Extending WeNMR e-Infrastructure outside Europe

Marco VERLATO[1]

Istituto Nazionale di Fisica Nucleare, Sezione di Padova
Via Marzolo 8, I-35131 Padova, Italy
E-mail: `Marco.Verlato@pd.infn.it`

WeNMR (Worldwide e-Infrastructure for Nuclear Magnetic Resonance and Structural Biology) is a EU FP7 project that brings together research teams in the Structural Biology area into a Virtual Research Community at a worldwide level, focusing on bio-molecular Nuclear Magnetic Resonance (NMR) and Small Angle X-ray Scattering (SAXS). This has been achieved through the implementation of a grid based e-Infrastructure, now fully integrated into EGI, aimed at providing the user community with a platform integrating and streamlining the computational approaches necessary for NMR and SAXS data analysis and structural modelling. Nowadays more than 20% of the users come from outside Europe. Therefore, in the last year the WeNMR team has worked to extend the e-Infrastructure with new resource centres from South Africa, Latin America, North America, Mediterranean and Asia-Pacific areas, with the goal of involving their NGIs to support their local NMR/SAXS community. In particular, a program to interoperate with the OSG grid in US in collaboration with their SBGrid Virtual Organisation has started, and the first tests of WeNMR job submission to OSG have been successfully carried out in the last months. WeNMR is also involved, through its collaboration with the EU FP7 project CHAIN (Coordination & Harmonisation of Advanced e-INfrastructures) in the effort to foster interoperability among EGI and those other grid infrastructures in China and India, which are not based on the EMI/gLite middleware.

---

[1]   Speaker

1.Introduction

During the last year the WeNMR operations team has worked to integrate into the WeNMR grid a number of resource centres from South Africa, Latin America, North America, Mediterranean and Asia-Pacific areas. The integration of one centre in South Africa (University of Cape Town), three centres in Latina America (UFRJ in Brasil, ULA in Venezuela and UNLP in Argentina) and one centre in Taiwan (ASGC of Taipei) has run quite smoothly being these centres deployed with the EMI/gLite middleware. In order to monitor the proper configuration and functioning of these grid sites WeNMR has leveraged on the VO-specific instances of the EGI operational tools deployed at one of its partners' site.

In North America the situation is different. WeNMR has been in contact since its beginning with the SBGrid Consortium, which is serving the Structural Biology user community in the US. SBGrid is also a VO of the Open Science Grid (OSG), the US national distributed computing grid for data-intensive research, which is interoperable with the EGI grid. At the beginning of 2011 it has been agreed between WeNMR and SBGrid that every US user registering with SBGrid VO is given the option to register with the enmr.eu VO too, allowing him to use the WeNMR services. Moreover, OSG, SBGrid and WeNMR representatives designed a technical plan aiming at setting up a test-bed for allowing the submission of WeNMR grid jobs towards the OSG resource centres supporting the SBGrid VO.

The first proof of concept has been achieved through the use of dedicated EMI/gLite middleware services components deployed in Europe and configured to submit WeNMR jobs to OSG. A further step, planned to enable the job submission from Europe making use of the Condor based GlideinWMS system services deployed in OSG and properly configured to receive WeNMR jobs, has been successfully tested in Spring 2012.

In the context of the collaboration with CHAIN, the project is involved in the effort to foster interoperability among EGI and those other grid infrastructures in China and India, which are not based on the EMI/gLite middleware. The adoption of the proposed standard based solutions to achieve a production level interoperability would also allow a further extension of the WeNMR grid for the benefit of the user community active in those countries.

2. The WeNMR project

WeNMR is a three years project started in November 2010 and building on the previous EU FP7 e-NMR project that starting since November 2007 has brought the bio-NMR community in the grid business and established an e-Infrastructure encompassing the computing centres of the NMR laboratories of the Johann Wolfgang Goethe University of Frankfurt in Germany, the University of Utrecht in The Netherlands and the University of Florence in Italy. During the e-NMR lifetime the computing resources of these three labs became part of their respective National Grid Initiatives (NGIs) and together with the resources provided by the fourth partner, the Italian National Institute of Nuclear Physics (INFN), and by a number of other NGIs supporting the enmr.eu Virtual Organisation (VO) formed the so called e-NMR grid. At the end of the e-NMR project (October 2010) the VO was supported by a number of central

grid services, partly dedicated and partly shared with EGI, and 18 resource centres providing a total of 9000 CPU-cores and 120 TB of disk storage space, shared with other production VOs.

The follow up project WeNMR was designed to achieve the following objectives:

- Operate and further develop a user-friendly e-Science gateway for the NMR and SAXS communities
- Establish a virtual research platform for (interaction with) the user community
- Provide support to software developers, users and other e-Infrastructure projects
- Foster the adoption and use of e-Infrastructure in a wide range of flanking disciplines within the life sciences
- Operate and consolidate the e-NMR grid infrastructure and to extend it to interoperate with other worldwide grid initiatives
- Develop a model to ensure sustainability of the project

In order to shield as much as possible the end user from the grid technicalities and all middleware related issues and commands, WeNMR chose to develop mainly web portals providing "protocolized" access to the grid. Their operation is facilitated when possible by using robot certificates. These portals are listed in the Virtual Research Community (VRC) WeNMR gateway shown in Figure 1 at www.wenmr.eu. About two dozen of portals are currently available and provide services for NMR and SAXS spectroscopists. Many of these run calculations on the grid.

Figure 1: The WeNMR VRC web portal at www.wenmr.eu

The computational services available can be divided into:

- Processing (MDD NMR)
- Assignment (Auto Assign, MARS, UNIO)
- Analysis (TALOS+, AnisoFIT, MaxOcc, RPF, iCing)
- Structure calculation (ASDP, CS-ROSETTA, CYANA, UNIO, Xplor-NIH)
- Molecular Dynamics (AMBER, GROMACS)
- Modelling (3D-DART, HADDOCK)
- Tools (Format converter, SHIFTX2, Antechamber, DisMeta, PREDITOR, RCI, UPLABEL)

The application software characteristics of those parts of calculations running on the grid are mainly CPU intensive and low I/O. Limited data transfer is needed to/from the computing nodes and usually no mass storage is required. The application software is deployed from remote (i.e. without the intervention of the local site manager) on the VO dedicated shared areas of the computing nodes.

3. Extension towards other grids worldwide

At March 21st 2012 the enmr.eu VO counted #395 registered users from 41 countries, #95 of them (24%) coming from 11 non-European countries. Of these, #57 were from USA. The overall geographical user distribution is shown in Figure 2.

The WeNMR user base has been since its beginning not limited to the European region. If we look for example in Figure 3 the geographical user distribution of one of the most popular modeling application, we may notice the large amount of users actually concentrated in USA, India and China. However, out of the about #2000 users registered with the HADDOCK portal only #125 have the grid access enabling them to run the more advanced tasks empowered by the use of grid resources. This is probably due to the still high barrier for grid adoption introduced by the need of X.509 PKI personal certificates, not easy to obtain and manage specially in the Life Science community institutions of several regions. Nevertheless it is quite clear that there is a strong potential user base for the WeNMR grid in all the continents which could benefit of the support of National Grid Initiatives like OSG in USA, GARUDA in India, CNGrid in China, SAGrid in South Africa, and of several EU FP7 international grid projects like CHAIN (Co-ordination and Harmonisation of Advanced e-Infrastructures), GISELA (Grid Initiatives for e-Science virtual communities in Europe and Latin America), EU-IndiaGrid (Sustainable e-Infrastructures across Europe and India) and EUMEDGRID-Support (Sustainability of e-Infrastructures across the Mediterranean).

Thus, during its first year of activity WeNMR has been represented in several workshops and conferences organized by the above projects where discussing about interoperability among different e-Infrastructures. Among the activities planned for the second project year in fact there was one dedicated to the worldwide grid extension outside Europe and divided in two tasks:

- Extension towards EMI/UMD based e-Infrastructures, especially the ones of South Africa, Latin America and Asia-Pacific that are all based on the EMI/gLite middleware

- Extension towards non-EMI/UMD based e-Infrastructures, as OSG in USA based on VDT (Globus, Condor, etc.), GARUDA in India based on Globus and CNGrid in China based on Grid Operating System (GOS) middleware.
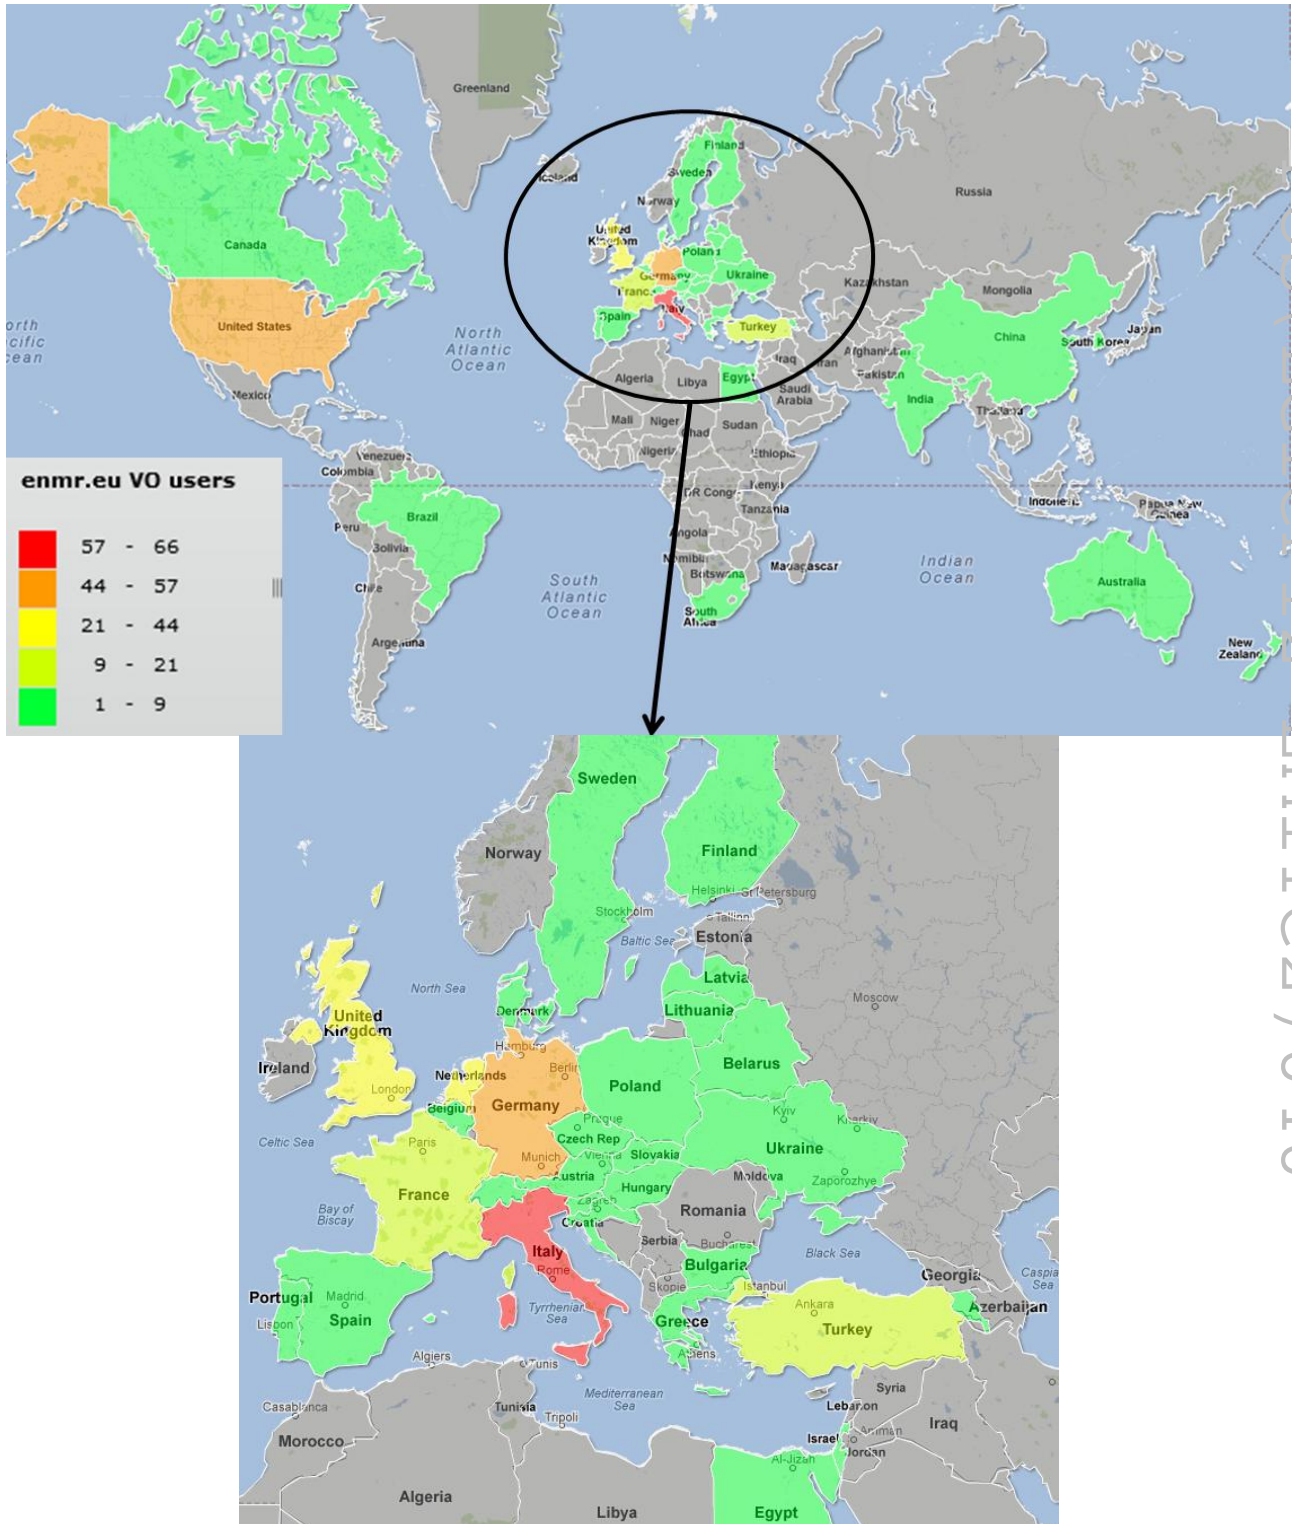

Figure 2: Geographical view of WeNMR user distribution (registered with enmr.eu VO)
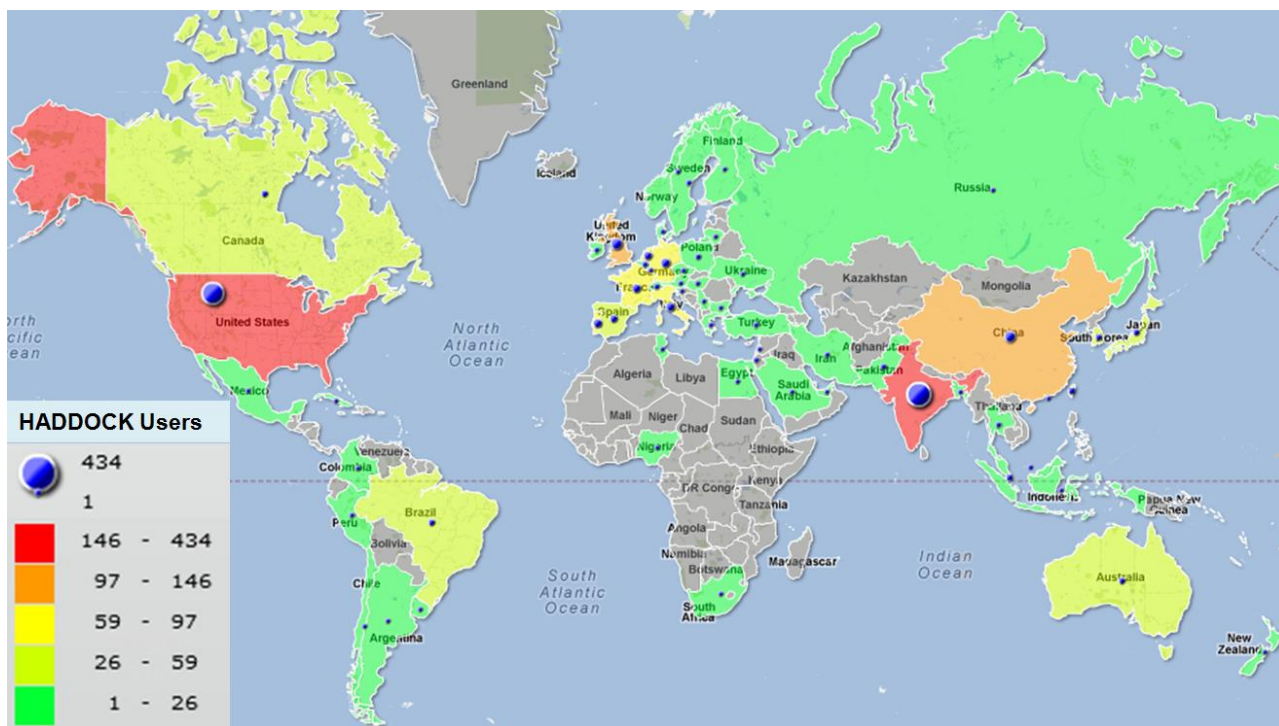
Figure 3: Geographical view of HADDOCK user distribution

The activities carried out within the above tasks will be described in the next sub-sections.

3.1 Extension towards EMI/UMD based e-Infrastructures

The WeNMR grid at March 2012 is based on the EMI/gLite middleware and is formed by 40 central grid services (part of them shared with other EGI VOs) and around 35.000 CPU-cores hosted by 27 resource centres of several European NGIs, plus 2.500 CPU-cores provided by 5 resource centres spread in South Africa, Latin America and Taiwan.

The extension of the WeNMR grid to include these resource centres outside Europe has started already during the past e-NMR project after an agreement achieved with the South African National Grid (SAGrid). The South African National Grid is a federation of Universities, national laboratories and the National Cyberinfrastructure programme founded in 2008. SAGrid has adopted the EMI/gLite middleware since the beginning of its operations and currently supports the enmr.eu VO with its computing centre at the University of Johannesburg, making available more than 200 CPU-cores shared with other national and international VOs. The operation and control of this site, initially supported by INFN, has been endorsed by the Africa & Arabia Regional Operation Centre (ROC) established in 2011 as a coordination and support point for all sites in the two Regions, with the support of many project and initiatives like ASREN (Arab States Research and Education Network), CHAIN, EPIKH (Exchange Programme to advance e-Infrastructure Know-How), EUMEDGRID-Support and SAGrid.

Starting from 2011 an agreement with the GISELA project has lead to the enablement of the enmr.eu VO in three resource centres of Latin America, namely the Universidad Nacional de La Plata in Argentina, Universidad de Los Andes in Venezuela and the Universidade Federal do

Rio de Janeiro in Brasil, providing a total of about 300 CPU-cores. GISELA inherits from the previous EU FP7 projects EELA and EELA-2 (E-science grid facility for Europe and Latin America) that initiated in 2006 a grid infrastructure based on the EMI/gLite middleware and now fully integrated into EGI operations.

In Spring 2011 a collaboration with the Academia Sinica Grid Computing (ASGC) of Taiwan was established and its resource centre is now supporting the enmr.eu VO by providing about 2000 CPU-cores. ASGC is since 2005 the WLCG Tier-1 and is conducting the Asia-Pacific regional e-Science collaborations, development and infrastructure operations. Given the large Structural Biology community present in the Asia-Pacific area and the leading role of ASGC in supporting e-Science, the WeNMR project has started the process of amending its work plan to include Academia Sinica, in particular its Institute of Biological Chemistry and Institute of Physics, as new partner of the WeNMR consortium, starting since Spring 2012. The new work plan foresees new activities in the area of Networking, Dissemination and Outreach (e.g. translation of documentation and publication of abstracts in Chinese language, coverage of major NMR, Biophysics and Structural Biology events in Asia-Pacific region); Services (e.g. allocation and operation of grid resources, dedicated queues for MD simulations with multi-threaded GROMACS version, training/FAQ/user support in Chinese language); Joint Research (e.g. developing protocols for joint usage of SAXS and NMR with docking, enhancing the working capacity of the SAXS beam line at the National Synchrotron Radiation Research Centre in Taiwan).

Figure 4 shows the current geographical distribution of the resource centres (in terms of CPU-cores) supporting mainly in opportunistic way the WeNMR project and its enmr.eu VO. Countries colored in green do not provide yet any resource to the WeNMR grid, while having a not negligible amount of users of the WeNMR services. Countries colored in blue have both users of WeNMR and resource centres supporting the VO.
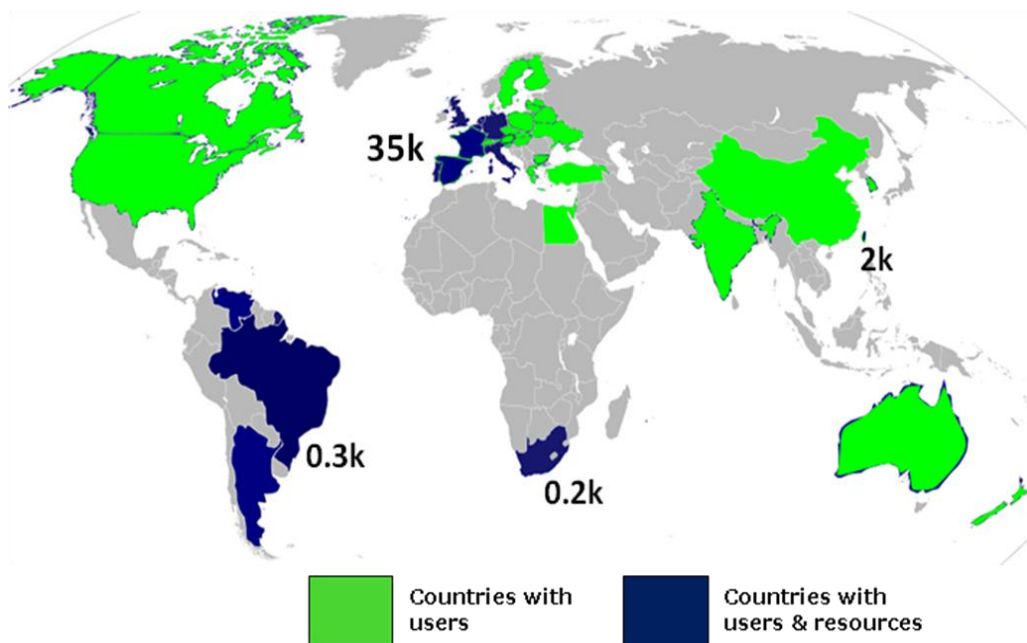
Figure 4: Geographical distribution of computing resources supporting the enmr.eu VO

The 5 resource centres outside Europe are fully monitored by the EGI operational tools, some of which have dedicated instances operated by WeNMR. This is the case of GStat and the VO-Nagios, for which a number of WeNMR application specific probes have been implemented and used to verify the proper functioning of all the resource centres supporting the enmr.eu VO.
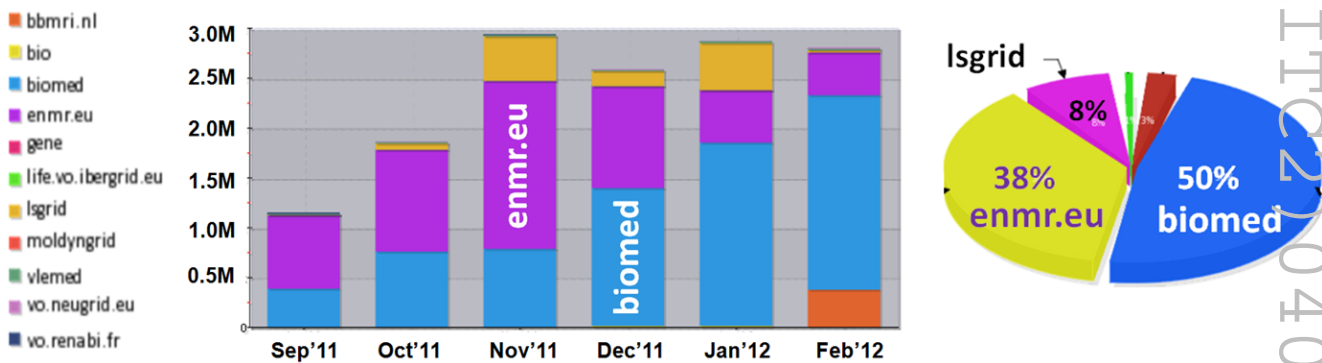

Figure 5: Normalised CPU time (kSI2K hours) vs time (left) and divided by VO (right) consumed by the VOs of the Life Science domain

The accounting of resource usage shown by the EGI Accounting Portal has demonstrated that the computing centres outside Europe contribute at the same level of the European centres in term of both number of jobs and CPU time. In the last six months until February 2012 about 170.000 jobs/month on average have been executed on the WeNMR grid, 7% of them on the outside Europe grid sites, and about 1 Million CPU-hours/month (normalized by kSI2K) on average have been consumed, 2% of them outside Europe. In the Figure 5 the normalized CPU

time consumed by the enmr.eu VO is compared with the other VOs competing in the Life Science domain.

3.2 Extension towards non-EMI/UMD based e-Infrastructures

3.2.1 Extension towards Open Science Grid

The Open Science Grid (OSG) e-Infrastructure operating in USA has a longstanding collaboration in place with the European grid initiatives dating back to the early 2000s with the EU-DataGrid project and later on with the EGEE project series and now EGI. In spite of OSG is based on a different middleware stack, the Virtual Data Toolkit (VDT) which is a collection of software components provided by Globus [1], Condor and other external projects, a full interoperability with EGI has been achieved already in the past in the context of the High Energy Physics (HEP) domain. In fact, many transnational VOs related to HEP experiments can run their application transparently on both OSG and EGI infrastructures without even notice where their jobs are physically executed.

For what concerns the Structural Biology domain, the collaboration between WeNMR and the US SBGrid Consortium dates back to the previous e-NMR project and has been well consolidated after the start of WeNMR. In particular, the registration with the enmr.eu VO is offered as an option to US researchers when they register with the SBGrid Science Portal. This has simplified a lot the access of US users to WeNMR services, and at March 21st #57 users owning a personal X.509 certificate issued by the DOEGrids CA were registered with the enmr.eu VO. Moreover, SBGrid turn out to be also a VO of OSG, that is enabled on more than 26 resource centres across US for a total of more than 60.000 CPU-cores available mainly in opportunistic way.

In a joint OSG-SBGrid-WeNMR-EGI meeting held in March 2011 it was decided to start implementing interoperability between OSG and EGI in the domain of Structural Biology. A number of key persons was appointed for each project to follow a well defined work plan that was detailed after the meeting. The plan established that WeNMR had to use the SBGrid VO for accessing the OSG resources. The enablement of enmr.eu VO on the OSG infrastructure was in fact not allowed by the OSG policy. The plan was designed in two phases. At first, the proof of concept that WeNMR application jobs can effectively run over OSG resources could be obtained quickly by using EMI/gLite grid services set up in Europe and appositely enabled with the SBGrid VO. After that, the second phase would have investigated the possibility to use the GlideinWMS service, the OSG recommended Condor based job submission system. The last one was the preferred solution in the perspective of going in production, because it allows the SBGrid Consortium to control the job flow from Europe and thus optimise the overall workload balance between WeNMR and SBGrid application jobs.

3.2.1.1   Phase 1: using EMI/gLite services

A set of EMI/gLite grid elements were properly set up for the initial tests at the INFN-Padova centre, that is one of the WeNMR partners. An existing gLite-UI and gLite-WMS were enabled with the SBGrid VO and a few testers registered with the SBGrid VOMS server hosted at Harvard Medical School in Boston, the leading partner of the SBGrid Consortium.

A dedicated LFC catalogue instance was also set up at INFN-Padova for the SBGrid VO, in order to test the EMI/gLite data management tools used to transfer files from/to the SRM based Storage Elements (SEs) of both OSG and EGI infrastructures. Data transfer from/to 10 OSG SEs with file registration on LFC were successful performed using the lcg-utils from both gLite-UI and OSG Worker Nodes (WNs). Properly configured OSG WNs in fact have in the path the EMI/gLite data management binaries and libraries.

After some "Hello World" test jobs submitted to OSG with a SBGrid VO proxy using the gLite-UI and gLite-WMS hosted at INFN-Padova, the software of one of the most popular and CPU time demanding WeNMR applications, CS-Rosetta, was properly installed from remote in the SBGrid VO shared area of 14 OSG Computing Elements (CEs) through the use of dedicated grid jobs or globus-job-run commands executing scripts directly in the CE head node. Finally, CS-Rosetta test jobs were successful submitted through gLite-WMS (with minimal changes of bash scripts) and run on OSG WNs.

3.2.1.2 Phase 2: using the GlideinWMS service

GlideinWMS is a Glidein Based WMS (Workload Management System) that works on top of Condor. Condor is a specialized workload management system for compute-intensive jobs developed by the Condor Research Project at the University of Wisconsin-Madison [2]. Glidein is a mechanism by which one or more grid resources (remote machines) temporarily join a local Condor pool. During the period of time when the added resource is part of the local pool, the resource is visible to users of the pool. After glidein, the user may submit jobs for execution on the added resource the same way that all Condor jobs are submitted.

The GlideinWMS is made of several services [3]:

- Users submit jobs to the User Pool Condor schedd process, running on a local Condor submit node.
- The GlideinWMS Frontend, managed by a VO, polls the user pool to make sure that there are enough glideins (worker nodes) to satisfy user jobs. It submits requests to the glidein Factory to submit glideins.
- The glidein Factory, managed by OSG, receives requests from the frontend(s) and submits a Condor startd wrapper to entry points (grid sites).
- The grid sites receive the jobs and starts a Condor startd process that joins the user pool. This shows up as a resource in the user pool.

Once setup, final users can submit regular Condor jobs to the local queue and the glidein Factory will provide the computing resources behind the scenes. From the final user point of view, the Condor pool just magically grows and shrinks as needed. The user need not worry about grid entry points, managing queues, or provisioning worker nodes.

The GlideinWMS services orchestration adopted by WeNMR is shown in Figure 6.
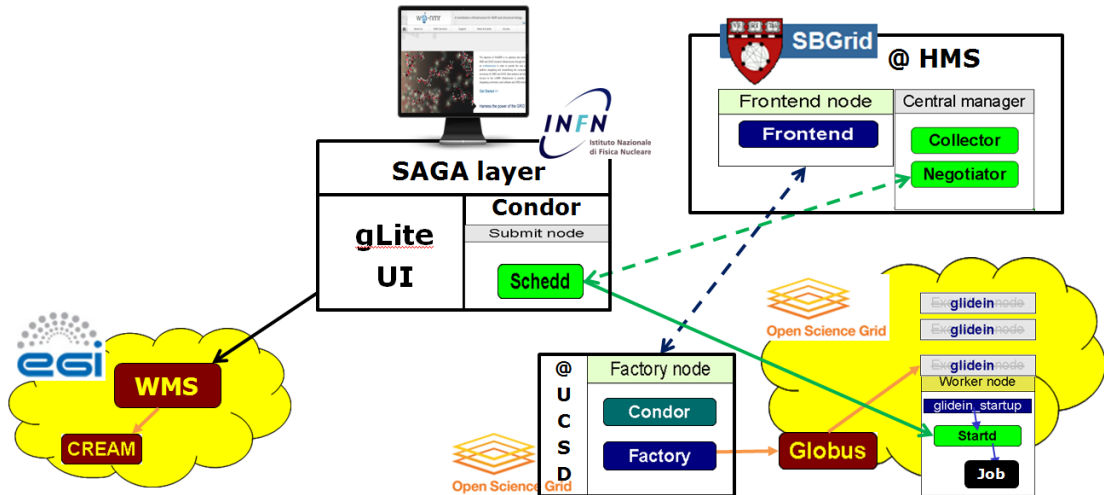
Figure 6: Interoperability schema adopted by WeNMR to enable job submission to OSG

A Condor submit node was installed in the same server hosting the gLite-UI at INFN-Padova. It was configured to work with the GlideinWMS Frontend node hosted at the Harvard Medical School (HMS) in Boston and operated by the SBGrid VO. The Frontend interacts with the Factory node hosted at University of California, San Diego (UCSD) and operated by OSG. While glideins are submitted by the Factory to the OSG nodes using a SBGrid VO proxy, WeNMR jobs can be submitted to the Condor pool using a enmr.eu VO proxy, as done when submitting to EGI via the gLite-WMS. This has two advantages: it allows interactions with the existing WeNMR LFC and SE services; it allows WeNMR jobs to be accounted by the Gratia system (the OSG accounting service) under the enmr.eu VO (and its groups/roles).

The WeNMR application software was installed on the OSG CEs (where not already present as result of the phase 1 tests) from remote via Condor jobs by selecting the desired site in the Condor submit file, or via globus-job-run command as in phase 1.

CS-Rosetta tests jobs were successfully submitted and executed in two ways:

1. Using the *condor_submit* command with a Condor submit file obtained by translating the EMI/gLite JDL (Job Description Language) file used in phase 1 and similar to the ones used in production by the WeNMR portal
2. Using the SAGA Python API implemented by the Louisiana State University team [4].

These testing jobs were run under a personal user account and were properly monitored by the Condor view tool available from the SBGrid Science Portal, as shown in Figure 7. However, the grid accounting system adopted by OSG is Gratia [5]. The Gratia Project (contributed as an OSG Activity) designs and deploys robust, scalable, trustable and dependable grid accounting, publishes an interface to the services and provides a reference implementation. The Gratia system consists of many probes operating on and uploading data from remote locations to a network of one or more collector-reporting systems. The Gratia accounting for the WeNMR test jobs was therefore enabled by deploying the Gratia probes on the Condor submit node of INFN-Padova. This required the installation of two software packages and the re-configuration of the Condor instance. A cron job then send the accounting records collected at the submit node to the

OSG Gratia web site gratiaweb.grid.iu.edu. Many interesting histograms are available from this web site. All of them can be queried by selecting User, VO and even VO group/role. An example of view is shown in Figure 8.
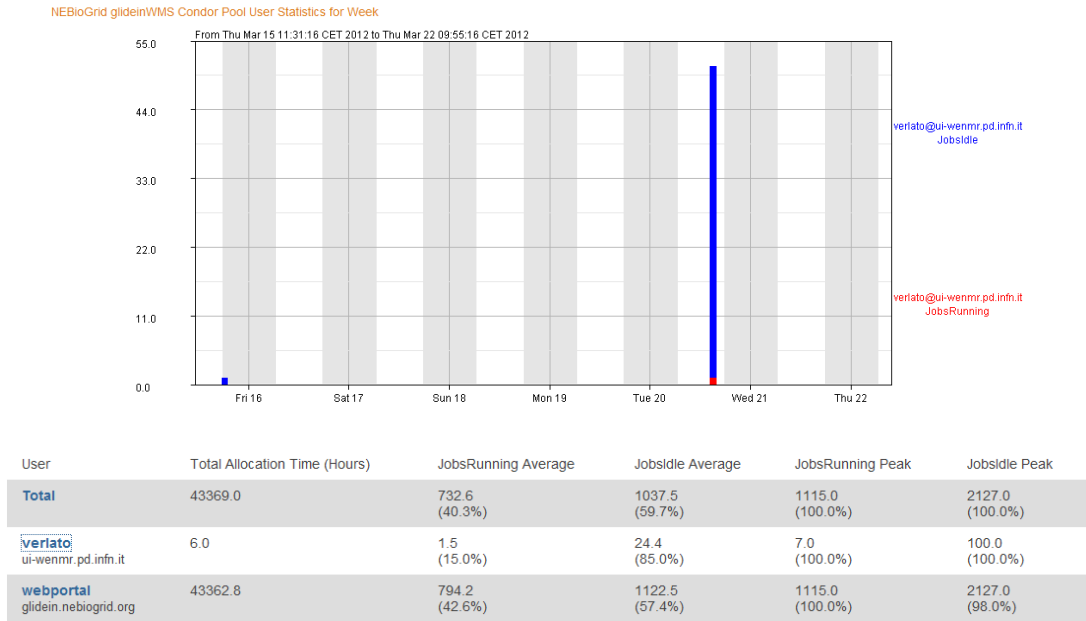

Figure 7: Testing jobs statistics as seen by the SBGrid Science Portal via the Condor view tool
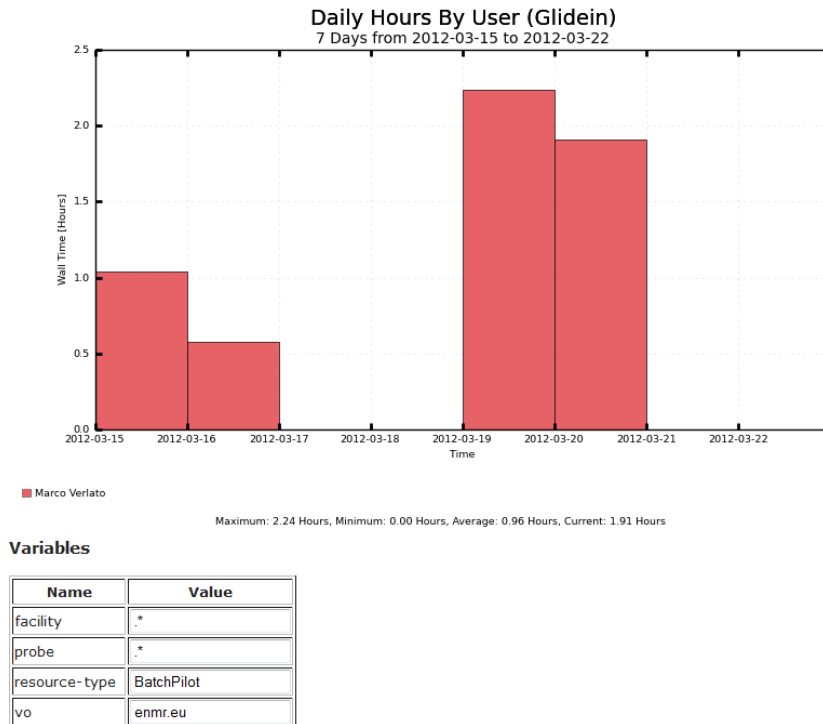

Figure 8: Testing jobs accounting by VO and user as seen by Gratia web site

12

SAGA (Simple API for Grid Application) is a standardized API for developing distributed applications that can run on grid and cloud infrastructure, that follows the specification from the Open Grid Forum (OGF). It has an emphasis on job handling and monitoring, file transfer and management as well as distributed orchestration mechanisms.

The advantage of using a SAGA layer on top of the gLite-UI and Condor clients is to abstract the application level scripts that submit the grid jobs and check their status from the underlying middleware flavours which are connected through a plug-in mechanism implemented by dedicated adaptors.

The open source SAGA implementation provided by the Louisiana State University supports the C++ and Python programming languages and provides backend support for many middleware distributions including the Globus Toolkit, Unicore, Condor and Amazon EC2. Unfortunately this SAGA implementation, while providing a Condor adaptor, still lacks of a EMI/gLite-WMS adaptor. On the other hand, another popular open source SAGA implementation supporting the Java programming language, JSAGA [6], has developed a stable EMI/gLite-WMS adaptor but still lacks of a Condor adaptor. Thus, WeNMR has established contacts with the contributors of both implementations to verify the possibility to develop one of the missing adaptors. The goal in fact is to go in production mode using the same robot proxy and the same SAGA based scripts behind the portal for handling job submission to both EGI and OSG grid respectively through the EMI/gLite-WMS and GlideinWMS services.

After the test jobs described above, a Condor submit node has been installed at the University of Utrecht hosting a number of WeNMR application portals (including the CS-Rosetta one) and further tests with real production jobs are underway.

3.2.2 Extension to India and China grids

The possibility to make the WeNMR grid interoperable with the grid infrastructures available in India and China leverages on the collaboration with the CHAIN project and its achievements. In fact last year CHAIN has proposed an interoperability plan involving four main middleware providers: EMI/UMD (adopted by EGI and PRACE infrastructures), GARUDA (adopted by GARUDA, the Indian National Grid), GOS (adopted by CNGrid, the China National Grid) and OurGrid (a peer-to-peer grid middleware adopted in Brasil by the OurGrid Community).

The project has analysed two interoperability levels:

1. Services level: interoperability between web services of different middlewares. It could be transparent to the users and they are thus not aware of using resources outside their usual infrastructure. It requires consistent work from middleware developers. More promising in the long term, more difficult to implement in short term.
2. Applications level: solution to submit specific applications to several infrastructures; can be obtained, for example, by using meta-schedulers or science gateways. In order to shield the users from the complexity of the middleware details it should implement mechanisms to choose automatically the target grid infrastructure. It requires simpler contribution from middleware developers.

In both cases the interoperability model will be based on a number of well known available OGF and OASIS standard specifications [7]:

- OGSA-BES, JSDL (job execution)
- SAGA, DRMAA (client API)
- GLUE (information system)
- SAML (security)
- GridFTP, SRM (data management)

Due to the time limitations imposed by the CHAIN project duration (end of 2012), it was decided to rely at first on the second interoperability level (Application level) for implementing in the short term a CHAIN interoperability test to demonstrate that:

- e-Infrastructures can be made interoperable among each other using standards (with the meaning of interoperability given above).
- VRC-specific applications can be submitted from anywhere and run everywhere

The current environment foresees a test/demo that is based on an implementation of a Science Gateway with SAGA plug-in to the different middleware flavors, in a similar manner as described in the previous section for the OSG case. WeNMR is one of the VRCs that has been selected by CHAIN to validate their interoperability model.

3.2.2.1   Extension towards GARUDA

GARUDA is India's first National Grid Initiative bringing together academic, scientific and research communities for developing their data and compute intensive applications with guaranteed QoS.
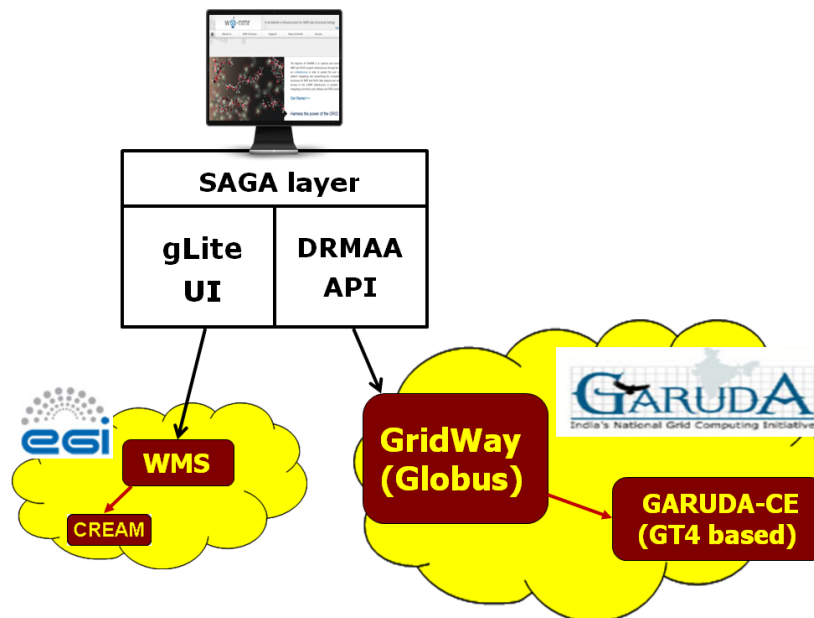

Figure 9: Interoperability schema investigated by CHAIN and WeNMR to enable job submission to GARUDA

GARUDA grid is an aggregation of resources comprising of computational nodes, mass storage and scientific instruments distributed across the country. Currently it provides around 6200 CPU-cores (70 TFlops) and 220 TB of mass storage distributed across 15 grid sites. The GARUDA middleware is based on the Web Services version of Globus 4.x, and the computing resources are managed by the GridWay metascheduler.

The relationship between GARUDA and the European grid initiatives has been strengthened since 2006 through the EU co-funded project series EU-IndiaGrid, in the context of which the interoperability between the two e-Infrastructure was fostered in several workshops, deeply investigated and experimented [8].

The interoperability approach under investigation within the CHAIN plan is based on the implementation of the DRMAA standard packaged with GridWay. The development of a SAGA adaptor for DRMAA is considered a possible solution to enable job submission and management to GARUDA grid, as shown in the simplified schema of Figure 9.

Contacts with the GARUDA management are underway to identify the human resources to dedicate to this effort and define a work plan in the short term.

3.2.2.2 Extension towards CNGrid

The People's Republic of China has funded since 2002 two grid infrastructures supported respectively by the Ministry of Science and Technology and the Ministry of Education, i.e. the China National Grid (CNGrid) and the China Education and Research Grid (ChinaGrid). CNGrid is the largest and the most active one in this region, counting more than 1400 users mainly from the industry simulation domain and biological information domain. It has aggregated computing power of more than 430 TFlops and 2.2 PB of mass storage distributed among 11 supercomputing centers including one located in HongKong SAR.
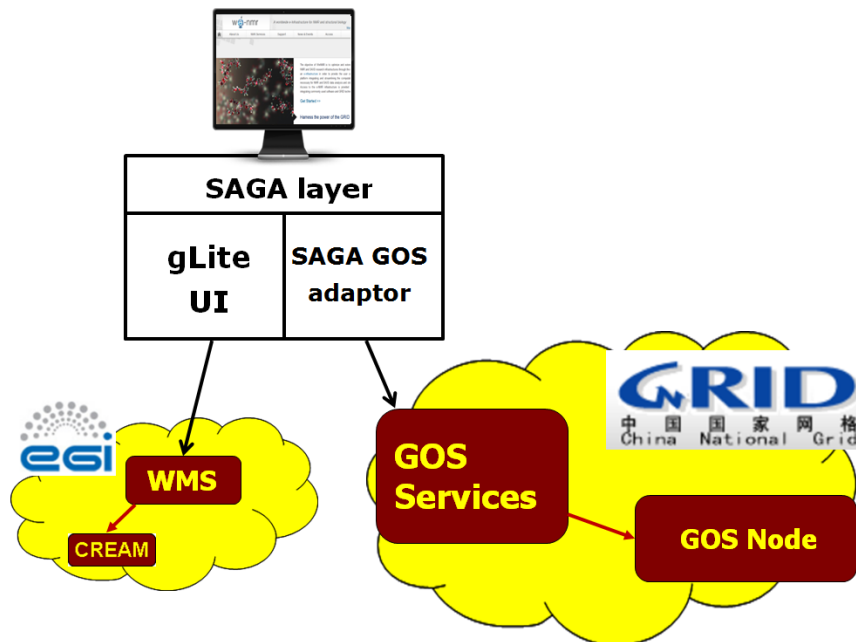

Figure 10: Interoperability schema investigated by CHAIN and WeNMR to enable job submission to CNGrid

Moreover, three PFlops-scale sites are going to be integrated soon. CNGrid uses its own developed middleware called CNGrid Grid Operating System (GOS). Actually at version 4.0, GOS is maintained from several institutions headed from Chinese Academy of Sciences. It is strongly based on Globus, but actually completely rewritten in Java with a quite different layout.

Similarly to the Indian case, the relationship between CNGrid and the European grid initiatives has benefit since 2006 of the EU co-funded project EUChinaGRID, aimed to support the interconnection of the existing European and Chinese grid infrastructures and enable their interoperability. One of the results of the project was the development of gateway prototype between EMI/gLite and GOS [9]. The gateway allowed to exchange jobs between the two infrastructures taking care of the differences related to the job description languages and the security mechanisms.

Building on this experience, CHAIN has discussed with CNGrid representatives about the opportunity of starting the implementation of a JSAGA adaptor to GOS, profiting of a grid school organized by CHAIN in China in April 2012. CNGrid gave a positive feedback, allocating human resources for this immediate effort, and also for the more long term perspective of implementing service level interoperability. It is therefore expected that the solution shown in the simplified schema of Figure 10 will be implemented on time before the CHAIN interoperability demonstration deadline, and thus also available for testing WeNMR applications on CNGrid infrastructure.

4. Usage model, policies and support

As pointed out in section 2, WeNMR chose since its beginning to develop mainly web portals providing "protocolized" access to the grid. Their web application portal architecture is divided in three layers:

1. the user web-based interfaces

2. the server-based request processing and workflow management scripts

3. the application execution scripts that are executed on remote grid sites where the application software has been pre-installed

An example of WeNMR web portal is available in a recently published article [10].

The user is only exposed to the first layer, so he might even never know on which grid infrastructure his jobs will be executed. He is completely hidden from the grid technicalities behind the portal, and is only required to get a personal X509 certificate from a IGTF compliant Certification Authority, upload it into his browser and register with the enmr.eu VO. More experienced users can always use their certificate e.g. to submit jobs to the EGI grid through the command line by using the EMI/gLite User Interface made available from the project as Virtual Machine, however the large majority of the users will login in an application portal and submit their tasks by interacting with the web-based interface.

The second layer is the one "middleware dependent", since it consists in a set of management scripts that handles the users' requests, validate the input files, prepare and submit the jobs to the grid. These scripts are generally bash or python scripts that create JDL files and use glite-wms-* client commands for managing the grid jobs, that are actually suited only to

interact with EMI/gLite based grid infrastructures as the ones described in section 3.1. In section 3.2.1 we demonstrate that properly replacing the JDL files with the corresponding Condor submit files and the glite-wms-* client commands with the corresponding condor_* client commands in the management scripts the WeNMR applications can properly run on the OSG grid. It is clear however that following this approach we'll have to rewrite the management scripts every time we want to plug behind our portal computing resources from grid infrastructures based on a different middleware flavor. We think instead that rewriting them once for all adopting the SAGA standard will easily allow us to access the other grids by exploiting the SAGA adaptors already available or in development in the context of the other international grid projects mentioned in section 3.2.2.

Having a world-wide infrastructure that encompasses so many sites scattered around the world necessarily requires to accurately define policies, procedures and support activities at different levels when interacting with the resource providers of the various grid infrastructures. WeNMR obtained before its start a number of formal letters of support from several world-wide grid initiatives (US-SBGrid, India-GARUDA, South Africa-SAGrid) and grid related projects (CHAIN, EELA2/GISELA, EUIndiaGrid, EUAsiaGrid) whose main goal is to foster the adoption of grid technologies for the benefit of the international scientific communities active in their geographic areas by promoting adequate support actions to address interoperability issues and define appropriate policies. Later on WeNMR became the first Virtual Research Community (VRC) to sign a Memorandum of Understanding (MoU) with EGI.eu, and signed also MoUs with CHAIN, EUMEDGRID-Support and EMI that defined in more detail a frame of collaboration between these projects and WeNMR. In particular the CHAIN project aims to coordinate and leverage recent efforts and results with a vision of a harmonised and optimised interaction model for e-Infrastructures by focusing on grid interfaces between Europe and the rest of the world. The project is defining a coherent operational and organisational model, where a number of EU countries/regions will possibly act, in collaboration with EGI.eu, as bridges/gateways to other Regions/Continents. The validation of this model is planned by supporting the extension and consolidation of worldwide virtual communities, which increasingly require distributed facilities (large instruments, distributed data and databases, digital repositories, etc.) across the regions for trans-continental research. CHAIN will also act as a worldwide policy-watch and coordination instrument, by exploring and proposing concrete steps for the coordination with other initiatives and studying the evolution of e-Infrastructures. The cooperation of WeNMR with CHAIN is therefore extremely valuable for achieving the world-wide extension of the WeNMR e-Infrastructure, and the WeNMR VRC is one of the world-wide scientific communities that has been chosen by CHAIN to validate its model.

For what concerns operational policies, WeNMR adheres to the policies and procedures approved and in use by the EGI community, and defined by the EGI policy groups [11]. The most relevant in this context are the Service Operations Security Policy, the VO Operations Policy and the Portal Policy. These policies are our reference guide when discussing and agreeing the procedures for integrating new resources from different grid infrastructures.

5. Conclusions

WeNMR grid has successfully extended towards EMI/gLite based infrastructures in South Africa, Latin America and Asia-Pacific area. More sites are expected to join soon specially from the Asia-Pacific area thanks to the new partnership with the Academia Sinica of Taiwan which is coordinating the grid activities in that region, and from the Mediterranean area thanks to the collaboration with the EUMEDGRID-Connect project and the Africa & Arabia ROC.

The very fruitful collaboration with the US SBGrid Consortium and the Open Science Grid has led to achieve a positive proof of concept of interoperability between WeNMR/EGI and SBGrid/OSG by enabling transparent WeNMR job submission, management, monitoring and accounting over the OSG grid exploiting the computing resources supported by the SBGrid Virtual Organisation. Move from proof of concept to production mode is expected in the short term.

More long term, the perspectives of extending WeNMR towards Indian and Chinese grids look very promising thanks to the collaboration with CHAIN project, that is planning an interoperability demonstration by the end of 2012 which is including the WeNMR applications as validation use case.

6. Acknowledgements

PoS(EGICF12-EMITC2)040

References

[1] I. Foster, "Globus Toolkit Version 4: Software for Service-Oriented Systems", IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp 2-13, 2006.

[2] Douglas Thain, Todd Tannenbaum, and Miron Livny, "Distributed Computing in Practice: The Condor Experience", Concurrency and Computation: Practice & Experience - Grid Performance, Volume 17 Issue 2-4, February 2005, pages 323-356

[3] D Bradley et al, "Scalability and interoperability within glideinWMS", 2010, J. Phys.: Conf. Ser. 219 062036

[4] www.saga-project.org

[5] Philippe Canal, "Gratia: New Challenges in Grid Accounting", 2011, J. Phys.: Conf. Ser. 331 062028

[6] grid.in2p3.fr/jsaga

[7] M. Riedel, E. Laure, et al., "Interoperation of World-Wide Production e-Science Infrastructures", Concurrency and Computation: Practice & Experience - A Special Issue from the Open Grid Forum, Volume 21 Issue 8, June 2009, pages 961-990.

[8] Balachandar R Amarnath, et al., "EU-IndiaGrid interoperability experiments using GridSeed tool", Multiagent and Grid Systems, Volume 6 Issue 3, August 2010, pages 261-269

[9] Bingheng Yan, et al., "Experiences with the EUChinaGrid Project - Implementing Interoperation between gLite and GOS", APSCC '07 Proceedings of the The 2nd IEEE Asia-Pacific Service Computing Conference, pages 224-231

[10] Marc van Dijk, Tsjerk A. Wassenaar, and Alexandre M.J.J. Bonvin, "A Flexible, Grid-Enabled Web Portal for GROMACS Molecular Dynamics Simulations", J. Chem. Theory Comput. 2012 8:3463-72, DOI: 10.1021/ct300102d

[11] www.egi.eu/about/policy/policies_procedures.html