use case-studies to illustrate the ease with which non-LHC communities (such as those engaged in medical research, or running Monte Carlo simulations on the Grid), have adopted Ganga as their chief job-submission tool.

In addition to providing a stable platform with which to conduct user analysis, the Ganga development team has deployed a range of supporting tools and interfaces. We will present developments of the job monitoring interface, a lightweight tool integrated into the Dashboard monitoring service, which allows users to track the status of their tasks submitted from the Ganga/Diane environments.

The Ganga team has also deployed a service capable of receiving application 'crash reports' (comprising a snapshot of the configuration, job parameters, input/output files and command history), and presenting this information in a form that can be viewed and downloaded by user-support teams. The Ganga client provides a simple method by which users can send such reports, but this service can also be exploited with the Ganga environment.

The error-reporting tool will be described, with specific reference to how it has been adopted by the CMS VO, a community who have their own task-management system in place of Ganga, yet who were able to make use of the web-based interfaces and underlying technology to deploy an error-reporting service for their users.

IMPACT

The end-users of grid computing resources demand that the tools they use are reliable, efficient and flexible enough to meet their needs. Most users, irrespective of the research community to which they belong, are generally not interested in developing grid-access tools, and nor should they be. Their role is to exploit the resources available as effectively as possible, and with minimum knowledge of how the underlying technologies function.

A wide range of grid-enabled tools have been developed which aim to shield the user from the complexity of distributed infrastructure technology. With this goal in mind, Ganga has been engineered to provide a homogeneous environment for processing data on a range of technology "back-ends", ranging in scale from a solitary user's laptop, up to the integrated resources of the Worldwide LHC Computing Grid.

Ganga is, by design, a tool that can be extended for a given user-community to add new functionality, whether that be enhancements for data-analysis routines, or integration of tools aimed at improving the end-user experience and support mechanisms.

CONCLUSION

The Ganga team has deployed new tools and core-functionality designed to enhance the experience of end-users. These developments include the advent of an interactive user-interface for task monitoring, and an intuitive mechanism for generating error-reports that can be accessed by user-support teams.

Future developments will see an improvement in the Ganga communication channels, including the launch of a blog to keep users informed of developments and provide usage tips.

Longer term goals include the development of advanced core-features, enhancement of the data-management interface and packaging Ganga for inclusion into popular Linux distributions.

http://ganga.web.cern.ch/ganga

Ulrik Egede u.egede@imperial.ac.uk Imperial College London • **Mike Kenyon** mkenyon@cern.ch CERN • **Jakub Moscicki** jakub.moscicki@cern.ch CERN • **Ivan Antoniev Dzhunov** ivan.antoniev.dzhunov@cern.ch University of Sofia • **Lukasz Kokoszkiewicz** lukasz.kokoszkiewicz@cern.ch CERN • **Laura Sargsyan** laura.sargsyan@cern.ch CERN/Yerevan Physics Institute • **Edward Karavakis** edward.karavakis@cern.ch CERN • **Julia Andreeva** julia.andreeva@cern.ch CERN

Cheminformatics platform for drug discovery application on grid

OVERVIEW

To extend e-science applications to public health, especially in Asian countries, the EUAsiaGrid in the past two years has established an in-silico high-throughput drug screening platform, GAP Virtual Screening Service (GVSS). Providing intensive computing power and effective data management, the production e-Infrastructure (such as EGI and EUAsiaGrid) enables opportunities for in-silico drug discovery on neglected and emerging diseases, for instance, Avian Influenza and Dengue Fever.

To extend the compound 3D structure database and compound information management, we will create 3D structure conformers of known compounds for in-silico screening and post structure-activity relation analysis.

DESCRIPTION

Cheminformatics is able to come into its own as a discipline and as an enabling science. The structure based molecular docking simulation is a common method for predicting potential interacting complexes of small molecules in protein binding sites. By promoting the integration of data, applications and workflows, it is providing scientists with broad access to information and methods. As cheminformatics systems become more prevalent, we can expect them to appreciably enhance our use of previous knowledge and our ability to generate new knowledge – significantly improving the drug discovery and development process.

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 ability to process large-scale molecular docking. Grid is an ideal environment, which can provide large-scale and on-demand resources, including computing and storage resources.

Inspired by the successful experiences on Avian Flu Data Challenges, ASGC developed the GVSS application package that incorporates the EGEE gLite middle-ware DIANE2 and AMGA. Therefore, ASGC coordinated the Dengue Fever Data Challenge via the EUAsiaGrid VO in June 2009. The objective is to use grid enabled high-throughput screening for structure-based computational methods to identify small molecule protease inhibitors. In addition to the ZINC CDI 300,000 compounds library we used, we established the new ZINC 800,000 ligands and Chembridge 300,000 ligands, a free database of available compounds library selected for virtual screening.

IMPACT

We introduce GVSS, a user-friendly graphical user interface 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, visualisation and analysis with this GUI, of course including the advanced refinement docking simulations. The end-users can finally have a real grid-enabled desktop utility for their daily research.

To improve the application domain and bridge the gap between experimental and in-silico endeavour, firstly, we will focus on common interest and also on the most important issues of public health. The activity, therefore, has two approaches aimed to accelerate and optimise the drug discovery process by computing method: firstly, to increase the collection of the compound library, so as to extend the present compound structure database; secondly, to execute in-silico to disease target, such as critical proteins in cell wall formation and viral replication pathway. Examples are: penicillin-binding proteins (PbP) for antibiotics development and non-structure protein (e.g. NS3 and NS5) in anti-flavivirus drug development. Validating the computer work with experimental assay is also the focus of this activity, which bridges e-science activity and bioscience.

CONCLUSION

Molecular docking simulation is a time-consuming process that searches exhaustively all correct conformations of a compound. However, the massive in-silico processes benefit from the high-throughput computing grid technology. Providing intensive computing power and effective data management, the production e-Infrastructure (EUAsia VO) enables opportunities for an in-silico drug discovery platform. We deployed this high-throughput in-silico massive molecular docking service benefiting from state-of-the-art grid technology to the activities for designing the cheminformatics platform on EUAsiaGrid infrastructure since 2010. Furthermore, these activities also facilitate more biomedical e-science applications, such as other diseases and compounds profiling.

Hsin-Yen Chen hychen@twgrid.org ASGC

VisIVO, visualisation in different grid environments, from gLite to desktop grid

OVERVIEW

The availability of distributed computing infrastructures has been growing rapidly in the last few years. Many scientific fields can now produce high-resolution numerical simulations with multi-dimensional datasets in the order of several petabytes. An essential part of modern scientific research is the necessity to employ computer graphics and scientific visualisation tools for appropriately displaying such datasets, so as to allow scientists to perform efficient visual discovery. We present VisIVO, a powerful environment for exploring highly complex multi-dimensional astrophysical datasets on the grid.

VisIVO provides an integrated suite of tools and services that can be used in many other scientific fields. VisIVO allows users to visualise meaningfully highly-complex, large-scale datasets and create movies of these visualisations based on distributed infrastructures. The deployment of VisIVO on the DG and gLite is carried out with the support of EDGI and EGI-InSPIRE projects.

DESCRIPTION

VisIVO supports high-performance, multi-dimensional visualisation of large-scale astrophysical datasets. Users can rapidly obtain meaningful visualisations while preserving full and intuitive control of the relevant parameters. VisIVO consists of VisiVO Desktop – a stand-alone application for interactive visualisation on standard PCs, VisIVO Server – a platform for high performance visualisation, VisIVO Web – a custom designed web portal supporting services based on the VisIVO Server functionality and VisIVOSmartphone – a web application allowing modern smart phones to exploit the VisIVOServer functionality.

We are currently in the process of porting VisIVO Server on grid environments. VisIVO Server consists of three core components: VisIVO Importer, VisIVO Filter and VisIVO Viewer respectively. To create customised views of 3D renderings from astrophysical data tables, a two-stage process is necessary. First, VisIVO Importer is used to convert user datasets into VisIVO Binary Tables (VBTs). Then, VisIVO Viewer is invoked for display. VisIVO Filters are collections of data processing modules able to extract interesting features in datasets for performing visual discovery.

Depending on the structure and size of datasets in consideration, the data exploration process could take several hours of CPU for creating customised views, and the production of movies could potentially last several days. For this reason an MPI parallelised version of VisIVO can play a fundamental role in increasing performance, e.g. it could be deployed automatically on nodes that are MPI aware. A central concept in our development is thus to produce unified code that can run as necessary either on serial nodes or in parallel by using HPC oriented grid nodes. Another important aspect, to obtain as high performance as possible, is the integration of VisIVO processes with grid nodes where GPUs are available. We have selected CUDA for implementing a range of computationally heavy modules.

IMPACT

The main characteristic of modern astronomical datasets is extremely large sizes, typically requiring storage in a distributed way. Gaining an insight into such datasets requires very powerful statistical and data analysis algorithms that can be fully realised, in a cost-effective manner, only through gateways to grid and cloud infrastructures. Forthcoming astronomical surveys are expected to collect petabytes of raw data resulting in massively large-scale distributed databases requiring that astronomers interact with them using very high performance data analysis and visualisation tools.

INAF-OACT is involved in numerous grid activities, e.g. the COMETA consortium that is developing a large grid infrastructure, not only for research but also for industrial applications. INAF is also a full partner of the Italian Grid Initiative (IGI) which regards science gateways as an extremely valuable service to be provided to the scientific community. IGI has expressed a strong interest in developing visualisation tools on the grid. IGI is also part of the European Grid Infrastructure together with many other European NGIs.

We are also working with science centres on a pilot experiment to identify novel ways for exploiting the functionality of VisIVO using smart phones for large-scale public engagement activities. The use of DG could significantly improve this experiment. Within the activities in the EDGI and EGI-InSPIRE projects we will port and install our portals in dedicated servers exploiting existing grid infrastructures. We envisage that using smart phones via these portals, potential users (researchers, citizen scientists, science centre visitors) would be able to exploit VisIVO functionality on-the-go, accessing archives of large-scale datasets and generating images or movies by exploiting the power of grid infrastructures.

CONCLUSION

Our future work will focus on fully porting VisIVO Server and VisIVO Web to grid environments, more specifically using desktop grids and gLite. The exploitation of grid technologies will enhance performance significantly by supplying the required computational capability. A very exciting possibility is the deployment of this technology for visual discovery in scientific disciplines apart from astrophysics. Furthermore, a pilot application built upon VisIVOSmartphone is underway for public engagement activities in astrophysics. Our ultimate vision is users (researchers, citizen scientists or even science centre visitors) being able to reserve resources for data analysis and visual discovery as required (possibly involving several production grids) and retrieving results on smart phones on-the-go irrespective of internet access or geographical location.

http://visivoweb.oact.inaf.it

UGO BECCIANI ugo.becciani@oact.inaf.it INAF

e-Infrastructures integration with gCube

OVERVIEW

Delivering an e-Infrastructure service to large organisations is a complex task that requires the integration of several technologies. The complexity of this service resides on:

1. very rich applications and data collections are maintained by a multitude of authoritative providers;
2. different problems require different execution paradigms: batch, map-reduce, synchronous call, message-queue, etc.;
3. key distributed computation middleware exist: gLite, Globus, and Unicore for grid-based wide resource sharing; Condor for site resource sharing; Hadoop and Cassandra for cluster resource sharing, etc.;
4. several standards exist in the same domain.

gCube, the D4Science empowering technology, offers solutions to abstract over differences in location, protocols, and models by scaling no less than the interfaced resources, by keeping failures partial and temporary, and by being autonomic reacting and recovering from a large number of potential issues.

DESCRIPTION

gCube is a large software framework designed to abstract over several technologies and offer them through a well-formed set of APIs. gCube consists of several packages offering:

- Access to several storage back-ends tailored for different needs. For example, it offers a storage server: for multiple-version software packages; for scientific data-sets stored as tables; for time series with an OLAP interface; for structured document objects; for geo-coded datasets compliant with OGC; and finally for storing files;
- Management of metadata in any format and schema that can be consumed by the same application in the same Virtual Organisation;
- A process execution engine to manage the execution of software elements in a distributed infrastructure under the coordination of a composite plan that defines the data dependencies among its actors. It supports several computational middlewares without performance compromises. A task can be designed as a workflow of invocation of different code components (services, binary executables, scripts, map-reduce jobs, etc.);
- A transformation engine to transform data among various manifestations. This engine is manifestation and transformation agnostic by offering an object-driven operation workflow. It is extensible through the addition of transformation-programme plugins;
- A Virtual Research Environment (VRE). Through VREs, groups of users have controlled access to distributed data, services, storage, and computational resources integrated under a personalised environment. VREs support cooperative activities such as: metadata cleaning, enrichment, and transformation by exploiting mapping schema, controlled vocabulary, thesauri, and ontology; processes refinement and show cases implementation; data assessment; expert users validation of products generated through data elaboration or simulation; sharing of data and process with other users.

IMPACT

gCube doesn't hide the infrastructure middleware. It is not another layer. Rather it turns infrastructures into a utility by offering single submission, monitoring, and access facilities. It offers a common framework to programming in the large and in the small. It allows concurrent exploitation of private virtualised resources organised in sites with resources provided by IaaS and PaaS cloud providers.

By using its set of facilities, several scientific applications have been implemented and delivered to the Fishery and Aquaculture Resource Management communities delivering the following applications (among others):

- A collaboration-oriented suite providing seamless access and organisation facilities on a rich array of objects (e.g. information objects, queries, files, templates, time-series). It offers mediation capabilities between the external world objects, systems and infrastructures (import/export/publishing) and it supports common file management features (drag and drop, contextual menu);
- A time-series framework offering tools to manage large datasets by supporting the complete time-series lifecycle (validation, curation, analysis, and reallocation). It offers tools to operate on multi-dimensional statistical data. It supports filtering, grouping, aggregation, union, mining, and plotting;
- An ecological niche modelling suite to predict the global distribution of marine species, that generates colour-coded species range maps using half-degree latitude and longitude blocks by interfacing several scientific species-databases and repository providers. It allows the extrapolation of known species occurrences to determine environmental envelopes (species' tolerances) and to predict future distributions by matching species tolerances against local environmental conditions (e.g. climate change and sea pollution).

CONCLUSION

gCube is currently deployed in the D4Science production infrastructure and provides to the Fishery and Aquaculture Resource Management communities a large number of APIs and highly specialised scientific applications running on an e-Infrastructure service which hides the complexity of the underlying heterogeneous set of middleware systems, standards, data types, metadata schemas, etc.

www.gcube-system.org

Pedro Andrade pedro.andrade@cern.ch CERN • **Pasquale Pagano** pasquale.pagano@isti.cnr.it CNR • **Andrea Manzi** andrea.manzi@cern.ch CERN

Distributed multi-scale computing: the MAPPER project

OVERVIEW

Today scientists and engineers are commonly faced with the challenge of modelling, predicting and controlling multi-scale systems which cross scientific disciplines and where several processes acting at different scales coexist and interact. Such multidisciplinary multi-scale models, when simulated in three dimensions, require large scale or even extreme scale computing capabilities. The MAPPER project develops computational strategies, software and services for distributed multi-scale simulations across disciplines, exploiting existing and evolving European e-Infrastructure.

DESCRIPTION

Driven by seven challenging applications from five representative scientific domains (fusion, clinical decision making, systems biology, nanoscience, engineering), MAPPER deploys a computational science environment for distributed multi-scale computing on and across European e-Infrastructures. By taking advantage of existing software and services, as delivered by EU and national projects, MAPPER will result in high quality components for today's e-Infrastructures. We develop tools, software and services that permit loosely and tightly coupled multi-scale computing in a user friendly and transparent way. We integrate our applications into the MAPPER environment, and demonstrate their enhanced capabilities.

MAPPER integrates heterogeneous infrastructures for programming and execution of multi-scale simulations. We reuse as much of the existing infrastructural and software solutions as possible. The MAPPER solution is developed on top of existing e-Infrastructures without the necessity to modify already deployed components. The functionality to be delivered is realised as extensions to existing e-Infrastructures. The integration is done using well defined APIs and standard based interfaces, thus reducing potential impact of changes on middleware level components.

IMPACT

MAPPER is driven by seven exemplar applications from five user communities (virtual physiological human, computational biology, fusion, hydrological engineering, nano material science), and these communities are specifically targeted. However, our solutions are generic and will enable distributed multi-scale computing for any multi-scale model fitting into our paradigm, and MAPPER therefore opens up to other user communities as well.

MAPPER partners have significant trans-Atlantic grid and HPC experience, and have been involved very actively in TeraGrid and with the US Department of Energy laboratories. We collaborate with the US TeraGrid to integrate infrastructures across the globe.

CONCLUSION

MAPPER started on 1 October 2010, and is now producing first results. It is expected that first demonstrations of Distributed Multi-scale Computing will be available during the conference. MAPPER will quickly move forward, and it expects to be able to offer a first Distributed Multi-scale Computing environment for external projects when entering its second project year.

The research leading to these results has received funding from the European Union Seventh Framework Programme under grant agreement n° 261507 (the MAPPER project).

www.mapper-project.eu

Alfons Hoekstra a.g.hoekstra@uva.nl University of Amsterdam

Theompi: a large MPI cluster on the grid for theoretical physics

OVERVIEW

We are reporting about the first experience using a large grid-enabled MPI-cluster and the problematic/advantage of its use for massive parallel applications.

In recent years EGEE/EGI has proven to be a robust and scalable infrastructure for sequential scientific computation, but hardly usable for parallel programming. The situation changed following the June 2010 proposal of the EGEE MPI WG (Working Group) for JDL attributes to select the slots needed for the execution and allocation of a parallel application. This recent evolution has convinced the INFN theoretical physics community to install a first large grid-enabled MPI cluster, named Theompi, taking advantage of an experimental implementation of the granularity attributes. We describe the porting of significant applications in theoretical physics, executed in this new environment and taking advantage of the new grid parallel attributes.

DESCRIPTION

The cluster has been installed at the INFN-Pisa site and is grid-enabled through a CREAM based computing element. The production version of the software has been enhanced and customised with an experimental patch to better support the parallel job execution on a multicore environment, according to the recommendations given by the EGEE MPI WG. This patch has been tested and deployed in collaboration with the gLite middleware developers.

In a preliminary operational period of the cluster a set of theoretical physics applications has been executed on the cluster in order to test the functionality of the new architecture for different kinds of parallel computational models, such as pure MPI, pure openMP and hybrid algorithms, taking into account also the CPU and memory affinity for the NUMA architecture of the Worker Nodes.

IMPACT

The INFN theoretical physics community consists of 700 FTE researchers distributed over 28 sites and involved in 60 research projects. In this community there is a widespread need for massive parallel applications on HPC systems with next neighbours high speed inter-connections (involving Lattice Quantum ChromoDynamics, Fluid Dynamics, Numerical Relativity) that in the past used small or medium-sized dedicated local clusters. The Theompi project aims to provide the community with a transparent and flexible mechanism to share the HPC resources.

CONCLUSION

The support of the new granularity attributes in the gLite middleware gives a real chance to use grid clusters for both MPI and multithreaded applications. We decided to start with a single large cluster, but the flexibility of this mechanism will allow the integration of other (old and new) clusters in a near future.

http://wiki.infn.it/cn/csn4/calcolo/csn4cluster/home

Roberto Alfieri roberto.alfieri@pr.infn.it INFN Parma – Parma University

CORAL – a relational abstraction layer for C++ or Python applications

OVERVIEW

The huge amount of experimental data from the LHC and the large processing capacity required for their analysis has imposed a new approach involving distributed analysis across several institutions. The non-homogeneity of policies and technologies in use at the different sites and during the different phases of the experiment lifetime has created one of the most important challenges of the LHC Computing Grid (LCG) project. In this context, a variety of different relational database technologies may need to be accessed by the C++ client applications used by the experiment for data processing and analysis. The Common Relational Abstraction Layer (CORAL) is a software package that was designed to simplify the development of such applications, by screening individual users from the database-specific C++ APIs and SQL flavours.

DESCRIPTION

CORAL is a C++ software package that supports data persistency for several relational database back-ends. It is one of three packages (CORAL, POOL and COOL) that are jointly developed by the CERN IT Department and the LHC experiments within the context of the LCG Persistency Framework project. The CORAL API consists of a set of abstract C++ interfaces that isolate the user code from the database implementation technology. CORAL supports several back-ends and deployment models, including local access to SQLite files, direct client access to Oracle and MySQL servers, and read-only access to Oracle through the Frontier/Squid and CoralServer/CoralServerProxy intermediate server/cache layers. Users are not required to possess a detailed knowledge of the SQL flavour specific to each back-end, as the SQL commands are executed by the relevant CORAL implementation libraries (which are loaded at run-time by a special plugin infrastructure, thus avoiding direct link-time dependencies of user applications against the low-level back-end libraries).

IMPACT

CORAL provides generic software libraries and tools that do not target specific data models and could therefore be used in any other scientific domain to access relational databases from C++ or Python applications.

CONCLUSION

The CORAL software is widely used for accessing from C++ and Python applications the data stored by the LHC experiments using a variety of relational database technologies (including Oracle, MySQL and SQLite). It provides generic software libraries and tools that do not specifically target the data models of the LHC experiments and could therefore be used in any other scientific domain.

Raffaello Trentadue raffaello.trentadue@cern.ch CERN • **Alexander Loth** alexander.loth@cern.ch CERN

Challenges in the adoption of the EGI paradigm for an e-Science/Tier 2 centre (ES-ATLAS-T2)

OVERVIEW

The High Energy Physics (HEP) community has reached great achievements during the last year, with the run of the Large Hadron Collider (LHC) again and obtaining first collisions on November 2009, and with a success run at 7 GeV at the beginning of 2010.

In this contribution we present the current production ready infrastructure of the IFIC institute member of the Spanish ATLAS Tier2 (ES-ATLAS-T2) and report its response at different stages, from data taking to experimental results, and we describe the facilities for data analysis (Tier 3) set up at IFIC.

We will also present the processes being implemented to adopt the new situation with the EGI transition in its different phases, with the key requirement of being the least disruptive possible on the production operations and on the day-by-day end-user work. As a member of the EGI project it will be also presented the contributions to several activities in order to provide better and stable software at the end infrastructure.

DESCRIPTION

The computing model of the ATLAS experiment at the LHC is based on a tiered hierarchy that ranges from Tier 0 (CERN) down to end-user's own resources (Tier 3). Levels 0 to 2 are well defined within the computing model and have been thoroughly tested in the last years. At Tier 0 takes place a first event reconstruction, the main purpose of the Tier 1s distributed all over the world is to reprove and to carry out analysis that needs to access to huge amounts of data and Tier 2s mainly take care of providing CPU and storage resources for the various physics groups' analysis and official Monte Carlo data production. Tier 3 sites are institution-level non-ATLAS funded that participate most frequently in support of the particular interests of local physicists.

In the case of ATLAS, the raw data coming from the detector is reconstructed to produce Event Summary Data (ESD) and Analysis Object Data (AOD) files. The ESD incorporates all of the information of the event reconstruction, and is mainly used for detector performance studies. The AOD contains only part of the information available in the ESD and is meant to be used for physics analyses. Data distribution and storage organisation is based on space tokens. According to the Computing Model, these space tokens are controlled through the Distributed Data Management (DDM) system, which is working 'a la grid', and they are associated to a path to a Storage Element (SE).

This model is supported by the LCG project that sums up the participant institutes to provide a coherent production infrastructure with three grid flavours including gLite, OSG, and ARC. This model is going to change at the beginning of 2011 and the WLCG will decide which of them has the blessing for adoption.

The transition from EGEE to EGI is being done and the important issues that affect ATLAS are:

- Support for middleware and tools.
- Infrastructure support according to required levels.
- Not disturbing the current operations and end-users.

IMPACT

This model and the IFIC prototype was previously presented in EGEE User Forums, and shown its validity for managing the ATLAS data, but with the finalisation of the EGEE project that was supporting the main software releases of gLite, new challenges and issues arise, including the risks of timeline releases and the availability of supporters. EGI-InSPIRE intends to continue supporting the production infrastructure, with a broader end-user community, but leaving middleware developments for outside providers, being the main one the releases from EMI (European Middleware Initiative).

It reflects the key tools that are being used like DDM (Distributed Data Management), and Ganga, and how are they supported through the Services for HEP Heavy User Community in EGI.

The main impact would be that this activity provides continued support for activities previously supported by EGEE while they transition to a sustainable support model within their own community or within the production infrastructure. The support model for High Energy Physics will be described in the document MS603 – Services for High Energy Physics and made available in the next version of the present deliverable. In this way, maybe the main impact will be the end of the lcg-CE to move to Cream. The proposal is that all sites supporting LHC experiments run CREAM and are no longer required to run LCG-CE for LHC.

CONCLUSION

The main conclusions would be:

1. LHC started again on November 2009 and successfully reached 7 TeV.
2. IFIC is part of Spanish Tier 2, and defined its Tier 3 to fulfil ATLAS requirements. Computing and Storage resources are in place according to 2010 pledges.
3. Common middleware services and operations are now supported by EGI IFIC. Users submit its analysis jobs where data is, replicating most used datasets to local storage.
4. Various tools are used, some of them supported by EGI, HEP being a Heavy User Community: Ganga, Panda, DDM or Dashboards xrootd for interactive analysis.

http://ific.uv.es/grid/e-science
https://twiki.ific.uv.es/twiki/bin/view/Atlas/GridComputing

Santiago Gonzalez de la Hoz santiago.gonzalez@ific.uv.es CSIC

Monitoring of the LHC computing activities during the first year of data taking

OVERVIEW

The Worldwide LHC Computing Grid (WLCG) provides the grid infrastructure used by the experiments of the Large Hadron Collider at CERN, which started taking data this year. The computing and storage resources made available to the LHC community are heterogeneous and distributed over more than a hundred research centres. The scale of WLCG computing is unprecedented; the LHC Virtual Organisations (VOs) alone run 100,000 concurrent jobs and the ATLAS VO can sustain an integrated data transfer rate of 3GB/s. Reliable monitoring of the LHC computing activities and the quality of the distributed infrastructure are a prerequisite for the success of the LHC data processing.

DESCRIPTION

The Experiment Dashboard system was developed in order to address the monitoring needs of the LHC experiments. It covers data transfer and job processing and works transparently across the various middleware flavours used by the LHC VOs. This presentation will describe the experience of using the system during the first year of LHC data-taking, focusing on the dashboard applications that monitor VO computing activities. Those applications that monitor the distributed infrastructure are the subject of a different presentation, 'Experiment Dashboard providing generic functionality for monitoring of the distributed infrastructure'. Though primarily the target user communities of the Experiment Dashboard are the LHC experiments, many of the Experiment Dashboard applications are generic and can be used outside the scope of the LHC. Special attention will be given to generic applications such as job monitoring, and the common mechanism that can be used by VO-specific Workload Management Systems (WMS) for reporting monitoring data.

IMPACT

The Experiment Dashboard system plays an important role in the computing operations of the LHC Virtual Organisations, in particular those of ATLAS and CMS, and is widely used by the LHC community. For example, the CMS VO's Dashboard server receives up to 5,000 unique visitors per month and serves more than 100,000 page impressions daily.

CONCLUSION

During the first year of the data taking the system coped well with growing load both in terms of the scale of the LHC computing activities and in terms of number of users. The Experiment Dashboard system became an essential component for the LHC computing operations. The variety of its applications covers the full range of the LHC computing activities. The system is being developed in a very close collaboration with the users. As a result, the Experiment Dashboard manages to respond well to the needs of the LHC experiments.

http://dashboard.cern.ch

Julia Andreeva julia.andreeva@cern.ch CERN • **Mattia Cinquilli** mattia.cinquilli@cern.ch CERN • **Marco Devesas Campos** marco.devesas.campos@cern.ch CERN • **Edward Karavakis** edward.karavakis@cern.ch CERN • **Elisa Lanciotti** elisa.lanciotti@cern.ch CERN • **Gerhild Maier** gerhild.maier@cern.ch CERN • **Sergey Mitsyn** svm@mammoth.jinr.ru CERN • **Michal Nowotka** michal.maciej.nowotka@cern.ch CERN • **William Ollivier** william.ollivier@cern.ch CERN • **Ricardo Rocha** ricardo.rocha@cern.ch CERN • **Timur Sadykov** timur.sadykov@cern.ch JINR • **Pablo Saiz** pablo.saiz@cern.ch CERN • **Laura Sargsyan** laura.sargsyan@cern.ch Yerevan Physics Institute • **Irina Sidorova** irina.sidorova@gmail.com JINR • **Jacobo Tarragon Cros** jacobo.tarragon.cros@cern.ch CERN • **David Tuckett** david.tuckett@cern.ch CERN

Optimisation of the job submission and data access in a LHC Tier2

OVERVIEW

Since LHC start-up, HEP users have increased their requests on the computing infrastructure in terms of performance and functionality, often going beyond the projected requirements of the experiments' computing model.

Several users want an interactive facility to test their code before submitting it to the grid. Usually such an interactive cluster is deployed on a small number of machines. In our experience this is not the best solution due to scalability limits and maintenance issues. Users also want to use the local batch submission to run their analysis quickly and reliably, using a controlled environment, on a small dataset produced by them.

The network topology and the storage infrastructure also needs to be set up in order to fulfil the I/O requirements of analysis jobs; this can be difficult especially within a big multi-VO site, where users can have very different use cases.

DESCRIPTION

Our work consists first of all in creating a recipe to provide the interactive facility using worker nodes; Torque interactive jobs have been a starting point to deploy such facility. We have also tested many different storage solutions in order to choose the one which best fulfilled all the user requirements. These storage tests have been performed on many different storage solutions, using different storage brands and different technologies (hardware RAID, software RAID based on the Linux kernel; RAID5, RAID6; fibre channel, external SAS) as one of the primary goals of the test was to assure that a heterogeneous storage infrastructure could be successfully built to achieve high performances. Lustre was chosen as it provides POSIX access, the best performances, and easy administration. As a result of those tests, our farm migrated from dCache to Lustre.

Finally we have successfully connected Lustre to the grid using consolidated solutions like StoRM, XROOTD, GridFTP, verifying the work in terms of performance and compliance with grid requirements.

IMPACT

Using Torque interactive jobs, the user connects to a front-end machine and runs a command to submit a job to the cluster, just like he does for batch local submission; the batch manager chooses one CPU to execute the job and returns an interactive shell. The user will keep that CPU until he releases the interactive jobs through a logout. Using the screen utility, the user can also interrupt his work, log out, and then log in back again preserving the session he had left.

Lustre has been mounted on all the nodes of the farm, so that users can access their data through grid jobs, local batch jobs and interactive jobs, simplifying users' life.

The new storage configuration has been tested using CMS analysis jobs, resulting in a very high CPU efficiency when compared with other storage solutions.

As a consequence of the design of the overall infrastructure, it's very easy for new users to start with local activities, and for new experiments to add new resources in terms of nodes and storage servers.

CONCLUSION

The number of users has considerably grown during the last year, involving several experiments, VOs and communities (CMS, Alice, Glast/Fermi, Pamela, Theophys, Magic V, computational chemistry, bio-medicine, bioinformatics, etc.) actively using the new infrastructure both locally and through EGI. This has also resulted in a considerable drop in the maintenance overhead of the site administrator.

Giacinto Donvito giacinto.donvito@ba.infn.it INFN

The EDGI infrastructure and its usage for the European Grid user communities

OVERVIEW

The EDGI (European Desktop Grid Initiative) FP7 project is aimed at deploying Desktop Grid (DG) and Cloud Computing services for the EGI research user communities that require large-scale distributed computing resources for multi-national projects. In order to achieve this goal EDGI develops middleware for extending Service Grids (SG) (e.g. ARC, gLite, UNICORE) with Desktop Grids (e.g. BOINC, XtremWeb) enhanced by Academic Cloud infrastructures based on Eucalyptus and OpenNebula. Software components of ARC, gLite, UNICORE, BOINC, XWHEP, Attic, 3GBridge, OpenNebula and Eucalyptus will be integrated into a SG→DG→Cloud platform for service provision and as a result EDGI will extend ARC, gLite and UNICORE grids with volunteer and institutional DG systems. In this way, the whole European e-science ecosystem will benefit from Desktop Grid extensions since parameter sweep applications can be directed from the expensive cluster and supercomputer resources to cheap desktop resources.

DESCRIPTION

The SG to DG bridging service which is being developed by the EDGI project leans on two key components: modified computing element on the Service Grid (gLite, ARC, UNICORE) to forward the job to a Desktop Grid site and the 3GBridge service on the Desktop Grid site which is able to receive and transform the job as a workunit for execution for the Desktop Grid (BOINC, XtremWeb) site.

The objective of automatically forwarding gLite jobs to Desktop Grid resources has been already achieved by the former EDGeS project that built a production infrastructure enabling the extension of gLite VOs with several volunteer and institutional Desktop Grids. However, this version had some scalability issues related to the gLite modified CE and to the fact that all data have been forwarded through the bridge components which became a bottleneck.

In EDGI one of the focuses is on eliminating scalability issues, thus first we changed gLite CE to CREAM CE, secondly we added URL pass through mechanism to the 3GBridge software in order to be able to forward the reference to files instead of the files themselves. Moreover, Attic (a P2P file system) has also been adapted to the EDGI bridging components to increase the efficiency of data transfer among the server and clients.

In order to seamlessly integrate cloud resources into the EDGI infrastructure, we have extended the bridge with a special, automatic Eucalyptus cloud handling mechanism, including instantiation and job execution. Due to the plugin framework of the 3GBridge this could easily be implemented by adding a new plugin without modifying the core software. With this extension jobs arriving to the bridge can be executed on cloud resources.

Based on the solution above, the first software release of the EDGI project has been finished. Currently, we are working on integrating our Desktop Grid infrastructure – containing tens/hundreds of thousands of computers – together with our bridging services into EGI.

IMPACT

Service (cluster) grids like EGI cannot always provide the required number of resources for many VOs. Therefore extending the capacity of these VOs with volunteer or institutional Desktop Grids would significantly increase the number of accessible computing resources that can particularly advantageously be exploited in case of parameter sweep applications.

EDGI is currently focusing on three different areas: The first main area is to create technical solutions and develop the necessary software extensions for ARC, UNICORE, Attic, 3GBridge, Application Repository, Monitoring, BOINC, XtremWeb, Eucalyptus and OpenNebula. With these developments we can enable the creation of a Desktop Grid in order to support service grids with a huge number of resources. The second main area is to continuously operate and maintain the EDGI infrastructure built by these software components and provided as a production service for all EGI User communities. The third main area is to extend the capacity of the infrastructure by institutional and volunteer desktop and cloud resources to continuously keep up with the requirements of computational capacity.

In order to provide sustainability for the project, EDGI together with its partner project called DEGISCO have recently established an organisation called International Desktop Grid Federation (IDGF) which aims to bring together people (technical experts, developers, operators, users or any interested people) from all areas related to desktop Grid computing and aims to share the common knowledge among its members.

CONCLUSION

The EDGI desktop grid and cloud infrastructure contains a large number (>100,000) of volunteer and institutional desktop resources. The EDGI infrastructure contains bridging services which are able to automatically forward jobs from service grids (currently gLite and ARC) to Desktop Grids (currently BOINC and XtremWeb). Moreover, cloud resources (currently Eucalyptus/Amazon) can also be integrated to expand the currently available reliable and non-reliable resources based on the latest developments.

EGI user communities can gain significant advantages from this huge pool of resources or if there is a need they can set up their own pool based on the EDGI software components. The solution brought by EDGI also ensures that EGI users do not need to change their well-known environment, since the infrastructure developers provided the access to the desktop grid and cloud resources through gLite and ARC user interfaces. Within a year the UNICORE interface will also be supported.

http://edgi-project.eu
http://desktopgridfederation.org

Jozsef Kovacs smith@sztaki.hu MTA SZTAKI

Cross-grids simulation scenarios with the GridSFEA framework

OVERVIEW

Interoperability of grids is a major concern for the scientific community. Whereas most efforts are focused on finding solutions for middleware interoperability, we propose an approach closer to applications and end-users. Our approach is based on extending the client capabilities of the grid simulation framework for engineering applications (GridSFEA), to support middleware such as Globus Toolkit (GT), UNICORE, and gLite. For demonstrating the application-level interoperability we present the cross grids computations of scenarios of numerical simulations. The prototype provides high-level end-user operations such as parameter investigations, handling of large jobs, handling of remote simulation data. The tools and the programming interface of the framework act as a mediator between users and applications on one side, and grids on the other side, converting the grid interactions into operations specific to a grid middleware. Experiences are gathered on RoGrid-NGI and on DEISA.

DESCRIPTION

The development of grid applications that keep pace with the rapid evolution of the grid middleware is still a challenging task. Moreover, we face today a great effort put in the direction of the interoperability of grids, of the development and adoption of common open standards. In this context, the main objective of the GridSFEA framework is to gridify simulation scenarios stemming from different fields of computational sciences and engineering, following standards, libraries, and best practices available in the grid community. The initial target middleware was the Globus Toolkit 4, currently client-side support for GT5, UNICORE, and gLite being available within GridSFEA.

The framework handles simulation scenarios on the grid, from their formulation as grid jobs, submission and monitoring of the status' execution, attachment to jobs of tasks for post-processing simulation results, to the manipulation of simulation checkpoints and scenario data (intra- or cross-grid transfers, downloads to end-user machines). Thus, high-level operations such as parameter investigations, management of long running simulations, and preview of remote results are made available to framework users. Scenarios are annotated; the respective metadata is collected at runtime by application wrappers, and it is registered with the grid services of the framework. These services are WSRF-based (GT4 specific). The client tools of GridSFEA are based on the Eclipse plugin technology. The main challenges we faced were: First, the integration of client-side libraries specific to different middleware proved to be difficult from both the software engineering and the technical perspectives (e.g. conflicting third-party libraries required). Second, the JSDL requests from users/clients are mapped to specifics of the middleware, entailing considerable work, thus only a partial implementation of the JSDL spec is currently supported.

IMPACT

The cross-grid set up employed for this work included the grids RoGrid-NGI (gLite), DEISA (GT4, GT5), and a test grid (GT4). The WSRF-based simulation services of GridSFEA were deployed on the test grid, application wrappers and simulation codes on all sites, the client tools of GridSFEA were installed on an end-user machine located outside these grids. A limitation of our set-up is the necessity of having at least one deployment of GT4 in the cross-grid that hosts our grid services. We computed fluid-structure interaction scenarios on the three grids, each job running within one site at the time. Continuation of jobs on other sites was handled by GridSFEA entirely. Thus, our prototype was capable of operating under the strict requirements of production grids the entire framework and simulation tools present in the user space only (excepting the simulation services). With our work, users can interact with different grids in the same way, with solely knowing a subset of JSDL. In this way, the framework is a concrete example of a tool providing application-level interoperability of grids.

We evaluated community libraries such as jGlobus, GAT, SAGA, DESHL, and CREAM and integrated them to some extent in GridSFEA. By using the Eclipse plugins as enabling technology, the reuse of plugins from and by other tools such as g-Eclipse is fostered. The export of parameter studies of GridSFEA as high-level operations in the WS-VLAM tool brings benefits to further user communities. G-Eclipse is the closest to our approach, but as far as we are aware of, it lacks operations such as migration of jobs, handling of scenarios, or simulation preview. GridWay is also capable of interacting with different middleware, but it is a middleware component and has a completely different goal, namely metascheduling.

CONCLUSION

GridSFEA has been successfully employed to gridify complex simulation scenarios. With the tools of the framework, the migration of scenarios on heterogeneous grids is made possible. With our approach, user applications can run on heterogeneous grids based on GT4, GT5, gLite, and Unicore. From the user's point of view, the grid interaction is the same regardless of the type of middleware in place. End-users of the grid benefit from high-level operations that hide the complexity of the different grids. Jobs are translated on the fly into other representations; data handling is done using GridFTP. We considered for the next step the investigation of the interoperability of GridSFEA with GridWay for providing flexible migration policies and for interactions with further middleware. This work was partly funded by the EC project "Infrastructure for Globus in Europe", contract nr. 261560/21.06.2010 and the Romanian National Council of Scientific Research, contract nr. RU-RP 10/2009.

http://acal.utcluj.ro

Ioan Lucian Muntean ioan.lucian.muntean@cs.utcluj.ro Technical University of Cluj-Napoca, Department of Computer Science

Integration, sharing and exploitation of national and international e-Infrastructures

OVERVIEW

The DEGISCO infrastructure support project inherited a production-level and hybrid Distributed Computing Infrastructure (DCI) from the successfully completed EDGeS project based on a generic grid-to-grid bridging technology.

The main aim of DEGISCO is the further extension and exploitation of the DCI in terms of:

- involved service, desktop, and volunteer grids,
- ported applications,
- supported user communities,
- number of volunteers.

The work is in progress with several achievements in strong collaborations worldwide with several ICPC partners, such as Russian Federation, China, and Brazil.

The operated DCI consists of more than 150,000 PCs from more than 15 global volunteer and local desktop grids, which have been connected to gLite based European VOs from EGI and SEE-Grid, and recently from national infrastructures, e.g. Russian Data-Intensive grid and TWGRID. More than 25 applications have been ported to the combined DCI using the EADM application development methodology.

DESCRIPTION

Traditional desktop grids consist of computers and other devices, including desktop PCs and notebooks that are used for general purposes but having unused computational (CPU/GPU) and storage capacities.

These DCIs can be formed inside institutes and universities (local desktop grids) or by citizens that voluntarily donate spare computing time to science (volunteer desktop grids). Both types of grids collect large numbers of underutilised resources and can offer them for scientific applications and users.

To be more useful for researchers, desktop grids have been integrated into scientific workflows on a regular basis; the elaborated generic bridge between desktop grids and traditional service grids together with the appropriate application development methodology and transparent access mechanisms/tools foster the convergence of distributed computing infrastructures.

The recently launched DEGISCO project transfers the knowledge concerning this combined DCI towards new communities by supporting the creation, integration, and operation of new desktop grids for e-Science worldwide.

As the result of the current activities, the project members provide best practices and well-organised assistance in application porting, as well as training about the grid and its usage.

Several popular and generic applications are already available: e.g. Autodock, Blender, MOPAC, or ISDEP, but several new applications have been ported as well, e.g. for solving optimisation problems.

In order to enhance further the exploitation of the infrastructure, DEGISCO is searching for potential new applications and users that could benefit from the research infrastructure. For the newcomers the International Desktop Grid Federation (IDGF) can provide efficient help; IDGF has been set up to exchange experiences about the usage of desktop grid technology, to expand scientific infrastructures, and in order to bring together grid operators, application developers, and other key players.

IMPACT

The project supports the creation of new desktop grids in ICPC (International Cooperation Partner Countries) countries and the connection of these grids to European DCIs and existing service grids in ICPC countries by employing 3G Bridge technology.

Moreover, building on the solid expertise of the DEGISCO partners, the project provides recommendations on best practices and defines joint roadmaps for ICPC countries and Europe. Thus, the presented work reinforces the global relevance and impact of European distributed infrastructures.

The well-established application related activities of DEGISCO help use the already more than two dozen ported applications on new connected DCI's in ICPC countries and support new applications.

The dissemination and training activities promote via various channels the interoperation between Desktop grid and Service grid infrastructures on a global scale, which leads to more awareness in the general public of computational science and distributed computing co-funded by the EC. As a result more citizens, students, and companies are expected to donate resources to scientific purposes.

CONCLUSION

The presented efforts of the DEGISCO project show the feasible ways and best practices for exploitation of the nationally and internationally integrated service/desktop/volunteer DCI particularly from the application developers' and users' point of view. In order to broaden the existing user communities, which are particularly from various areas of bioscience and physics, several steps have been performed; best practices and enhanced support services are provided by the project, and the application developers and users can join and benefit from the International Desktop Grid Federation.

In the case of the DEGISCO project the role of ported applications is extremely crucial and two-folded; new generic applications with high social impact would attract both volunteers (more resource providers for the infrastructure) and scientists (more users with more research results). Therefore, the key element of the exploitation plan in DEGISCO is to reach and keep these potential communities.

www.degisco.eu
http://desktopgridfederation.org

Robert Lovas rlovas@sztaki.hu MTA SZTAKI

Multidisciplinary approach for computing in Emilia Romagna (Italy)

OVERVIEW

A multidisciplinary context has been set up in order to build a common distributed computing infrastructure and support the applications of each partner. Comput-ER (Computing in Emilia Romagna) is based mainly on commodity farms but new GPU-based farms for parallel applications are under testing. Comput-ER resources are accessed through the gLite middleware. These resources are either based on real hardware systems, or are dynamically provisioned via virtual machines using the Worker Nodes on Demand Service (WNoDeS) grid/cloud virtualisation system.

DESCRIPTION

Experience has been gained with applications in the fields of genome annotation, mid-ocean ridge processes, ensemble methods for ocean forecasting and molecular systems computation with clusters of grid-accessible GPU-enabled systems. These applications are very CPU demanding and data-driven thus using sequential computation requires months of CPU time. The peculiar features offered by a grid environment, such as a large number of storage and computation resources and advanced high level job sub mission services, allow to reduce consistently the amount of computation time spreading the jobs on different resource sites.

The key solution is based on their execution as loose parallel applications, splitting the input data in several parts and using the parametric job submission feature of the gLite WMS to optimise the management of the computational tasks.

IMPACT

In particular in the field of genome annotation we instrumented an application for extended and robust protein sequence annotation over conservative non-hierarchical clusters based on the Bologna Annotation Resource v 3.0. Tests have been done with 150 genomes splitting the data in 8,000 pieces and implementing a loose parallel computation over the Grid. The above computation has been done on 200 nodes and took two weeks instead of months if running in a single computer.

Concerning mid-ocean ridge modelling, we simulate mantle flow dynamics to study mantle rock serpentinisation, a process central to a variety of chemical exchanges between solid earth, hydrosphere and biosphere. Sub-lithospheric mantle flow and mantle thermal structure are obtained from the solution of Stokes and heat equations by semi-analytical-pseudo-spectral and finite difference techniques, respectively. The grid infrastructure allowed submitting in a single instance n x m jobs, achieving n x m solutions at approximately the time of one single run.

Ensemble method technique has been used to quantify the forecast uncertainty in short-term ocean forecasting systems. In this study, we explore the short-term ensemble forecast variance generated by perturbing the initial conditions. Grid allowed us to perform several ensemble forecast experiments with 1,000 members: they are completed within five hours of wall clock time after their submission, and the ensemble variance peaks at the meso-scale.

Molecular dynamics simulations were performed using NAMD v2.7, a parallel molecular dynamics code, to study the trans-membrane lipid translocation processes. These processes are the basis of important properties and functions of cell membranes. In this study, the dynamics of a lipid bilayer, under the effect of two electric field intensities is investigated. Up to a ten-fold decrease in total computation time has been achieved by using a system with three GPU instead of traditional CPUs .

CONCLUSION

This paper describes how the above applications have been modified in order to be executed in a loose parallel way and shows the advantages obtained by using most of the available computing resources of a distributed infrastructure and the benefit of the knowledge-sharing from a multidisciplinary scientific computing environment.

Paolo Veronesi paolo.veronesi@cnaf.infn.it INFN-CNAF (National Institute of Nuclear Physics – National Computer Science Center), Bologna Italy • **Marco Bencivenni** marco.bencivenni@cnaf.infn.it INFN-CNAF (National Institute of Nuclear Physics – National Computer Science Center), Bologna Italy • **Daniele Cesini** daniele.cesini@cnaf.infn.it INFN-CNAF (National Institute of Nuclear Physics – National Computer Science Center), Bologna Italy • **Davide Salomoni** davide.salomoni@cnaf.infn.it INFN-CNAF (National Institute of Nuclear Physics – National Computer Science Center), Bologna Italy • **Antonia Ghiselli** antonia.ghiselli@cnaf.infn.it INFN-CNAF (National Institute of Nuclear Physics – National Computer Science Center), Bologna Italy • **Damiano Piovesan** piovesan@biocomp.unibo.it Biocomputing Group, University of Bologna • **Pier Luigi Martelli** gigi@biocomp.unibo.it Biocomputing Group, University of Bologna • **Piero Fariselli** piero@biocomp.unibo.it Biocomputing Group, University of Bologna • **Ivan Rossi** ivan@biodec.com Biocomputing Group, University of Bologna • **Daniele Guerzoni** guerzoni@biocomp.unibo.it Biocomputing Group, University of Bologna • **Rita Casadio** casadio@biocomp.unibo.it Biocomputing Group, University of Bologna • **Giovanni Bortoluzzi** giovanni.bortoluzzi@ismar.cnr.it CNR/ISMAR (National Research Council-Institute of Marine Sciences), Bologna Italy • **Marco Ligi** marco.ligi@bo.ismar.cnr.it CNR/ISMAR (National Research Council-Institute of Marine Sciences), Bologna Italy • **Stefano Ottani** stefano.ottani@isof.cnr.it CNR/ISOF (National Research Council-Institute for the Organic Synthesis and Photoreactivity), Bologna Italy • **Alessandro Venturini** alessandro.venturini@isof.cnr.it CNR/ISOF (National Research Council-Institute for the Organic Synthesis and Photoreactivity), Bologna Italy • **Nadia Pinardi** n.pinardi@sincem.unibo.it INGV (National Institute Geophysics and Volcanology),Bologna Italy • **Enrico Scocimarro** scoccimarro@bo.ingv.it INGV (National Institute Geophysics and Volcanology), Bologna Italy • **Srdjan Dobricic** dobricic@bo.ingv.it CMCC (Euro-Mediterranean Centre for Climate Change), Bologna Italy • **Antonio Navarra** navarra@bo.ingv.it CMCC (Euro-Mediterranean Centre for Climate Change), Bologna Italy

InSilicoLab – grid environment for supporting numerical experiment in chemistry

OVERVIEW

Although the capacity and technical advancement of computational resources offered by the European computing centres is constantly improving, the progress in development of high-level tools facilitating access to those resources is still slow. In this paper, we present our approach to such a high-level, customised environment, called InSilicoLab – designed especially for chemists who want to use the grid infrastructure.

DESCRIPTION

InSilicoLab is a web portal for the scientists from the Computational Chemistry domain who want to take advantage of Grid Resources. The portal allows them to create and manage scientific experiments and their results using terms specific to their domain of science. The multi-layer architecture of InSilicoLab supports this domain-specific interaction by introducing a ‘scientific’ layer, responsible for communication with the users, while a separate layer is reserved for communication with computational and storage resources. These two layers are interconnected via a third – intermediate – layer, which processes and translates the objects and methods defined by the user to parallel grid jobs and/or calls to other services and tools required by the user. This separation of the domain- and grid-specific layers and concepts related to them makes the InSilicoLab portal more intuitive to users, enabling them to focus on their work instead of on the technical details of configuring the computation, and, thus, making their work more efficient.

IMPACT

Computational Chemistry is one of the three science domains that utilise grid resources most – along with high energy physics and life sciences. While many computing centres offer a variety of scientific packages for chemists, the use of these software suites is hindered by the lack of intuitive interfaces to them. InSilicoLab, not only provides the users from the Computational Chemistry domain with easy access to Grid by the means of job management, but, more importantly, supports scientific experiment planning and evolution. While the first functionality is already offered by grid portals, and some aspects of the latter are covered by specialised tools invented for Computational Chemistry (like WebMO [www.webmo.net] or ECCE [http://ecce.pnl.gov]), only expensive commercial User Interfaces (like Accelrys Material Studio [http://accelrys.com/products/materials-studio/index.html]) join these features. Still, these proprietary tools are mostly desktop applications – which makes them inaccessible from outside a local computer and hinders collaboration.

CONCLUSION

The InSilicoLab portal facilitates in-silico experiments performed with the help of the most popular Computational Chemistry packages – Gaussian [www.gaussian.com], GAMESS [www.msg.ameslab.gov/GAMESS] and Turbomole [www.cosmologic.de].
The tool is available to any user with a valid Grid certificate registered in Gaussian VO [http://egee.grid.cyfronet.pl/Applications/gaussian-vo] or vo.plgrid.pl (a Virtual Organisation maintained by the Polish NGI – PL-Grid [www.plgrid.pl]). It can be accessed at:

http://insilicolab.grid.cyfronet.pl

Mariusz Sterzel m.sterzel@cyfronet.pl CYFRONET • **Joanna Kocot** j.kocot@cyfronet.pl CYFRONET • **Tomasz Szepieniec** t.szepieniec@cyfronet.pl CYFRONET • **Daniel Harężlak** d.harezlak@cyfronet.pl CYFRONET • **Klemens Noga** k.noga@cyfronet.pl CYFRONET

The GRelC Project: main service, functionalities and relevant use cases

OVERVIEW

Grids encourage and promote the publication, sharing and integration of scientific data, distributed across Virtual Organisations (VO). The complexity of data management within a grid environment comes from the distribution, heterogeneity and number of data sources. In the last ten years there was a strong interest in grid-database access and management topics. Moreover tools and services able to access relational databases in grid are also strongly required in the ES, LF, A&A User Communities. Within the proposed talk we will present in detail the Grid Relational Catalog (GRelC) Project, an integrated environment for grid database management, highlighting the vision/approach, architecture, components, services and technological issues. The most relevant use cases (in particular in the earth science and environmental domains) will be described in detail.

DESCRIPTION

The key topic of this talk is the GRelC Service.

The GRelC service is a GSI/VOMS enabled web service addressing extreme performance, interoperability and security. It efficiently, securely and transparently manages databases on the grid across VOs, with regard to emerging and consolidated grid standards and specifications as well as production grid middleware. It provides a uniform access interface, in grid, both to access and integrate relational (Mysql, Postgresql, SQLite) and non-relational data sources (XML DB engines such as eXist, XIndice and libxml2 based documents).

The GRelC service provides:

- basic functionalities (query submission, grid-db management, user/VO/ACL management, etc.) to access and manage grid-databases;
- efficient delivery mechanisms leveraging streaming, chunking, prefetching, etc. to retrieve data from databases in grid providing high level of performance (in terms of query response time, number of concurrent accesses, etc.);
- additional functionalities such as asynchronous queries,
- a data Grid Portal (GRelC Portal) to ease the access, management and integration of grid-databases, as well as user/VO/ACL management, etc.

The GRelC middleware has been included into the EGEE RESPECT Programme (Recommended External Software Packages for EGEE CommuniTies) since it works well in concert with the EGEE gLite software by expanding the functionality of the grid infrastructure (w.r.t. database management in grid). The GRelC service is currently adopted as grid metadata management service in the Climate-G testbed to enable geographical data sharing, search and discovery activities. Moreover it is currently used at the Euro-Mediterranean Centre for Climate Change to manage climate metadata across the Italian CMCC data grid infrastructure through the CMCC Data Distribution Centre portal.

IMPACT

A grid-database service is fundamental for distributed data-oriented infrastructures, production grids, grid testbeds, since it enables the management of crucial information. Some examples concern biological sequences in the bioinformatics domain, spatial metadata information in the earth science and environmental domains, patient-related information into a Health Information System (HIS), astrophysics databases, etc. The talk presents the GRelC service highlighting also how it provides cross-VO capabilities. Due to the service nature, architecture and functionalities common use cases can be defined across different disciplines. This helps to identify common requirements, formalise exploitation patterns paving the way towards sustainability.

CONCLUSION

The GRelC service is a grid-database management service that can be exploited both at VO and site level. This service is currently successfully deployed and positively evaluated by end-users in the Earth Science and Environmental contexts (e.g. CMCC and Climate-G). Moreover it is also available for tutorial purposes in the GILDA t-Infrastructure.

www.grelc.unile.it

Sandro FIORE sandro.fiore@unisalento.it University of Salento • **Giovanni ALOISIO** giovanni.aloisio@unisalento.it University of Salento

Data Management

Data management is an area of constant evolution: both experiment and user activity always increase and oblige Virtual Organisations to ensure the future scalability of their systems by automating manual operations, optimising the usage of available resources and adapting to evolving middleware, more performant infrastructure and new technologies.

This session will present work from the Data Management projects of ATLAS and CMS, two of the LHC experiments distributing multi-Petabyte data volumes across affiliated computing sites in the Worldwide LHC Computing Grid and EGI sites.

The session will give a general overview of ATLAS's Distributed Data Management and the architecture and operational strategies that have been decided in order to provide a fault-tolerant, highly intelligent system. The presentation will explain how different independent subcomponents are used to fully profit from the deployed resources and make the best usage of space at the sites in an automated fashion.

Next, the presenters will describe the effort in both experiments dedicated to improving the data transfers and monitoring the throughput rates between sites. Both collaborations base their file transfers on common underlying gLite middleware and pursue the optimal usage of their network and evolution of the distribution topologies in their respective Computing Models. The tendency is to move to a less hierarchic data transfer model where all sites can communicate with each other. An insight into the instrumented tools and steps followed to ensure the commissioning of the links will be given.

Jamie Shiers Head of the Experiment Support Group, IT Department, CERN

An insight into the ATLAS Distributed Data Management

OVERVIEW

ATLAS, one of the four LHC experiments, fully relies on the usage of grid computing for offline data distribution, processing and analysis. The ATLAS Distributed Data Management is the project built on top of the WLCG middleware and is responsible for the replication, access and bookkeeping of the multi-PB ATLAS data across more than 100 distributed grid sites. It enforces data management policies decided on by the collaboration and defined in the ATLAS computing model. It has been in production since 2004 and to date is considered one of the largest open data management environments and an example of a global multi-grid hybrid system. This contribution will give an overview of the architecture, operational and deployment strategies of this highly automated system, as well as details about different subsystems and monitoring solutions that could be of interest for other communities.

DESCRIPTION

ATLAS Distributed Data Management is the system that manages the experiment's detector, simulated and user data while enforcing the policies defined in the Computing Model. It provides functionality for data placement, deletion, bookkeeping and access on a hierarchic grid model composed of around 100 sites with heterogeneous storage technologies, services and protocols.

The system is currently managing 50 PB of data corresponding to almost 200 million files and is being used to achieve aggregated throughput rates far beyond the initial requirement of 2 GB/s, having reached a throughput peak of over 10 GB/s. To ensure further scalability, the core of the system has been designed as a set of independent agents which work around a global database, the Central Catalogues. This architecture allows distributed, fault-tolerant services that interact with the grid resources and checkpoint centrally the status of the requests.

Given the volume of files managed, it is critical to reduce the load on operations by providing compact monitoring information, automatically curing inconsistencies and automate systems as far as possible.

The talk will give an overview of the general architecture and furthermore explain in detail several independent subcomponents targeted towards resource optimisation, such as storage space accounting, the automatic site cleaning agent and the centralised site exclusion service and how these can be extended to interact between each other. The ideas and concepts presented will provide inspiration for any Virtual Organisation that is currently planning to move their data to the grid or working on improvements to their usage of grid, network and storage resources.

IMPACT

Given the scale of the project, ATLAS Distributed Data Management can be considered an important example of distributed data management on the grid. Although the implementation is experiment specific, it is based on common underlying middleware and is composed of independent subcomponents, auxiliary systems and monitoring solutions that can be of inspiration and re-implemented to fit the needs of any grid community.

Relatively simple subcomponents such as the storage space accounting, which was initially implemented to account the used, free and total storage space of grid storage endpoints was extended to allow the breakdown of the used space by different metadata, have been claimed to be needed by different communities inside and outside the WLCG community. The communication of the storage space accounting with a centralised site exclusion service or an automatic site cleaning agent are the next steps, in order to avoid replicating data to full sites or to trigger the cleaning of full sites respectively. In the case of ATLAS, these simple tools have helped to significantly reduce manual operations and improve day to day operations.

CONCLUSION

Data management is an area of constant expansion: both experiment and user activity keep increasing. The ATLAS DDM team oblige to ensure the future scalability of the system by: automating manual operations, optimising the usage of available resources and adapting to evolving middleware, higher performant infrastructure and new technologies.

To mention one example, DDM is currently replacing its messaging infrastructure, based on open protocols such as http, by more recent, asynchronous message queuing technologies, which should enhance the independence of subcomponents and therefore the overall fault-tolerance.

Another area of development, given the evolution of the network throughput, is the development of the hierarchic data distribution policies towards a more dynamic, relaxed Computing Model where gradually regional boundaries are relieved.

http://bourricot.cern.ch/dq2/accounting/global_view/30

Fernando Harald Barreiro Megino fernando.harald.barreiro.megino@cern.ch CERN • **Simone Campana** simone.campana@cern.ch CERN • **Vincent Garonne** vincent.garonne@cern.ch CERN • **Andrii Tykhonov** andrii.tykhonov@ijs.si Jozef Stefan Institut

Improving CMS data transfers among its distributed computing facilities

OVERVIEW

CMS computing needs reliable, stable and fast connections among multi-tiered computing infrastructures. PhEDEx provides a data management layer composed of a series of collaborating agents, which manage data replication at each distributed site. It uses the File Transfer Services (FTS), a low level data movement service responsible for moving sets of files from one site to another, while allowing participating sites to control the network resource usage. FTS servers are provided by Tier 0 and Tier 1 centres, and need to be setup according to the grid site's policies, including all the virtual organisations making use of the grid resources at the site, and properly dimensioned to satisfy all the requirements for them. Managing the service efficiently requires knowledge of CMS's needs, of all the transfer workflows that need to be handled, along with the potential benefits and conflicts to other virtual organisations (VO) using the same FTS transfer managers.

DESCRIPTION

This contribution deals with a complete revision of all FTS servers used by CMS, customising the topologies and improving their set up in order to keep CMS data transfers to the desired levels in a reliable and robust way.

We use the FTS Monitor, a web-based monitoring system developed at the CC-IN2P3 Tier 1, which provides a graphical view of the FTS activity. This service retrieves data directly from the FTS back-end database to generate summary statistics and to provide detailed reports about transfer activities. The FTS Monitor's web pages display channel configuration, statistics about transfers in the last 14 days on each channel, and detailed information on all jobs submitted in the last 24 hours, including the status and throughput of each individual transfer.

Each transfer detail is published in machine-readable XML format, which is parsed daily. The FTS Monitors each CMS Tier 1 to keep the history of all transfers and important values, such as transfer rates per file and per stream, SRM response times, FTS channel congestions, etc. A wealth of information that is extremely useful to spot issues and debug problems.

IMPACT

This global study is collecting statistics on each individual transfer for the whole CMS distributed fabric. This has never been done before. In fact, we do not know of any other VO that has conducted a study like this. The impact is the improvement of data transfers for CMS by either solving FTS channel bottlenecks or identifying configuration problems at the sites, or even on individual transfer links.

As an example, the study has already been useful to propose a policy to set up the FTS at PIC Tier 1 for PIC to Tier 2 transfers. The impact of the change was the increase by a factor 2 in the overall PIC to Tier 2 transfers, as well as a clear improvement on the data transfer qualities, by means of implementing dedicated FTS cloud channels for low and high throughput transfer connections, based on results collected by this tool.

By the time of the conference, these results will be available on a web page containing all the relevant information and plots, at the disposal of the operation teams as well as sites administrators.

CONCLUSION

With the help of such statistics the central team of distributed data transfers of CMS is running a campaign to spot issues and highlight them to the sites. Regular usage in operations will, of course, give more feedback on which are the most relevant statistics to gather. Some new statistics are already planned to be included. At the moment FTS Monitor Parser is gathering data only for the CMS VO, but in the future it can be opened to other VOs. In fact, we would like other VOs to be engaged with us in a global WLCG data transfer study, and we are collaborating with FTS Monitor developers, so more information can be published.

Josep Molina jflix@pic.es PIC/CIEMAT

Evolution of the ATLAS data placement model

OVERVIEW

The ATLAS experiment at the LHC is fully relying on the usage of grid computing for its offline data placement, processing and analysis. For data placement the ATLAS computing model defines a set of policies, which establish a hierarchical tier organisation according to the network topology, which was laid out for data distribution. However, since the original creation of the computing model, network capabilities have significantly increased and it is convenient to gradually start relaxing some of the imposed boundaries. This talk will focus on the work that is being carried out in the ATLAS Distributed Data Management project in order to evaluate more dynamic constraints and provide the necessary framework for commissioning network links.

DESCRIPTION

The ATLAS Distributed Data Management is the project built on top of the WLCG middleware and is responsible for the replication, access and bookkeeping of the multi-Petabyte ATLAS data across the grid while enforcing the policies defined in the ATLAS Computing Model. Following this model ATLAS sites are grouped into ten clouds by geographical and organisational reasons. Each cloud is formed by one Tier 1 that must provide a high level of service and is responsible for data storage and reprocessing and several Tier 2s and Tier 3s, which are used for analysis and Monte Carlo production and depend directly on the Tier 1.

Network connections between the Tier 0 (CERN) and the Tier 1s are guaranteed by the Optical Private Network and inside a cloud are generally performant, while inter-cloud links are less guaranteed. Thus, Tier 1s usually act as the access point to get data in and out of the cloud and direct cross-cloud communication between Tier 2s is generally avoided. However, since networking capabilities have significantly evolved, the computing model is moving towards more dynamic data management policies and the cloud boundaries are gradually being reduced in a controlled fashion.

Consequently, the data distribution framework in the ATLAS Distributed Data Management project has been instrumented to measure the durations of gLite File Transfer Service (FTS) transfers between sites and store them in an Oracle database. The transfer statistics will be used as feedback to optimise the source selection and choose between multi-hop transfers through the Tier 1s or direct cross-cloud transfers. The statistics are also visualised in a dynamic web page in order to monitor the throughput performance of the network links. In parallel, an ad-hoc load generator will trigger transfers on the complete mesh of ATLAS sites and will provide the information for a first attempt of link commissioning.

IMPACT

New data brokering models such as PanDA Dynamic Data Placement have been recently introduced by ATLAS. The idea behind this model is to replicate datasets from any Tier 1 to any Tier 2 only after they have exceeded a popularity threshold, thereby eliminating the replication of unpopular datasets. This work is a first attempt in the ATLAS Distributed Data Management project to optimise these cross-cloud transfers, improve the network usage and provide the necessary statistics needed for link commissioning activity with the final goal of reducing cloud boundaries.

A similar initiative exists in the CMS experiment, another of the LHC experiments. Their approach is somewhat different, as instead of measuring the transfer events, the statistics are collected by parsing html files retrieved from each one of the gLite File Transfer Service servers. Also the statistics are not fed back into the system in order to optimise transfers, but are only displayed for link commissioning.

CONCLUSION

The collected transfer statistics are known directly by the FTS servers, but are not made available through an programmatic interface. Having access to them is not an ATLAS particular interest, but can be of use to any grid-based community. This contribution wishes to open the discussion between the different user communities with the ultimate goal of moving in the future towards a common, central solution.

http://bourricot.cern.ch/dq2/ftsmon

Fernando Harald Barreiro Megino fernando.harald.barreiro.megino@cern.ch CERN • **Simone Campana** simone.campana@cern.ch CERN • **Andrii Tykhonov** andrii.tykhonov@ijs.si Jozef Stefan Institut

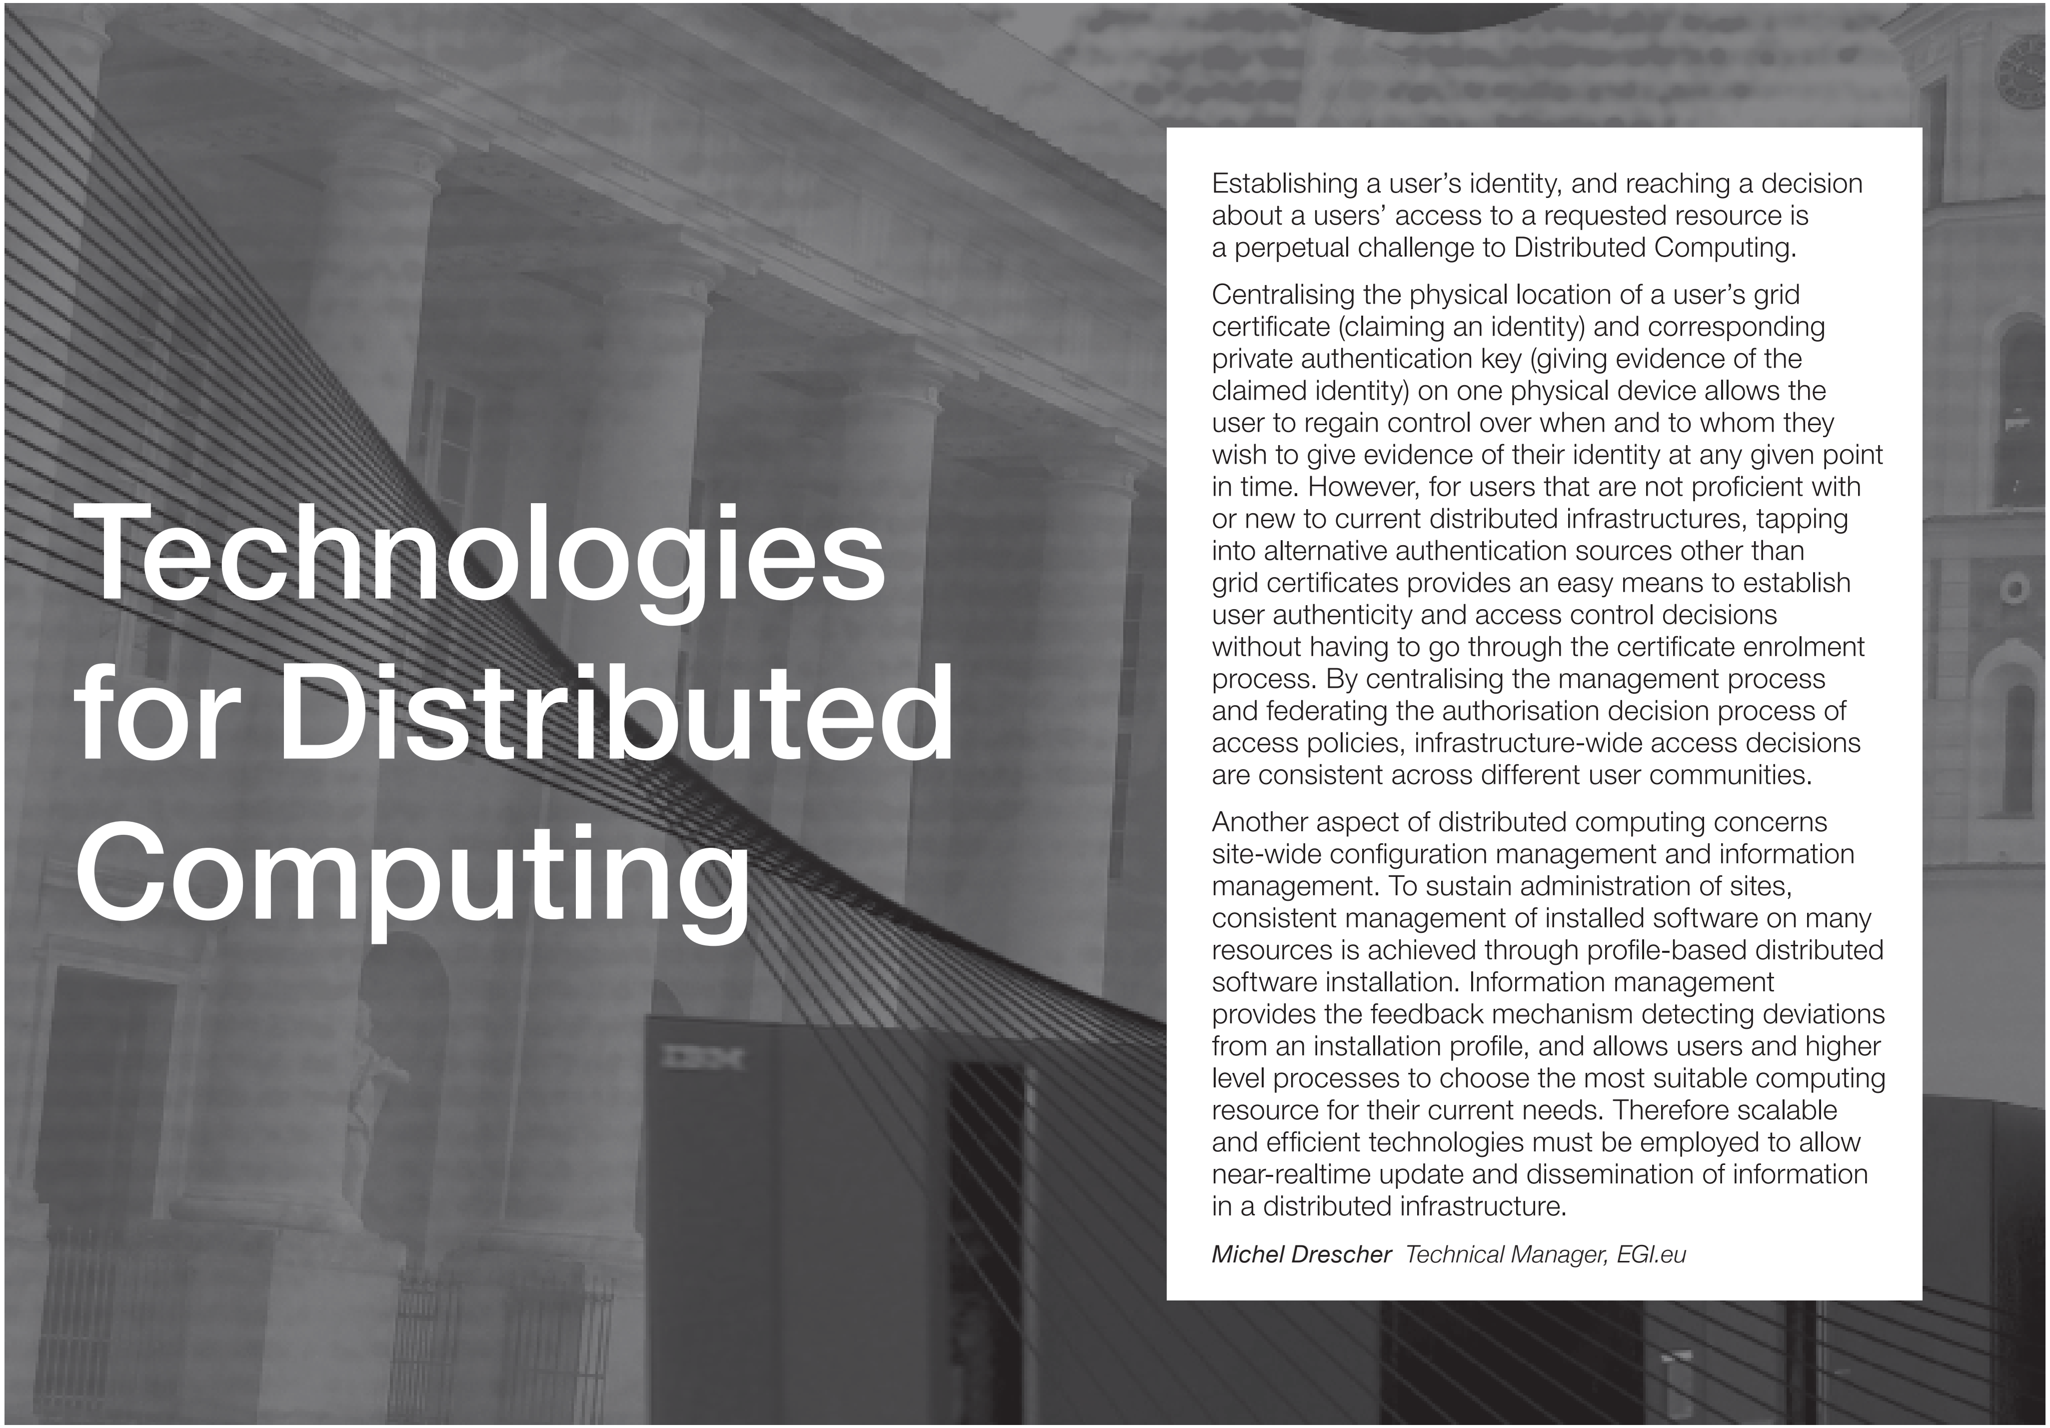

Technologies for Distributed Computing

Establishing a user's identity, and reaching a decision about a users' access to a requested resource is a perpetual challenge to Distributed Computing.

Centralising the physical location of a user's grid certificate (claiming an identity) and corresponding private authentication key (giving evidence of the claimed identity) on one physical device allows the user to regain control over when and to whom they wish to give evidence of their identity at any given point in time. However, for users that are not proficient with or new to current distributed infrastructures, tapping into alternative authentication sources other than grid certificates provides an easy means to establish user authenticity and access control decisions without having to go through the certificate enrolment process. By centralising the management process and federating the authorisation decision process of access policies, infrastructure-wide access decisions are consistent across different user communities.

Another aspect of distributed computing concerns site-wide configuration management and information management. To sustain administration of sites, consistent management of installed software on many resources is achieved through profile-based distributed software installation. Information management provides the feedback mechanism detecting deviations from an installation profile, and allows users and higher level processes to choose the most suitable computing resource for their current needs. Therefore scalable and efficient technologies must be employed to allow near-realtime update and dissemination of information in a distributed infrastructure.

Michel Drescher Technical Manager, EGI.eu

A new 'lightweight' crypto library for supporting an advanced grid authentication process with smart cards

OVERVIEW

Many of the existing grid middleware, and in particular gLite, rely only on the adoption of a Public Key Infrastructure (PKI) of digital certificates for user authentication, and these credentials must be present on each User Interface (UI) that is used by users to access the computational and storage resources. Distributing the certificate's private key on multiple locations is considered a security breach, as the certificate may be subjected to possible fraudulent use by non-authorised people logged into the UI. Furthermore, there is lack of support for other authentication mechanisms such as smart cards even if this hardware is available and can help in keeping these certificates safe and avoid any fraudulent use. The public part of an X.509 certificate stored on this hardware can usually be accessed by users, applications, portals and/or science gateways but the corresponding private key can never be copied off the smart card.

DESCRIPTION

In this contribution we describe our work in the design and implementation of a new grid authentication method based on the use of digital certificates stored on smart cards. The solution we propose extends the native Sun PKCS#11 cryptographic APIs with the Bouncy Castle and the cog-jGlobus APIs library in order to implement a new 'lightweight' crypto utility which can be used by users to access the digital certificates stored on a smart card and generate a VOMS proxy. The Bouncy Castle APIs were used to generate a X.509 version 3 certificate by reading the credentials available in the smart card, while the support with the cog-jGlobus was introduced to establish a secure connection with the VOMS server and add the Attributes Certificate to the original grid proxy. In this first implementation, the library runs with the Aladdin e-Token PRO 32K directly plugged into a remote 64-bit UI based on Scientific Linux 5 where the Aladdin's e-Token PKI Client software (pkiclient-full-4.55-34) was previously installed. This software enables e-Token USB operations and the implementation of e-Token PKI-based solutions. The library has been successfully tested with both personal user's and robot certificates and used to generate grid proxies in the new e-Collaboration environment based on Liferay and GENIUS/EnginFrame technologies. It is of course not restricted to any grid portals and/or science gateways. Concerning the credentials stored into the smart card, we used robot certificates issued by the INFN Certification Authority (CA) because they are the type of certificates that best match the use case of these credentials in our environment. Nevertheless, nothing prevents using personal certificates instead if their planned usage is consistent with the restrictions the CAs put on it.

IMPACT

The work carried out and reported in this contribution is particularly relevant for several user communities, applications, grid portals and/or Science Gateways developers. The library allows using the credentials stored into smart cards for generating VOMS-compliant proxies thus enhancing the Java technology to deal with private and public keys. The benefits introduced in this work are far-reaching. The new crypto library can be used, for instance, to help developers to design and develop science gateways for several scientific communities and provide, especially for non-expert users, a transparent and easy access to e-Infrastructures. Last but not least, the introduction of smart cards for storing digital credentials can massively improve the security of grid infrastructures since they protect sensitive information from malicious applications. As we mentioned before, the crypto library has been successfully tested on Aladdin e-Token smart cards only. No compatibility tests have been performed on smart cards of other vendors, but these can be done on request.

CONCLUSION

The Java SE platform provides developers with a large set of security APIs, algorithms, tools and protocols. Among them, we can point out the native Sun PKCS#11 cryptographic tokens which have been used in this work together with the Bouncy Castle and the cog-jGlobus Java APIs to implement a new module for the gLite grid middleware which allows us to sign the proxy certificate using the digital credentials stored in a smart card. The open source solution described in this work can be used by users, applications, grid portals and/or science gateways developers to generate VOMS proxies using Java APIs starting from the credentials stored in e-Token smart cards. The new library provides an asset in raising grid awareness and encourages broader use by a wider number of potential users.

Giuseppe La Rocca giuseppe.larocca@ct.infn.it INFN Catania

Federated access to grids

OVERVIEW

Grids have established their own PKI-based Authentication & Authorisation Infrastructure (AAI), which is not linked to any existing identity-management systems. The result is not always well received by users. Even if the situation is changing and grids are aiming at utilising identity federations, there still remain crucial aspects that prevent the smooth usage of federated identities in grids.

One basic issue is the way of obtaining a credential to routinely access grids, which should be simple to use, yet secure enough. Utilisation of federated identities is highly desired by the users and several solutions have been adopted recently, e.g. with the Terena TCS certification authority. Another useful approach is transparent credential conversion, which allows obtaining grid credentials without the users' explicit intervention. We will demonstrate how the federated Moonshot infrastructure can provide such a conversion, opening grids to a larger user base.

DESCRIPTION

Moonshot is a set of technologies for providing federated access to applications. At a technical level, federation decouples management of credentials within an organisation from authentication proofs between organisations. Many existing technologies such as Security Assertion Markup Language (SAML), RADIUS, and Diameter support federation. However these existing technologies are focused on a single application domain. SAML provides federation for the web and web services, RADIUS and Diameter provide federation for network access. Moonshot integrates RADIUS federation into most application protocols, whose users can then leverage from federated access. At the same time, SAML is fully supported and provides rich attributes to describe federated subjects. The Moonshot infrastructure builds upon components that are well-known and understood, and have been in use for a long time. Unlike other federation middlewares (e.g. Shibboleth), Moonshot is not tied with the web environment and can be utilised by non-web applications. The Moonshot architecture is based on open standards and the Moonshot community actively contributes to the ABFAB working group of the IETF.

Moonshot is a good fit for the grid because many grid components, particularly grid middleware components, do not use web technologies. We are therefore seeking ways to find out how grids can benefit from the Moonshot infrastructure. In this contribution, we will demonstrate the design and current achievements focusing on users running jobs on the grids using a Moonshot-generated id. As the result, the users will be able to use their federated identity to launch and control jobs. Finally, we will also demonstrate the flexibility of Moonshot by demonstrating how a particular grid service can be plugged into the Moonshot infrastructure so that users could use the service without having to obtain a certificate first. For this purpose, we will show how the gLite Logging & Bookkeeping service can be adapted to natively support Moonshot.

IMPACT

Current identity federations are limited in the number and type of applications supported. The most prevalent federations are closely tied with the web (e.g., Shibboleth) or network access (Eduroam). However, none of these federations are suitable for facilitating access to the grid. The Moonshot architecture, on the other hand, is open for any type of services regardless the service protocols and/or interfaces they use.

Having Moonshot supported by the grid environment would increase the number of potential users and also make their management easier. In particular, we propose an integration of the Moonshot and PKI so that users can easily obtain credentials needed to join the grid, based on their federated identities. Identity federations also maintain sets of attributes assigned to the users, which may also be valuable for the grid, too. For instance, the resource providers and/or VO managers could specify access control rules based on these attributes (e.g. users' affiliation). The big advantage of the federation model is that the attributes are kept current and therefore can be relied upon by their consumers.

CONCLUSION

A federated identity is a very convenient means for authentication of users and passing additional information about them. Current identity federations are not generally enough to fit the grid environment, though. The aim of this contribution is to introduce the Moonshot project, which fills in the gap between existing identity federations and applications that cannot be integrated with them due to completely different underlying technologies.

Moonshot can make the access to the grid easier for the users, since they could mostly use the credentials they use also for other services. Also the resource providers could benefit from this integration and rely on the users' attributes propagated from their identity providers.

While the Moonshot project is quite young and the work is still in progress, we plan to demonstrate its benefits for the grid community to provide a notion of its potential.

www.project-moonshot.org

Daniel Kouril kouril@ics.muni.cz CESNET • **Sam Hartman** hartmans@painless-security.com Painless Security • **Josh Howlett** josh.howlett@ja.net JANET • **Jens Jensen** jens.jensen@stfc.ac.uk STFC Rutherford Appleton Laboratory • **Michal Prochazka** michalp@ics.muni.cz CESNET

Argus, the EMI Authorisation Service

OVERVIEW

Authorisation across EMI middleware stacks is currently not homogeneous, and components often have their own mechanisms of handling it. To address this inconsistency and to unify the authorisation, the Argus Authorisation Service was chosen as the EMI solution. It renders consistent authorisation decisions for distributed services (e.g., user interfaces, portals, computing elements, storage elements). The service is based on the XACML standard, and uses authorisation policies to determine if a user is allowed or denied to perform a certain action on a particular service. The presentation introduces Argus to site and system administrators. Different deployment scenarios are presented, as well as tools that Argus provides to manage authorisation policies. The roadmap for integration of Argus in other EMI components is presented together with an update on the status of common authorisation profiles standardised in EMI.

DESCRIPTION

The Argus Authorisation Service is composed of three main components. The Policy Administration Point (PAP) provides the tools to author authorisation policies, organise them in the local repository and configure policy distribution among remote PAPs. The Policy Decision Point (PDP) implements the authorisation engine, and is responsible for the evaluation of the authorisation requests against the XACML policies retrieved from the PAP. The Policy Enforcement Point Server (PEP Server) ensures the integrity and consistency of the authorisation requests received from the PEP clients. Lightweight PEP client libraries are also provided to ease the integration and interoperability with other EMI services or components. An infrastructure-wide user banning mechanism is a major security feature currently missing in existing deployments despite its importance. This feature can be effectively implemented leveraging the modularity of Argus that allows linking together PAP services deployed at different levels of the infrastructure. Leveraging the Argus policy distribution mechanism, banning lists can be defined at various levels (European, NGI, institutional etc.) and sites can import such lists so that local authorisation mechanisms are aware of malicious users. In order to preserve site autonomy, Argus allows to override policies imported from remote PAPs with local ones, leaving the control of local resources authorisation in the site administrators' hands. Authoring XACML policies in XACML itself is not straightforward: XML per se is perceived by many users as difficult to read, and editing can be prone to error. Argus's solution to this issue is the Simplified Policy Language (SPL), which facilitates the authoring of policies. Site administrators can write policies in the SPL and import them in the PAP. Policies are then transparently converted in XACML and stored in the local repository. Command line tools are also provided to add and manipulate commonly used authorisation policies.

IMPACT

The Argus Authorisation Service is based on XACML; using a standard for authorisation eases the integration with other grid components, and provides interoperability across EMI middleware stacks. Interfacing Argus into Compute Elements as well as Storage Elements will enable maintenance of the authorisation policies in one service for an entire site. The chaining of different PAPs at the international as well as national level will lead to effective global banning of malicious users during incident handling.

A special focus on the site and system administrator will be given in this presentation that will allow him/her to better understand the advantages of deploying Argus at his/her site and the roadmap of integration into other services over the lifetime of the EMI project. Argus service deployment and concrete examples of authorisation policies will be discussed.

CONCLUSION

The Argus Authorisation Service is the solution to render and enforce consistent authorisation decisions across the different EMI middleware stacks. Using a standard XACML based authorisation service facilitates integration and interoperability between the different EMI services and eases the authorisation management at individual sites.

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

Valery Tschopp valery.tschopp@switch.ch SWITCH • **Andrea Ceccanti** andrea.ceccanti@cnaf.infn.it INFN-CNAF

Site integral management with Puppet

OVERVIEW

After the first LHC phase, experiments and sites focused efforts to deliver the demanding WLGC metrics. Now, sites are facing a new challenge to consolidate and automate the computing services to reach a functioning steady-state. The installation and post-installation mechanics are two of the most important points for the computing centres, in order to avoid initial installation and configuration problems. At Port d'Informació Científica (PIC), the adopted solution is to steer all post-install and dynamic post-configuration using Puppet.

Puppet uses a master entity to easily define profiles that get propagated around the cluster, hence fulfilling the necessities of post-install configurations, after the raw installation, and ensuring the persistence of the profile and the defined services once they have been completely installed.

DESCRIPTION

Managing hundreds of nodes is one of the challenges of any computing centre, and PIC is no exception. There are many tools able to cope with the problem; ranging from the very simple to the very complicated. We did a study of some of the available tools (quattor, cfengine, puppet). Our requirements include the ability to do incremental configuration (no need to bootstrap the service to make it manageable by the tool), simplicity in the description language for the configurations and in the system itself, ease of extension of the properties/capabilities of the system, a rich community for assistance and development, and open-source software.

We found in Puppet the balanced trade-off between simplicity and flexibility that was the most fitting for our requirements. The Puppet approach to system management is simplistic, non-intrusive and incremental; Puppet doesn't try to control every aspect of the configuration but only the ones you are interested in. Our sysadmins were able to build complex configurations in a short time due to the easy learning curve.

IMPACT

Puppet allows you to administrate a whole site from a central service, easing a lot of potential reconfiguration issues or speeding up disaster recovery procedures.

Having a centralised management of grid services profiles results in a very easy scaling method for adding resources on the fly: push the install button and let Puppet do all the rest of the work until the service is ready for production.

Potentially Puppet doesn't have to be a site specific tool but a master service able to do the same work on distributed resources or sites. This advantage can be considerable in centrally managing several sites that share a common middleware infrastructure, which is in line with the current strategy of middleware implementation.

CONCLUSION

Having a steered configuration system like Puppet is a huge advantage from the administration point of view in easing the daily work of sysadmins.

Having the whole services configuration of a site under Puppet provides an invaluable confidence, ensuring a streamlined way for service deployment.

Unifying the work model is also an important feature; all services are defined using the same pattern in a common location. Therefore removing the constraint of service-responsible by having a common language of administration transversal to all sysadmins.

Xavier Espinal espinal@pic.es Port d'Informació Científica (PIC) • **Arnau Bria** arnaubria@pic.es Port d'Informació Científica (PIC)

A new scheme for efficient distribution of information in grid infrastructures

OVERVIEW

The gLite grid middleware utilises BDII as an information system in a layered approach. In our work we have developed an alternative path for distribution of the information, which enables considerable decrease in the amount of data that is exchanged. We exploit the fact that much of the information that is exchanged does not change over time by introducing a versioning scheme. In this way, most of the traffic consists of data in a compressed diff-like format.

DESCRIPTION

We implemented an alternative path for distribution of information that is currently available via the gLite BDII service. The new lightweight middleware components are based on the well-known properties of the information that is currently transferred between the various layers in production grids based on gLite, and follow a simple protocol to ensure consistency.

Since the information that a site-level information system displays changes frequently, the problem is how to distribute in an efficient way the changes between successive 'versions' towards the nodes that require them (mostly top-level BDIIs). Since the changes are usually limited to certain attributes and the appearance or disappearance of whole sections of information is related to rare events, like failing or new nodes, we use the routines from the 'diffutils' package to compute the difference between successive 'versions' of the information that a site displays.

Moreover, we compress the result, which decreases even further the amount of information that needs to be transferred. A client, which can be co-collocated with a top-level BDII, can query our service over the http protocol, supplying a hash of the information that it has at a given point. If this information is sufficiently recent, within the last 1–2 minutes, then only the gzipped difference is transferred. The client proceeds to decompress the patch and applies it to its copy of the information, obtaining the same result as when all the information is downloaded. We have implemented a testbed deployment of the new system, where we mimic the infosystems of the European Grid Infrastructure sites in real time, which enabled us to perform a series of benchmarks and to analyse the performance of the new scheme.

IMPACT

The new scheme can be deployed in combination with the current BDII solution. It reduces drastically (by orders of magnitude) the amount of network bandwidth being used, which will be important to accommodate for further increases in the number of sites in the infrastructure and the complexity of information about them. It also decreases the latency in the movement of data between the layers.

CONCLUSION

In this work we present the architecture of the system and show the quantitative results to support its advantages.

Our system design and implementation are sufficiently generic, thus allowing applications in other areas, where the distribution of information follows the same patterns.

Emanouil Atanassov emanouil@parallel.bas.bg IICT-BAS

Coarse-grained interoperability of heterogeneous grid workflow systems

OVERVIEW

Several grid workflow management systems (WfMS) have emerged in the last decade. These systems were developed by different scientific communities for various purposes, and many applications have been ported to the grid using them. Reusing workflows in different WfMS and building new experiments by utilising existing workflows as building blocks is currently not supported. WfMS typically have their own user interfaces or APIs, description languages, provenance strategies, and enactment engines, which are not standard and do not interoperate.

The European SHIWA project aims to leverage existing solutions and enable cross-workflow and inter-workflow exploitation of DCIs by applying coarse- and fine-grained strategies. The coarse-grained approach treats workflow engines as distributed black box systems, while the fine-grained approach addresses language interoperability by defining an intermediate representation. This contribution concentrates on results of the coarse-grained approach.

DESCRIPTION

The coarse-grained interoperability work inside the SHIWA project refers to the nesting of different workflow systems to achieve interoperability of execution frameworks. The solution integrates different workflow engines to a grid application repository and submitter service, called GEMLCA, and allows grid based workflow systems to access non-native workflow engines if they have to execute embedded workflows.

The current approach facilitates the embedding of Kepler, Taverna, Triana, GWES and Moteur workflows into P-GRADE workflows. GEMLCA is capable to either invoke pre-deployed workflow engines through a variety of grid middleware solutions (including GT2, GT4 and g-Lite), or to submit and execute the engines on resources of infrastructures built with these types of middleware.

The first version of the SHIWA Simulation Platform (SSP) uses the NGS P-GRADE portal and resources of the Westfocus VO of the UK NGS to execute embedded workflows. The platform is currently being tested and utilised by the SHIWA user community. The first considered workflows implement pilot applications from different areas: neuroimaging (Charité Berlin, D-Grid, GWES); DNA sequence alignment (AMC Amsterdam, DutchGrid, MOTEUR); simulation of chemical reactions (Charité Berlin, EGI, MOTEUR) and medical simulation (CNRS Creatis, EGI, MOTEUR). The existing workflows had to be adapted to increase their portability respective to credentials needed to access the work-flow description and DCI resources for enactment; flexibility and neutrality of invocation and management of input data and output results; data location and transfer protocols; description language; and implicit functionality of the workflow system; etc. These workflows are now available at the SHIWA repository and can be used to compose more complex "meta-workflows".

IMPACT

Coarse-grained workflow interoperability enables the coordinated execution of workflows created in different workflow systems, from a host workflow. The created 'meta-workflows' may span multiple heterogeneous grid infrastructures.

Workflow interoperability allows workflow sharing to support and foster the adoption of common research methodologies, improves efficiency and reliability of research by reusing these common methodologies, increases the lifetime of workflows, and reduces development time for new workflows.

The SHIWA project develops a coarse-grained interoperability solution that enables the publishing, searching and sharing of workflow engines and workflows, in a repository. Using the repository, users can select already existing workflows or workflow components and embed them into their native workflow. The developed solution supports the above functionalities without requiring users to understand the nature and technical details of the workflow system that hosts the embedded workflow. The embedded workflow is considered as a black-box, and is represented by its input and output parameters and the executing workflow engine only.

The new services developed and deployed by SHIWA allow different scientific communities to exchange applications, workflows, and data resources regardless of the workflow system in which they are used. Therefore, it enables the development of such inter-disciplinary and inter-organisational workflow applications that were not possible before.

The results of the project are directly utilised by selected user communities from the areas of medical imaging, bioinformatics and chemistry. The communities provide pilot application workflows that serve as benchmarks for the developed platform and form the basis of a larger user community that will utilise SHIWA services in the future.

CONCLUSION

The experience gained so far in this project has indicated that the reusability of work-flows can be improved by following simple guidelines to reduce dependencies on the specific environment adopted by the user community that develops and runs them. However some dependencies on the description language, workflow system, credential handling and file access protocols cannot be controlled at the workflow level, so they need to be solved at the SHIWA platform level.

The SHIWA project is currently completing the implementation of SSP including a coarse-grained interoperability mechanism. Future work consists of enacting the pilot workflows in different workflow systems using the SSP and preparing other (more complex) applications for publication on the SSP.

www.shiwa-workflow.eu

Tamas Kiss kisst@westminster.ac.uk University of Westminster • **Tamas Kukla** t.kukla@my.westminster.ac.uk University of Westminster • **Dagmar Krefting** dagmar.krefting@charite.de Charité – Universitätsmedizin Berlin • **Vladimir Korkhov** v.korkhov@amc.uva.nl Academic Medical Center of the University of Amsterdam • **Gabor Terstyanszky** g.z.terstyanszky@westminster.ac.uk University of Westminster • **Noam Weingarten** n.weingarten01@westminster.ac.uk University of Westminster • **Peter Kacsuk** kacsuk@sztaki.hu MTA SZTAKI • **Gabor Kecskemeti** kecskemeti@sztaki.hu MTA SZTAKI • **Johan Montagnat** johan@i3s.unice.fr CNRS I3S Lab, UMR 6070 • **Ketan Maheshwari** ketan@i3s.unice.fr CNRS I3S Lab, UMR 6070 • **Silvia D. Olabarriaga** s.d.olabarriaga@amc.uva.nl Academic Medical Center of the University of Amsterdam • **Tristan Glatard** glatard@creatis.insa-lyon.fr CREATIS CNRS UMR 5220

EGI.eu

Science Park 105
1098 XG Amsterdam
The Netherlands

www.egi.eu
contact@egi.eu

The EGI-InSPIRE (INFSO-RI-261323) and EMI (INFSO-RI-261611) projects are partially funded by the European Commission.

ISBN 978 90 816927 1 7