2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing

Exploring the Potential of Using Multiple e-Science Infrastructures with Emerging Open Standards-based e-Health Research Tools

M. Riedel, B. Schuller, M. Rambadt,
M.S. Memon, A.S. Memon,
A. Streit, Th. Lippert
Jülich Supercomputing Centre
Forschungszentrum Jülich
Jülich, Germany
m.riedel@fz-juelich.de

S.J. Zasada, S. Manos, P.V. Coveney
Department of Chemistry
University College London
London, UK

Felix Wolf
Laboratory of Parallel Computing
German Research School for Simulation Sciences
Aachen, Germany

Dieter Kranzlmüller
Ludwig Maximilians University Munich
Munich, Germany

***Abstract* — E-health makes use of information and communication methods and the latest e-research tools to support the understanding of body functions. E-scientists in this field take already advantage of one single infrastructure to perform computationally-intensive investigations of the human body that tend to consider each of the constituent parts separately without taking into account the multiple important interactions between them. But these important interactions imply an increasing complexity of applications that embrace multiple physical models (i.e. multi-physics) and consider a larger range of scales (i.e. multi-scale) thus creating a steadily growing demand for interoperable infrastructures that allow for new innovative application types of jointly using different infrastructures for one application. But interoperable infrastructures are still not seamlessly provided and we argue that this is due to the absence of a realistically implementable infrastructure interoperability reference model that is based on lessons learned from e-science usage. Therefore, the goal of this paper is to explore the potential of using multiple infrastructures for one scientific goal with a particular focus on e-health. Since e-scientists gain more interest in using multiple infrastructures there is a clear demand for interoperability between them to enable a use with one e-research tool. The paper highlights work in the context of an e-Health bloodflow application while the reference model is applicable to other e-science applications as well.**

e-Health, HPC, HTC, Reference Model, Interoperability

I. INTRODUCTION

e-Health makes use of information and communication methods and tools to support the understanding of body functions. It thus plays a significant role in improving the health of the global population. While e-Health is a large scientific research field, we focus in this contribution on recent work towards the *Virtual Physiological Human (*VPH*)* [1]. In this context, e-Scientists take already advantage of single e-science infrastructures such as Grids to perform computationally-intensive investigations of the human body that tend to consider each of the constituent parts separately without taking into account the multiple important interactions between them. These subdivisions make it impossible to investigate the systemic nature in which the body functions. Many e-science applications in this area are limited by the computational power provided in their respective e-science infrastructures, because the power has to be shared with numerous other applications in science and engineering. In contrast, the VPH vision is a methodological and technological framework that enables collaborative investigations of the human body as a unique complex system. This initiative is part of the larger international Physiome Project that *clearly raises the demand in its roadmap [2] for making world-wide e-science infrastructures interoperable so that it becomes suitable for collaborative research* in order to tackle the computational intensive challenges that the simulation of the VPH addresses. The VPH community seeks to serve the development and integration of multi-scale models, which have different computational requirements ranging from single processor desktop machines to the largest supercomputers available in *different kinds of production Grids* today.

We have shown in earlier work [3] that interoperability of Grid infrastructures, among others, can be considered as one approach to perform e-science. Hence, world-wide interoperable Grids are not a unique demand of the VPH community and are emerging in many different scientific research fields. One particular example that shows the effectiveness of jointly using different kinds of production Grid infrastructures is the use of seamless access to both *High Throughput Computing (*HTC*)* driven Grids and infrastructures driven by *High Performance Computing (*HPC*)* needs. This is attractive since HTC resources are cheaper and do not provide a good interconnection between cpus/cores while HPC resources are costly but provide good interconnections between cpus/cores and are capable of running large-scale applications. For instance, we have

978-0-7695-4039-9/10 $26.00 © 2010 IEEE
DOI 10.1109/CCGRID.2010.96

shown in earlier work [4] how e-Health scientists of the WISDOM project benefit from the interoperability between the HTC-based EGEE Grid [5] and the HPC-driven DEISA infrastructure [6]. Here, the seamless access to different kinds of e-science infrastructures enables a cheaper and faster drug discovery. Apart from the beneficial access to different kinds of systems, Rodero et al. proved in [7] that *infrastructure interoperability can improve the overall system performance and significantly enhance the resource utilization* in interoperable Grids. In this paper we outline the landscape of today's production Grids and point to the problems why interoperable infrastructures are not seamlessly provided. We argue that this is due to the absence of a realistically implementable *infrastructure interoperability reference model (*IIRM*)*, which we have already applied to different use cases from different scientific domains [8]. This paper highlights one specific IIRM instance that is used to enable multi-site blood-flow applications that contribute as one step towards the VPH.

The remainder of this paper is structured as follows. After the introduction in Section 1, the scene is set in Section 2 where we survey the state-of-the-art e-science infrastructures with a particular focus on interoperability challenges. Section 3 presents our proposed IIRM design and its core elements. Section 4 gives insights to one particular e-Health application that leverages the implementation of the IIRM between various Grid infrastructures and highlights challenges encountered in interoperability and computational steering. Finally, after surveying related work in Section 5, we present our conclusion in Section 6.

II. State-of-the-Art e-Science Infrastructures

Today, scientists regard computational techniques as the third pillar alongside experiment and theory [3]. As an addition to these three pillars that are the foundations for science, the term *enhanced science (e-science)*, sometimes also called electronic science evolved in the last couple of years. We base our considerations on the description of e-science, which has been given by John Taylor [9] and that is defined as follows: *e-science is about global collaboration in key areas of science and the next generation infrastructure that will enable it*. Often, this definition has been extended in several ways to include emerging technologies. For instance, recently, by the addition of dynamic deployment features using virtualization techniques of so-called clouds to the required features of next generation infrastructures. Nevertheless the definition is still valid and we keep this mature definition as a base for our discussions.

The above mentioned next generation infrastructures can be considered as a solid basement for the three pillars of traditional scientific computing and thus enable e-science that can be seen as a roof that stands for the collaboration in key areas of science. At the time of writing, these next generation infrastructures are implemented as Grids, often named as e-science infrastructures in the rather scientific domain. But over the years, various types of e-science infrastructures evolved that can be basically *classified into different categories* according to their services in general and their offered resource-types in particular. The first category is represented by HPC-driven infrastructures since Grids of this type integrate mostly large-scale clusters or supercomputers to enable *massively parallel applications*. Two famous examples in this category are TeraGrid in the US and DEISA in Europe. In contrast, the second category is represented by HTC-driven infrastructures since Grids in this category are mostly tuned to support *farming applications (aka embarrassingly or nicely parallel jobs)* that does not require a good interconnection between single cpus/cores. Known infrastructures of this type are EGEE in Europe, OSG in US, or NorduGrid in the nordic regions.

There are also some infrastructures that we consider as hybrid infrastructures since they provide access to a limited set of large-scale facilities while still providing also access to pc pools and smaller clusters. Infrastructures of this category are, among others, the National Grid Service (NGS) [10] of the UK, or the German national Grid DGrid. In theory it is basically hard to define clear boundaries for this classification. In practice, however, the boundaries of these categories are fundamentally different, especially when the resource type is considered as well as the overall usage policies. We also learned from this categorization that these different e-science infrastructures will remain since it breaks basically down to the physical hardware (i.e. cpu/core interconnection) or usage paradigms such as HPC and HTC concepts.

The major problem of these different categories lead to the fact that there are a well-known set of world-wide non-interoperable *Grid islands* majorly funded through public sources today. A deeper investigation [11] into these Grid islands reveals that each e-science infrastructure runs its own technology. Even within one category, different middleware technologies exist. While TeraGrid uses Globus, DEISA is based on UNICORE. EGEE has deployed gLite, while OSG uses VDT, which is largely based on Condor, and NorduGrid uses ARC. The hybrid category is a bit different, since NGS runs the OMII-UK software stack that includes different software technologies (i.e. GridSAM) and D-Grid became much known by its concept of deploying UNICORE, Globus and gLite in parallel.

At the time of writing all these different Grid middleware systems are as a whole not interoperable with each other. We argue that this is due to the absence of a realistically implementable infrastructure reference model in Grids. As a consequence, today's e-science infrastructures still struggle to provide e-scientists with a common access to all different types of Grids in order to fully leverage the power of globally existing resources by one particular e-science project.

III. The Infrastructure Interoperability Reference Model based on Improved Standards

Taking the review of the state-of-the-art e-science infrastructures into account, we still observe a slow adoption of the Open Grid Services Architecture (OGSA) initially defined by Foster et al. [12]. While OGSA represents a good architectural blueprint for Grids in general, we argue that the scope of OGSA is too broad to be realistically implementable in today's production e-science infrastructures in particular. This is clearly indicated by the complexity of the latest published OGSA-Roadmap document [13] and its hugh amount of services and profiles that raise economic considerations of how all these components should be maintained within Grid middleware distributions in the future. The absence of a realistically implementable reference model is diametric to the fundamental design principles of software engineering and has thus lead to numerous different architectures of production Grids and their deployed middleware systems.

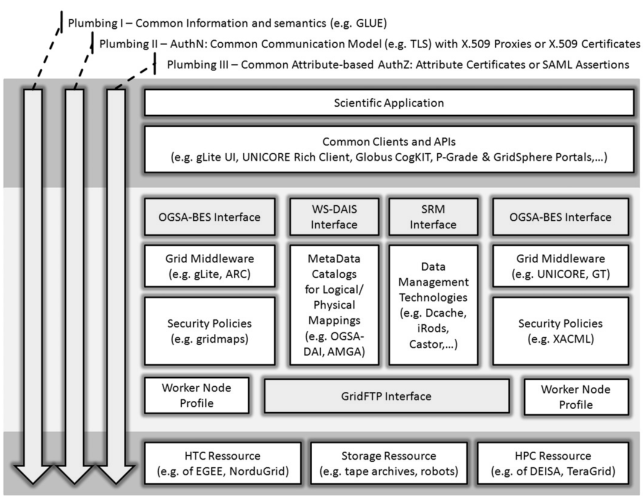

Figure 1. Core building blocks of the Infrastructure Interoperability Reference Model.

Although one goal of OGSA is to facilitate the interoperability of different Grid technologies and infrastructures in e-science and e-business, we state that the requirements for interoperability have to be specified much more precisely than within OGSA. Therefore, we focused on earlier work [14] on the definition of a set of requirements based on lessons learned obtained from world-wide interoperability efforts between production Grid infrastructures gained within the OGF Grid Interoperation Now (GIN) community group [11]. We further organized the International Grid Interoperability and Interoperation Workshops (IGIIW) 2007/2008 [15] at the e-Science conference series to evaluate numerous world-wide interoperation approaches. Based on these requirements and lessons learned, we presented earlier [16] and defined more recently the *infrastructure interoperability reference model (*IIRM*)* that is closer oriented towards the interoperability of production e-science infrastructures than OGSA and is already used in numerous different interoperability use cases that in turn proves its feasibility [8].

It is important to notice that this reference model does not aim to replace OGSA but rather trim it down in functionality by dropping several parts of it and refining other parts that are mostly relevant to interoperability of production Grids today. Thus, most of the core building blocks of the IIRM shown in Figure 1 are already deployed on production Grids and thus we consider our approach as a bottom-up approach compared to OGSA standardization efforts that can be seen as top-down approach according to its roadmap. But the fundamental idea of the IIRM is not only to review already deployed open standard implementations on production infrastructures, but also to identify their weaknesses, provide missing links between them as well as certain refinements and tunings as illustrated in Figure 2.

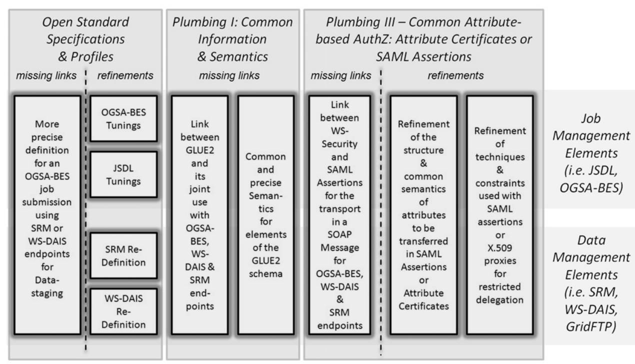

Figure 2. Overview of the key elements of the IIRM that are based on refinements and missing links of open standard specifications gained by production Grid interoperation.

As shown in Figure 1, our approach is thus based on (refined and tuned) open standards such as the Storage Resource Manager (SRM) [17] and WS-DAIS specification[18] that both using GridFTP [19] underneath for transport. In terms of job management we base our approach on the OGSA-Basic Execution Services (OGSA-BES) [20] specification to enable interoperability.

Nevertheless, we argue that the standardization of these specifications basically in isolation from each other is not suitable for an adoption within production e-science infrastructures. Also, many specifications clearly state that security considerations and information exposure aspects are out of scope. In contrast, we thus further argue that security and information exposure aspects have to be closely taken into account in order to achieve a full interoperable solution within production Grid infrastructures that lead to the definition of different *plumbings*. As shown in Figure 1 and 2, our reference model actually defines three different plumbings from the security and information domain to indicate that these rather *'vertical standards'* have a high impact on the rather *'horizontally used standards'* to achieve full interoperability between production Grid infrastructures. In other words, the compliance checking of numerous specifications is just a precondition for production interoperability since the plumbings have to be applied in parallel in order to achieve real production-level interoperability.

In this context, we are well aware of the approach of different profiles that are part of the OGSA roadmap. But investigations in the broader GIN community revealed [11] that profiles such as the HPC Basic Profile [21] does not satisfy the needs of todays production e-science infrastructures, for instance, in the context of attribute-based authorizations required for the member management of virtual organizations (VOs). Nevertheless, profiles are a good step towards the right direction to address the joint use of open standards instead of standardizing them in isolation from each other. But the rather module-driven standards lead to many different profiles providing more flexibility for Grid deployments, but significantly decreasing the chance of interoperable Grids. In this context the OGSA roadmap does not clearly address how all standards should work together since the relationship details between them are missing.

In contrast to the OGSA in general, and its roadmap in particular, our aim with the definition of the IIRM is to specifically take production experience from world-wide interoperation efforts into account. An overview of the key elements of the IIRM is shown in Figure 2 that is based on refinements and missing links of open standard specifications gained by production Grid interoperation experience. Left out for simplicity in Figure 2 is plumbing II (cp. Figure 1) since the protocols and interfaces of it does not have to be refined (either full X.509 certificates or X.509 proxy certificates). Note that we provide no specification versions since the most standards are basically either tuned or refined.

We achieved the design of the core building blocks of the IIRM and its key elements by investigating open standards-based technologies as a whole ecosystem in production setups. Thus we have been able to gather the significant standards for production Grids, and identified the missing links and reviewed specifications in order to perform certain refinements underpinned by lessons learned out of production Grid interoperation efforts in GIN. One example are tunings of the OGSA-BES specification in terms of data push use cases where end-users would like to create an activity without starting it and then perform data staging manually. Once the data staging has been performed, the activity should be started. This functionality, among others required by production Grid use cases, are not reflected in the current specification and thus considered as 'tunings'. We also fed them back to the standardization processes.

In fact, since we have shown implementations of the IIRM in numerous successful Grid interoperation activities, we have given the IIRM as an input to the standardization process of OGF via the Production Grid Infrastructure (PGI) working group [22], which members are representatives of numerous different world-wide production Grids.

Finally, a complete definition (i.e. schemas, specifications, etc.) of the IIRM in general and its key elements in particular is clearly out of scope of this paper due to page restriction and our focus on one IIRM use case in e-Health. Please refer to [8] for more information about IIRM elements.

IV. E-HEALTH SIMULATIONS ON INTEROPERABLE GRID INFRASTRUCTURES BASED ON THE REFERENCE MODEL

According to Fettke et al. [23], a reference model can be evaluated with numerous empirical perspectives in general and case studies in particular. In order to proof the feasibility of our proposed IIRM design, we have applied its components to two different real world case studies in the field of e-Health that both require interoperability of e-science infrastructures. The first is described by Riedel et al. in [4] enabling a cheaper and faster drug discovery by using EGEE and DEISA. The second case study, which is the focus of this paper, is related to a pre-production setup towards the VPH that use the hybrid infrastructure NGS and the HPC-driven infrastructures DEISA and TeraGrid to get access to a broader variety of HPC systems at the same time.

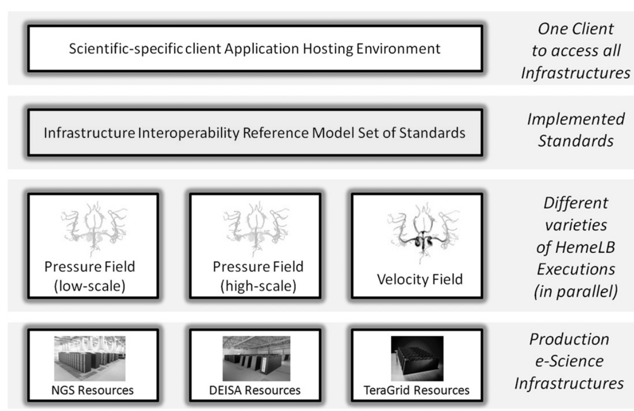

Figure 3. Core building blocks of the Infrastructure Interoperability Reference Model.

In general, the VPH roadmap [2] raised the demand for having access to an even more amount of computational resources in interoperable infrastructures, but we review one particular use case here in order to understand the computational-intensiveness. The scientific applications being considered here offer the possibility of performing patient specific virtual experiments to study the effects of course of treatment in silico, without danger to the patient in question. In the field of patient specific medical simulation there are a wide variety of different studies that can run with different levels of scale in parallel on a wide variety of computing systems provided by numerous different e-science infrastructures. This leads to the demand of interoperable e-science infrastructures that use, for instance as shown in Figure 3, the HemeLB [24] code on numerous different types of computer hardware architectures provided by different kinds of infrastructures.

The fundamental goal of our particular pre-production setup in general and the HemeLB application in particular is to contribute to the research field of cardiovascular diseases, which are the cause of a large number of deaths in the developed world. The problems of patients are often due to anomalous blood flow behavior in the neighborhood of bifurcations and aneurysm within the brain. In this context, cerebral blood flow behavior plays a crucial role in the understanding, diagnosis, and treatment of this disease. However, the simulations with HemeLB raise a demand for a large amount of computing resources offering different scales, because simulating a whole brain flow is computational-intensive and patient-specific. Thus, the central goal of the IIRM implementation use case illustrated in Figure 4 is to satisfy this demand by providing a number of technologies that enable simulations of patient-specific blood flow behavior using the large amount of computer power available on interoperable production e-science infrastructures.

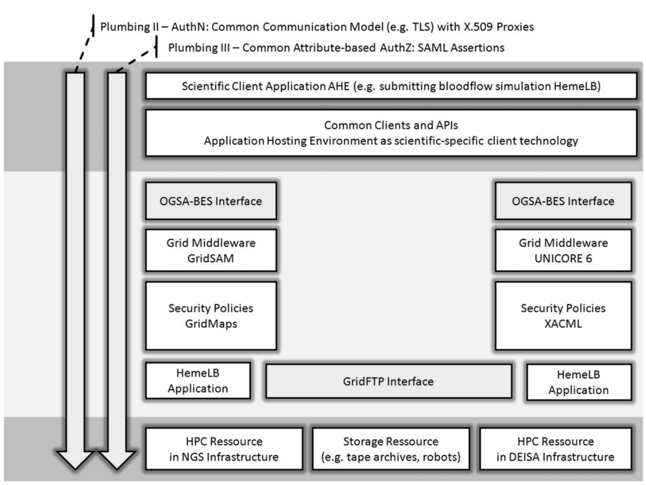

Figure 4. IIRM implementation providing seamless access to NGS and DEISA resources

We have been collaborating with members of the GENIUS project, which is mainly concerned with performing neurovascular blood flow simulations in support of clinical neurosurgery. This e-science project uses the lattice-Boltzmann code HemeLB [24] that is designed to simulate fluid flow in the sparse topologies of the patient brains. In this context, the particular simulation models are typically derived from patient-specific brain x-ray angiography scans that are in turn used as input to the specific HemeLB simulation runs.

In more detail, a separate segmentation tool typically runs on resources local to the clinician in order to segment the imaging data to generate a model for the HemeLB simulation. In other words, the HemeLB simulation doesn't use MRI scan data directly, but indirectly via the segmentation tool and its output, which is available on the Grid for its use on computational Grid resources that run HemeLB. As a consequence, our IIRM-based infrastructure setup actually requires the possibility for large file transfers that are able to transport the x-ray-based data to the computing resources across existing e-science infrastructures boundaries. The latter is specifically important to circumvent a duplicate storage of large patient image datasets in each different e-science infrastructure that is used with HemeLB. The IIRM satisfies this demand by providing GridFTP, which has been enabled under UNICORE 6 for our particular setup as shown in Figure 4. In order to use GridFTP, UNICORE 6 was also enhanced to support X.509 proxies during OGSA-BES-based job submissions and it's data-staging.

However, Figure 3 and 4 indicates another key requirement of our e-Health application that is related to the seamless access of different infrastructures. The e-scientists typically use their own specific client that in our case is named as the *Application Hosting Environment (*AHE*)* [25]. It is an open source research tool based on open standards, which furnishes a facile mechanism for accessing federeated Grid resources in standards and non-standard ways if necessary. The fundamental goal of this tool is to allow clinicians to seamlessly interact with a large amount of computational power available in different e-science infrastructures even from within their operation theatre.

In more detail, the AHE is a middleware that realizes the concept of Grid application virtualization. The key aspect of this concept is to simplify the use of Grid resources and infrastructures by focusing on the Grid applications instead of plain known Grid middleware usage. It introduces a layer of Web services between the user and the various flavors of Grid middleware systems thus hiding much of the complexity of the Grid infrastructures. The AHE achieves Grid application virtualization by providing an abstract interface to a given scientific application deployed on the Grid. Features of AHE include an uniform environment that users can use to reserve time on Grid resources in advance (if supported by the corresponding underlying middleware), launch cross-site parallel applications and also interact with them through real time steering and visualizations.

As shown in Figure 4, the particular IIRM design implementation also supports the use of the scientific-specific client AHE, because we added an OGSA-BES client to its features. This client in turn is using this standard to submit the computational Grid jobs that execute HemeLB to the OGSA-BES interface of UNICORE 6 deployed on DEISA and to the OGSA-BES implementation of OMII-UK named as GridSAM deployed on the NGS. In terms of security, we used the plumbings II of the IIRM design with X.509 proxies for authentication and to enable also GridFTP transfers required to move the patient-specific input data for HemeLB to the respective resources within TeraGrid an DEISA. In terms of authorization, we used the IIRM plumbing III with signed SAML assertions, which convey attributes of end-users such as different roles and are later used in the middleware technologies to achieve attribute-based authorization. Hence, the IIRM design implementation fundamentally supported the interoperability between the used e-science infrastructures.

Finally, the specific HemeLB use case also implies the requirement to have interactive access to resources of these infrastructures as well in order to perform computational steering in real-time. In this context, computational steering refers to the change of application parameters on the fly during the application execution on one of the resources within an e-science infrastructure. Hence, the goal of these real time visualization and computational steering is to allow clinicians to interact with the simulations as they run in order to review the possible effects of various surgical investigations performed on different e-science infrastructures. Figure 5 illustrates the encountered challenges in using the BlueGene/P in Juelich that is part of DEISA for computational steering, which needs a bi-directional connection realized on sockets. We learned that not many people have worked with opening sockets on computing cores and thus, as a side-effect, our work contributed also to the understanding of the BlueGene/P technology in terms of mapping sockets of compute cores to i/o nodes. Finally, we used SSH forwarding mechanisms in order to have the image shown at the steering client-side. Hence, the data is directly rendered on the compute cores and then forwarded via the i/o node to the login nodes that runs the steering clients. The steering client then uses an X-window connection in order to be shown at the corresponding remote site where the scientific-specific Grid client AHE is running.

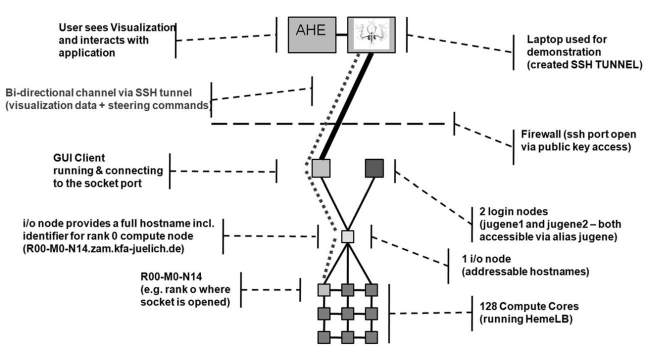

Figure 5. Setup on the BlueGene/P in Jülich to enable computational steering in DEISA.

V. Related Work

As already mentioned in the introduction of this paper, the rather high-level OGSA initially defined by Foster et al. [12] defines an architecture model taking many requirements from e-science and e-business into account and develops rather top-down standards to satisfy them. A more detailed comparison between OGSA and the IIRM identifies basically four major differences. First, the IIRM is less complex in order to address the rather long standardization processes in standards bodies. To provide an example, the OGSA-BES standard took 36 iterations to be completed as a full standard over the duration of approximately two years. We only take a subset of currently available standards. Second, OGSA is not specific enough with respect to e-science infrastructures since it takes a huge requirement set into account. Most notably, it does not address the missing links between several specifications that have been provided by the IIRM for a specific subset of open standards.

We argue that the standardization process according to the OGSA roadmap is difficult since the job standards do not address data issues while the data standards do not address job issues, and additionally, both job and data standards declare that security is out of scope. We argue that this produces rather isolated standards, instead of well-linked standards that have been defined by the IIRM design by also specifying the missing links between important standards (i.e. SRM-based data staging with OGSA-BES secured via the specified security profiles).

As a third difference, we argue that the many components that are basically part of the OGSA roadmap raise economic issues. At the time of writing, many e-science infrastructures struggle to have a sustainable strategy in terms of maintenance cost and thus commercial options are considered as one way. In other words, we argue that a full grown OGSA conform technology with all dedicated services might be actually too expensive to be maintained by middleware providers over a long period. IIRM, in contrast, only focuses on a small subset and can be thus considered to be an economic version of OGSA. Finally, the fourth difference relies in the bottom-up approach of the IIRM compared to OGSA. This means we know already from the production experience that the standards defined by the IIRM are working while a full grown OGSA conform technology has still to perform a lot of test runs before many of these standards will be deployed and used in production Grids.

Apart from the reference model there is also related work in the field of Grid application virtualization known as the Simple Api for Grid Applications (SAGA) [26] standard. In a similar manner like the AHE, the SAGA open standard can be used to focus on the Grid application itself while using the wide variety of Grid middleware flavors transparently underneath. In principle, the SAGA API can be even embedded into the AHE to realize a standard-based abstract application interface.

VI. Conclusions

In this paper, we raised the demand for an infrastructure interoperability reference model to promote interoperability between today's production e-science infrastructures. We have shown that the elements of this reference model are based on experiences and lessons learned from many worldwide interoperability projects. We can conclude that the IIRM represents a trimmed down version of OGSA in terms of functionality and complexity, but still providing the most significant features used in e-science infrastructures over the last couple of years. We have shown the core building blocks of the IIRM and many of them are already deployed on the infrastructures and only minor additions (i.e. missing links) have to be done in order to achieve interoperable e-science infrastructures.

The particular focus on this paper was an e-Health application that used one IIRM implementation to perform bloodflow simulation across different kinds of e-science infrastructures. We have also shown pre-production setups of this IIRM implementations as life demonstrations at the Supercomputing 2008, but clearly we are in the process of going towards production via the DEISA virtual community "VPH" and beyond. We have also used the IIRM in numerous other scientific use cases in different research fields such as life sciences (i.e. drug discovery) and fusion science. However, we can conclude that the IIRM implementation satisfy the demand of the scientists by providing access to multiple interoperable infrastructures and thus provides assurances that an IIRM-based client such as the AHE can access their required multiple infrastructures by still using single sign-on and the same security credentials.

A closer look also reveals that we have identified several improvements of GLUE2 and JSDL respectively within our reference model. In fact, we have worked on a better support of providing shape characteristics and network topologies (mesh, etc.) during job submission that was basically raised as a demand while working with the hemeLB application and others. These improvements of open standards enable a better efficient use of resources such as the BlueGene/P and others and are thus one of the key benefits of our proposed reference model based on improved open standards. Other areas of improvements related to the work in this paper are precise core/task mappings that can also make a difference in several use case runs.

Since our evaluation use cases have been very successful, we have given the IIRM as an input to OGF by creating a GIN spin-off activity named as the PGI working group. By chairing this group, our goal is to standardize the IIRM elements and plumbings to assure that also numerous other Grids can benefit from our propose IIRM design. This will basically contribute to the vision of having an *interoperable united federation of world-wide e-science infrastructures* in the future offering standardized PGI-compliant access.

Finally, history of computer science shows that often complex architectures were less used than their trimmed down versions. For instance, the complex SGML is less used than its smaller version XML. Also, the original ISO/OSI reference model consisted of seven layers, while its much more successful trimmed down version TCP reference model become the de-facto standard in networking. We argue that the same principles can be applied with OGSA by defining a more limited, but more usable reference model such as our proposed IIRM.

Acknowledgment

We thank the OGF GIN and PGI groups. This work is partly funded via the DEISA-II project funded by the EC in FP7 under grant agreement RI-222919.

References

[1] Website, "Virtual Physiological Human (VPH)," http://www.europhysiome.org.

[2] S. Consortium, "Seeding the europhysiome: A roadmap to the virtual physiological human (vph)," http://www.europhysiome.org/roadmap.

[3] M. Riedel, D. Kranzlmueller, *et al.*, "Classification of Different Approaches for e-Science Applications in Next Generation Computing Infrastructures," in *Proceedings of the Int. e-Science Conference, Indianapolis, USA*, 2008.

[4] M. Riedel *et al.*, "Improving e-Science with Interoperability of the e-Infrastructures EGEE and DEISA," in *Proceedings of the MIPRO*, 2007.

[5] Website, "EGEE," http://www.eu-egee.org/.

[6] P. Website, "DEISA," http://www.deisa.org/.

[7] I. Rodero *et al.*, "Modeling and Evaluating Interoperable Grid Systems," in *Proc. of the TeraGrid 2008 Conference, Virginia, USA*, 2008.

[8] M. Riedel *et al.*, "Research Advances by using Interoperable e-Science Infrastructures - The Infrastructure Interoperability Reference Model applied in e-Science," in *Journal of Cluster Computing, SI Recent Research Advances in e-Science*, 2009.

[9] J. Taylor, "e-Science Definition," http://www.escience. clrc.ac.uk.

[10] Website, "National Grid Service," http://www.gridsupport. ac.uk.

[11] M. Riedel, E. Laure, *et al.*, "Interoperation of World-Wide Production e-Science Infrastructures," in *Journal on Concurrency and Comp.: Practice and Experience*, 2008.

[12] I. Foster *et al.*, *Open Grid Services Architecture (OGSA), Version 1.5*. Open Grid Forum Grid Final Document Nr. 80, 2006.

[13] C. Jordan *et al.*, *Defining the Grid: A Roadmap for OGSA Standards v.1.1*. Open Grid Forum Grid Final Document Nr. 123, 2008.

[14] M. Riedel *et al.*, "Experiences and Requirements for Interoperability between HTC- and HPC-driven e-science Infrastructures," in *Proceedings of the Workshop HPC for e-Science, Korea e-Science All Hands Meeting*, 2008.

[15] Website, "International Grid Interoperability and Interoperation Workshop (igiiw)," http://www.fz-juelich.de/jsc/igiiw.

[16] M. Riedel, "Grid Interoperability Now – Towards an Interoperability Reference Model," 2008, http://indico.hep.lu.se/conferenceDisplay.py?confId=616.

[17] A. Sim *et al.*, *The Storage Resource Manager Interface Specification Version 2.2*. Open Grid Forum Grid Final Document Nr. 129, 2008.

[18] M. Antonioletti *et al.*, *Web Services Data Access and Integration - The Core Specification, Version 1.0*. Open Grid Forum Grid Final Document Nr. 74, 2006.

[19] I. Mandrichenko *et al.*, *GridFTP v2 Protocol Description*. Open Grid Forum Grid Final Document Nr. 47, 2005.

[20] I. Foster *et al.*, *OGSA Basic Execution Service Version 1.0*. Open Grid Forum Grid Final Document Nr. 108, 2007.

[21] B. Dillaway *et al.*, *HPC Basic Profile, Version 1.0*. Open Grid Forum Grid Final Document Nr. 114, 2007.

[22] Website, "OGF - Production Grid Infrastructure Working Group (PGI-WG)," http://forge.ogf.org/sf/projects/pgi-wg.

[23] P. Fettke, "Multiperspective Evaluation of Reference Models - Towards a Framework," in *Conceptual Modeling for Novel Application Domains*. Springer.

[24] S. Manos *et al.*, "Patient specific whole cerebral blood flow simulation: A future role in surgical treatment for neurovascular pathologies," in *Proc. of the TeraGrid 2008 Conference, Virginia, USA*, 2008.

[25] S. J. Zasada and P. V. Coveney, "Virtualizing Access to Scientific Applications with the Application Hosting Environment", Computer Physics Communications, 180, (12), 2513-2525, (2009), DOI: 10.1016/j.cpc.2009.06.008.

[26] T. Goodale, S. Jha, H. Kaiser, T. Kielmann, P. Kleijer, A. Merzky, J. Shalf, C. Smith, „A Simple API for Grid Applications (SAGA)“, http://www.ogf.org/documents/GFD.90.pdf