On Importance of Service Level Management in Grids

Tomasz Szepieniec[1], Joanna Kocot[1], Thomas Schaaf[2], Owen Appleton[3], Matti Heikkurinen[3], Adam S.Z. Belloum[4], Joan Serrat-Fernández[5], and Martin Metzker[2]

[1] ACC Cyfronet AGH, Krakow
[2] Ludwig-Maximilians-Universitaet, Munich
[3] Emergence Tech Limited, London
[4] University of Amsterdam
[5] Universitat Politcnica de Catalunya, Barcelona

Abstract. The recent years saw an evolution of Grid technologies from early ideas to production deployments. At the same time, the expectations for Grids shifted from idealistic hopes — buoyed by the successes of the initial testbeds — to disillusionment with available implementations when applied to large-scale general purpose computing. In this paper, we argue that a mature e-Infrastructure aiming to bridge the gaps between visions and realities cannot be delivered without introducing Service Level Management (SLM). To support this thesis, we present an analysis of the Grid foundations and definitions that shows that SLM-related ideas were incorporated in them from the beginning. Next, we describe how implementing SLM in Grids could improve the usability and user-experience of the infrastructure – both for its customers and service providers. We also present a selection of real-life Grid application scenarios that are important for the research communities supported by the Grid, but cannot be efficiently supported without the SLM process in place. In addition, the paper contains introduction to SLM, a discussion on what introducing SLM to Grids might mean in practice, and what were the current efforts already applied in this field.

Keywords: SLM, service delivery, Grids, ITIL.

1 Introduction

Since the 1990's, when the term 'Grid' was coined, Grids have changed from early prototype implementations to production infrastructures. However, despite maturing considerably during this time, Grids still suffer from the lack of service management solutions that would be suited to an infrastructure of the size and user base of the current Grid. The maturing Grid technologies need to incorporate understanding of the business models of the users and service providers. When possible, they should be composed from standard business solutions that support service management and delivery.

M. Alexander et al. (Eds.): Euro-Par 2011 Workshops, Part II, LNCS 7156, pp. 64–75, 2012.

In parallel with the evolution of the infrastructures, the understanding of what the Grid should be was subject to a change that was by no means less significant and rapid. The users of the Grids, as well as the specialists from the distributed computing domain, were at first fascinated with the potential of the Grid technology. However, they gradually became disappointed with what was really offered to them, and shifted towards new trends and paradigms (with the same elevated hopes). This change is somewhat alarming, as these new technologies, will be (or already are) facing the same problems [1] related to provision of computational and storage resources to the users. There is a danger of repeating the vicious enthusiasm-disillusionment cycle, as long as the users are looking for "miracle cures" and there are over-optimistic proponents of untested solutions.

The authors believe that the Grid technologies are in fact mature enough to meet most of the user needs. However, the quality of service provision and management needs much more attention. A professional service management approach is the key to engaging with users and improving their satisfaction services. It also acts as a tool for capturing and transferring requirements and best practices that can be used for more informed evaluation of new e-Infrastructure services (and more efficient uptake, eventually). This process becomes even more critical when an e-Infrastructure intends to serve more and more demanding and complex projects. For customers engaged in such initiatives, warranties related to resource provisioning and service level are crucial. The common *de facto* assumption of the e-Infrastructure service providers, which sees any vague, qualitative service level or best-effort operation (beyond what is provided by the software itself) as sufficient, is no longer valid. Grid Computing and other e-Infrastructures must follow similar paths towards maturity as the general solutions available for IT services. The realization of this goal can be sped up by basing it on documents such as ITIL [2], which provide best practices for implementation and management of processes important to contemporary Grids.

In this paper, we argue that mature and competitive e-Infrastructures implementing the Grid ideas cannot be delivered efficiently without implementing processes of Service Level Management (SLM). To prove this thesis, we provide a range of arguments – starting from the elements of Grid theory in Sec. 4, through the analysis of benefits for Grid customers and providers in Sec. 5 that SLM might bring, to actual scenarios of real computations using Grids in Sec. 6. Additionally, we give a short introduction to SLM in Sec. 2, and to ideas how SLM can be applied to Grids in Sec. 3. Related works and implementation are presented in Sec. 7 and in Sec. 8.

2 Background: On the Relevance of SLM

The efficient delivery of high-quality IT services — especially in a constantly changing environment, with ever-growing customer and user demands — poses a major challenge for the IT service providers. To rise to this challenge, more and more (commercial) providers are adopting IT service management (ITSM) processes as described by the IT Infrastructure Library (ITIL) [2][1], or ISO/IEC 20000 [3][2].

ITSM can be regarded as a set of organisational capabilities and processes required by an IT (service) provider to keep his utility and warranty promises / commitments. In this context, *Service Level Management* and *Service Delivery Management* are the most important sub-disciplines:

Service Level Management (SLM) describes the processes of:
- defining a catalogue of IT service offerings;
- specifying services and service components, including their dependencies and available service level options;
- negotiating and signing Service Level Agreements (SLAs) with customers, underpinning SLAs with internal Operational Level Agreements (OLAs) and suitable contracts with external suppliers,
- monitoring and reporting on the fulfilment of SLAs as well as (early) notifications of SLA violations.

Service Delivery Management (SDM) provides guidelines for managing the delivery of SLA-aware IT services through their lifecycle including:
- planning of details of service delivery;
- monitoring and reporting on capacity, availability, continuity, and security;
- managing changes and releases in a controlled manner;
- maintaining accurate information on the infrastructure and its configuration;
- handling incidents and user requests, and resolving and avoiding problems.

Following a process approach in the implementation of Service Level Management and Service Delivery Management means providing a clear definition of tasks, activities and procedures. This must be supported by unambiguous delegation of responsibilities, identification of all interfaces, as well as steps to ensure adequate documentation, traceability and repeatability of all processes.

The main focus of this paper is Service Level Management, since it forms the foundation for effective Service Delivery Management. In general, SLM is a vital part of customer-oriented provision of high-quality IT services. It is important for achieving an improved relationship between an IT service providers and their customers, as well as for aligning "what the IT people do" to "what the business requires". In the relationship management domain SLM provides common understanding of expectations, mutual responsibilities and potential constraints between different domains.

Various approaches for supporting effective SLM have evolved from research and practice, mostly focused on business IT, throughout the last decade. Still, it should be noted that SLM in general is evolving beyond the "traditional" IT service provisioning scenarios – hence, introducing it to such infrastructures as Grids can be seen as part of a natural progression.

[1] A set of handbooks describing good practices for ITSM.

[2] An international standard for ITSM which features a process framework, which is in many aspects aligned with ITIL.

3 Model of SLM for Grids

Before specifying how the Service Level Management principles can be applied to Grid infrastructures, the main actors and relations between them have to be identified and described. The main actors considered in an SLM model of a Grid infrastructure are:

- Virtual Organisation (VO) is a set of individuals and organisations (i.e. users) that cooperate by sharing resources according to formal or informal contract which defines the rules of cooperation. We understand that a Virtual Organisation is the customer of a Grid Initiative.
- Grid Initiative (GI) is an approved body that provides Grid computing services or represents Grid providers in a region, country or group of countries. Grid Initiatives may be organised in larger bodies, creating a hierarchical structure, with primary GIs federated in secondary GIs etc. In Europe, for example, the primary GIs are created at national level – forming National Grid Initiatives (NGIs), and are federated in the European Grid Initiative (EGI.eu). The infrastructure and middleware supporting the GIs on all levels constitute a Grid. The GI is a single point of contact for a VO, representing the Grid as a whole. The added value of a GI may range from a simple aggregation (GI as "mediator") to full integration (GI as "service provider") of the underlying resources.
- Site is an infrastructure provider that offers computing and storage infrastructure available through Grid protocols; they usually do not provide the whole set of technical services to support Grid Computing.
- An External Partner/Supplier is supporting any of the above mentioned primary actors in the fulfilment of their duties.

The SLM model for Grids assumes that not all the actors interact with each other directly, and that the interaction between the parties can be formalised using a set of agreements. The model for these interactions is presented in Fig. 1. It allows the relationships to form hierarchical SLA&OLA framework that is compatible with a general IT SM approach. However, the model was designed to be applicable to different types of GIs in terms of amount of extra warranties added on higher level of OLAs or SLAs. The model allows the following interactions:

- VO – GI: The GI is responsible for provision of a Grid service to the VO. Formalisation of such relationship is done through a Service Level Agreement (SLA). SLA describes the Grid service, documents Service Level Targets, and specifies the responsibilities of the GI and the VO.
- GI – Site: The Site is responsible for delivering services to the GI customers (VOs). These relationships are formally described with an Operation Level Agreement (OLA). The OLA framework within a GI supports the fulfilment of the targets agreed in the SLAs between the GI and its VOs. Hence, OLAs may be established for one of the two purposes: in order to support one or more specific existing or intended SLAs, or as a general and/or preparatory basis for establishing new services/SLAs.

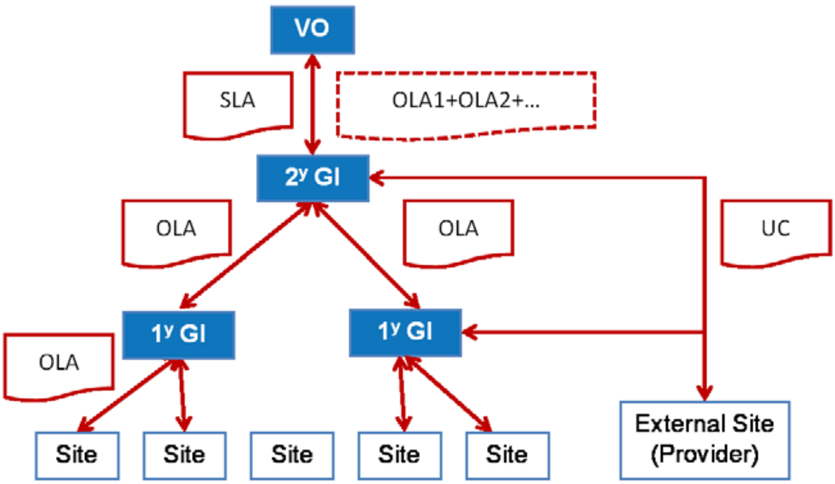

Fig. 1. OLA/SLAs defining relations between actors in Grids. 1y GI and 2y GI stand for Primary and Secondary GI, respectively.

- Primary GI – Secondary GI: The nature of this relation is similar to the interaction between GI and a Site, and may be described with the same formalisms (OLAs).
- External Partner/Supplier – GI: The relationship between External Partner/Supplier and a GI (primary or secondary) or any other actor is formalised through an Underpinning Contract (UC). As UCs are formal contracts with external bodies they may contain references to general terms and conditions or specification of commercial and legal details.

4 Elements of SLM in Grid Theory

The term Grid was introduced to describe a federated infrastructure, providing computing resources to its users. Ian Foster, considered as one of the original authors of the Grid concept, required delivering *nontrivial qualities of service* as one of three main characteristics of the Grid in his most commonly cited definition of Grid [4]. This feature was also explained as *various qualities of service* which are set-up to *meet complex user demands*. Even if Forster's definitions do not explicitly mention a need of negotiation and signing of an agreement – an SLA, it is obvious that the quality of service needs to be described using measurable metrics.

On the other hand, Plaszczak and Weiner [5] claim that one of three main advantages of Grid Computing is *on-demand provisioning* as opposed to classical resource provisioning realised by purchasing and installing hardware and software. If this process is to be reliable to the user, a kind of warranty that the resources provisioned are available when they are needed, is crucial. Therefore, such a warranty has to be a subject of an agreement between the provider and the customer.

A similar conclusion can be drawn from [6], where the authors introduce a distinction between customers with specified expectations and customers that can accept any (unspecified) quality of service (QoS). According to the authors, the former needed to be serviced by so-called “commercial” Grids, that require an SLA framework. The latter are just a limited class of users and applications – which means that without SLM, Grid technologies shall become niche technologies of very limited usage.

However, the history of Grid Computing, apart from these ambitious theories, provides also an explanation why the current infrastructures provide their resources only on best effort basis. The first implementations of Grid-like technologies were built by voluntary computing based on desktop machines like the Seti@Home Project[3], in which only best-effort QoS approaches were possible. Many people still believe that Grid Computing is simply voluntary computing and it will always remain of low QoS[4]. We consider this view a stereotype. One can note that technologies aimed at federating resources are orthogonal to single resource reliability. It applies to federated resources with low reliability (voluntary computing), as well as to resources with high reliability (professionally managed computer centres) – both types of resources have their own groups of users. Obviously, reliability and other QoS parameters of federated infrastructures are strictly related to the same parameters with which single services are provided and it means that providing federated resources with high level of quality is possible.

5 Actors Perspective

In this section we analyse how each actor of the SLM model for Grids would benefit from introducing Service Level Management solutions. We will also assess the cost of such operation.

5.1 Customer Perspective

The groups most directly interested in the quality of the Grid services are their users. Here we focus on issues typically raised by them, which can be solved by introducing SLM to the Grid infrastructure. This analysis is partially based on a survey the authors performed on users of different Grid infrastructures and the Virtual Organisation managers gathered at EGI User Forum 2011[5].

Any users activity in the infrastructure is usually a stage of a scientific plan, project or experiment that needs to be accomplished in a limited time. This is strictly related to a need for a warranty of availability of certain resources, fulfilling certain requirements (parameters) within the requested time. Therefore,

[3] `http://setiathome.berkeley.edu/`

[4] This observation was confirmed recently by a survey performed on participants of the International Supercomputing Conference ISC’11.

[5] `http://uf2011.egi.eu/`

a crucial benefit of introducing the SLM mechanisms to Grid is the possibility of planning ahead the activities that require services. This requirement was confirmed by the aforementioned survey results – the users perceived "no or poor warranty of obtaining resources in reasonable (finite) time" to be the second[6] most discouraging issue with Grid technology. Also, "improving warranty" was the most desired improvement suggested by the questioned users.

Naturally, the introduction of any kind of warranty introduces additional managerial overhead for users who need to apply and negotiate for such warranties. The balance between benefits and costs of additional effort seems important success factor of the SLM deployment. In our survey, 60% respondents indicated that they are ready to invest in more strict and complex procedures in exchange for improvements in the Grid quality issues.

The motivation for improvement in management of the Grid services can be also drawn from how the users evaluate the infrastructures they had experience with. It is significant, that the quality of the resource provision and management in most infrastructures, according to their users, is considerably lower than the quality of the resources themselves. The disparity of the average grade spans from 0,52 to 1,35 in 1-5 scale, for the larger, international infrastructures. While, for national infrastructure the same parameters are perceived as slightly better.

All this agree with a fundamental psychological fact: user's satisfaction strongly depends on predictability of resource characteristics.

5.2 Sites Perspective

Usually, sites (or resource providers in general) tend to be reluctant in adopting SLM, as they see themselves as the side that is forced to promise and give warranties. However, deeper analysis shows some important benefits for them too, coming from the adoption of SLM.

Primarily, in the SLM process, a provider obtains detailed specification of the user needs usually some time in advance. This gives them the opportunity to better manage the resource provision and perform capacity planning – by allowing the providers to better estimate the parameters of the resources that would be needed by customers. That, in consequence, leads to optimising the resource costs. Based on the known requirements, the provider can better handle prioritization also in terms of executing internal policies and preferences. What is more, introducing SLM facilitates (or enables) accounting, which, especially in academia, usually requires justification in terms of reported results of scientific research. With SLM the previous result reports may serve in negotiating new SLAs.

In further perspective, SLM stimulates and strengthens the relations with customers, which naturally results in evolution of the maturity of the providers, who better know customer needs and can assess their satisfaction. SLM may serve also to improve the communication with the customers to keep the providers better informed about the user needs, distribute offers and provide means for marketing solutions. The latter actions are now usually neglected by the resource providers,

[6] The first were "technical difficulties" – which are out of scope of this document.

especially academic computer centres, although their R&D departments require close collaboration with the users.

The main cost of introducing SLM from the resource providers point of view is the additional effort of maintaining SLM-related processes which include negotiations, reconfiguration of resources, usage monitoring, and accounting.

5.3 Grid Infrastructure Provider Perspective

The Grid Infrastructure can be perceived as a *virtual resource provider*, since, what it offers, are resources operated by other (physical) resource providers. So, the benefits and cost of adopting SLM are similar to a provider's described above. However, the perspective of GI is broader, as it can handle many resource providers.

In ITIL all the activities should be focused to deliver more *value* to the users. Value comes from two elements: resources itself and the warranties. In terms of resources GI usually does not provide anything that cannot be delivered by the sites. However, by maintaining OLAs with sites, the GI can deliver more warranty than any specific task. In that sense, a GI that implements SLM can provide more value to customers, and, in this way, be competitive among the Grid infrastructures. Otherwise, the GI's role is limited to providing technical solutions for integration.

6 Application Scenarios That Require SLM

6.1 Large Collaboration Case

From the beginning of the realisation of the Grid concepts, the key Grid customers were large-scale projects with worldwide collaboration. For such projects, implementation of at least some processes from Service Level Management seems unavoidable. Giving the example of the main customers of the European Grid Initiative (EGI) [7], we show how SLM was necessary for them and how it was realised.

The most representative example of a large European project using Grids is Large Hadron Collider (LHC) built in CERN. LHC is the largest research device worldwide, gathering thousands of researchers in four different experiments. Each of these experiments requires petabytes of storage space and thousands of CPU cores to process the data. The data produced by the experiments are produced continuously while the LHC is running. Therefore, there should be enough resources capable of handling large data amounts and throughput supplied, both, on the short- and on the long- term. This includes computational and storage resources, as well as network facilities. Thus, the long-term goals require special focus on infrastructure planning.

The LHC way of defining contracts related to resources was to launch the process of *Memorandum of Understanding (MoU)*preparation and signing. The process was extremely hard and problematic - it required many face-to-face meetings and took several months. MoUs had to be signed by each Resource Provider. It was possible to agree only on very general metrics related to capacity of resources provided in the long term. The process of MoUs signing was planned to be a single action. However, they required fulfilling other quality metrics defined by OLAs acknowledged by the sites entering the Grid infrastructure. These metrics were not related to any specific customer.

Even with these simple means, the result of signing of the MoUs was a considerable increase of job success rate (a factor describing the amount of tasks submitted to the Grid that could be completed normally) [9].

6.2 Data Challenges

Other important scenarios for Grid Computing are experiments that need to mobilise large resources for a relatively short period of time. Usually, this kind of experiments, also known as *data challenges*, are planned and co-located with a public event or a tight schedule of some research - the International Telecommunication Union Regional Radio Conference in 2006 may serve here as an example. During this event, representatives from 120 countries negotiated a radio-frequencies plan[7]. The conference lasted about one month and required weekly major revisions of the global plan and daily over-night refinements for certain regions. Both processes were computationally intensive and needed to be completed in a defined period of time, as their results were needed to continue negotiations.

Scenarios like this clearly show a need for warranty that the required resources are available on time and in sufficient amount. Technically, in such well described and scheduled computations, a Service Level Agreement would specify resource reservations. Reservation schedule can be subject of negotiations before signing an SLA.

6.3 Urgent Computing

Urgent computing is a class of applications that typically request large number of computing power at a specific time. Early Warning Systems (EWS) are good example of such applications. New generations of EWS framework, such as the one targeted by the UrbanFlood project[8], are used to extend existing ones with new internet and sensor networks technologies. The EWS targeted in the UrbanFlood project run as an internet service, able to host multiple EWSs, corresponding to various environmental issues and belonging to different organizations and authorities. In such systems, data streams from sensors need to be processed in order to analyze a current status of the monitored systems, make a prediction, validate a model, or recommend an action. Sophisticated simulation models that are computationally intensive are used to process the collected data.

[7] `https://twiki.cern.ch/twiki/pub/ArdaGrid/ITUConferenceIndex/C5-May2006-RRC06-2.ppt`

[8] `http://www.urbanflood.eu`

In case of emergency, the processing of the data needs to be delivered in real time in which case urgent computations might be triggered automatically and require a rapid access to a large computing power.

Even though Grids have the potential infrastructure in term of available computing power to handle such emergency cases, there is no guarantee that the required computing power will be available at the time it is needed, as the current state of Grid administration is based on best effort and the queuing time of the jobs submitted to the Grid varies from a couple of minutes to a couple of hours and sometimes more. Reservation of computing resources is currently the only way to guarantee the availability of computing resource in Grid but is not applicable in the case of urgent computing as it is not known in advance when the resources will be needed. Furthermore, in some emergency scenarios, such as flooding, the Grid infrastructure could be hampered and, thus, Grid sites become isolated and unreachable. For this reason it would be ideal to use resources close to the sensor data so as to limit the point of failures. In major emergencies it would also be ideal to have mutable resources where computing resources stops whatever they are doing and focus all their attention to the emergency.

The solution of urgent computing cases lies clearly in the area of SLM. Additionally, it shows how different SLAs influence each other, by e.g. including an option of killing other jobs based on one SLA by the others in case of emergency. It also shows that SLA should be implemented in the configuration of the site and linked with proper authorization procedure [8]. Those are valid research tasks.

7 Related Works

The Extended Telecommunication Operations Map (e-TOM) [10] is a reference framework promoted by the Telemanagement Forum for processes to be conducted within the network operators and service providers. e-TOM is hierarchical, in the sense that processes are grouped in categories or levels. Among others, it puts special emphasis on service delivery and service level management. In the following paragraphs we summarize the structure of that framework in respect to SLM.

Service Level Management is covered by a level-2 process called Service Quality management. This, in turn, is decomposed into seven level-3 processes. The latter processes are meant to monitor, analyse, improve and report the service quality. In addition, in case of service degradation, these processes track and manage the service quality resolution.

SLAs, OLAs and SLSs are defined within a model (informal model) in [11]. Each of these concepts is related to each other and to a set of actors and metrics to allow determining potential performance degradations according to what is established in the above contracts.

Management of SLA is also of particular relevance within e-TOM. A closely related process in that field is Customer QoS/SLA Management, which is a level-2 process that is decomposed in other six level-3 processes aimed to the assessment and report the SLA fulfilment to the customer. Also, these level-3 processes cover the lifecycle of the process of management and resolution of eventual SLA violations.

Performance management is not only considered in the relationship between the service provider and their customers, but also between the service provider and their partners/suppliers. In this sense, it is worth to mention the level-2 process called Supplier/Partner Performance Management. It decomposes further into five level-3 processes covering aspects like the performance assessment, its reporting and eventual actions to undertake in case that the contracted quality drops below established thresholds. The performance of the service to be provided by a supplier or a partner is also collected in SLAs (Supplier/Partner SLA).

8 Examples of Implementation of SLM Elements in Grids

Although there was no coordinated effort to introduce Service Level Managements in the main European Grid Initiatives (the gSLM project is the first one), there were several attempts to implement some of its aspects to the infrastructures, mainly (but not only) at the national level.

The example of such are projects SLA@SOI[9] and SLA4D-Grid[10]. The first project is concerned mainly with Service-Oriented infrastructures and aimed on industrial use cases. Its main concern is assuring predictability and dependability for the business partners. These features are achieved by introducing an SLA framework for automatic SLA negotiation and management, which may not be possible in such large infrastructures as Grids. The aim of the second project is to design and implement a Service Level Agreement layer in the middleware stack of the German national Grid initiative D-Grid. The targets of the SLAs in the project are warranties of the quality of service and fulfilment of the negotiated business conditions. The SLA4D-Grid project focuses on the tools for automatic SLA creation and negotiation, offering also support for monitoring and accounting, it does not, however, provide a model for an integrated SLA framework enabling interaction with other Grid infrastructures than D-Grid. Important part is effort to standardise SLA negotiation protocols based WS-Agreement [12].

Within PL-Grid Project[11], a project called Grid Resource Bazaar [13] has developed a framework for Service Level Agreement negotiation, designed for resource allocation in the Polish NGI. In this framework, an NGI can act as a mediator between user groups and sites. Users can apply for resources, specifying several optional metrics that are later translated into computational and storage resource configurations. Sites can optionally delegate negotiations of some SLAs to NGI, based on special types of Operational Level Agreement. The process is maintained using a specialized Web-based platform that facilitates complexity management.

[9] `http://www.sla-at-soi.eu/`
[10] `http://www.sla4d-grid.de/`
[11] `http://www.plgrid.pl/en`

9 Summary

In this paper authors advocate introducing Service Level Management to Grid Computing. Extensive analysis shows the need of such solutions and the benefits they could provide. However, it is also clear that adoption of the SLM processes in federated infrastructures is challenging, and requires considerable effort to deliver and maintain. Today, IT infrastructure services are becoming closer to other industrial standards in their approach, in order to meet business requirements of their customers. The Grid infrastructures cannot ignore this trend, lest they will lose their users as a consequence of poor levels of user-satisfaction.

Acknowledgments. This work was funded by EU FP7 gSLM. Tomasz Szepieniec thanks PL-Grid Project for support of this work.

References

1. Schwiegelshohn, U., et al.: Perspectives on Grid computing. Future Generation Grid Computing 26(8), 1104–1115 (2010)
2. Taylor, S., Lloyd, V., Rudd, C.: ITIL v3 - Service Design, Crown, UK (2007)
3. ISO/IEC 20000-1:2011 IT Service Management System Standard
4. Foster, I.: What is the Grid: A Three Point Checklist, Grid Today, July 20 (2002)
5. Plaszczak, P., Wellner, R.: Grid computing: the savvy manager's guide. Elsevier (2006) ISBN: 978-0-12-742503-0
6. Leff, A., Rayfield, J.T., Dias, D.M.: Service-Level Agreements and Commercial Grids. IEEE Internet Computing 7(4), 44–50 (2003)
7. Candiello, A., Cresti, D., Ferrari, T., et al.: A Business Model for the Establishment of the European Grid Infrastructure. In: The Proc. of the 17th Int. Conference on Computing in High Energy and Nuclear Physics (CHEP 2009), Prague (March 2009)
8. Kryza, B., Dutka, L., Slota, R., Kitowski, J.: Dynamic VO Establishment in Distributed Heterogeneous Business Environments. In: Allen, G., Nabrzyski, J., Seidel, E., van Albada, G.D., Dongarra, J., Sloot, P.M.A. (eds.) ICCS 2009, Part II. LNCS, vol. 5545, pp. 709–718. Springer, Heidelberg (2009)
9. Moscicki, J., Lamanna, M., Bubak, M., Sloot, P.: Processing moldable tasks on the Grid: late job binding with lightweight User-level Overlay. Accepted for pub. in FGCS (2011)
10. Business Process Framework, Release 8.0, GB921, TMForum (June 2009)
11. SLA Management Handbook, Release 3.0, GB917. TMForum (May 2010)
12. WS-Agreement Specification, version 1.0. Approved by OGF, `http://forge.gridforum.org/`
13. Szepieniec, T., Tomanek, M., Twarog, T.: Grid Resource Bazaar. In: Cracow Grid Workshop 2009 Proc., Krakow (2010)