PL-Grid: Foundations and Perspectives of National Computing Infrastructure

Jacek Kitowski[1,2], Michał Turała[2], Kazimierz Wiatr[2], and Łukasz Dutka[2]

[1] AGH University of Science and Technology, Faculty of Electrical Engineering, Automatics, Computer Science and Electronics, Department of Computer Science, al. Mickiewicza 30, 30-059 Kraków, Poland

[2] AGH University of Science and Technology, ACC Cyfronet AGH, ul. Nawojki 11, 30-950 Kraków, Poland

Abstract. The Polish Grid Initiative commenced in 2009 as part of the PL-Grid project funded within the framework of the Innovative Economy Operational Programme. The main objective is to provision a persistent heterogeneous computational platform for the Polish scientific community with a unified interface, enabling easy access to the distributed large-scale computing infrastructure. The project establishes a countrywide computing platform which supports scientific research through integration of experimental data and results of advanced computer simulations carried out by geographically distributed computer infrastructure and teams. The solutions applied in setting up this e-infrastructure facilitate integration with other, similar platforms around the world.

In this chapter the foundations of the PL-Grid architecture, a brief history of the project and its most important objectives are presented.

Keywords: Grid, e-Infrastructure, PL-Grid Consortium.

1 Introduction

Over the past several years we have witnessed rapid increases in the dependence of scientific research on access to large-scale computing systems. Theoretical sciences rely on computing to evaluate hypotheses in many areas, including mathematics, economics, astrophysics, chemistry, climatology, physics, etc. Computerized simulations enable researchers to test their ideas and draw initial conclusions before moving on to the much more expensive and time-consuming experimentation phases. Moreover, experimental and natural sciences, including climatology, meteorology, high energy physics, and astrophysics generate enormous amounts of data which needs to be processed on several levels and distributed to research teams all over the world.

Many of the largest research initiatives are driven by international institutions such as CERN (European Organization for Nuclear Research), EMBL (European Molecular Biology Laboratory) or ITER (International Thermonuclear Experimental Reactor) which can share their results with geographically distributed research teams, increasing their demand on computational power and storage resources.

M. Bubak, T. Szepieniec, and K. Wiatr (Eds.): PL-Grid 2011, LNCS 7136, pp. 1–14, 2012.

Fortunately, at the same time we have observed unsurpassed growth in the computational power and storage capabilities of high-end computing solutions, ranging from simple clusters through cluster platforms with advanced interconnects to custom-built supercomputers with peak performance in the PFlop range (10^{15} floating-point operations per second). In order to efficiently harness such computational power special software solutions are necessary. Research scientists need to be able to actually use the resources in a transparent manner, without having to deal with the issues of scalability, distribution, data access and security.

Currently the most popular paradigm for provisioning distributed computing resources for research is grid computing – a term coined in 1999 [1,2]. The initial idea was to develop a global computational infrastructure which would be as readily accessible as the power grid and whose users would form Virtual Organizations (VO) based not on their geographical location but on their respective research domains. Over the last decade, several such infrastructures have been developed all over the world – most notably the LHC Computing Grid [3], TeraGrid [4], EGEE [5] and DEISA [6]. Most grid deployments are based on one of the following frameworks: Globus [7], UNICORE [8], Condor [9] or ARC [10], all of which evolved in parallel and provide comprehensive sets of components and features, allowing users to run their computational jobs on federated computing resources. The emergence of grids facilitates integration and collaboration of research teams in the US and Europe, at least in the scope of large computing centers.

Unfortunately, although major Polish HPCs have indeed participated in grid development projects such as CrossGrid [11] or EGEE, Polish scientists often had to rely solely on their local computing resources or personally apply for access to external infrastructures (for instance through EGEE channels). This strongly discouraged the Polish research community from conducting interdisciplinary computational research with other institutes in Poland and produced a situation where different research teams were affiliated with different European grid domains and developed software for different grid systems (e.g. gLite [12] or UNICORE).

Eventually, in 2007, the major computing centers in Poland (see Fig. 1) – specifically the Academic Computer Centre Cyfronet AGH (ACC Cyfronet AGH) in Kraków, the Poznań Supercomputing and Networking Center (PSNC), the Interdisciplinary Centre for Mathematical and Computational Modeling (ICM) in Warsaw, the Gdańsk Academic Computer Centre (CI TASK) and the Wrocław Centre for Networking and Supercomputing (WCNS) – came together to form the PL-Grid Consortium [13] with the stated goal of "providing the Polish scientific community with an IT platform based on grid computer clusters, enabling e-science research in various fields".

The foundation of the consortium was a long-term agreement expressing the will to integrate national distributed computing resources in order to provide the Polish research community with on-demand access to a sustainable computing platform and participate in international grid initiatives and e-science activities under the umbrella of a federated national infrastructure.

Fig. 1. PL-Grid Consortium

The PL-Grid Consortium prepared a project (called PL-Grid) which obtained funding from the European Regional Development Fund as part of the Innovative Economy Program [14] and commenced in March 2009. The total project budget was 21M€ and project completion was scheduled for March 2012, with the aggregated offer for the community consisting of 215 TFlops of computational power and 2.5 PB of storage resources. The overall goal of this project was to deploy a persistent national large-scale distributed computing infrastructure for Polish scientists, allowing them to conduct interdisciplinary research on a national scale, and giving them transparent access to international grid resources via affiliated international grid infrastructures.

As the project enters its final stage, the PL-Grid Consortium carries on its activities in the area of grids and clouds for scientific community and has prepared a follow-up project called PL-Grid Plus, funded in October 2011 by the same institution with a total budget of 18M€. This project will last until the end of 2014.

In addition, consortium partners have been participating in other national and European projects extending the infrastructure and available middleware while constantly enhancing the maturity and functionality of the resources provided to Polish scientific teams.

2 Objectives

The main objectives of the project were drawn from the experience of all consortium partners based on their activities in international grid initiatives as well as the requirements expressed by users of local computing resources. The list of expected properties of the Polish grid infrastructure included the following:

– Multiscale,
– Multidisciplinary,
– Heterogeneity,
– Unification of access,
– Capability of international cooperation,
– SLA support and computational grant tools,
– Sustainability,
– Training and user support.

Multiscale. The PL-Grid Consortium believes that the national grid platform must be easy enough for users not to discourage them even if they run relatively simple computations. Basing on experience from projects which were tailored for large-scale experiments and suffered from the lack of user variety, our intention from the start was to make the infrastructure as user friendly as possible, with full support even for the smallest, entry-level computations. Moreover, in practice, the overhead the users must "pay" when migrating their applications to the grid sometimes is a barrier, as most experiments start from small-scale test cases and then evolve into more advanced scenarios. The PL-Grid infrastructure simplifies running the smallest jobs using direct access to PBS systems installed at particular locations, while encouraging the users to progressively migrate their applications to scalable grid job implementations in order to get the full benefit of the infrastructure.

Multidisciplinary. One of the fundamental aims of PL-Grid was to build a computational platform for multiple scientific domains. By analysing the structure and heterogeneity of Polish scientific communities we discovered that the platform could not be bound to, or otherwise favor, any particular scientific domain, but must be flexible enough to support different communities with various software requirements, different collaboration practices and varying security needs. This challenge led to the concept of domain grids, as planned in PL-Grid Plus, which allow creation and grouping of scientists and resources based on their specializations (see Fig. 2).

Heterogeneity. The multidisciplinary aim described above determines another very important challenge which PL-Grid must face – infrastructural heterogeneity. Provisioning a flexible infrastructure, with support for various scientific domains, requires the project to maintain a set of differing computing architectures and storage solutions. Thus, while the majority of PL-Grid resources are based on x86 computing clusters (mainly due to their universality) we also integrate additional, more specialized platforms and systems based on such technologies as GPGPUs and software enhanced clusters by Versatile SMP (vSMP). Although integration of such systems with the wider grid infrastructure requires substantial effort, we believe it will benefit many users who deal with atypical problems – which can be solved more efficiently using custom hardware.

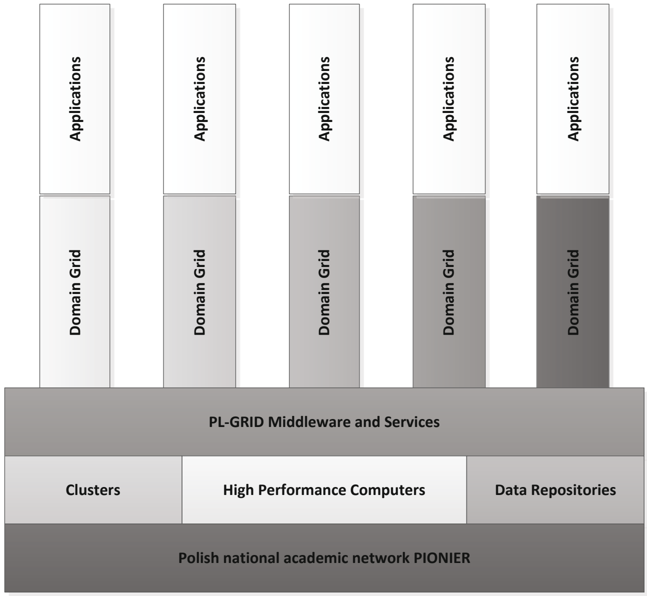

Fig. 2. Layered view of the PL-Grid infrastructure

Unification of Access. Due to the fact that many of our users have already used some grid infrastructure in the past and are accustomed to different types of interactions with the grid, another major challenge was to make the migration of their applications to the national grid platform as natural as possible. In practice, this resulted in the need to support two main grid middleware stacks, i.e. gLite and UNICORE. Furthermore, a special portal has been developed, enabling users to acquire credentials with which they can access both grid middleware components.

SLA Support and Computational Grant Tools. Another important requirement from the management perspective is ensuring that each of the computing centers which contribute resources to the PL-Grid infrastructure complies with a certain Service Level Agreement (SLA) in terms of the amount of resources and their availability. On the other hand, the resources centers must be able to monitor how their resources are used and that only authorized users obtain access and run jobs on the infrastructure, respecting their assigned quotas.

Training and User Support. In spite of substantial effort put into making the PL-Grid infrastructure as user-friendly as possible, it is still necessary to provide training for current and future system users. We need to not only show how to access and use the system, but also how to exploit the full potential of the hardware and software capabilities provided by our infrastructure in order to allow the users to create and run scalable and sustainable in-silico experiments.

Moreover, even advanced and experienced users sometimes require expert support on emerging technical on administrative issues, and thus PL-Grid provides persistent helpdesk support through the user portal.

Availability and Sustainability. In order to support the entire Polish scientific community with a distributed computing infrastructure it is crucial that the service is provided 24/7/365, with minimal downtime. Although most of the infrastructure is distributed (thus minimizing the likelihood of total system failure), we nevertheless do our best to ensure that the system remains at its full potential all the time. Hence, the average infrastructure availability is 95%-98%, depending on the period. Another issue here is that the development and deployment of an infrastructure of such a scale cannot be justified just for duration of a single project and PL-Grid Consortium partners have made long-term commitments to support the hardware infrastructure, software and users, as well as to promote continuous evolution by extending the hardware and software made available to the Polish scientific community.

Capability for International Cooperation. Last but not least, the PL-Grid infrastructure aims to be a key player in the international grid community, giving our users access to a much more extensive set of varied computational resources. The goal is to integrate our infrastructure with such initiatives as EGI and other European grids, PRACE and, where possible, US infrastructures.

3 Infrastructure

At the moment PL-Grid provides 230 TFlops of computational infrastructure and 2500 TB of storage to the Polish scientific community. Over the coming 3 years the consortium plans to extend these resources to 700 TFlops and 6000 TB of storage.

Currently the total number of cores is about 25 000 and the number of jobs ran per quarter is constantly growing, having reached $3 * 10^6$ in Q2 2011 (see Fig. 3). The number of registered users is approximately 800 (see Fig. 4).

To fulfill the aims of the project and the expectations of consortium members regarding availability, datacenters were redesigned to the multi-megawatt scale. Each server room is equipped with highly efficient megawatt-scale air conditioning systems and sustainable power supplies backed by large Uninterrupted Power Systems and diesel power generators, giving the possibility to operate the computational infrastructure for many hours despite any mains power failures.

The described infrastructure consists of:

- Intel-based Xeon clusters with 4- or 6-core processors and Infiniband connectivity,
- AMD-based clusters with 8- or 12-core processors and Infiniband connectivity,
- vSMP integrated machines with extensive memory,
- GPGPU accelerated hosts running NVidia Tesla cards.

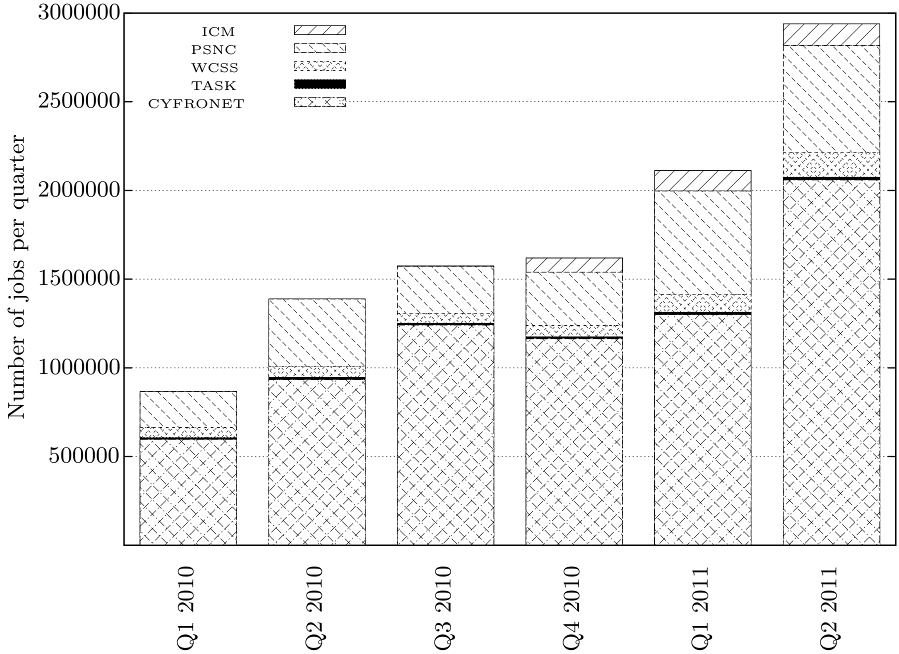

Fig. 3. Number of jobs executed in PL-Grid infrastructure per quarter

In addition to the computational infrastructure we provide large-scale storage using fiber channel connectivity and (in most cases) the Lustre distributed file system.

4 Technological Advancements

The requirements of the project and its end users, i.e. the Polish scientific community, have resulted in the overall layered architecture of the PL-Grid platform, which is presented in Fig. 5. In addition to deploying the basic hardware infrastructure and standard grid middleware stacks, we have designed several new features and components with the aim of improving user experience and extending the scope of possible applications which can be developed by the research communities.

4.1 Middleware Plurality

One of the main technological challenges was to provide unified access to grid infrastructures managed by both gLite and UNICORE middleware stacks. Although this issue has been addressed (to some extent) by other projects in the past, we could not find a solution which would satisfy all our needs and thus we decided to extend the QosCosGrid [23] system for scheduling jobs across

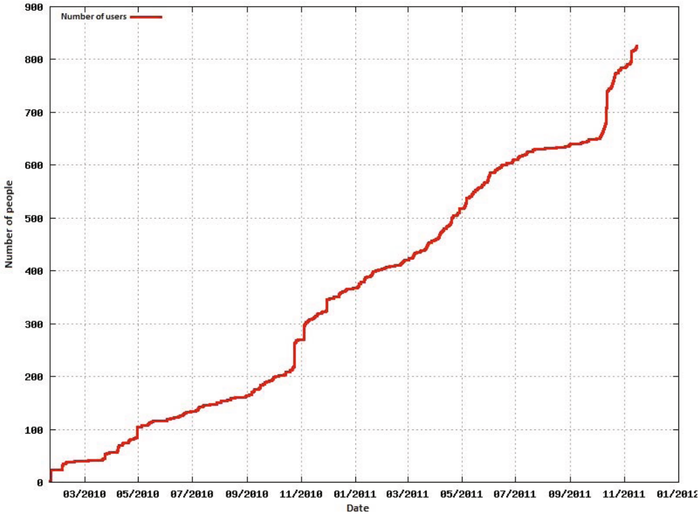

Fig. 4. Number of registered PL-Grid users over time

grid middleware domains. The need to support multiple middleware suites is dictated by users' applications. Some users have to use gLite because they collaborate with international partners who already use this middleware. On the other hand, UNICORE users tend to prefer this system due to the ability to schedule computational jobs in a workflow-like manner. Moreover, QCG (developed in PL-Grid) offers easy-to-use mapping, execution and monitoring capabilities for a variety of applications, such as parameter sweeps, workflows, MPI or MPI-OpenMP hybrids.

4.2 SLA and Grants

In order to enforce the expected quality of service of our platform and allow users to describe their hardware and software requirements, PL-Grid provides a collaborative SLA negotiation tool called Bazaar, which simplifies the process of establishing a compromise between user expectations regarding the infrastructure and the actual capability which can be assigned to handle a particular user request. Bazaar is compatible with our computational grant model and takes into account information about user quotas and current infrastructure status, as well as the history of resource provisioning and consumption, enabling system managers to evaluate how different QoS parameters are respected by providers and users.

4.3 New User Interfaces

PL-Grid provides computational power for users representing various scientific domains. In many of these users have very specific ways of conducting their experimentation, which, in some specific cases, involves demands for new

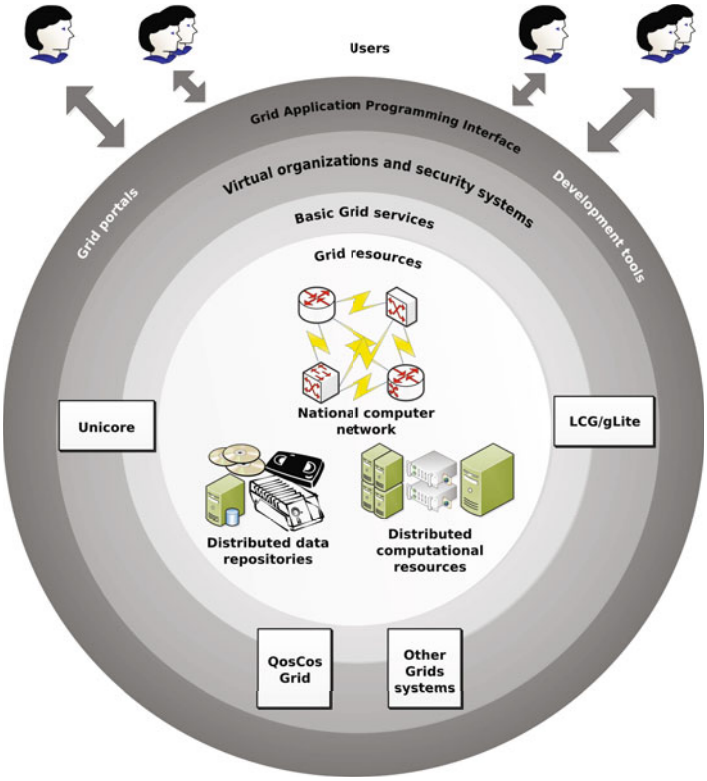

Fig. 5. Structure of the PL-Grid Computing Platform

user interfaces. One of the interesting outgrowths of PL-Grid users' requests is the GridSpace2 Experiment Workbench [18,19]. It is a novel virtual laboratory framework which enables researchers to conduct virtual experiments on grid-based resources. A web portal has been developed for this tool and serves as the main access point to all the mechanisms of the virtual laboratory from any workstation equipped with a Web browser. The portal encapsulates the so-called Experiment Workbench layer which provides tools to create experiments, collaborate, communicate, share resources and run large-scale computations in a user-friendly manner.

Many users deem command-line grid interfaces as too unwieldy and obfuscated. Thus, a graphical user interface, called the Migrating Desktop [20,21], has evolved in PL-Grid. This advanced GUI, similar to a window-based operating system, hides the complexity of the grid middleware and makes access to grid resources easy and transparent. It allows the user to run applications and tools, manage data and store personal settings independently of the location or terminal type. The Migrating Desktop offers a flexible personalized roaming work environment, scalability and portability, a rich set of tools, a single sign-on mechanism and support for multiple grid infrastructures.

4.4 Security

The dynamically evolving scene of security threats requires the project to provide its system administrators with targeted security management tools. One such tool is ACARM-ng [22] – an extensible plugin-based alert correlation framework. It consists of abstractions for correlation, reporting, reaction, gathering data from multiple sources and data storage. Real-time reporting is supported, which means that alerts can be reported while still being correlated. A web-based administrative interface presents the gathered and correlated systemwide data in a consistent way.

Another tool developed for PL-Grid security management is SARA (System for Automated Reporting and Administration). The system combines information on vulnerabilities stored (with the help of the standards mentioned above) in the National Vulnerability Database. It makes management of well-known vulnerabilities easier and more efficient.

5 International Computational Arenas

Collaboration with existing and future high-performance computing initiatives is a very important goal for the PL-Grid Consortium, as it allows the Polish scientific community to collaborate and conduct research in truly international teams and gain access to special hardware and software resources available abroad.

In March 2010 an international organization called the European Grid Initiative, with headquarters in Amsterdam, was launched with the goal to oversee the operation and development of the pan-European grid infrastructure. About 40 European National Grid Initiatives (NGIs), three European Research International Organizations (EIROs) and several grid organizations from outside of Europe supported this initiative. In this context the PL-Grid platform is one of the NGI members.

As of today EGI.eu integrates and supervises more than 200 thousand logical CPUs (cores) and about 100 PB of disk and 80 PB of tape storage; more than 13 thousand users are organized in more than 180 VOs, out of which about 30 are active – about 1 million jobs are executed daily, mainly related to particle physics but also representing the domains of archeology, astronomy, astrophysics, civil protection, computational chemistry, earth sciences, finance, fusion, geophysics, life sciences, multimedia, material sciences, etc.

PL-Grid is one of the key players in the EGI infrastructure, currently ranked 4th (see Fig. 6). PL-Grid was actually the first NGI in the EGI initiative and the PL-Grid Consortium has devised many procedures for migrating from an EGEE ROC to National Grid Initiative status, thus paving the way for other countries. PL-Grid operates a regional Service Availability Monitoring system with the ability to dispatch Regional Technical Support teams to assist service administrators in investigating problems. In addition, we are responsible for our grid configuration management, i.e. certification of new sites in Poland, adding services and downtimes. Recently we have contributed to the EGI UNICORE

Integration Task Force, helping NGIs willing to contribute their UNICORE resources to EGI. PL-Grid is strongly involved in the EGI security team called CSIRT (Computer Security Incident Response Team). Our responsibility is to coordinate security incidents with all Polish sites, including fixes for vulnerabilities reported by EGI as well as those found by our own staff.

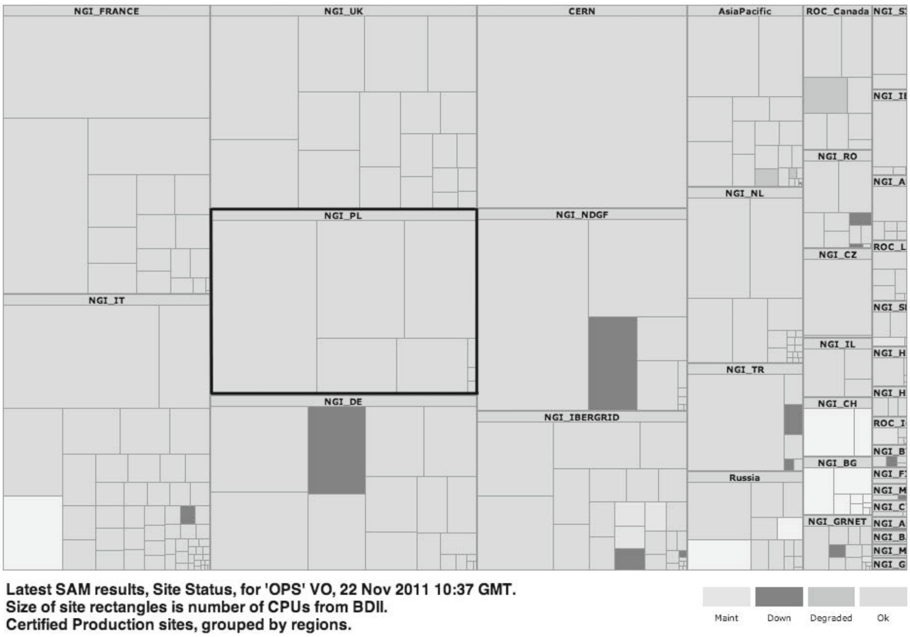

Fig. 6. PL-Grid site named *NGI_PL* size in comparison to others NGIs registered in EGI. Size of the boxes represents the computational power. The PL-GRID box has been highlighted.

Additionaly, several activities related to middleware development have been undertaken in cooperation with other international partners. For instance, the advanced reservation and co-allocation of distributed computing resources provided by QosCosGrid drew attention from other supercomputing centres from USA and Europe. Consequently, the QCG community has grown and today includes users from INRIA, the University of Dortmund, the Leibniz Supercomputing Centre, SARA, the Queensland University and the Louisiana State University.

6 Future Work

PL-Grid is a stable and sustainable distributed computing platform operating on a national level. However, the increasing user requirements and expectations force us to keep provisioning the best possible technological solutions. Given the current performance level of computing hardware and the observed rate of progress, the high-performance computing community should reach the exascale performance range in the next 5-7 years. Furthermore, current trends in distributed computing, especially in business scenarios, favor simpler and more scalable infrastructural models such as cloud computing. Thus, the designers and developers of the PL-Grid infrastructure are actively participating in several projects focusing on technological advancement in both hardware and software areas. The most interesting solutions will be integrated with our infrastructure in the future.

Over the next three years, we will focus however mainly on supporting and building true distributed user communities called domain grids. This approach is motivated by the need to bring together scientists from all over Poland and attract them to our computing infrastructure, encourage collaboration and sharing of results and improve the impact of research performed in Poland as well as on the international stage. We wish to involve single users and entire user communities, both existing and emerging, in the process of refining the infrastructure in order to meet their needs. Since there is no universal solution which solves all the domain-specific problems we have decided to invite community representatives and involve them in the process of making the grid easier and more efficient to use.

Near-term technological refinements will include forays into the cloud computing world, making our infrastructure more flexible from the system administration point of view while still providing all the grid features usually lacking in cloud platforms.

7 Conclusions

The PL-Grid Polish National Grid Initiative was initiated almost 5 years ago. Over time it has received support from several projects, which extended its hardware and software capabilities. At the moment the infrastructure is both powerful and mature. The quality, availability and versatility of tools offered to users increases its impact, as evidenced by the constant growth in the number of registered participants (see Fig. 4).

All of the objectives stated in the design phase already have been fulfilled to some extent. We provide a flexible infrastructure, supporting users from all domains of science, including high energy physics, chemistry and linguistics.

The availability indicators of the infrastructure, especially when taking into account its complexity, remain very high. At present users can run their experiments on different platforms, including x86-based clusters, vSMP with up to 6TB of memory and GPGPU-accelerated clusters.

The extensive effort of more than 100 people working on practical applications of modern computational technologies has made the Polish NGI a key player in the European grid community, and we believe that multi-level heterogeneity is the key to the project's unquestionable success.

References

1. Foster, I., Kesselman, C. (eds.): The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufmann, San Francisco (1999)
2. Foster, I., Kesselman, C.: The Grid 2: Blueprint for a New Computing Infrastructure. Morgan Kaufmann (2004)
3. LCG Project, `http://lcg.web.cern.ch/lcg/`
4. Bollen, J., Fox, G., Singhal, P.R.: How and where the TeraGrid supercomputing infrastructure benefits science. J. Informetrics 5, 114–121 (2011)
5. EGEE Project, `http://www.eu-egee.org/`
6. Gentzsch, W.: DEISA, the Distributed European Infrastructure for Supercomputing Applications. In: Gentzsch, W., Grandinetti, L., Joubert, G.R. (eds.), vol. 18, pp. 141–156. IOS Press (2008)
7. Foster, I.: Globus Toolkit Version 4: Software for Service-Oriented Systems. In: Jin, H., Reed, D., Jiang, W. (eds.) NPC 2005. LNCS, vol. 3779, pp. 2–13. Springer, Heidelberg (2005), doi:10.1007/11577188_2
8. Streit, A., Bala, P., Beck-Ratzka, A., Benedyczak, K., Bergmann, S., Breu, R., Daivandy, J.M., Demuth, B., Eifer, A., Giesler, A., Hagemeier, B., Holl, S., Huber, V., Lamla, N., Mallmann, D., Memon, A.S., Memon, M.S., Rambadt, M., Riedel, M., Romberg, M., Schuller, B., Schlauch, T., Schreiber, A., Soddemann, T., Ziegler, W.: Unicore 6 – Recent and Future Advancements. Annales des Telecommunications 65, 757–762 (2010)
9. Thain, D., Tannenbaum, T., Livny, M.: Distributed computing in practice: the Condor experience. Concurrency – Practice and Experience 17(2-4), 323–356 (2005)
10. Eerola, P., Ekelöf, T., Ellert, M., Grønager, M., Hansen, J.R., Haug, S., Kleist, J., Konstantinov, A., Kónya, B., Ould-Saada, F., Smirnova, O., Szalai, F., Wäänänen, A.: Roadmap for the ARC Grid Middleware. In: Kågström, B., Elmroth, E., Dongarra, J., Waśniewski, J. (eds.) PARA 2006. LNCS, vol. 4699, pp. 471–479. Springer, Heidelberg (2007)
11. Bubak, M., Malawski, M., Zając, K.: The CrossGrid Architecture: Applications, Tools, and Grid Services. In: Fernández Rivera, F., Bubak, M., Gómez Tato, A., Doallo, R. (eds.) Across Grids 2003. LNCS, vol. 2970, pp. 309–316. Springer, Heidelberg (2004)
12. Laure, E., Grandi, C., Fisher, S., Frohner, A., Kunszt, P., Krenek, A., Mulmo, O., Pacini, F., Prelz, F., White, J., Barroso, M., Buncic, P., Byrom, R., Cornwall, L., Craig, M., Di Meglio, A., Djaoui, A., Giacomini, F., Hahkala, J., Hemmer, F., Hicks, S., Edlund, A., Maraschini, A., Middleton, R., Sgaravatto, M., Steenbakkers, M., Walk, J., Wilsonm, A.: Programming the Grid with gLite. Computational Methods in Science and Technology (2006)
13. PL-Grid Project, `http://www.plgrid.pl/`
14. Innovative Economy Operational Programme, `http://www.poig.gov.pl/english/`
15. International Organization of Standardization (ISO), Information technology – Service Management – Part 1: Specification (ISO/IEC 20000-1:2005)
16. EGI Consortium, `http://www.egi.eu/`
17. Szepieniec, T., et al.: On Importance of Service Level Management in Grids. Accepted for Workshop: Managing and Delivering Grid Services 2011 (MDGS 2011) (2011)

18. Bubak, M., Malawski, M., Gubała, T., Kasztelnik, M., Nowakowski, P., Harężlak, D., Bartyński, T., Kocot, J., Ciepiela, E., Funika, W., Król, D., Baliś, B., Assel, M., Tirado Ramos, A.: Virtual Laboratory for Collaborative Applications. In: Cannataro, M. (ed.) Handbook of Research on Computational GridTechnologies for Life Sciences, Biomedicine and Healthcare, Information Science Reference, ch. XXVII. IGI Global (2009) ISBN: 978-1-60566-374-6
19. GridSpace2 Experiment Workbench, https://gs2.cyfronet.pl/
20. Owsiak, M., Palak, B., Plociennik, M.: Graphical Framework for Grid Interactive and Parallel Applications. Computing and Informatics, 223–232 (2008)
21. Kupczyk, M., Lichwala, R., Meyer, N., Palak, B., Plociennik, M., Wolniewicz, P.: Applications on demand as the exploitation of the Migrating Desktop. Future Generation Comp. Syst., 37–44 (2005)
22. ACARM-ng, http://www.acarm.wcss.wroc.pl
23. QosCosGrid Middleware, http://www.qoscosgrid.org/