Deliverable D2.2

Assessment of the e-NMR infrastructure and GOC report

SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures

INFRA-2007-1.2.2 - Deployment of eInfrastructures for scientific communities

Grant agreement for: Combination of Collaborative projects & Coordination and support actions

Proposal/Contract no.: 213010 – e-nmr

Project full title: Deploying and unifying the NMR e-Infrastructure in System Biology

Project coordinator: Prof. Dr. Harald Schwalbe

Project website: http://www.enmr.eu/

Period covered: from **01-11-2007 to 30-4-2009**

Editor: Marco Verlato (INFN)

Table of contents

1. Introduction

1.1. Purpose

This document is the project deliverable D2.2 due by Month 18. It aims at reporting the activity of the Grid Operation Centre (GOC) task T2.2 assigned to the Work Package 2 (WP2) of the e-NMR project. It is intended to be a first assessment of the e-NMR grid infrastructure usage.

1.2. Document organisation

The document is organised as follows:
Section 1 contains the purpose of the document, its references and a glossary of terms and acronyms.
Section 2 summarizes the content of the document.
Section 3 describes the current status of grid services and hardware deployed on the e-NMR grid, updating the one presented in D2.1 document at PM 12.
Section 4 shows the usage of the e-NMR grid in the last six months.
Section 5 describes the operational procedures and tools adopted to manage, control and monitor the e-NMR grid.
Finally, Section 6 tracks the conclusions and Section 7 contains the annexes.

1.3. References

[R1]	Public link not available	*e-NMR Annex I*
[R2]	https://igi.cnaf.infn.it/	*Italian Grid Infrastructure portal*
[R3]	https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf	*Usage of Glue Schema v1.3 for WLCG Installed Capacity information*
[R4]	http://haddock.chem.uu.nl/enmr/eNMR-portal.html	*e-NMR Applications web portal*
[R5]	http://www.jspg.org/wiki/VO_Portal_Policy	*VO Portal Policy – Version 3.0 beta-1 draft*
[R6]	https://emds.cern.ch/document/723928/1	*Operations Metrics Defined, EGEE-II Milestone MSA1.1*
[R7]	cic.gridops.org	*EGEE Operations Portal*
[R8]	sam-docs.web.cern.ch	*SAM web page*
[R9]	gridice.forge.cnaf.infn.it	*GridICE web page*
[R10]	goc.grid.sinica.edu.tw/gstat	*GStat web page*
[R11]	twiki.cnaf.infn.it/cgi-bin/twiki/view/WMSMonitor/	*WMSMonitor web page*
[R12]	https://twiki.cern.ch/twiki/bin/view/LCG/GridView	*GridView web page*
[R13]	doi.ieeecomputersociety.org/10.1109/GRID.2003.1261716	*"An Economy-based Accounting Infrastructure for the DataGrid", proceedings of the 4th International Workshop on Grid Computing (GRID2003)*
[R14]	http://www.iop.org/EJ/article/1742-6596/119/5/052012/jpconf8_119_052012.pdf	*"HLRmon: a Role-based Grid Accounting Report Web Tool", Journal of Physics: Conference Series 119 (2008) 052012*
[R15]	igrelease.forge.cnaf.infn.it	*INFNGRID Release Wiki*
[R16]	https://edms.cern.ch/document/860386/0.5	*EGEE-II Service Level Description between ROCs and Sites*
[R17]	http://www.sagrid.ac.za	*South African National Compute Grid web portal*
[R18]	www.enmr.eu	*e-NMR project web portal*
[R19]	http://en.scientificcommons.org/984536	*"Global Grid User Support: The Model and Experience in LHC Computing Grid", proceedings of Computing in High Energy and Nuclear Physics (CHEP06), Mumbai, India, February 2006*
[R20]	https://edms.cern.ch/document/927171/1	*Operations Automation Strategy, EGEE-III Milestone MSA1.1*

1.4. Terminology

This subsection provides the definitions of terms, acronyms, and abbreviations required to properly interpret this document.

Term	Definition
BCBR	Bijvoet Centre for Molecular Research, University of Utrecht, The Netherlands
BDII	Berkeley Database Information Index
BMRZ	Centre for Biomolecular Magnetic Resonance, Goethe University, Frankfurt, Germany
CA	Certification Authority
CE	Computing Element
CIRMMP	Interuniversity Consortium for Magnetic Resonance on Metalloproteins, Florence, Italy
DGAS	Distributed Grid Accounting System
DoW	Description of Work
EGEE	Enabling Grids for e-Science
EGI	European Grid Initiative
GGUS	Global Grid User Support
gLite	Codename of the middleware software suite developed by EGEE
GOC	Grid Operation Centre
HLR	Home Location Registry (component of DGAS system)
IGI	Italian Grid Infrastructure
INFN	National Institute of Nuclear Physics, Italy
KSI2K	KiloSpecInt2000: computer benchmark for CPU's integer processing power
LB	Logging & Bookkeeping service
LCG	LHC Computing Grid
LFC	LCG File Catalogue
NGI	National Grid Initiative
NMR	Nuclear Magnetic Resonance
PM	Project Month
RA	Registration Authority
ROC	Regional Operation Centre
SAM	Service Availability Monitoring
SE	Storage Element
SL	Scientific Linux
UI	User Interface
VO	Virtual Organization
VOMS	Virtual Organisation Membership Service
WMS	Workload Management System
WP	Work Package

2. Executive summary

As stated in the DoW [**R1**], the main objective of the task T2.2 is

"to set up an operational structure for the management, control and support of the e-NMR infrastructure. This will be constituted by a set of tools to proactively monitoring the e-Infrastructure, checking the status and accounting the use of the central grid services and of the remote computing and storage resources. These tools will be chosen among the ones already in use within EGEE (e.g. GridICE, DGAS, SAM, GStat, RTM), in order to make the integration of the operations of the e-NMR infrastructure into EGEE smoother.
[…]
An important parameter to measure the success of the project via the GOC will be the number of members enrolled in the e-NMR VO, including both researchers who port their application(s) to the e-NMR GRID as well as new users of the e-NMR platform. A reasonable target for the project is to attain 120 members at its end, including also non-European members."

This document thus describes the operational tools put in place to establish a Grid Operation Centre (GOC) for the e-NMR grid, and evaluates the use of the infrastructure in terms of number of users, number of jobs run, CPU time used, storage capacity used, integration level with EGEE, effectiveness of GOC procedures and user support.

The number of registered users with enmr.eu VO has reached the value of 65. As a comparison, the number of users registered with the EGEE global VO "biomed" is currently 197. Notice however that biomed VO is covering since the beginning of EGEE project a broad scientific applications area divided into the three sectors of medical imaging, bioinformatics and drug discovery, and counting actually about twenty regularly used applications.
The planned training events and workshops demonstrating the e-NMR application platform organised in the second half of the project are expected to attract new users for enmr.eu VO, allowing meeting the target of 120 VO members declared in the DoW.
Anyway, nearly 150000 jobs and 12 years of normalized CPU elapsed time have been consumed on the grid by e-NMR applications in the last six months.

Most of the deployment and operational tools designed for the EGEE and the IGI production grid infrastructures have been successfully adopted by the e-NMR grid, making easy in the short term the process of integrating it with the EGEE coordinated operational procedures.
In particular, the CIRMMP grid site already managed to join EGEE grid through the IGI organisation, which is managing the Italian ROC of EGEE.
Moreover, Northern European ROC and Germany/Switzerland ROC have ensured to provide support for including in EGEE respectively BCBR and BMRZ grid sites.

The participation to 4th EGEE User Forum in March 2009 has raised great interest, among others, of the South African National Compute Grid, which has committed to provide support to the e-NMR project and to share part of its newly deployed resources with e-NMR grid.
The e-NMR GOC will take care of coordinating and supporting this activity in the next months.

3. Current grid deployment status

The resources deployed for building the e-NMR grid infrastructure have been widely described in the document D2.1 delivered at PM 12. Only a summary updated at PM 18 will be given here.
The enmr.eu VO has been created at the beginning of the project and registered with the EGEE Operations Portal as global VO under the Life Sciences discipline. In the rest of the document the term e-NMR grid will refer to the set of grid services and resources supporting the enmr.eu VO.

3.1. Core grid services

CIRMMP and INFN centres are hosting since the beginning of the project the central core services for e-NMR grid. While services hosted by CIRMMP are fully dedicated to e-NMR, the services hosted by INFN are shared with the Italian Grid Infrastructure (IGI) [**R2**], which in turn is part of the EGEE grid infrastructure.
e-NMR grid is in fact a mixed infrastructure, composed by services and resources which are in part shared with other VOs of EGEE, and in part neither included in any National Grid Initiative (NGI) nor in the EGEE production grid.
Currently, central core grid services span all the areas of the gLite architecture: job management, information system, security, data management, monitoring and accounting.
GStat service is hosted by Academia Sinica Grid Computing (ASGC) of Taipei, Taiwan, the coordinator of EGEE Asia-Pacific ROC, which courtesy agreed to add an instance for e-NMR grid on its EGEE-wide server. A failover instance of GStat is provided by INFN.
The full list of the core grid services available to e-NMR is given in table 1.

Site	Hosting machine	Service	Scope
Job management services			
CIRMMP	wms-enmr.cerm.unifi.it	WMS	e-NMR
INFN	prod-wms-01.pd.infn.it	WMS	EGEE
INFN	glite-rb-00.cnaf.infn.it	WMS	EGEE
CIRMMP	lb-enmr.cerm.unifi.it	LB	e-NMR
INFN	prod-lb-01.pd.infn.it	LB	EGEE
INFN	lb009.cnaf.infn.it	LB	EGEE
Information system services			
CIRMMP	bdii-enmr.cerm.unifi.it	Top-BDII	e-NMR
INFN	egee-bdii.cnaf.infn.it	Top-BDII	EGEE
Security services			
INFN	voms2.cnaf.infn.it	VOMS	EGEE
INFN	voms-02.pd.infn.it	VOMS	EGEE
Data management services			
INFN	lfcserver.cnaf.infn.it	LFC	EGEE
Monitoring services			
CIRMMP	gridice-enmr.cerm.unifi.it	GridICE	e-NMR
ASGC/INFN	gstat.gridops.org.gstat/e-nmr	Gstat	e-NMR
CIRMMP	ui-enmr.cerm.unifi.it	SAM	e-NMR

INFN	eu-india-01.pd.infn.it	WMSMonitor	e-NMR
Accounting services			
CIRMMP	dgas-enmr.cerm.unifi.it	HLR	e-NMR
INFN	dgas.cnaf.infn.it	HLRmon	EGEE

Table 1: Central grid services deployed for e-NMR

3.2. Local grid services

Computing and storage resources of e-NMR grid are geographically distributed across the four partners' institutions: BCBR in Utrecht, BMRZ in Frankfurt, CIRMMP in Florence and INFN in a number of centres in Italy. An additional site managed by the CNR, member of IGI, supported enmr.eu VO on best effort basis.

The overall picture is shown in Table 2, which collect data taken from GStat.
Computing and storage resources of BCBR, BMRZ and CIRMMP are fully dedicated to enmr.eu VO, and correspond currently to 236 CPUs for a total computing power of 473 KSI2K, as computed according to [**R3**], and 2.87 TBs of disk space available for grid storage. Additional resources of IGI added 522 CPUs, for a total computing power of 1016 KSI2K, and 26.45 TBs of disk space available for grid storage, shared with others regional and global VOs.

The main difference with respect to PM 12 was the addition of 12 CPUs at the NMR centres, and of the best effort support to enmr.eu VO from INFN-LNL-2 and INFN-TRIESTE grid sites, making available additional 228 CPUs and 20.43 TBs of storage capabilities.

In April 2009 BiG Grid, the Dutch e-Science Grid, officially agreed to support enmr.eu VO. As a consequence, a large cluster farm with 1600 CPUs at Philips in Eindhoven is expected to join the e-NMR grid in the short term.

The number of CPUs is a very dynamic metric. CPUs are continuously replaced by newer, more powerful ones. Nowadays many systems have multi-core architectures, and in this case it is actually the number of the cores that is counted. Throughout the document cores are referred to as CPUs.
e-NMR grid is very heterogeneous as far as processor architectures are concerned. It ranges from old Intel Xeon machines of 0.9 KSI2K/core to most modern Xeon quad-core E5430 machines of 2.8 KSIK/core. Many AMD Opteron machines ranging from 1.3 to 2.0 KSI2K/core are also present.
Both x86-64 and i386 instruction set architectures are supported, even if most of the applications run 32-bit code.

The current middleware release deployed on e-NMR grid is the INFNGRID gLite 3.1/SL4 Update 42/43/44 (24/4/2009).

	Month 12			Month 18		
Site	**CPUs**	**KSI2K**	**TBs**	**CPUs**	**KSI2K**	**TBs**
BCBR	176	314	1.81	176	314	1.81
BMRZ	16	40	0.89	28	70	0.89
CIRMMP	32	89	0.17	32	89	0.17
Total dedicated	**194**	**443**	**2.87**	**236**	**473**	**2.87**
INFN-PADOVA	58	60	5.65	86	144	5.65
INFN-CATANIA	204	399	-	204	399	-
CNR-ILC-PISA	4	3	0.37	4	3	0.37
INFN-LNL2	-	-	-	132	356	-
INFN-TRIESTE	-	-	-	96	114	20.43
Total shared	**266**	**462**	**6.02**	**522**	**1016**	**26.45**
TOTAL	**460**	**905**	**8.89**	**758**	**1489**	**29.32**

Table 2: Computing and storage resources forming the e-NMR grid

4. Resource Usage

The enmr.eu VO was set up in December 2007, just one month after the official start of the e-NMR project. It was at first populated with people working for the project at the partners' institutes, as soon as these managed to contact their local Registration Authority (RA) of the corresponding national Certification Authority (CA), or became themselves RA.
At the end of the first year, October 2008, the VO counted 26 people registered, all of them internal to the Consortium. After the milestone M2.1: "e-NMR grid operational" was achieved at PM 12, the registration was opened to research groups of the bio-NMR community outside the Consortium. Currently the VO counts 65 members. It includes people from Slovenia, UK, US and New Zealand.

11 applications were deployed on the e-NMR grid at PM 12, covering NMR data processing (NMRPipe, PROSA), NMR data analysis (CNS, CYANA, XPLOR-NIH, HADDOCK, CSRosetta), molecular dynamics simulations (GROMACS, AMBER) and automated assignment (CYANA, MAPPER). Two of them (XPLOR-NIH and HADDOCK) were accessible via a user friendly web portal developed by CIRMMP and BCBR partners.
At PM 18 a new application with access via portal has been available, CcpNmr FormatConverter, and access via portal has been also enabled for CYANA, CSRosetta and AMBER. Furthermore all of the applications portals were made available from a single web page hosted at BCBR [**R4**]. From the portal is now possible, for example, to start a XPLOR-NIH calculation, then take its output and refine it with AMBER.
Table 3 summarizes the evolution of the number of VO members, applications and portal development in the last 6 months.

The access rights to the portals were guaranteed by requiring:

1) a X509 personal certificate issued by a IGTF accredited CA loaded into the users' browser
2) the personal certificate DN had to be registered with the enmr.eu VO

Behind the portal a business logic software component performed the needed operations to allow job submission over the e-NMR grid. Actual job submission was implemented making use of X509 robot certificate issued per application by NIKHEF and INFN CAs, according with the VO Portal Policy draft documented in [**R5**] by the Joint Security Policy Group (JSPG) of EGEE.
A custom User Tracking System was implemented by portal developers in order to keep track of applications usage by each user, associating the DN of the users' certificate read at portal access time with the identifiers of grid jobs submitted with the robot on behalf of the user. Of course the usual grid accounting tools will see only the aggregate application usage, while the portal administrator will be always able to see the usage records of each user.
Following the recommendations of the EU officers, a plan for replacing robot-based job management with one based on users' certificate has been set up after the DoW amendment carried out in March, and is currently under development. The release of the new implementation is expected by June 2009.

Month 12			Month 18		
# VO members	**Application**	**Access via portal**	**# VO members**	**Application**	**Access via portal**
26	NMRPipe		65	NMRPipe	
	PROSA			PROSA	
	INFIT			INFIT	
	CNS			CNS	
	CYANA			CYANA	✓
	XPLOR-NIH	✓		XPLOR-NIH	✓
	HADDOCK	✓		HADDOCK	✓
	CSRosetta			CSRosetta	✓
	GROMACS			GROMACS	
	MAPPER			MAPPER	
	AMBER			AMBER	✓
				CcpNmr FormatConverter	✓

Table 3: Evolution of the number of enmr.eu VO members, applications and their availability through web portal

The DGAS system, described in the next section, has been used to account the usage of the e-NMR grid by enmr.eu VO members. The period considered is in the interval from the 1st of November 2008, official opening of the e-NMR grid, to 30th of April (end of PM 18). Figure 4-1 shows the number of jobs run per site in the given time interval.

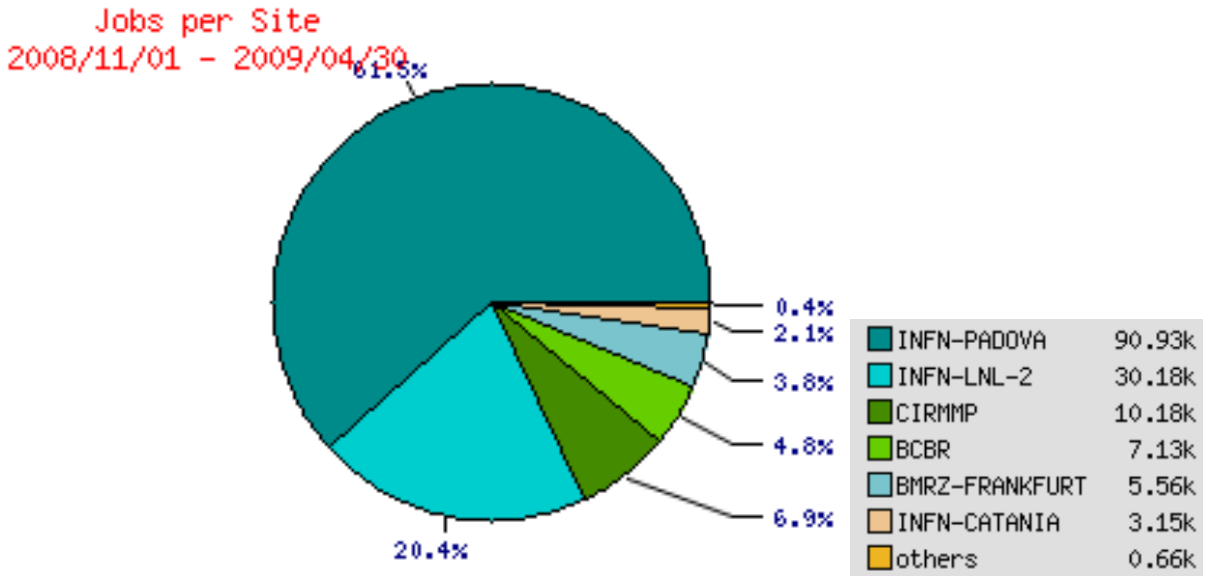

Figure 4-1: Number of job run by each grid site (147768 in total)

A total of 147768 jobs were run in the period, 82% of them in two INFN grid sites.
A deeper analysis can be performed by the GOC manager, who has access to the Distinguished Names (DNs) of the proxy certificates used to submit grid jobs. This sensitive information can be retrieved only by authorized people, through a mechanism explained in section 5.3.2, and is shown partially in figure 4-2. Here the users DNs have been hidden to ensure privacy, but the ones corresponding to robot certificates.

Users activity summary

User	Site	VO	Job Number	CPU Time [s]	Wall Time [s]
Robot: grid client - [redacted]	INFN-PADOVA	enmr.eu	89595	49694641	73161087
Robot: grid client - [redacted]	INFN-LNL-2	enmr.eu	29730	64782645	66248756
Robot: grid client - [redacted]	CIRMMP	enmr.eu	8348	8710531	10725718
Robot: grid client - [redacted]	BCBR	enmr.eu	6561	16282010	28070901
Robot: grid client - [redacted]	BMRZ-FRANKFURT	enmr.eu	4761	8221167	12921132
Robot: grid client - [redacted]	INFN-CATANIA	enmr.eu	2810	1756013	3558934
[redacted]	INFN-PADOVA	enmr.eu	900	980	44929
[redacted]	CIRMMP	enmr.eu	674	134769	171974
Robot: WEB-UI - [redacted]	CIRMMP	enmr.eu	585	496773	559998
[redacted]	CNR-ILC-PISA	enmr.eu	542	547	7985
[redacted]	BMRZ-FRANKFURT	enmr.eu	342	17301	20875
[redacted]	CIRMMP	enmr.eu	337	24188	28197
[redacted]	INFN-LNL-2	enmr.eu	288	16166	19497
[redacted]	INFN-PADOVA	enmr.eu	171	11672	17734
Robot: WEB-UI - [redacted]	BCBR	enmr.eu	159	157069	300467
[redacted]	BMRZ-FRANKFURT	enmr.eu	141	6874	111563
[redacted]	INFN-CATANIA	enmr.eu	134	11771	17555
[redacted]	BMRZ-FRANKFURT	enmr.eu	131	248307	665425
[redacted]	INFN-CATANIA	enmr.eu	122	2454	7032
[redacted]	INFN-PADOVA	enmr.eu	117	9827	12263
[redacted]	BCBR	enmr.eu	112	4897	8683
[redacted]	CNR-ILC-PISA	enmr.eu	100	3715	5561
[redacted]	CIRMMP	enmr.eu	98	4140	5259
[redacted]	BMRZ-FRANKFURT	enmr.eu	84	3200	91613
[redacted]	BCBR	enmr.eu	77	1601	12242
[redacted]	BCBR	enmr.eu	63	405	309615
[redacted]	BCBR	enmr.eu	60	3580	7222

Figure 4-2: Users activity summary ordered by number of jobs submitted to a grid site by a given user. Only the robot related part of the certificate DNs are shown here, while single user DNs have been hidden.

The first 6 lines of figure 4-2 show that the 96% of the jobs (141805 over 147768) were actually submitted to the entire e-NMR grid through the grid-enabled web portal hosted at BCBR, making use of its robot certificate.
The internal User Tracking System of the portal has shown that these jobs correspond to about 200 HADDOCK runs (a typical run submits to the grid 250 jobs, but for guru users the number can increase up to 15000), plus about 10 CSRosetta test runs (each run in this case submits from 1000 to 2500 jobs).

Figure 4-3 shows the daily CPU and Wall clock time (in days) consumed by the same grid jobs, summing up a total of 4.9 CPU.years and 6.4 WT.years.
Notice the big peaks after the 16th of April. Despite these correspond only to about the 17% of the total jobs (25k over 148k), these account respectively for the 73% (3.6 over 4.9 years) and 63% (4 over 6.4 years) of the total consumed CPU time and Wall time.

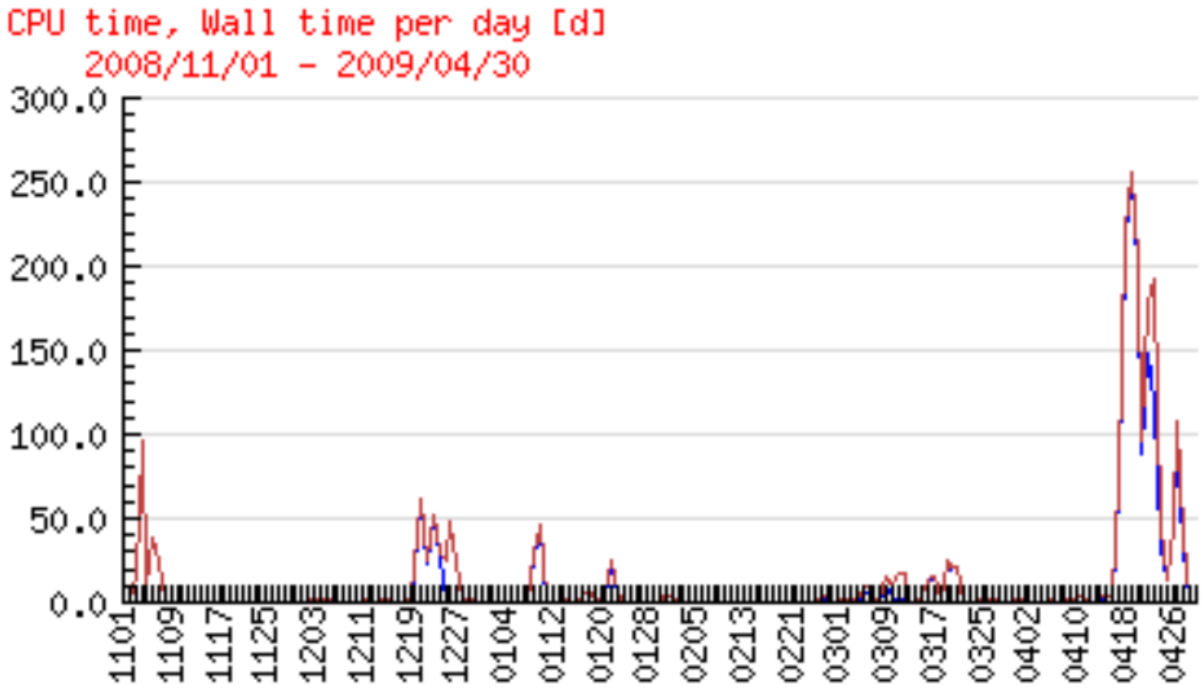

Figure 4-3: CPU-days (blue line) and WT-days (red line) consumed per day (4.9 and 6.4 years in total)

Figure 4-4 shows the normalized Wall time per grid site, i.e. the Wall clock time consumed on the computing nodes in KSI2K.days units, according with the "size.2" metric defined in [**R6**].
A total of 12.2 KSI2K.years of normalized Wall time was consumed in the reported period. In this case about 32% of the computing was provided by the NMR dedicated centres of BCBR, BMRZ and CIRMMP.

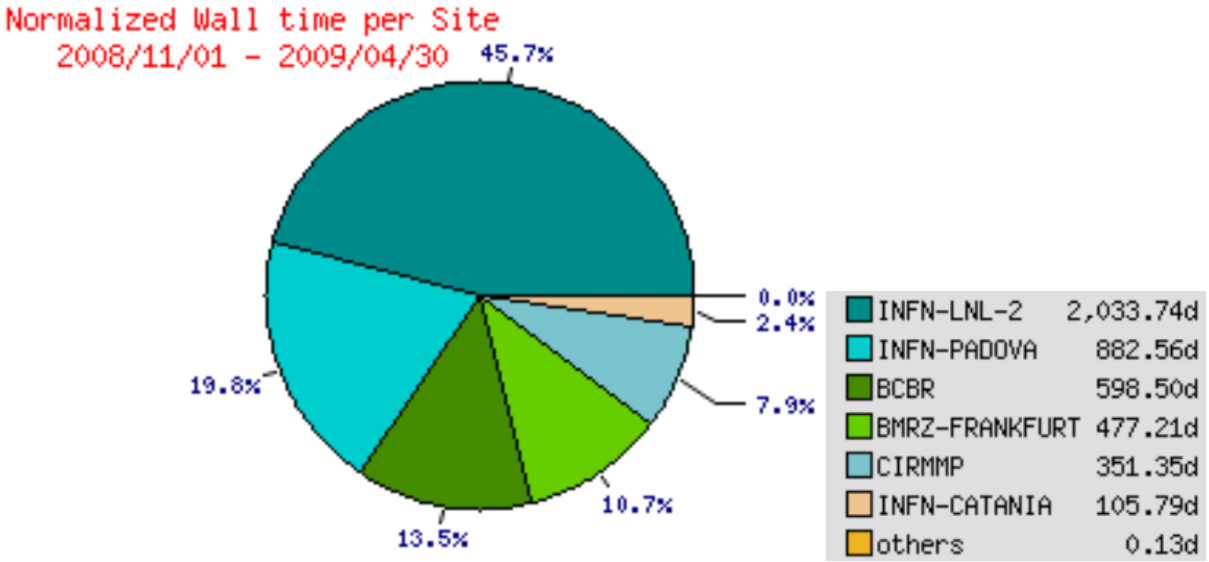

Figure 4-4: Normalized wall clock time consumed in each grid site (12.2 KSI2K.years in total)

5. Grid Operation Centre

The term Grid Operations Centre (GOC) indicates an operational organisation for the management, control and support of a grid infrastructure.
The e-NMR grid in its initial phase can be considered a small scale clone of the EGEE production grid, comprising a few resource centres deploying gLite middleware and supporting only the enmr.eu VO and a higher priority VO for operational testing. Since the beginning however it includes a number of resource centres part of the EGEE production grid which agreed to support the enmr.eu VO. These centres within the EGEE organisation are managed by the Italian Regional Operations Centre (ROC). e-NMR grid is therefore a "mixed" grid formed by a number of EGEE grid sites plus other few grid sites completely unknown to EGEE.
One of the goals of the e-NMR project is however to fully integrate e-NMR grid into the EGEE production grid, or its future evolution within the EGI model.
Thus, it has been a quite natural choice to set up for the e-NMR grid a small scale GOC based on a sub-set of the current core operational tools adopted by EGEE project to manage its production grid operations.

These tools fall into five main areas:

1) **Provision of information about resources.**
The static and dynamic information concerning both grid sites belonging to the different ROCs and Virtual Organisations supported in EGEE are collected in three main repositories: the GOCDB, the Operations Portal and the Information System (BDII). e-NMR GOC make use of Operations Portal for what concerns the static information about the enmr.eu VO, and of the top-BDII hosted at CIRMMP site as Information System for the whole e-NMR grid. It is clear that as soon as the e-NMR sites not yet known by EGEE will join their relevant ROC to be included into the EGEE production grid their static information will be registered with the GOCDB.

2) **Grid monitoring and reporting.**
A number of different tools are used for the oversight of EGEE grid infrastructure such as GStat, GridView and SAM. Other tools are used at regional and site level, such as WMSMonitor, Nagios and GridICE.
For e-NMR GOC we decided to deploy a standalone instance of SAM, an instance of GridICE, an instance of GStat and an instance of WMSMonitor.

3) **Grid accounting and reporting.**
A EGEE-wide repository for accounting data is provided to collect and aggregate data sent from grid sites. Different accounting systems such as APEL and DGAS are used to collect the data and send them to the central system. Portals such as the EGEE Accounting Portal and HRLmon display this data and produce reports.
For e-NMR GOC the DGAS system and the HLRmon visualization tool have been deployed.

4) **Resource and user support.**
In EGEE a global ticketing system is provided by GGUS that links together regional ticketing systems. Site managers and users are free to send support requests either to their

regional support or global GGUS. Proper escalation between regional and GGUS systems on both directions is managed by the support teams handling the tickets.
Given the current small size of e-NMR grid, and being grid technicalities mainly hidden to the final users behind friendly application portals, the set up of a dedicated ticketing system instance was not considered a priority. Instead, for the e-NMR GOC a closed mailing list with archive functionality for answering to resource administrators' and software developers' support requests has been set up. A first level support team at INFN formed with personnel funded by e-NMR project is in charge to handle the requests and re-direct them to GGUS or to Italian ROC support system if needed.

5) **Follow-up of alarms created by monitoring systems.**
In EGEE the Operations Dashboard provides a mechanism for a grid operator to create tickets for alarms raised by monitoring systems such as SAM.
For the e-NMR GOC, the same support team of INFN is in charge to daily check the status of the grid using directly the GStat and GridICE tools, and to promptly react to alarms raised by the standalone instance of SAM.

In the following sub-sections we'll analyse the various aspects of the GOC established to oversee the e-NMR grid.

5.1. Provision of information about resources

The EGEE Operations Portal [**R7**] has been used to store and make publicly available all of the information about enmr.eu VO: name, discipline, status, scope, description, homepage, enrolment URL, acceptable use policy (AUP), VOMS server details, VOMS supported groups and roles, contacts for VO administrators, managers, deputies and experts, specific procedures to be followed by sites to support the VO.
A snapshot of this information is shown in figure 5-1.

Static and dynamic information about e-NMR grid resources are of course collected and published by the gLite information system, whose core component is the top-BDII. It is at the base of the monitoring tools such as SAM, GStat, GridICE that will be described in more details in the next sub-section.
In particular, the top-BDII instance run by CIRMMP centre collects information from all of the e-NMR grid sites, while the instance run by INFN has EGEE scope, so it does include only those e-NMR sites also part of the EGEE production grid. For these sites of course other relevant static information is stored on the GOCDB too.

Very recently, in March 2009, CIRMMP grid site joined the EGEE production grid after following the procedure implemented by the Italian ROC for adding new resources to IGI grid infrastructure. This procedure foresees site's data registration on the GOCDB by the ROC manager, certification of the proper functioning and configuration of the grid site via basic testing, and finally full inclusion of the site into regional and EGEE grid.
Figure 5-2 shows the GOCDB view of the CIRMMP site. To be registered with GOCDB as a certified site in production guarantees its full integration with EGEE grid and its operations. For these sites the whole set of operational tools listed above are available.

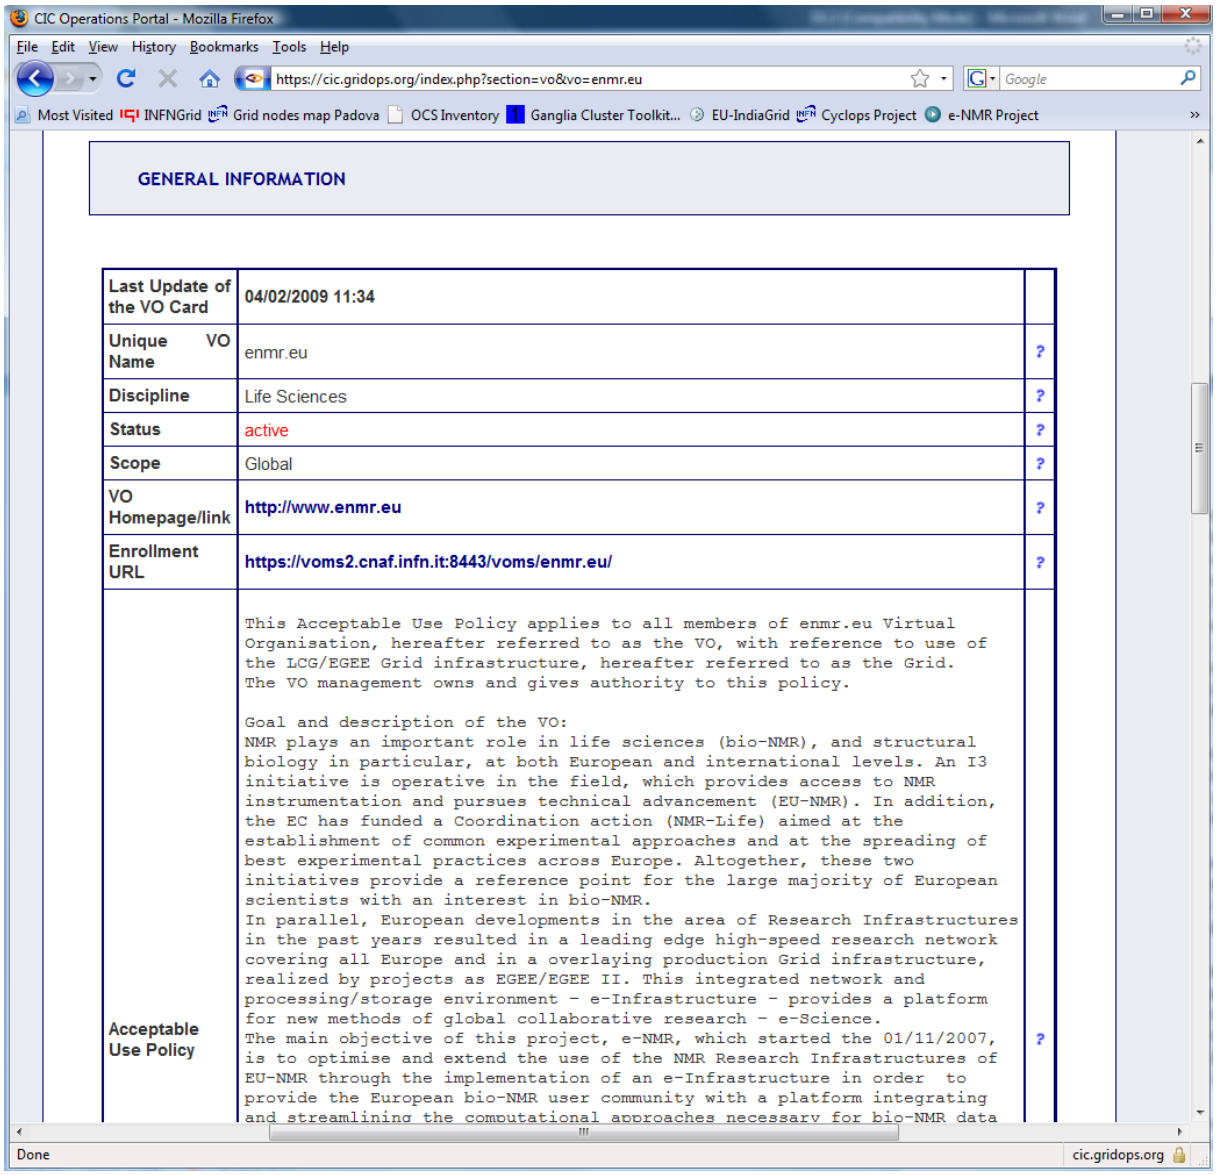

Figure 5-1: enmr.eu VO information card at cic.gridops.org: the EGEE Operations Portal

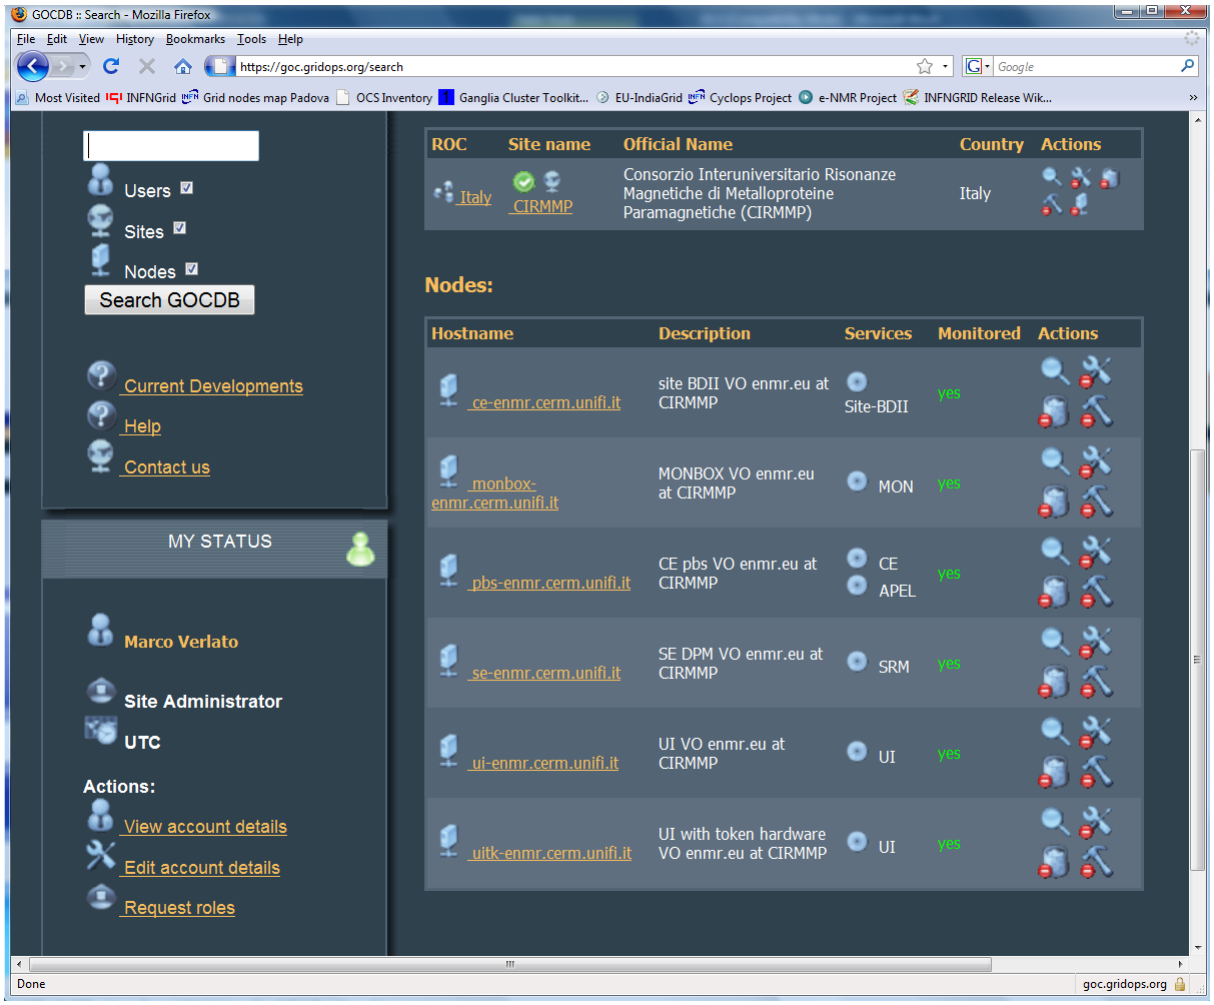

Figure 5-2: CIRMMP resource centre summary information stored in the GOCDB

5.2. Grid monitoring

A set of tools have been designed to monitor proactively the operational state of the grid and its performances, with the goal of initiating corrective actions to remedy problems arising with either core infrastructure or grid resources. Most of the tools were implemented by the Services Activities of the EGEE project during its three two-year phases. Independent instance of them were set up by the e-NMR WP2 team, as described in the following sub-sections.

5.2.1. SAM

The Service Availability Monitoring framework (SAM) [**R8**] aims to provide a site independent, centralized and uniform monitoring tool for all grid services. It is a generalized framework to monitor all EGEE services like CE, SE, BDII, WMS, LB, LFC, etc. It is designed for high-level grid operations and is being used in the validation of sites and services with calculation of availability metrics in the EGEE infrastructure.
Information on the operational state is given by the results of the SAM tests. There are SAM tests for different grid services. Figure 5-3 shows the SAM interface where the test results for different services can be selected. In addition, the region and VO, for which the tests are displayed, can be selected. In the standalone version the region is set as “Default”, since no

GOCDB information is available. In our deployment a site is required to accept the test jobs that are run under the "infngrid" VO. This VO is used for SAM jobs to do not interfere with application jobs. It is deployed and has highest priority on all of the e-NMR grid sites. Currently the tests are run every 3 hours.

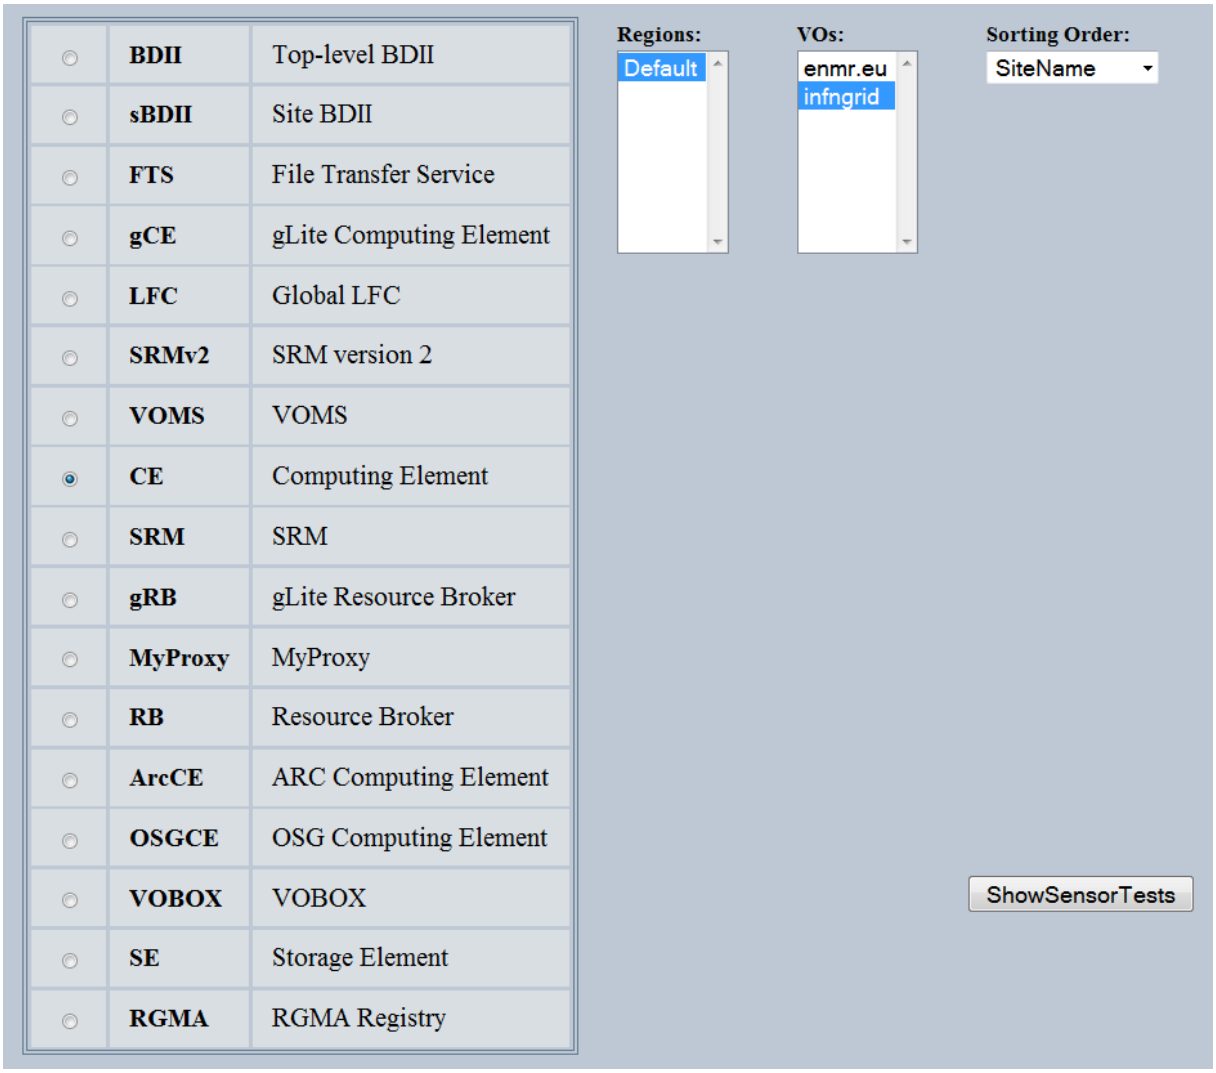

Figure 5-3: SAM GUI for selecting test results of a given grid element

The SAM testing and monitoring framework is made up of a number of components:

- lcg-sam-jobwrapper - SAM jobwrapper
- lcg-sam-client - SAM Client
- lcg-sam-client-sensors - SAM Client Sensors
- lcg-sam-server-db - SAM Database
- lcg-sam-server-portal - SAM Environment Web Portal
- lcg-sam-server-ws - SAME Query Web Service

and the figure 5-4 gives an idea of its working schema.
Notice that Standalone SAM version does not requires sites' registration on GOCDB, while taking all of the needed information from the top-BDII.

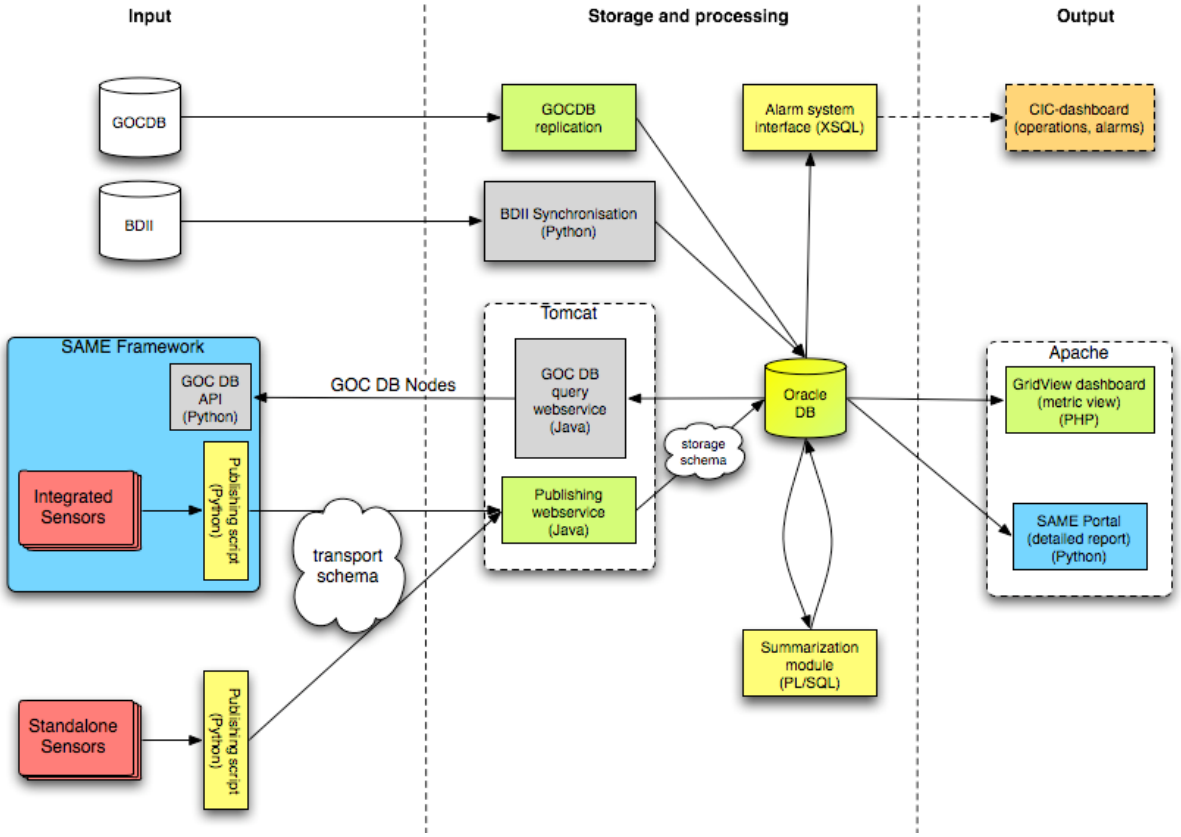

Figure 5-4: SAM architecture

A lot of effort was spent to set up a standalone SAM server as to simplify the sanity-check of the e-NMR grid resources.
The installation and configuration was not as simple as expected.
We needed to make a lot of debugging and made many iterations of information exchange with the SAM developer team to achieve the result.
As a matter of facts, we also opened some tickets in the Savannah-CERN SAM support unit.
It seems that most of the problems came from the controlled environment at CERN that the developer team assumed as a reference for all the other installations.
In details, we encountered the following difficulties installing the Standalone SAM service:

1) Error while deploying SAM database.
Create_probe_interface sql script generates errors during DB configuration (GGUS ticket no 29212). The problem is now reported in Savannah (https://savannah.cern.ch/bugs/index.php?31941) and then later fixed with release of new version of lcg-sam-server-db.
2) Dependency on python m2crypto libraries.
Dependencies were not announced with previous release of SAM web portal (GGUS ticket 29041).
3)Tomcat templates written in tomcat 5.5.
Tomcat templates were written in tomcat v5.5 whereas the repository still had older version of tomcat (v5) giving error while querying to SAM web services (https://savannah.cern.ch/bugs/?33911)
4) Unable to see detailed errors for the SAM test.
It was not possible to see detailed errors in SAM web portal for the SAM test. SAM web portal RPM were without dn.txt file (Savannah bug already reported 33033).

The e-NMR SAM portal is now running on the host ui-enmr.cerm.unifi.it and reachable at

https://ui-enmr.cerm.unifi.it/sam/sam.py. A snapshot is reported here in figure 5-5.

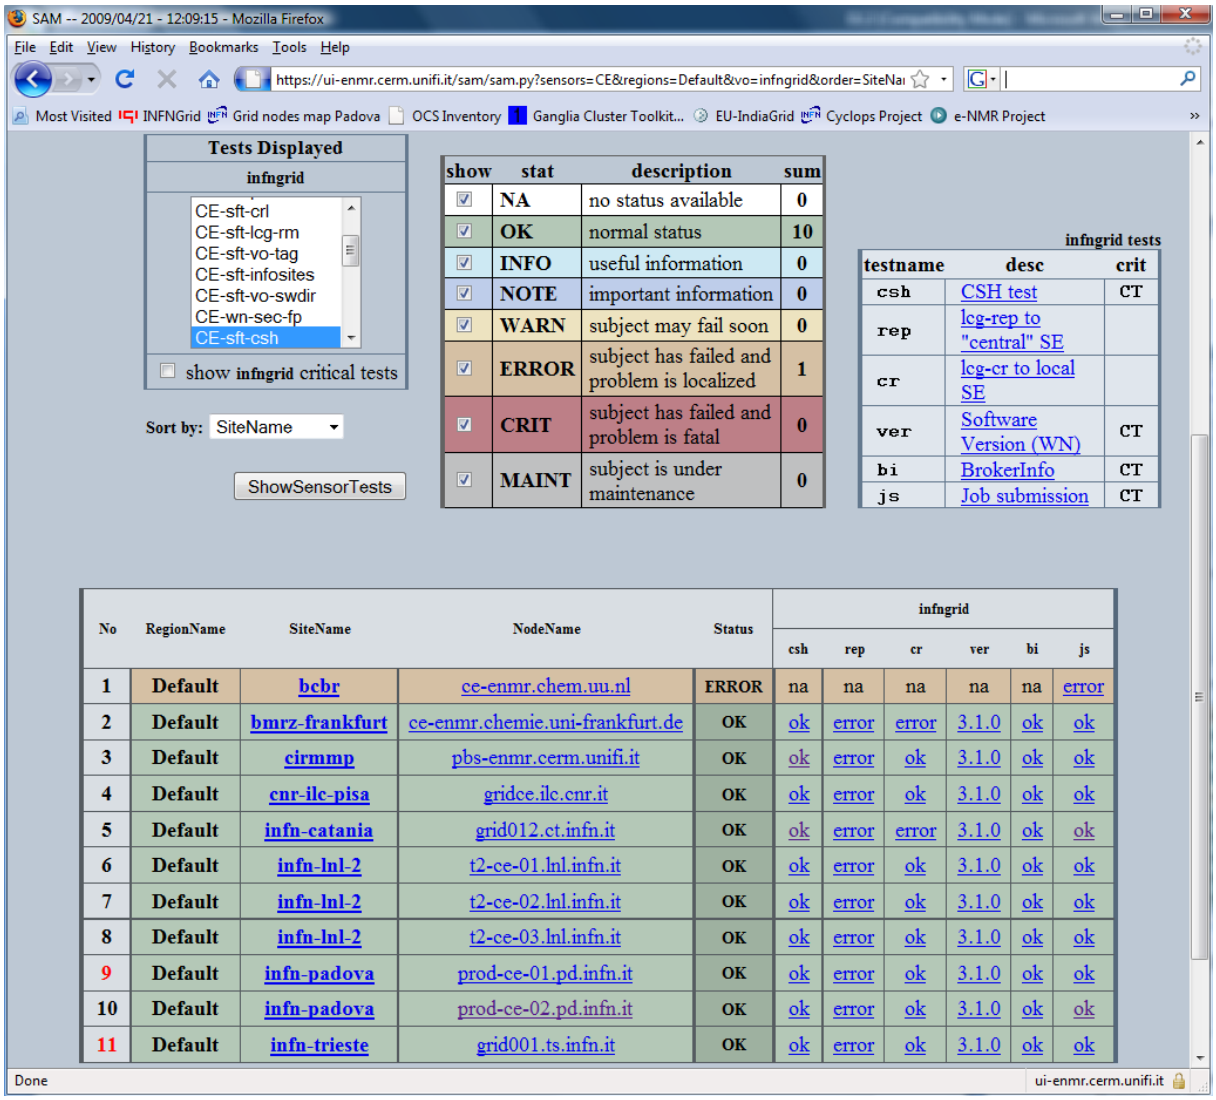

Figure 5-5: SAM test results summary for e-NMR CEs

5.2.2. GridICE

GridICE [**R9**] is a distributed monitoring tool designed for grid systems developed by INFN. It promotes the adoption of de-facto standard Grid Information Service interfaces, protocols and data models. Further, different aggregations and partitions of monitoring data are provided based on the specific needs of different users' categories (VO, GOC, Site). Being able to start from summary views and to drill down to details, it is possible to verify the composition of virtual pools or to sketch the sources of problems. A complete history of monitoring data is also maintained to deal with the need for retrospective analysis.

Its main futures are:

- designed to support the grid operations and VO management activities
- automatic discovery of new resources to be monitored via the Grid Information Service
- powerful and complete web-based interface for data presentation
- notification service

- complete set of monitored metrics, from host-related to grid service related characteristics
- supports and extends the GLUE Schema
- support for the following batch systems: OpenPBS, Torque, LSF
- each view of the web-based interface offers the same data in XML format
- support for customized graph generation
- integrated with network-related infrastructure for monitoring the connectivity of a Grid
- maintain history of metrics and support a per-attribute threshold based mechanism for data reduction
- open source

The GridICE architecture is shown in figure 5-6.

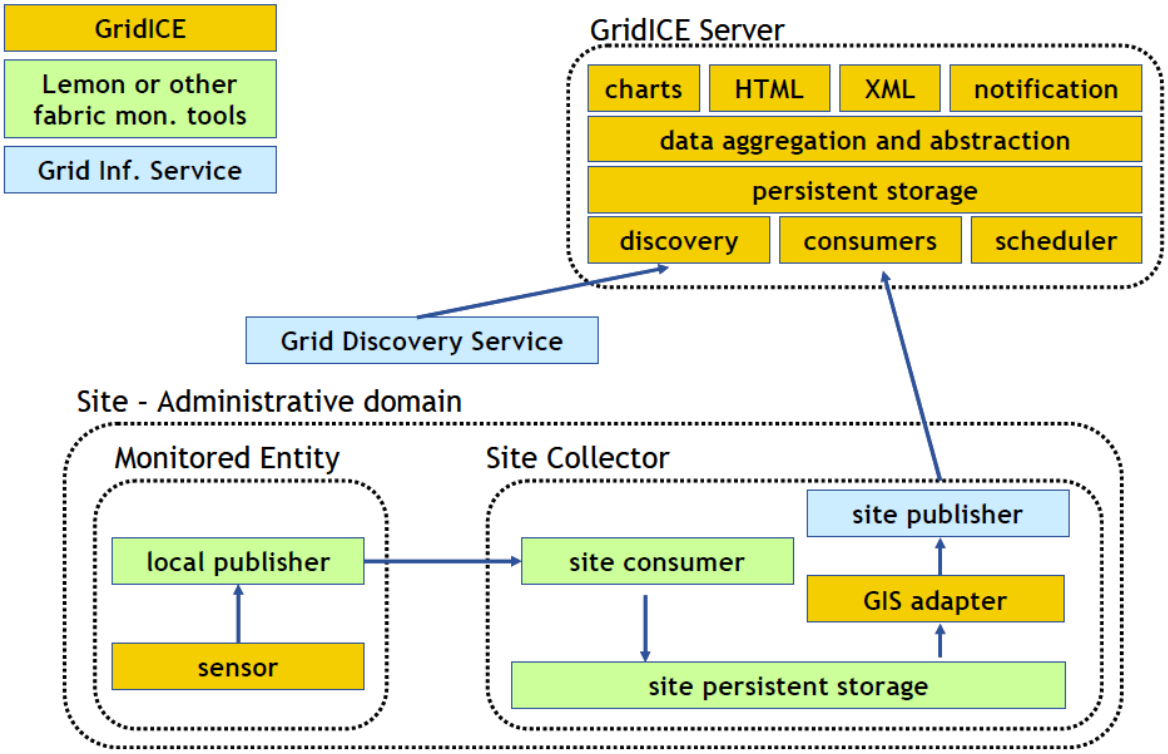

Figure 5-6: GridICE architecture

An instance of GridICE for e-NMR is now running on gridice-enmr.cerm.unifi.it and reachable at http://gridice-enmr.cerm.unifi.it/site/site.php
It is used by site managers to check the status of their resources, and by grid operators of the support team to get the overall status of the e-NMR grid together with the SAM and GStat views. It is often used also at conferences and dissemination events to provide a quick overview of e-NMR grid.
Here below some snapshots of GridICE views: figure 5-7 shows the sites' summary view, figure 5-8 shows the status of CEs available for enmr.eu VO.

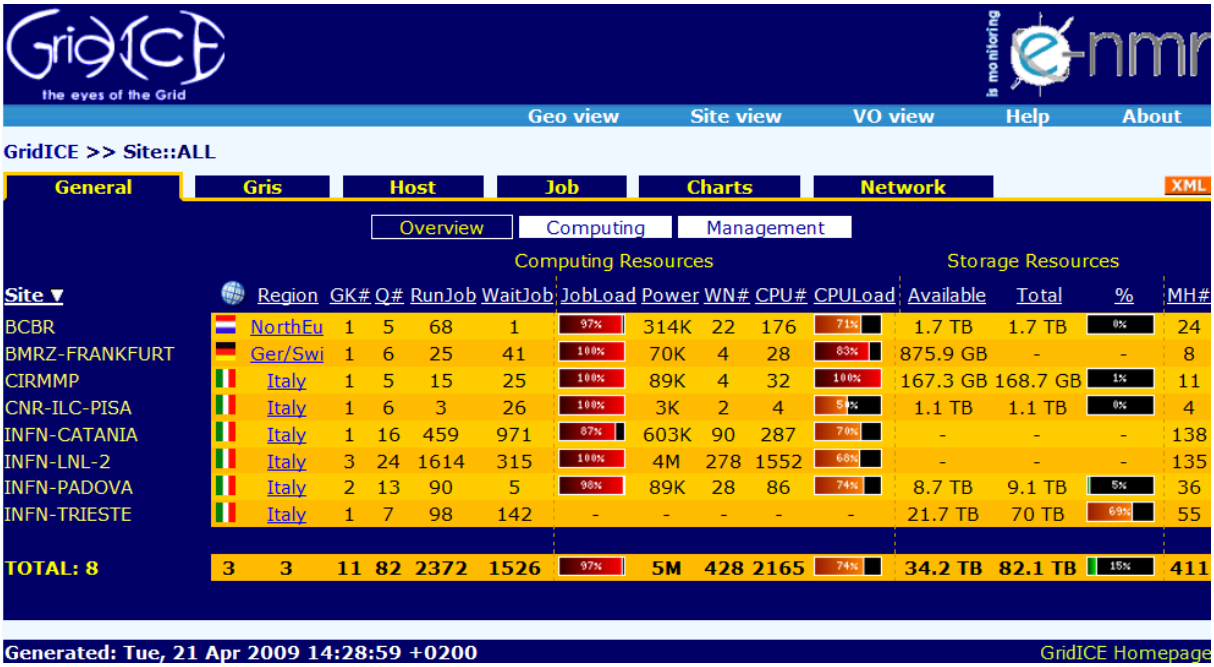

Figure 5-7: All sites' summary view of e-NMR grid

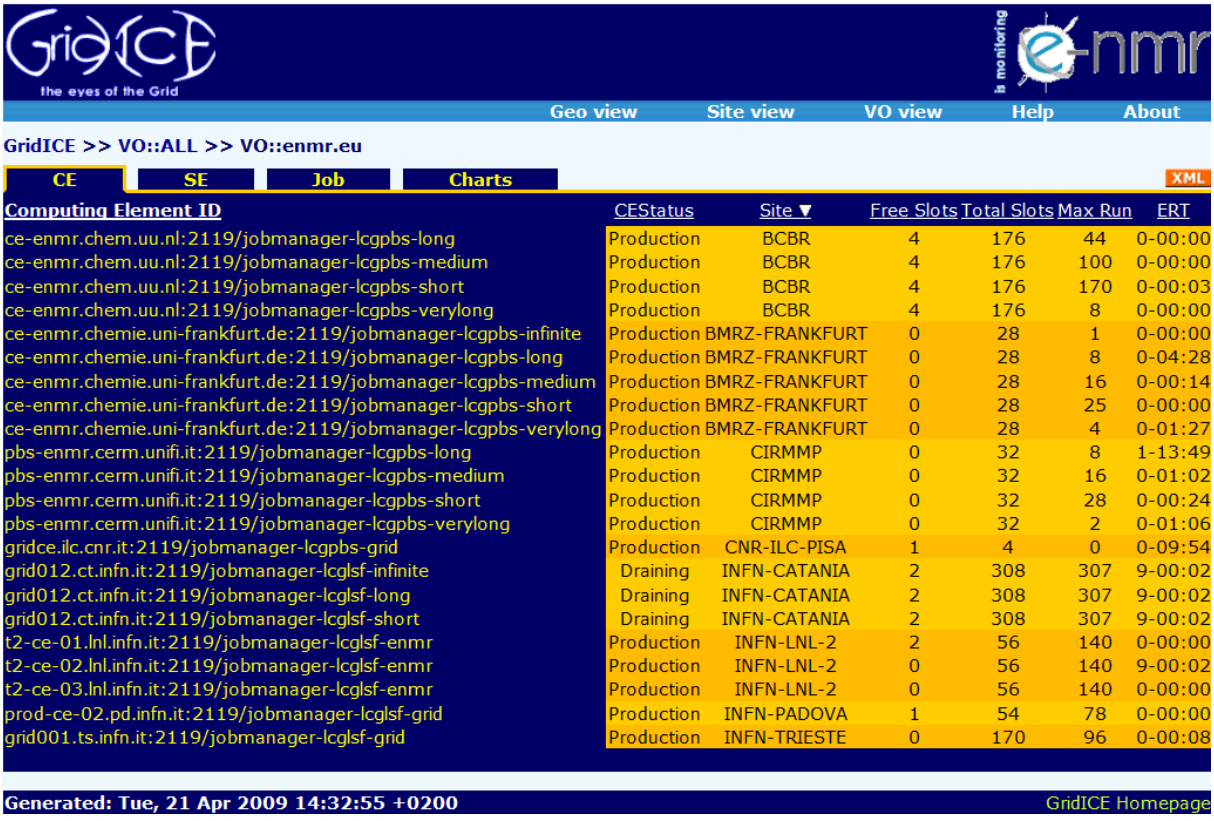

Figure 5-8: Detailed view of CEs supporting enmr.eu VO

5.2.3. GStat

GStat [**R10**] is an application developed by Academia Sinica Grid Computing (ASGC) centre of Taiwan, designed to monitor and check the health of EGEE compatible Information Systems (i.e. LDAP services with data conforming to the GLUE schema). GStat's primary goal is to detect faults, verify the validity and display useful data from the Information System. GStat tests the Information System approximately every 30 minutes. The test does not rely on any submitted job, but rather on queries to site GIISes/BDIIs. This is done to gather information and perform so called sanity checks to point out any potential problems with individual sites. The test covers the following areas:

- Site and service information: provide information about the site, services, software and VOs supported at that site.
- Usage information: provides the statistics on job slots, jobs, storage space.
- Information integrity: checks if the Information system is publishing data that meets specific syntax and value rules.

GStat runs on a single server. From this server, GStat executes queries, process the results and generates static HTML reports. Execution of queries and result processing are accomplished by agents and filters respectively. Currently there are two servers running GStat, at ASGC centre and at INFN-CNAF centre in Italy. Both servers can be accessed from the alias http://gstat.gridops.org/gstat/ .

GStat agents are responsible for making queries and collecting raw data for further analysis by filter components. Filters execute test logic and generate processed data which is in turn used to create a web based test reports. GStat currently stores numeric data in RRD databases with data reduction. Additional historical results can be found in daily snap shot archives. The configuration of GStat heavily depends on the data found in the GOCDB. GStat queries the GOCDB for the site's GIIS contact string, nodes information and other basic site information. However, it can also be used for monitoring sites not yet registered into GOCDB. In that case in fact the relevant information will be taken directly from the configuration file of the top-BDII acting as Information System of the grid that has to be monitored.

Thanks to this feature, there was no need to install a dedicated server for having a GStat instance to monitor the e-NMR grid. It was enough to contact ASGC site managers, who are hosting the master GStat server for the whole EGEE grid, and make them available the link of the top-BDII configuration file collecting all e-NMR grid sites GIISes:
http://beta.cerm.unifi.it/enmr-all-sites.conf
This allowed them to provide automatically an instance of GStat for e-NMR grid, which is reachable at http://goc.grid.sinica.edu.tw/gstat/e-nmr/ .
Its summary page view is shown in figure 5-9 here below.

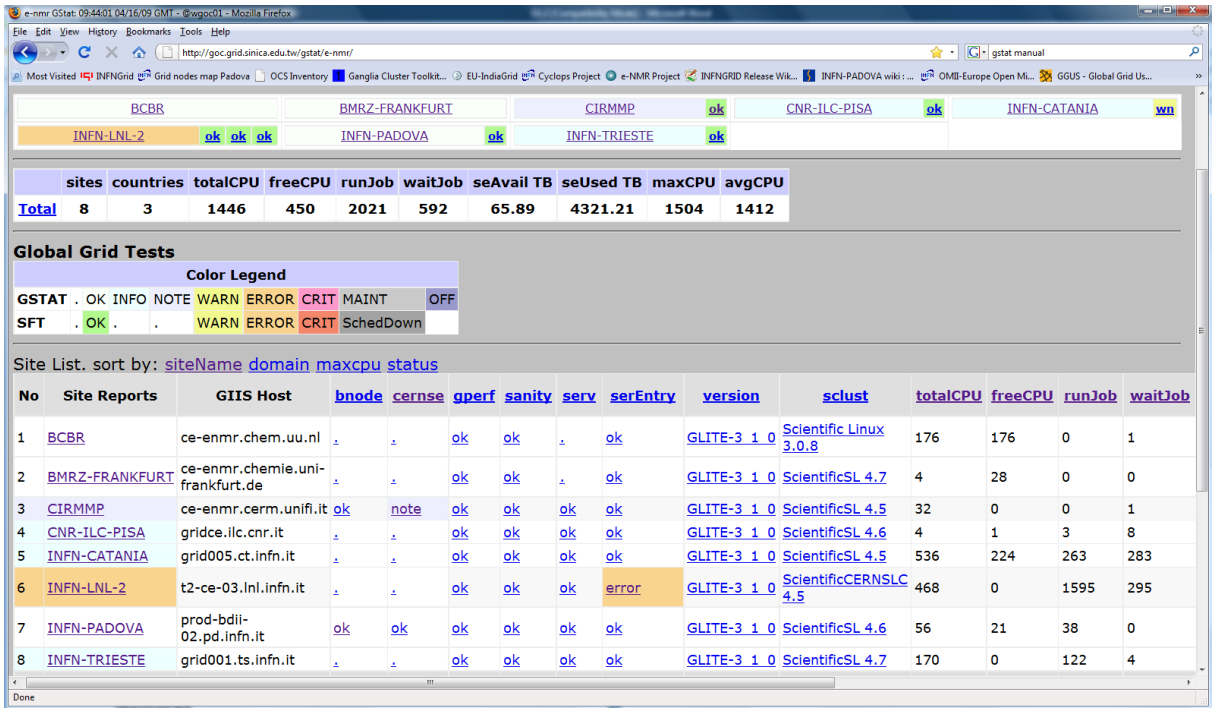

BCBR | BMRZ-FRANKFURT | CIRMMP ok | CNR-ILC-PISA ok | INFN-CATANIA wn
INFN-LNL-2 ok ok ok | INFN-PADOVA ok | INFN-TRIESTE ok

	sites	countries	totalCPU	freeCPU	runJob	waitJob	seAvail TB	seUsed TB	maxCPU	avgCPU
Total	8	3	1446	450	2021	592	65.89	4321.21	1504	1412

Global Grid Tests

Color Legend

GSTAT	. OK	INFO	NOTE	WARN	ERROR	CRIT	MAINT	OFF
SFT	. OK	.	.	WARN	ERROR	CRIT	SchedDown	

Site List. sort by: siteName domain maxcpu status

No	Site Reports	GIIS Host	bnode	cernse	gperf	sanity	serv	serEntry	version	sclust	totalCPU	freeCPU	runJob	waitJob
1	BCBR	ce-enmr.chem.uu.nl	.	.	ok	ok	.	ok	GLITE-3_1_0	Scientific Linux 3.0.8	176	176	0	1
2	BMRZ-FRANKFURT	ce-enmr.chemie.uni-frankfurt.de	.	.	ok	ok	.	ok	GLITE-3_1_0	ScientificSL 4.7	4	28	0	0
3	CIRMMP	ce-enmr.cerm.unifi.it	ok	note	ok	ok	ok	ok	GLITE-3_1_0	ScientificSL 4.5	32	0	0	1
4	CNR-ILC-PISA	gridce.ilc.cnr.it	.	.	ok	ok	ok	ok	GLITE-3_1_0	ScientificSL 4.6	4	1	3	8
5	INFN-CATANIA	grid005.ct.infn.it	.	.	ok	ok	ok	ok	GLITE-3_1_0	ScientificSL 4.5	536	224	263	283
6	INFN-LNL-2	t2-ce-03.lnl.infn.it	.	.	ok	ok	ok	error	GLITE-3_1_0	ScientificCERNSLC 4.5	468	0	1595	295
7	INFN-PADOVA	prod-bdii-02.pd.infn.it	ok	ok	ok	ok	ok	ok	GLITE-3_1_0	ScientificSL 4.6	56	21	38	0
8	INFN-TRIESTE	grid001.ts.infn.it	.	.	ok	ok	ok	ok	GLITE-3_1_0	ScientificSL 4.7	170	0	122	4

Figure 5-9: GStat view of the e-NMR grid

Notice that, being the GStat production server connected anyway with the GOCDB, for those sites of e-NMR part of EGEE grid is also published the information coming from SAM tests operated by EGEE (look at the green boxes in the picture).

5.2.4. WMSMonitor

WMSMonitor [**R11**] is a monitoring tool developed by INFN collecting information from a cluster of distributed WMS instances about service status and Job Flow. Collected information is aggregated and presented exploiting two main keys: single WMS instance information and VO information.
These two information aggregation keys lead to two parallel branches of metrics presentation in the GUI:

- WMS view focusing on WMS cluster status and performance monitoring. In the first branch we find a main page reporting an overview of most important status variables for all WMS instances monitored through which we can access a Specific WMSLB Instance Detailed data page with detailed data for each monitored WMS instance including Job flow rates between internal components other than system status indicator measurements.
- VOView providing Job flow statistics using the VOs as aggregation criterion.

Starting from WMSMonitor version 2.0, clicking on "WMS view" or "VO View" tabs, a window menu let you navigate through all available WMSMonitor web pages for each WMS instance or VO respectively. A valid X509 personal certificate and Flash software is required to be installed on the browser to access the web interface.

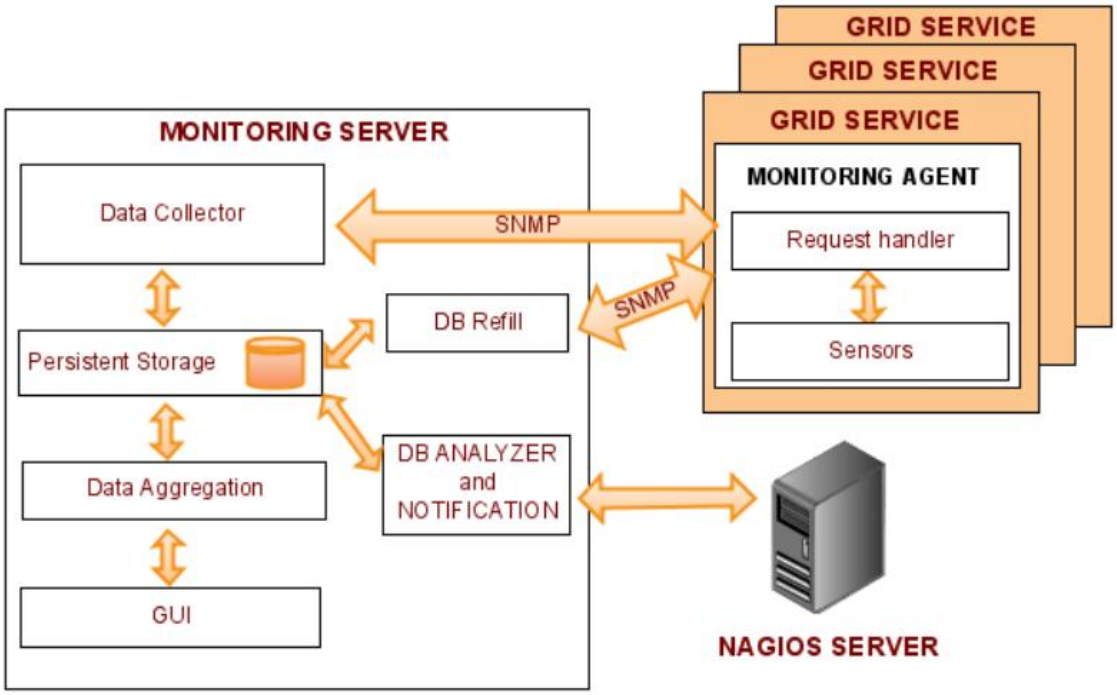

Figure 5-10: Schema of WMSMonitor architecture

WMSMonitor, whose architecture is shown in figure 5-10, is composed by the following main components:

- WMSMonitor Server: WMSMonitor runs on a dedicated server machine with a Mysql database to store data, PHP, Apache and Python on it.
- WMSLB Instances: a cron on the WMSMonitor server collects data from each WMS/LB instance executing sensors on them. These sensors send data to the WMSMonitor server using snmp. Job flow rates are calculated using direct queries to the LB database.
- WEB Publishing: exploiting PHP based functions; data are retrieved from a local Mysql DB, aggregated and published on a graphical user interface through a (secure) http protocol. Charts are implemented exploiting "Open Flash Chart", an open source Flash based set of libraries.
- DB Analyzer: a Python daemon that checks the latest status of the cluster and sends notification to a Nagios server providing mail and/or sms alert notifications.
- DB Re-filler: it can happen for various reasons that data for some jobs are not available on the LB database when the sensors are triggered on the LB servers, resulting in a loss of data in the WMSMonitor Server DB. The data collector does its best trying to fix the situation, but if data does not reach the LB server in a reasonable time (about ten hours by default, a configurable timeout) the collector gives up and a manual reconstruction is needed. This utility performs such a reconstruction.

An instance of WMSMonitor has been set up at INFN-Padova centre, and is reachable at https://eu-india-01.pd.infn.it:50080/wmsmon/main .

Figure 5-11 shows the detailed page for the WMS run at CIRMMP taken the 17th of April 2009. The text boxes show the status, load and statistics for each component of WMS and LB services in the last 15 minutes, while the histograms summarize the job flow and final states achieved daily in the last two weeks.

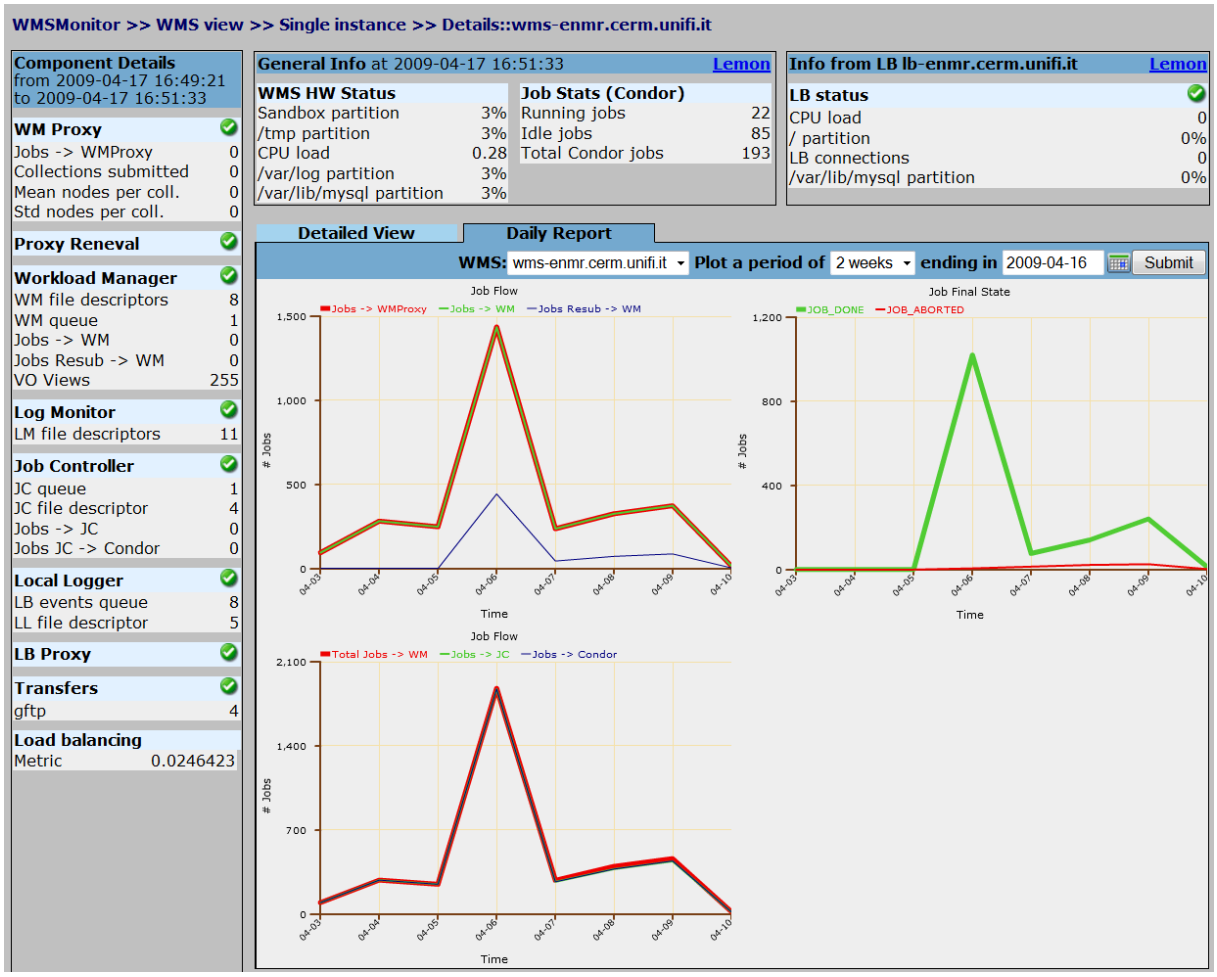

Figure 5-11: WMSMonitor statistics and job flow graphs for e-NMR WMS/LB services.

5.2.5. GridView

GridView [**R12**] is a visualization tool developed by the Bhabha Atomic Research Centre (BARC) of Mumbai, India, that provides a high level view of various functional and performance aspects of the EGEE infrastructure. GridView works by collecting various sources of raw monitoring information, called the 'raw data', from grid services at different sites. It analyzes this data and computes various derived metrics. The sources of raw data for GridView and the set of derived metrics are the following:

- GridFTP service logs: average throughput and aggregate data movement across sites computed site-wise and VO wise for different periodicities
- Resource Broker logs: site wise and VO wise job statistics for different job states and computation of average turnaround time and job success rates
- FTS service logs: average throughput and aggregate file transfer activities across sites, computed site-wise and VO wise including identification of 7 day and 12 hour peak transfers
- SAM test results: computation of service and site availability and reliability for different periodicities

Regarding visualization of the data, GridView provides a web based user interface to display these computed data by means of graphs and charts. A report generating module is available that can generate monthly site availability and reliability reports.
We decided to do not install a dedicated instance of GridView for e-NMR grid, being the others monitoring tools described above considered sufficient for an adequate GOC activity. Nevertheless the main EGEE-wide instance reachable at http://gridview.cern.ch/GRIDVIEW can be exploited to monitor those e-NMR grid sites already part of EGEE. Of particular interest are actually the plots showing the availability of grid services computed on the basis of the SAM test results. Service availability metrics are defined in [**R6**], and are measured as the fraction of time the service is properly working, i.e. it successful pass the corresponding SAM test.
Figure 5-12 shows the global site availability in the month after the 17th of March 2009 for the 6 out the 8 grid sites currently forming the e-NMR grid.

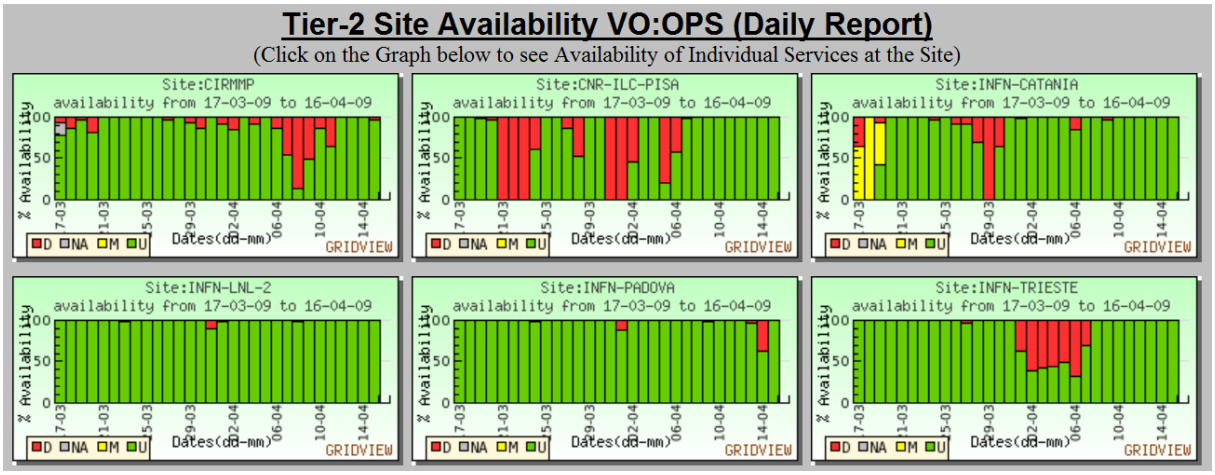

Figure 5-12: Daily availability for e-NMR sites belonging to EGEE grid

Figure 5-13 shows the availability of each grid service hosted at CIRMMP in its first month of operations after having joined the IGI / EGEE production grid the 17th of March 2009. This tool is typically used to enforce the Service Level Agreement signed by the local grid site with its coordinating operations centre, as described e.g. in the EGEE-II Service Level Description document [**R16**].

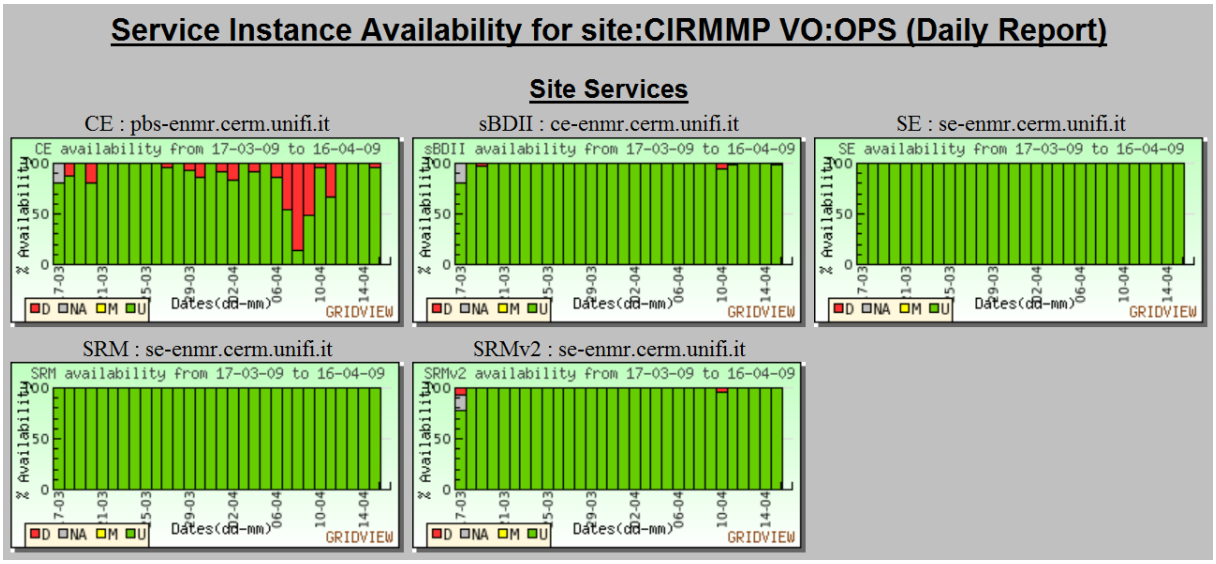

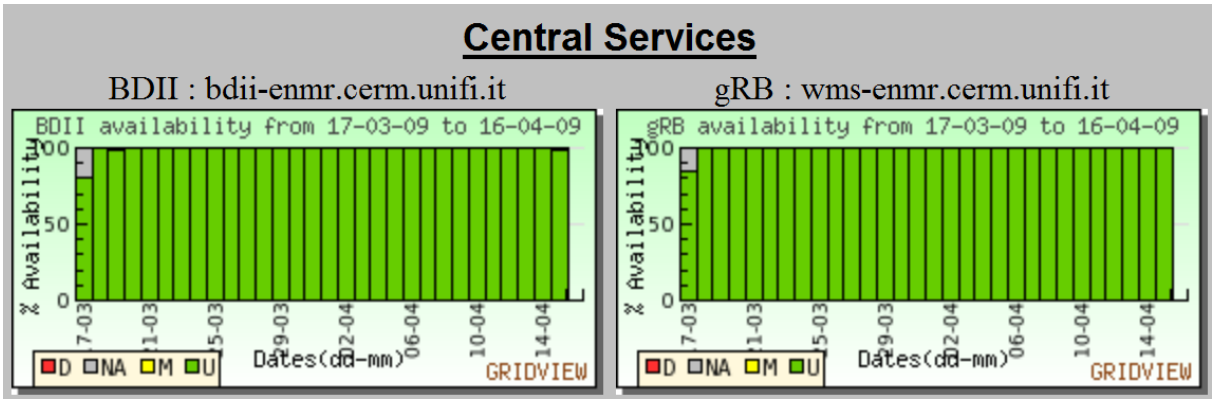

Figure 5-13: Grid services availability for CIRMMP site

5.3. Grid accounting

5.3.1. DGAS

The Distributed Grid Accounting System (DGAS) [**R13**], formerly called DataGrid Accounting System, was originally developed by INFN within the EU Datagrid Project (until March 2004) and later maintained and re-engineered within the EGEE project. The DGAS developers' team also participate in the OMII-Europe project for designing an OGF Resource Usage Service (RUS) interface for DGAS.
Currently the DGAS system is deployed on the IGI production grid, and is one of the four accounting systems available for grid infrastructures, the others being APEL (used in EGEE), SGAS (used within the North European countries), and Gratia (used in the Open Science Grid US project). All of the four systems are interfaced with the EGEE Accounting Portal, a graphical front-end to accounting data. DGAS has also its own graphical front-end called HLRmon.
The Purpose of DGAS is to implement Resource Usage Metering, Accounting and Account Balancing (through resource pricing) in a fully distributed grid environment. It is conceived to be distributed, secure and extensible.
The system is designed in order for Usage Metering, Accounting and Account Balancing to be independent layers, as shown in figure 5-14.

Usage Metering is done by lightweight sensors (DGAS Gianduia) installed on the Computing Elements. These sensors parse PBS/LSF/Condor event logs to built Usage Records that can be passed to the accounting layer (by DGAS Pushd). For a reliable accounting of resource usage (essential for billing) it is important that the collected data is unequivocally associated to the unique grid ID of the user (certificate subject/DN), the resource (CE ID) as well as the job (global job ID). Furthermore, usage records, when sent to the accounting layer, are signed using the user's certificate proxy in order to make sure that usage records are always associated to the user that submitted a job, thus preventing the accounting system from being vulnerable to spoofed usage records.

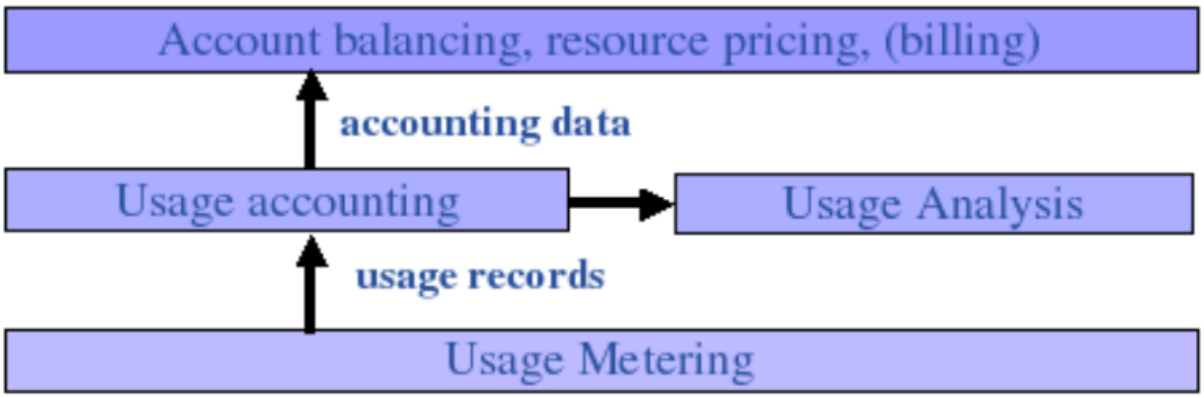

Figure 5-14: Layers of DGAS architecture

The usage of grid resources by grid users is registered in appropriate servers, called Home Location Registers (HLRs) that manage both user and resource accounts.
In order to achieve scalability, accounting records can be stored on an arbitrary number of independent HLRs.
Each HLR keeps the records of all grid jobs submitted or executed by each of its registered users or resources, thus being able to furnish usage information with many granularity levels: per user or resource; per group of users or resources; per VO.
Accounting requires accurate usage metering, but not necessarily resource pricing and billing. In fact, the upper layer of figure 5-14 is optional (the DGAS metering and accounting infrastructure can work without it) and may be used to implement economic models based on a market for grid resources, that may aid in balancing demand and supply as well as balancing the incoming workload. It requires the deployment of an additional component called Price Authority (PA).
Several important issues are addressed by DGAS:

- Privacy:
  - ✓ all communications are encrypted;
  - ✓ only authorized access to accounting data (users, resource administrators, VO administrators) is allowed
- Security/Reliability:
  - ✓ Authorization of the accounting/payment transaction (indirectly) by the user.
  - ✓ Usage records transmitted from CE to HLR using the user's proxy certificate.
  - ✓ Usage records stored by both User HLR and Resource HLR.
  - ✓ User HLRs accept usage records only for registered users.
  - ✓ Resource HLRs accept transactions only for registered resources and only from trusted User HLRs.
  - ✓ Usage Record transmissions and transactions between HLR are asynchronous and in case of failures (e.g. temporary network problems) are retried.
- Scalability:
  - ✓ Decentralized infrastructure with an arbitrary number of HLRs/PAs.

Figure 5-15 shows the current DGAS servers deployment adopted by e-NMR grid. It exploits a number of HLR servers providing the DGAS accounting for the IGI production grid.
A second level HLR hosted at INFN-Torino does collect Usage Records from all of the IGI site-level HLRs hosted by grid sites supporting the enmr.eu VO. Among them, CIRMMP site-level HLR, dgas-enmr.cerm.unifi.it, does collect the Usage Records from CIRMMP, BCBR and BMRZ CEs.

The visualization tool HLRmon described in the next sub-section performs queries to the second level HRL.

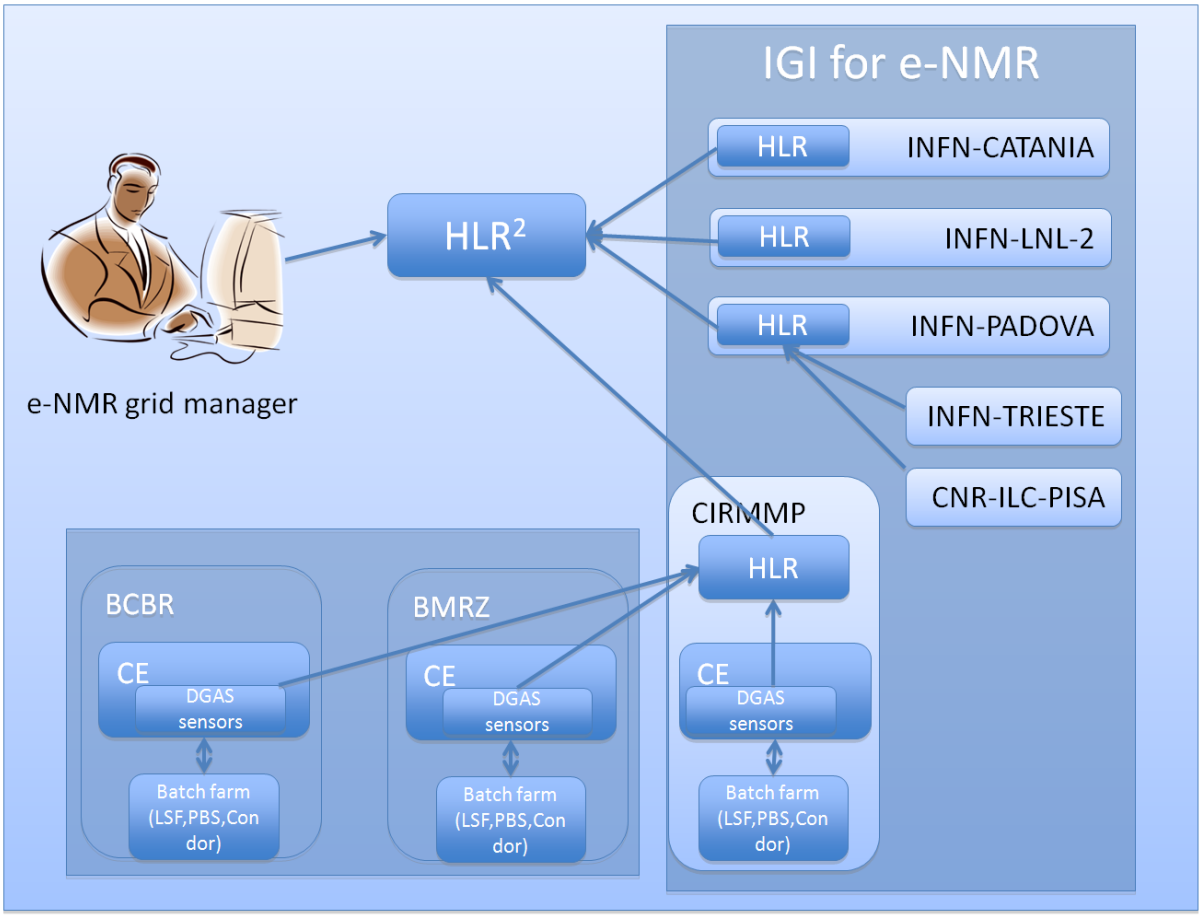

Figure 5-15: Overview of DGAS components deployed for serving e-NMR grid

5.3.2. HLRmon

HLRmon, the visual web component of DGAS, enables access to accounting information for managers, administrators and users [**R14**].
HLRmon gathers the accounting information via the DGAS command line tools to query the HLR (Home Location Register) accounting database and displays a variety of accounting information such as Jobs per VO, Wall time per VO, CPU time per day, etc.
Information is presented in a graphical or tabular form in order to satisfy different user needs.
Due to its ability to authenticate clients through certificate and related authorization rights, it can a-priori restrict the selectable items offered to the web user, so that sensitive information will only be provided to specifically authorized people.
Access to the accounting information, provided by HLRmon, is granted to users whose identity is authenticated through a valid digital certificate, which must be released by a trusted Certification Authority and installed in the client's browser.
HLRmon matches the user's identity with his access membership privileges and preferences. This allows one to directly offer activity reports for the proper subset of information related to group membership, and persistently keep track of user's preferences (such as date interval or other selections) through subsequent visits.

The following four different groups have been defined:

- ROC Manager: can access all available information about activity for every Site, VO and VO user.
- Site Manager: can access all available information about activity for every VO and VO user but only at sites where the Site Manager role is granted for the HLRmon user.
- VO Manager: can access available information about activity of VOs and VO users across all sites in the ROC, but only on VOs where the VO Manager role is granted for the HLRmon user.
- VO user: can access available information about VO user activity only for VOs where the VO user role is granted for the HLRmon user. This is the only role that cannot access other information activity but its own.

An instance of HLRmon for e-NMR grid has been set up and is reachable at https://dgas.cnaf.infn.it/hlrmon_enmr .
It is used by the e-NMR grid manager to produce resource usage reports like the ones shown in the previous section 4.

5.4. Middleware deployment and introduction of new resources

The INFNGRID Middleware Release (IG Release) is the reference distribution in IGI. IG Release is based on the gLite Middleware Release and integrates it with additional middleware components of interest to support both regional and international VOs in Italy. It has been chosen for e-NMR grid in order to simplify the deployment of gLite middleware at the three NMR laboratories which were completely unaware of grids at the beginning of the project. The choice had two advantages: the enablement of DGAS accounting, which is not included in the gLite original release; the guaranteed support of IG Release Team, mainly hosted at INFN-Padova, the leading INFN centre involved in the e-NMR project.

In a typical grid site, installing a large number of computers with the same characteristics (e.g. a set of Worker Nodes) is a common issue. In this case the IG Release Team provides the site with an automatic method to setup the farm, where each node has the same basic configuration, with the aim to simplify the grid middleware deployment. The procedure is provided through documents [**R15**] describing how to setup a server for RedHat OS installation via Kickstart using PXE technology and to manage a local YUM repository for OS and middleware. To manage the package repository a customized version of MREPO, a tool to mirror RPM repositories, is used, while the OS installation is based on normal Kickstart files. The tool YAIM is used to install/upgrade the middleware release.

The gLite middleware stack is based on a few core components, which define a given main release. If these core components change, it requires the whole stack to be rebuilt against them. In the case where a core component changes a new release will have to be made. This will result in a new baseline with which all components must work. The current main release is gLite 3.1. There is however a continuous release process where subsystems and components can be updated individually. Thus in gLite we have a baseline release to which updates are continuously added. This happens on average every two weeks. IG Release team typically tests, packs and distributes the updates in a combined IG Update about every month.

Having completed the procedures described in the documentation, the site administrators can ask the e-NMR WP2 “resource and user support” team to include their site on the Information

System and submit certification jobs to verify that resources have been properly configured. This can be done without interfere with the normal grid workload using a dedicated queue set up to allow access only to a testing purpose VO (currently the "infngrid" VO).
BCBR, BMRZ and CIRMMP sites followed these procedures to be included into e-NMR, as recalled in section 2.

In March 2009 the CIRMMP site joined IGI grid, the Italian ROC of EGEE, following a procedure analogous to the one described above, which includes the signature of the Service Level Description document [**R16**] with IGI, establishing a minimum of service level provisioning and behaviour rules acceptance.

At the 4th EGEE User Forum Conference held in Catania, Italy, the 2-6 March 2009, contacts were taken with Benelux ROC and Germany/Switzerland ROC managers, in order to prepare the joining to EGEE of respectively BCBR and BMRZ grid sites.
Both ROCs provided the e-NMR management with letters of support for the project, and are planning to include these sites in their relevant NGIs before the end of the project.
In particular, at the end of April 2009, BCBR site became officially supported by BiG Grid, the Dutch e-Science Grid, and initiated the process of becoming an EGEE site.
This commitment is of course of great importance for ensuring sustainability of the infrastructure after the end of EGEE-III and e-NMR projects.

At the same conference the South African National Compute Grid (SAGrid) [**R17**] manager has shown great interest in e-NMR project. SAGrid is deploying production high-performance grid computing services to South African research institutes and universities using the gLite stack from EGEE. It is currently supporting four scientific domains: Life Sciences, High Energy Physics, Nuclear Medicine (radiotherapy), and Astronomy/Astrophysics.
The interest into e-NMR was formalized with a letter of support (in annex), which stated the support of the Bioscience Unit of the South African Council of Scientific and Industrial Research (CSIR), and the offer of sharing with enmr.eu VO computing resources from the CSIR Cluster Computing Centre.

SAGrid is deploying on its sites the IG Release, so their integration into e-NMR grid is expected to be quite straightforward. A 7 steps work plan was agreed with the SAGrid Operations group as follows:

1) virtual site installation at University of the Free State (UFS) in Bloemfontein or at University of Cape Town (UCT), test integration into e-NMR central services
2) repeat installation at CSIR centre, validate configuration
3) applications installation, tag and testing on virtual site
4) repeat and validation at C4, final application configuration
5) run pilot jobs to test infrastructure (at this point we will have an evaluation of the procedure to discuss major issues)
6) CSIR centre into e-NMR production grid
7) repeat at any site which wants to support e-NMR VO

We expect that some of the most stable sites of SAGrid will join the e-NMR grid in summer 2009.

5.5. Resource and user support

Resource and user support is of paramount importance for any reliable grid infrastructure and one of the main goals of a GOC. Operational support to both grid users and site managers has to be assured. Grid sites not properly configured or unattended are often responsible of job's abortion, heavily damaging the task's execution efficiency. Being typically user's expectations from grid technologies quite high, low Quality of Service is often leading to frustration. User support requests have to be promptly addressed indeed.

A wiki page has been set up by INFN partner at the beginning of the project with two guides containing a quick description and the main links for more detailed documentation: the "Quick gLite Deployment Guide" and the "Quick gLite User Guide". They are accessible following the link under "wiki" menu of the project's web portal [**R18**].

A team of e-NMR funded personnel has been set up at INFN with the task of checking daily the status of e-NMR grid using the monitoring tools described in the previous sections. Due to the limited number of people involved in grid operations in the first half of the project, support requests from site administrators to grid experts are handled through the appositely set up electronic mailing list enmr-sitemanagers@lists.infn.it . This is a closed mailing list which includes all of the site managers of the three NMR centres. It was set up in January 2009, and has an activity of about 30 mails per month. The e-NMR/INFN support team promptly notifies via the mailing list e-NMR site managers about local grid issues or malfunctioning, in order to minimize the impact of badly working sites on job efficiency. The site affected by the problem is indicated on the subject of the e-mail, but the e-mail it is sent to the full list to make all site managers aware of the kinds of problems which can arise on e-NMR grid, so they can learn each others. The list has an archive allowing the tracking of problems for further reference.
For those sites of e-NMR which belong to EGEE too (all located in Italy until now), the IGI ticketing system is used for communication with their site managers. When sites from others countries will join the e-NMR grid, the use of the Global Grid User Support (GGUS) system [**R19**] will be the natural choice for communication.
The e-NMR/INFN team is also taking care of coordinating and ensuring the proper implementation at the sites of the periodic IG middleware updates happening nearly every month.

Concerning user support, the e-NMR/INFN support team has among its duties also the one of answering support requests coming for grid users.
More in general, any question or feedback about the services offered by e-NMR project can be asked through the "Contact" menu of the project's portal, shown in figure 5-16.
It has to be noticed however that, being the e-NMR end-user typically using the grid via the application portals, most if not all of the grid complexity is hidden to him, so that most of his requests are filtered out by the application portal providers.

5-16: User support page at www.enmr.eu project portal: for each area the link opens a web form to be filled with the question, that is then submitted to the relevant key person/team.

6. Conclusions

The present document has reported about the WP2 task T2.2 activities in the first half of the e-NMR project lifetime.
The current status of the e-NMR grid infrastructure, the use of its computing resources, and the organisation of its Grid Operations Centre have been illustrated, six months after the official opening of e-NMR platform to bio-NMR community outside the project consortium.

It is known that during EGEE-III phase the operational tools that currently rely on central deployment and operations will have to be adapted to allow the regional operation of tools by NGIs. It is in fact one of the goals of EGEE-III to enable smooth transition towards a European Grid Initiative (EGI) which would come into effect in 2010. This will be achieved following well defined operations automation strategy described in [**R20**].
WP2 task T2.2 will closely follow the activities of the Operations Automation Team of EGEE-III/SA1 group in order to extend the current e-NMR GOC with new or modified components, whenever required.

7. Annexes

7.1. Letter of support from SAGrid

Meraka Institute
(African Advanced Institute for Information
and Communications Technology),
Council for Scientific and Industrial Research
PO Box 395
Pretoria 0001
South Africa

Letter of Interest in Joining the e-NMR Network

Date: 12 March 2009

To: Prof. Dr. Harald Schwalbe
Coordinator, e-NMR
Zentrum für Biomolekulare
Magnetische Resonanz
Goethe Universitaet Frankfurt
Max-von-Laue-Str. 7
N160, 314, 3. Stock
D-60438 Frankfurt

Dear Professor Schwalbe,

At a recent meeting with your colleagues at the EGEE User Forum in Catania, we discussed the possibilities for joining the EU FP7 e-NMR network, in terms of collaboration at a scientific, as well as computing resource level.

As coordinator of the South African National Grid (SAGrid), I hereby state my full support to the e-NMR project. The scientific merit of the work supported by the e-NMR project would is highly attractive and we consider that scientific collaboration using the grid and NMR infrastructure would yield several important advances in the field of structural biology. The project would be potentially extremely useful to the

scientific endeavour of the Biosciences Unit of the Council for Scientific and Industrial Research (CSIR), represented by Dr. Colin Kenyon.

The grid computing infrastructure in South Africa is compatible with that employed by e-NMR. In the event that scientists from South Africa join the e-NMR Virtual Organisation, we would like to contribute computing resources to the collaboration via the middleware, at those sites in the SAGrid federation which choose to support the VO. Most notably, we would offer the computing resources of the CSIR Cluster Computing Centre and encourage further sites in the federation to provide storage and computing resources as required.

Yours sincerely,

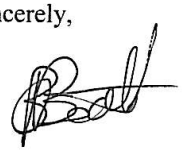

Bruce Becker
SAGrid National Coordinator,
Meraka Institute

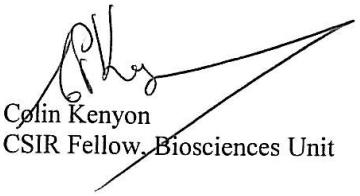

Colin Kenyon
CSIR Fellow, Biosciences Unit

Albert Gazendam
HPC research group leader
Meraka Institute, CSIR