University of Amsterdam

UvA-DARE (Digital Academic Repository)

gSLM: The Initial Steps for the Specification of a Service Management Standard for Federated e-Infrastructures

Serrat, J.; Szepieniec, T.; Belloum, A.; Rubio-Loyola, J.; Appleton, O.; Schaaf, T.; Kocot, J.

Link to publication

gSLM: The Initial Steps for the Specification of a Service Management Standard for Federated e-Infrastructures

Joan Serrat[1], Tomasz Szepieniec[2], Adam Belloum[3], Javier Rubio-Loyola[4], Owen Appleton[5], Thomas Schaaf[6], Joanna Kocot[2],

[1] Universitat Politècnica de Catalunya, Barcelona, Spain
serrat@tsc.upc.edu
[2] ACC Cyfronet AGH, Krakow, Poland
{t.szepieniec, ymkocot}@cyfronet.pl
[3] University of Amsterdam, Amsterdam, The Netherlands
a.s.z.belloum@uva.nl
[4] Cinvestav, Tamaulipas, Mexico
jrubio@tamps.cinvestav.mx
[5] Emergence Tech., Norwich, UK
owen@emergence-tech.com
[6] Ludwig-Maximilians-Universität, Munich, Germany
schaaf@nm.ifi.lmu.de

Abstract. This paper presents a methodology used to create a site independent assessment process of the capabilities of service management systems in federated e-infrastructures that can contribute to introduce or improve service management in these application domains. Based on ISO/IEC 20000 concepts it consist of an actors and relationships model, a set of management processes with their corresponding requirements and a capability model that all converge in an easy to use assessment tool. The methodology has been evaluated through its adoption in an existing federated e-infrastructure.

Keywords: IT Service Management; e-Infrastructures; Actors Model; Capability Model; Management Processes.

1 Introduction

Federated computing services such as Grid, Cloud and other e-infrastructures are increasingly important to scientific and technical research. Grid Computing started as an approach to realize the virtualization of distributed resources. Initially the services offered by Grids and federated infrastructures were best effort-based and limited to the scientific community but when Grid adopted the service-oriented paradigm it started to attract more and more commercial applications. However, much stricter guarantees are needed to attract these applications to use Grid-based systems.

Cloud Computing emerged as a much more commercial oriented approach provided that resources were owned by a unique service provider who could apply

commercial IT service management techniques. Nevertheless, as soon as the advantages of federated service delivery and management in the Cloud domain were realized, similar challenges to the Grid domain had to be faced.

Federated e-infrastructure scenarios are characterized by two main facts. First, there is no single central authority with control over service delivery, service levels and management processes. Second, there is no hierarchic service chain with clear distribution of responsibilities. In addition, most of today's Grid infrastructures are set on the grounds that service commitments and dedication of resources to the Grid are on a voluntary basis and in general there are no formal agreements on service functionality and quality. For these reasons, well known established IT service management frameworks like ITIL [1], ISO/IEC20000 [2] or others are not straightforward applicable.

Although the need of managing the services offered by e-infrastructures is widely accepted, there is not a common understanding on how to achieve the goal. For that reason, in September 2010 the EU-funded project gSLM [3] was started aiming a framework for the introduction and development of IT service management in such e-infrastructures. Since the beginning, gSLM understood that rather than proposing the implementation of service management in full at once, a grading of its level of introduction would be preferable in order to show the process as a set of achievable steps from no service management at all to commercial-level service management compliant with systems such as ITIL or ISO/IEC20000. The key was the use of the capability and maturity model concepts. A maturity model is an abstraction of a system, which stratifies the maturity of the system into an easily manageable set of levels, based on the capabilities of each of its constituting processes. Each level is accompanied by a description of diagnostic features of that level of maturity. It helps building a common understanding of maturity such that they can be improved, leading to the improvement of the overall system as well.

The capability/maturity model adopted in gSLM had to be framed in an actors and business models as well as in a set of well-defined service management processes that constitute the baseline to determine which functional aspects to consider. Finally, a set of requirements to be fulfilled by an e-infrastructure to achieve a given level of capability in each of the different functional aspects was issued. All in all this was part of the contribution of gSLM that after its two years duration, ended with a set of results that were all grouped in a methodology what we will call the "gSLM approach". This project has had a follow-up called FedSM [4] that is still ongoing and that will end with a standard for service management [5]. Hence gSLM was the precursor setting the grounds for such a standard.

This paper presents the gSLM approach with some details on its components and the evaluation that took place within the project. After this introduction, Section II is devoted to present a snapshot of the state of the art in service management of federated e-infrastructures. Section III presents the methodology that constitutes the core of our contribution. Section IV enters at the detail that the space restrictions are allowing in the supporting components of the gSLM approach. Here we describe the proposed actors' model, management processes, and the requirements that an infrastructure must fulfil to reach given capability levels. Section V presents a summary of the results of applying our approach in a federated Grid environment. Finally, the paper ends with Section VI presenting our concluding remarks.

2 Snapshot of Service Level Management Maturity in e-Infrastructures

As the starting point encountered at the beginning of the gSLM project we summarise the status, in terms of service management development, existing in six well known e-infrastructures. This is by no means an exhaustive list of e-infrastructures but a representative set of the most important ones.

The Open Science Grid (OSG) [6] is a multi-disciplinary partnership in USA to federate local, regional, community and national cyber-infrastructures to meet the needs of research and academic communities. Offered services are subject to formal Service Level Agreements (SLA). They are published on a Wiki page with public access, so that both, users and providers are able to browse them. The services are organised in a kind of a catalogue, and every new service is required to follow a path of SLA specification – external services included. In addition, there are change management processes in place. Therefore we could qualify the maturity of OSG as quite high in the full range of aspects considered.

The European Grid Infrastructure (EGI) [7] is a series of efforts launched in 2010 to provide access to high-throughput computing resources across Europe using grid computing techniques. EGI worked-out the Operational Level Agreement (OLA) framework, which defines relationships between sites, national grid infrastructures and the central EGI. Those OLAs are without SLA context so they define general service level for all the users. Considering SLAs, which should define relations between a customer and provider, we have to say that the level of maturity is quite lower on this field. In fact, EGI has not fully determined its role in the process of delivering services for grid users. The heterogeneity in service management adoption depending of the specific functional aspects considered yields to the conclusion that EGI maturity has to be ranked from low to high.

The Swiss National Grid (SwiNG) [8] was launched in 2006 to coordinate grid services in Switzerland. The level of formalisation and automation of the processes and procedures differs very much, depending on the area of interest – from semi-automated procedures with clear contact points and responsibilities (these are mainly related to security), to less formal (the rest of the processes). The infrastructure is aware of the need to manage the reliability of the services provision, but does not implement accordant mechanisms at the moment. Only some services provided within the infrastructure are monitored with automated monitoring tools, which are used also for accounting purposes. From the service level management perspective, SwiNG is then a quite low mature e-infrastructure.

The Nordic DataGrid Facility (NDGF) [9] is essentially a production grid facility that leverages existing national computational and storage resources and grid infrastructures in Scandinavian countries. The contractual framework is limited to OLAs between NDGF and sites. However, the OLA coverage is still incomplete, as some sites are considered operational in the infrastructure without signed OLA, and the agreements may not be consistent. The incident reporting is mainly automated, and the escalation path is more-or-less known, but not formally defined and includes informal contact means. Even when there are written agreements between the infrastructure and the sites, the enforcement is moderate and does not seem to follow

a defined path. NDGF is also a low mature infrastructure comparable in that sense to SwiNG.

The PRACE RI [10] is a pan-European infrastructure responsible for High-Performance Computing resources and services for public research. PRACE uses SLAs (OLAs) to specify and control commitments of the sites it incorporates. PRACE has also SLAs with its users that usually define the user's responsibilities and means of reporting incidents, however, the services are always considered best-effort without real warranties. The infrastructure offers a service catalogue, called "Application and tool catalogue", with distinguished service levels, however, nothing can be enforced on the sites and no escalation procedures exist. At the operation level, there is an escalation procedure and a kind of help-desk, taking care that the user's incident reports are transferred to the party on whose side the problem is supposed to lay. PRACE is relatively low mature in some management functional aspects but quite mature in others. Like EGI his maturity spans through several levels.

Amazon Web Services (AWS) [11] offers a complete set of infrastructure and application services materializing the cloud computing paradigm for this corporation. Each of the Amazon's services is described with an SLA with non-trivial metrics. The performance of the provided services is thoroughly measured and monitored by dedicated tools. All the procedures are well-defined and documented. The automation of all necessary processes makes the service provision management easier, and, thus, allows focusing more on the user needs and expectations – following a more user-centric approach. The overall maturity level of the Amazon Web Services may be evaluated as the highest found in all the above considered e-infrastructures.

Summarising, the above six considered e-infrastructures exhibit a large span of different levels of service management maturity. We note that heterogeneous infrastructures like Grids, which have an academic scientific background, lag behind commercial approaches like Amazon Web Services, because efforts are made mostly within Grid sites, while inter-organisational approaches to SLM are hardly developed. More difficult would be to quantify the levels of maturity and much more difficult would be to reach the same scores when evaluated from different actors or perspectives. In this context the identifiable contribution of the gSLM project is to bring an unambiguous and repeatable way to quantify the level of maturity of e-infrastructures irrespectively of its nature and purpose. This is a cornerstone to advance in a coherent way to reach ever increasing levels of SLM maturity, which is cost effective and at a pace tailored to each e-infrastructure organization.

3 Objective and Methodological Approach

gSLM aims to establish a baseline that can be adopted by federated e-infrastructures in order to determine their maturity level and future improvement plans in a unified way, independently of technology, business model and geographical distribution of its federation members. This is a fundamental tool for federations in order to understand, under a standard point of view, how they have to evolve in order to reach given levels of service delivery and service level management. Our perspective should be as broad

as possible in order to cover the most important service lifecycle management activities and not be concerned only with the very restricted niche of SLAs as the majority of efforts have done up to now.

3.1 Actors, Processes and Requirements

ISO/IEC 20000 is a much more simplified framework than the broadly accepted de-facto standard ITIL v3 in the commercial IT service management domain. In that sense the former is a lighter and more generic approach. This fact and also that our target service management system is not grounded in a conventional IT service management scenario, where a single authority acts as unique service provider, were the main reasons to adopt ISO/IEC 20000 as our starting point. But ISO/IEC 20000 is not a solution per-se. Instead it consist of a collection of management processes that are defined along generic requirements that have to be fulfilled in each of them. The challenge to solve our problem in the context of federated e-infrastructures consisted in finding an appropriate business model that made this set of processes applicable and specialize these processes in terms of appropriate and concrete requirements. Here the term appropriate is understood as satisfying at a time simplicity to allow for its implementation in existing infrastructures and completeness to ensure a good manageability.

Adopting the ISO/IEC 20000 framework doesn't mean that we have to consider all the processes herein defined or that we cannot use any other not explicitly recognized. The set of processes that were under the gSLM scope were those in direct support of the business model adopted by the federated e-infrastructure. Then, in case of a business model where the federator were a simple mediator or service broker, its responsibility and duties in respect to the service users would be much less than in case it were a full service integrator, adopting the role of a full fledge service provider. In summary the first challenge was to select the relevant set of management processes.

Having decided which processes to consider the next step would be the specialization of each one in terms of requirements to be fulfilled in the implementation of each of these processes in our problem domain. Really, ISO 20000 processes are generic in nature both in terms of the activities entailed as well as in terms of the subjects and objects to which they apply. Subjects and objects would be in fact the actors that would execute and receive the actions of the corresponding activities. The challenge at this point was to decide the granularity level to decompose in activities each of the above selected processes and at the same time identifying the actors that would be involved and their mutual relationships.

Facing the above challenges was done under a pragmatic perspective. Based in the consortium expertise, which covered various players in the federated e-infrastructure arena, and after an analysis of the trends of an amount of representative cases that included those presented in Section 2 [6-11], we decided to implicitly adopt a business model of a *full service integrator*. That is, we would assume the existence of a federator entity entrusted to take the responsibility of delivering the services and managing the corresponding service levels, acting as a single contact point between

users and infrastructure suppliers. That said, a proposal was detailed with the aim to be refined after an evaluation process in a federated e-infrastructure environment.

3.2 Capability Model

Once the requirements structured around preselected processes were defined, the question was how to establish ways to quantify up to which extend a given e-infrastructure was supporting a determined requirement. The number of quantification levels should be necessary greater than two (YES/NO) in order to better capture the real status of the infrastructure. To this aim we adopted a solution based on the concept of *Capability*. This term is expressing the quality of implementation within a single process. Capability is measured in levels. This means that for each capability level we have to specify the conditions to be fulfilled to grant that level. Therefore we will be able to say that the e-infrastructure "x" is supporting the management requirement "y" with capability level "n" that at the end was one the objective pursued in the gSLM project.

Capability levels for all the management requirements can then be stablished as targets to be reached in order to have a service management system characterized by a determined level of maturity. In other words, maturity levels characterize a whole service management system as capability levels characterize a single requirement or process. In gSLM we adopted capability levels per single requirement in each of the selected processes. Capability and maturity models have been traditionally used in assessment and governance frameworks like ISO/IEC 15504 [12], CMMI [13] and COBIT [14]. A six-level model inspired in these frameworks was precisely the adopted one for our purpose. The challenge now was to refine this six level generic model for each requirement defined in the previous steps. This phase of our methodology would end with a description, per requirement and capability level, of the conditions that if fulfilled would grant that capability.

Last, as the process of ranking an e-infrastructure could be tedious and prone to errors we developed a simple application that by means of NO/YES answers checks the fulfillment of the requirements catalog and depicts the corresponding capability levels using different colored marks. Nevertheless, as we have pointed out the conclusions, we realized through the evaluation of this framework that the use of the tool has to be done by really knowledgeable people to avoid misinterpretations and hence false assessments.

4 Detailed Components of the gSLM Approach

In this section we elaborate on the design elements of the methodological approach outlined earlier.

4.1 Actors and Relationships Model

Three principal actors are envisaged, namely Virtual Organization, Federated Infrastructure and Site. These actors and part of their main relationships are graphically depicted in Figure 1.
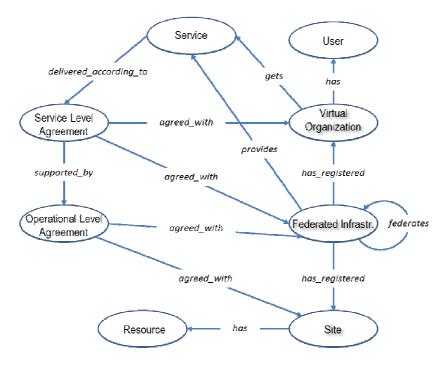

Fig. 1. Primary actors and part of their main relationships within the gSLM model

A Virtual Organization (VO) consists of a group of people that are the users of a given service. The VO represents collectively all this set of users in respect to the interactions with the entity that will provide the services.

A Federated Infrastructure (FI) is an approved body that provides e-infrastructure-based services in a region, country or group of countries. FIs may be organized in larger bodies, creating a hierarchical structure, with primary FIs federated in secondary FIs, etc. The role of a FI in the value chain can range from a simple mediator between VOs and Sites to more complex responsibility to ensure that services are provided with a pre-agreed quality, thus acting like a service provider.

A Site is an administrative independent domain that brings the resources to be federated with resources provided by other Sites and offered as a service through the participation of a FI. From the perspective of our model, Sites are then suppliers of FIs and therefore they are not directly dealing with VOs.

4.2 Processes Considered

According to ISO/IEC 20000 and also taking into account the need to have a clear understanding and management of the service portfolio, the following fourteen processes were envisaged

Table 1. Management processes

Name	Abbreviation
Service Portfolio Management	SPM
Service Level Management	SLM
Service Reporting	SR
Service Continuity & Availability Management	SCAM
Capacity Management	CapM
Information Security Management	ISM
Customer Relationship Management	CRM
Supplier Relationship Management	SuppM
Incident & Service Request Management	ISRM
Problem Management	ProbM
Configuration Management	ConfM
Change Management	ChM
Release & Deployment Management	RDM
Continual Service Improvement	CSI

The semantics and scope of each of these processes is therefore the same as defined in ISO/IEC 20000.

4.3 Specification of Requirements

A minimal set of requirements for an effective service management system (SMS) were defined keeping in mind the double aim of providing enough structure for services to be well managed on the one hand, but keeping the number small enough to have a realistic expectation that they can be implemented. This requirements include key activities for IT service management (ITSM) processes as well.

Requirements were classified in two categories, namely General Requirements and Process Specific Requirements. Both types shall be taken into account by a Federator as relevant criteria for implementing a service management system (SMS).

General Requirements are related to the top management responsibilities, documentation creation and maintenance, the scope of the service management and the planning, implementation and monitoring/reviewing of the service management system. For example, concerning the top management responsibilities, these General Requirements specify that top management shall provide evidence of its commitment to planning, implementing, operating, monitoring, reviewing, and improving the service management system (SMS) and services. In addition it concretes as duties of the top management the assignment of a management representative, the definition

and communication of goals, the definition of a general service management policy and the conduction of management reviews at planned intervals.

Process Specific Requirements were defined for each one of the fourteen considered processes. As an example we present these requirements for the Service Portfolio Management and Service Level Management processes

Table 2. Requirements of the Service Portfolio Management process

Process	Requirement
Service Portfolio Management	PR1.1: A service portfolio shall be maintained. All services shall be specified in the service portfolio.
	PR1.2: Design and transition of new or changed services shall be planned. Plans shall consider goals, acceptance criteria, timescales, responsibilities, new or changed technology, new or changed SLAs, testing and communication.

Table 3. Requirements of the Service Level Management process

Process	Requirement
Service Level Management	PR2.1: Services to be delivered shall be agreed with customers. SLAs shall include agreed service targets.
	PR2.2: A service catalogue shall be maintained.
	PR2.3: Services and SLAs shall be reviewed at planned intervals.
	PR2.4: Service performance shall be monitored against service targets.
	PR2.5: For supporting services or service components provided by Federation members, OLAs shall be agreed.

A total of 61 process specific requirements were identified that means and average of about 4 requirements per process in line with the above mentioned idea of simplicity. The exhaustive list of requirements can be obtained in [15].

4.4 Capability Model

To evaluate and categorize the level of implementation of the process specific requirements standard requirements, a process Capability Model has been established. The proposed Capability Model is composed of six Capability Levels that will characterize the status of Service Management implementation in each process and it will serve as basis for maturity assessments

The six Capability Levels are defined as follows.

Non Existent. In this capability level Service Management is nonexistent. The service provider is not aware of the tasks necessary to provide the service. In a practical world, as soon as a process is defined this level is surpassed.

Ad-hoc. The service provider is aware of the tasks needed to provide the service but the execution is undocumented, uncontrolled and reactive. Success in accomplishing the task is not guaranteed and it is likely to depend on individual efforts.

Repeatable. The service provider has a solid understanding of activities to be performed and that, most of the time, will lead to repeatable results or outputs. Tasks are realised mostly intuitively since the documentation and recording are poor.

Defined. The service provider has clearly defined and documented procedures, roles and responsibilities, but process effectiveness and efficiency are not measured nor reported.

Managed & Measured. The service provider has clearly defined and documented procedures, roles and responsibilities. Process effectiveness and efficiency are measured and reported, although, this information isn't necessarily leading to process change and improvement.

Efficient/Optimal. For the regarded process Service Management System is fully implemented, meaning that Capability Level Managed & Measured is enhanced by using information to systematically improve the regarded process.

Although the highest capability levels are of singular importance in commercial IT service management, in the context of federated infrastructures and more in particular in Grid computing services, where service management has been dealt with other perspectives, we understood that reaching more basic levels would be enough in terms of manageability and simplicity of implementation. For these reasons it was decided to develop only the first four of the six levels above.

On the other hand, to make use if these capability levels it was necessary to specify each one for every one of the process specific requirements. That would result in a granularity that would make it a useful tool for the process administrators at the time of evaluating the SMS. As an example, the following table provides the concretion of the Ad-hoc, Repeatable and Defined levels for the requirements of the Service Portfolio Management process. Note that the Non-Existent level is trivial and therefore doesn't need any concretion.

Table 4. Requirements of Service Portfolio Management process with capability descriptions

Process	Requirement	Capability Level	Description
Service Portfolio	PR2.1	Ad-hoc	The organization is aware and can generally describe the service offer

Management			(past, current, planned) and related information through some (undefined) format.
		Repeatable	There is a clear understanding of the service offer (past, current, planned), relationship to support activities and other related information that connect to the value creation capabilities. This list is maintained on an informal basis.
		Defined	There is a defined and documented procedure for maintaining the service portfolio, clearly specifying differentiated service offerings and links to strategic objectives. This is shared across the organization.
	PR2.2	Ad-hoc	Transition to new services or changes to current services are managed with no structured approach or control. The quality of the output depends on the individual efforts.
		Repeatable	There is an understanding of a structured approach for the transition to new services or changes to current services that is applied routinely, but this is not documented. Acceptance criteria, timescales are managed at a qualitative level.
		Defined	There is a clearly documented procedure to handle service transition or change that includes all the mentioned attributes.

5   Results of the Model Adoption in a Real e-Infrastructure

In this section we present the outcomes of gSLM approach applied to a real e-infrastructure. For this purpose we choose the Polish National Grid Initiative (PL-Grid) [16], which maintains grid infrastructure provided by 5 main computing centers in Poland offering around 25k cores and serving about 1500 individual scientists. PL-Grid is an example of federated organization with centralized control of services supported by distributed computer centers, which not only delivers resources but owns them and decides about their usage. In the last two years PL-Grid made a substantial effort to establish an SLA/OLA framework [17]. Therefore the gSLM assessment method was interesting for The PL-Grid staff to check the current SMS

solutions for completeness and define a way to improve them. On the other hand it was for us an excellent opportunity to evaluate our assessment methodology and improve it through the observations and concerns of the operations people. For that reason the evaluation process was carried out in close collaboration between the gSLM consortium and the PL-Grid staff.

Prior to the execution of the assessment, it was assumed that the role of the Federated Infrastructure (FI) is adopted by the PL-Grid. VOs and Sites were also easily mapped to actual users' groups and computing centers. In addition, the definition of existing operation activities at PL-Grid were mapped to activities of our model to assure that there was a similar understanding about what was conceived in the gSLM model and by the operations staff of PL-Grid. The matching was facilitated by the fact that the gSLM approach is based on very simple atomic management activities.

At the moment of writing this paper only seven out of the fourteen processes have been evaluated. The reason to proceed that way is precisely to get feedback from the evaluation of an initial set of processes. In particular, the processes included in this first phase were SPM, SLM, SR, CRM, SRM, ISRM and CSI. In summary, as explained in detail hereafter, we concluded that the gSLM framework is matching the reality and trends in the PL-Grid federation and therefore that the gSLM methodology fulfils its design goals. The following tables present the evaluation results of the Service Portfolio Management (SPM) and Service Level Management (SLM) processes. It is interesting to note the comments within the rationale of each ranking.

Table 5. PLGrid's Service Portfolio Management process evaluation

Process	Requirement	Self-assessment Score	Rationale for Score & Comments
Service Portfolio Management	PR1.1	Repeatable	PLGrid service portfolio is published on the website for users. It may need some reorganization to distinguish value creating capabilities and support activities. Services are defined on such abstract level that we do not expect to add or retire services. We do not have a documented procedure for maintaining the service portfolio.
	PR1.2	Defined	Changes at the service portfolio are not expected. However new technical service components are being added following well-defined procedures requiring service availability monitoring, establishing a support team, security and operational audits. These are mainly new scientific applications, web based domain-specific portals and tools.

Table 6. PLGrid's Service Level Management process evaluation

Process	Requirement	Self-assessment Score	Rationale for Score & Comments
Service Level Management	PR2.1	Defined	Applying, negotiating and signing an SLA is implemented using PLGrid User Portal and Bazaar tool. Any user (or group) willing to use PLG resources needs to apply for a computational grant specifying at least two metrics which are reflected in our portfolio: amount of wall-time and storage space. Other predefined metrics are possible regarding technical aspects of use of resources.
	PR2.2	Defined	Compute and storage services are described together with additional metrics defining technical aspect of use (some of them still being defined). Availability of scientific applications is listed on a regularly updated web page.
	PR2.3	Repeatable	Each SLA is periodically reported by users who can provide their comments about use of resources. Compliance with SLA targets is checked. Appropriate actions taken for violated SLAs. Some processes are covered by documented procedures. We do not have regular service reviews.
	PR2.4	Repeatable	The data for confronting sites with their agreed offerings exists however there is no procedure to follow. Technical service availability is monitored on daily basis, alarms are generated in case of failures.
	PR2.5	Defined	In fact, each SLA is broken into a set of agreements between Customer and actual resource provider (sub-SLAs) which PLGrid Operations Centre monitors and endorses.

Similarly, all the other requirements were evaluated. In total there were 30 requirements out of which 4 were ranked as *Ad-hoc*, 13 as *Repeatable* and 13 as well

as *Defined*. Establishing as global maturity level target having all these evaluated processes at capability level 3 we can conclude that the actual maturity level of PL-Grid is medium to high. In addition having a closer look at the four requirements that have lower marks we observe the following. In no particular order, the first one is PR3.2, in the SR process, which states that *the description of each service report shall include its identity, purpose, audience, frequency and content*. PLGrid produces service reports but these ones do not contain all these mentioned fields. The second requirements ranked at ad-hoc level is PR7.6, belonging to the CRM process, which states that *customer satisfaction shall be managed*. The reason for this low ranking here is because although users can express their overall satisfaction on PLGrid and appropriate reaction can be triggered, there are no means to measure user satisfaction. The third requirement ranked as ad-hoc is PR8.4, belonging to SRM process, stating that supplier performance shall be monitored. The reason here is because there is no specific activity related to performance monitoring, although individuals of the PLGrid operation staff often exchange their impression with suppliers. Finally, the fourth low ranked requirement was PR9.8, belonging to ISRM process, which states that *there shall be a definition of a major incident and major incidents shall be classified and managed according to a documented procedure*. The reason for the ad-hoc level is because no notion of a major incident exists at PLGrid although problems affecting many users or causing severe consequences are handled with a special care including extensive information and quicker response times.

The requirements and capability levels associated to each were found clear by the operations staff but not without a quite detailed explanation of the meaning and scope of each one along the process of evaluation, which was provided by personnel of the gSLM consortium. This means that it is almost mandatory the organization of training workshops for the operations personnel on the above described methodology and concepts, before them can be competent to carry out a self-assessment of the capabilities and hence to quantify the maturity of the service management system of an e-infrastructure. This is also a key aspect to assure that these people will find themselves comfortable with the introduction and deployment of that methodology in their service management systems.

6 Conclusions

To the best of our knowledge the contribution of this paper consists of being the first in formalizing the adaptation of IT service management best practices, in turn well established in single-administrative-domain IT service delivery, to the federated e-infrastructure contexts. Our approach is based in an actors and relationships model with well-defined roles and activities that were adopted from the ISO/IEC 20000 framework and conveniently adapted to the above mentioned federation domains. In the implementation presented in the paper we instantiated that model for the case where the e-infrastructure federator (the FI as called throughout the paper) adopts the role of a service integrator, but other business models are possible as well, thus making in practice that approach flexible enough to adapted to any real application context.

Although one of the components of our contribution is an online assessment tool, our contribution has not to be seen as an automatic black box that gives a result when a set of environmental conditions are fed. Instead, it has to be understood as a guidance methodology that requires careful judgment and expertise if the outcome of its application is pretended to be meaningful. This has been evidenced by the evaluation of the methodology we carried out in the PL-Grid, in close collaboration with its operations staff. Lesson learned from that evaluation were indeed incorporated in our methodology through several refinement loops. We believe that similar conclusions could be derived from the adoption of the methodology in other federations.

Our future plans include the evolution of the gSLM methodology through the inclusion of feedback from other stakeholders and we are already working together with new partners in the context of a recently launched project. Our goal is to make network and service management and instrument that leverage federated e-infrastructures in the industry and the research community. In that sense we can claim that at the time this paper was completed, PL-Grid is starting with the first changes according to an improvement roadmap designed by means of the gSLM approach here presented.

References

[1] Information Technology Infrastructure Library (ITIL) 2nd edition, s.l. Stationery Office, 2011.

[2] ISO/IEC 20000-1:2011 IT Service Management System (SMS) Standard

[3] 3Service Delivery and Service Level Management in Grid Infrastructures – the gSLM project. http://www.gslm.eu, 2014

[4] Implementing Service Management in Federated Infrastructures. http://www.fedsm.eu/, 2014

[5] Standards for Lightweight IT Service Management. http://www.fedsm.eu/fitsm, 2014

[6] The Open Science Grid. https://www.opensciencegrid.org/bin/view, 2014

[7] The European Grid Infrastructure. http://www.egi.eu, 2014

[8] The Swiss National Grid association. http://www.swing-grid.ch, 2014

[9] The Nordic DataGrid Facility. http://www.ndgf.org/ndgfweb/home.html, 2014

[10] The Partnership for Advanced Computing in Europe. http://www.prace-ri.eu , 2014

[11] Amazon Web Services. http://aws.amazon.com, 2014

[12] ISO/IEC 15504:2004 Information technology — Process assessment

[13] CMMI v1.3. http://www.sei.cmu.edu/cmmi/index.cfm

[14] IT Governance Institute. Control Objectives for Information and Related Technology. 4.1

[15] D3.2: Minimum requirements for service management in Federated e-Infrastructures. http://www.fedsm.eu

[16] Polish National Grid. http://www.plgrid.pl/en, 2012

[17] Tomasz Szepieniec at al. Implementation of Service Level Management in PL-Grid Infrastructure. Lecture Notes in Computer Science Volume 7136, pp 171-181, 2012