

EGI-InSPIRE

APPDB ACTIVITIES SUMMARY AND WORKPLAN (Y2A)

Document identifier:	AppDBWorkplanY2A.doc
Date:	07/05/2011
Activity:	NA3
Lead Partner:	IASA/GRNET
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/26

Abstract

This document presents the objectives and responsibilities of the Database of applications, tools, and their developers (AppDB) subtask, developed in the framework of EGI TNA3.4 Technical Services. It starts with a summary of the activities developed in the second six months (1-Nov-2010 to 30-Apr-2011) of the EGI project, followed by a work plan to be accomplished in the following six months (1-May-2011 to 31-Oct-2011).

I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	W. V. Karageorgos	IASA,GRNET/NA3	7/05/2011
Reviewed by	Moderator: Reviewers:		
Approved by	AMB & PMB		

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	28/04/2011	First draft	W. V. Karageorgos / IASA,GRNET M. Chatziangellou / IASA
2	06/05/2011	Second draft	W. V. Karageorgos / IASA,GRNET N. Zarife / IASA N. Theofilopoulos / IASA
3	07/05/2011	First complete version	W. V. Karageorgos / IASA,GRNET A. Nakos / IASA,GRNET N. Zarife / IASA N. Theofilopoulos / IASA
4	07/05/2011	Review of first version	M. Chatziangellou / IASA A. Nakos / IASA,GRNET

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed:

<https://wiki.egi.eu/wiki/Procedures>



VI. TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>.



VII. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed – both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting ‘grids’ of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today’s production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.



The production infrastructure supports Virtual Research Communities (VRCs) – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

VIII. EXECUTIVE SUMMARY

This document presents work progress and future plans, in terms of objectives and responsibilities, for the EGI Applications Database (AppDB) sub-task, developed in the framework of EGI TNA3.4 Technical Services. AppDB is the descendant of the EGEE Applications Registry [R 2] portal, which was initially developed by the IASA regional coordination team during the course of the EGEE-III project. It provides a catalogue of applications that have been ported, or are being ported, within the infrastructure [R 1]. As such it enables new communities to discover and reuse EGI applications, thus avoiding duplication of effort. By the reuse of ported applications one of the main barriers of grid adoption is eliminated.

During the second six months of the project, about 86% of the work items described in [R 3] have been completed. The remainder is still work-in-progress due to issues involving third parties. Moreover, additional effort has been invested in work items which arose during this period from new requirements capturing, mainly through the dedicated GGUS support unit.

Next term work plans fall under the key areas of Quality of Information, Information Retrieval, Dissemination, Compatibility, and Architecture.



TABLE OF CONTENTS

1 INTRODUCTION	7
2 ACTIVITY SUMMARY	8
2.1 Service integration.....	8
2.2 Features and enhancements	9
3 WORKPLAN	10
3.1 Architecture	11
3.1.1 S1: Migrate DBMS to PostgreSQL.....	11
3.1.2 S2: Audit and re-factor existing codebase.....	11
3.2 Information Retrieval	13
3.2.1 S3: Re-implement searching/filtering mechanism.....	13
3.2.2 S4: Applications Entry Tagging Mechanism.....	14
3.3 Quality of Information	14
3.3.1 S5: User Comment/Ranking System.....	14
3.3.2 S6: Entry Problem Reporting System.....	14
3.3.3 S7: Broken Link Detection Notification System.....	14
3.3.4 S8: Application Revocation Mechanism.....	14
3.3.5 S9: Detect and promote or remove non-finished apps based on status.....	15
3.3.6 S10: Application/tool name uniqueness validation mechanism.....	15
3.4 Dissemination	15
3.4.1 S11: Notification Services (email/RSS).....	15
3.5 Cross-browser Compatibility	15
3.5.1 S12: Investigation and estimation of Internet Explorer 9 (IE9) compatibility.....	15
3.6 EGI User Support Platform	16
3.6.1 S13: Provide write-access through API.....	16
3.6.2 S14: Develop new/extend existing gadget with write access to the AppDB.....	16
4 CONCLUSION	17
5 REFERENCES	18



1 INTRODUCTION

The AppDB sub-task is developed in the context of the EGI TNA3.4 Technical Services activity. It stores tailor-made computing tools for scientists and NGI User Support Teams to use. The applications and tools filed in AppDB are finished products, ready to be used on the European Grid Infrastructure. By storing pre-made applications and tools, the AppDB aims to

- ⤴ alleviate the need for scientists and NGI User Support Teams to spend time developing their own software
- ⤴ avoid duplication of effort across the EGI community

In the following sections, we will summarize the progress that has been achieved during the past six-month period, according to the existing workplan [R 3], and propose a new workplan for the next six-month period (1-May-2011 to 31-Oct-2011). Main development effort will be focused on the following key areas:

- ⤴ Quality of information
- ⤴ Information retrieval
- ⤴ Notification / Dissemination
- ⤴ Cross-browser compatibility
- ⤴ EGI User Support Platform
- ⤴ Architecture



2 ACTIVITY SUMMARY

The activities of the past six months were mainly focused on service integration. As such, effort was allocated towards interfacing the AppDB with the EGI Operations Portal and the GOCDB on one hand, and on the other, on providing a web-API for other services to interface with the AppDB. A web-gadget was also provided as a proof-of-concept of what can be done through the said web-API. The rest of the activities mainly concerned new features or enhancements to the GUI, and addressing bugs and requests from the GGUS unit and the EGI RT system respectively.

2.1 *Service integration*

Integration with the EGI Operations Portal [R 4] has been successfully completed, thus providing VO data through the AppDB, with respect to application and tool entries. Data is retrieved via an XML export facility of the Operations Portal, and is cached on the AppDB server. If more than one hour stale at the time of a relevant user request, then the data is refreshed. Relevant server-side Active Record objects have been added to the object model, in order to manipulate the VO data and index references in the AppDB Relational Database Management System (RDBMS). It should be mentioned that the Operations Portal development team is working on a SOAP web-service based on the Lavoisier framework [R 6], in order to deliver VO and NGI data. At the time of integration, this service was not readily available, but switching to it, instead of using the XML export functionality is to be evaluated at a later point, after the service has been successfully deployed.

Similar effort has been allocated to NGI data integration through the GOCDB PI [R 7], meaning that the relevant RBMS structures and server-side Active Record objects have been prepared. Nevertheless, integration has been stalled, due to the fact that the GOCDB stores obsolete EGEE data (e.g. ROC entries) along with new EGI data (NGI entries), without providing any kind of discrimination. Moreover, country relations to NGIs are not provided either, thus making the mapping to the AppDB data model impossible without ad-hoc data injections. The GOCDB development team has been contacted regarding these issues, but has expressed no intent of implementing the needed modifications in the near future, due to excess workload. The alternative of using the Operations Portal Lavoisier web-service instead will be evaluated once the service is available, in cases it better matches the AppDB data model requirements.

Furthermore, large effort has been allocated towards providing and properly documenting a prototype RESTful web-API, so that other services or portals in general may interface with the AppDB. The said API provides unauthenticated read-only data, in detail about applications and tools, related publications, and researcher profiles, as well as in reference about VOs, middlewares, countries, and (sub)disciplines. Documentation about the API can be found in the EGI NA3 wiki pages [R 8], and XML schemata are provided from the dedicated FQDN which provides the API itself [R 9]. Currently, the API is in beta status, and plans about extending it with authenticated read/write access are in the works.

Finally, a web-gadget (portlet) has been developed, as a proof-of-concept, of what can be achieved



through the web-API. This web-gadget is already in use by various regionalized sites [R 10], and an on-line editor is available through the AppDB site [R 11]. Through the said web-gadget, third parties can provide on-line embedded detailed information to filtered entries about applications, pertaining for example to a certain VO, or country. The possibility of providing more gadgets, or extending the existing one with more functionality, such as write-access -- provided the API is also thus extended -- is scheduled to be investigated.

2.2 Features and enhancements

The rest of the terms effort, development-wise, has been mainly attributed to work related to new features and enhancements of the portal. More than 15 GGUS and RT tickets have been addressed, along with the issues that were communicated over e-mail. Examples of such actions include a preliminary reporting tool for managers, major speed improvements of the portal, support for non-UMD middlewares, enhancements to the data export functionality, addition of more details in application and people entries, and overall usability of the portal's user interface. Of such actions that were defined in the previous workplan, the one pertaining to IE8 compatibility is still work-in-progress, since the constant addition and modification of features greatly impaired the ability to thoroughly test. With the public release of IE9, new tests have been planned, and compatibility is expected to be better.

3 WORKPLAN

As declared by the UCST, a new release of each of the TNA3.4 services is expected every 6 months. New releases must be prepared according to development plans endorsed by the UCB, based on the requirements collected from user communities. The requirements collection for the AppDB has been completed and documented in [R 5]. A summary of the workplan for the next 6 months (26 weeks) is presented in Table 1.

One issue that has to be mentioned is that in Table 1, items in red (also marked with a dagger) are sprints that our team will try to accomplish in the current six months work-cycle, by involving more than one engineer at a time, but this can not be guaranteed. This discrimination is made due to the fact that the overall effort required to follow the extended workplan, exceeds the upcoming six months timeframe by 8 weeks, given one person is allocated at any given time.

Sprint	Task	Duration	
		Weeks	Persons
Architecture			
S1	Migrate DBMS to PostgreSQL	4	1
S2	Audit and re-factor existing codebase	4	1
Information Retrieval			
S3	Re-implement searching/filtering mechanism	3	1
S4	Applications entry tagging mechanism	2	1
Quality of Information			
S5	User comment/ranking system	3	1
S6	Entry problem reporting system	1.5	1
S7	Broken link detection notification system	1.5	1
S8	Application revocation mechanism (only for Managers)	2	1
S9	Detect and promote or remove non-finished apps based on status	1	1
S10	Application/tool name uniqueness validation mechanism	1	1
Dissemination			
S11	Notification services (email/RSS)	2	1
Cross-browser Compatibility			
S12	Investigation and estimation of Internet Explorer 9 (IE9) compatibility	1	1
EGI User Support Platform			
S13	Provide write-access through API	4 [†]	1
S14	Develop new/extend existing gadget with write access to the AppDB	4 [†]	1
Total		26 (34[†])	

Table 1 AppDB six-month workplan

Ongoing Effort

Apart from the items mentioned in the above workplan, there should also be mention about the effort that will be allocated to tasks such as bug-fixing, support for new requirements from the users, and maintenance of the project in general. This effort is estimated to be c. 25% of the term, i.e. about 6.5 person-weeks, and will take place in parallel to the rest of the work-items, by allocating more than one persons in average.

3.1 Architecture

3.1.1 S1: Migrate DBMS to PostgreSQL

The existing RDBMS of choice, namely MySQL with an InnoDB engine, suffers from well-know limits with regards to scalability, performance, and query language expressiveness. During the first year of the project, where design and development was done through RAD processes, this was a good choice due to the simplicity of set-up, portability, and small learning curve, since most professionals are already familiar with it. Nevertheless, as the projects matures, and requirements rise, this choice fails to meet the augmenting demands for power.

On the other hand, a more advanced, yet open RDBMS such as PostgreSQL, now seems to be more suitable for the project's needs. Its superior query language expressiveness, paired with better performance and advanced features can deliver greater data transformation abilities inside the data tier. This automatically translates to better optimized code within the existing codebase architecture, and thus a better overall performance in general.

In order to migrate to the new RDBMS, the server-side Active Record object model must be adapted to account for SQL discrepancies between the two RDBMS's, despite the framework's abstraction layer and use of the same data model. Moreover, in-house code such as stored procedures, views, triggers, etc. will have to be re-written to a great extent, since said discrepancies are great, with the de facto lack of an abstraction layer. Finally, members of the project that are not familiar with PostgreSQL, will have to face a steep learning curve. All of the above lead to a much optimistic 4 week sprint for the migration process.

3.1.2 S2: Audit and re-factor existing codebase

Currently the project is based on a tree-tier client-server architecture, where web content is delivered to the presentation tier through a server-side Model-View-Controller (MVC) design. In this scheme, controllers in the business tier handle user requests by asking for data from objects of the Active Record model in the data tier , which they then pass on to the views in the presentation tier for rendering. Rendering is mainly done through server-side code (PHP), with minor post-presentation client-side handling through Javascript, mainly due to the extensive use of AJAX (Image 1).

This combination of architecture and design patterns has many well-know benefits, such as good modularity. For instance, migrating to a much different RDBMS, as described in the previous paragraph, would be rather unfeasible without it. It also serves very well in terms of abstraction in the presentation tier, providing for ease in rendering changes, yet it can only serve one kind of client

type transparently -- in this case, the portal's GUI.

This has worked well so far, however, with the introduction of the web-gadget, and much talk about more gadgets, or even iOS clients, such as mobile phones or tablets, this design suffers from duplication of effort in the presentation tier. In plain words, each client type must provide most of the rendering code from scratch. However, introducing a second, client-side MVC design inside the presentation tier would provide the framework to integrate all client types into a unified rendering codebase. This is readily feasible by exploiting the newly deployed RESTful API. With the proposed scheme, a user request is again handled by a controller, who this time simply provides the view with a template setup with little or no calls to any model. The view then renders the content asynchronously in the client through AJAX, by calling the API instead, which in a second iteration performs the full server-side MVC cycle to deliver the relevant data (Image 2).

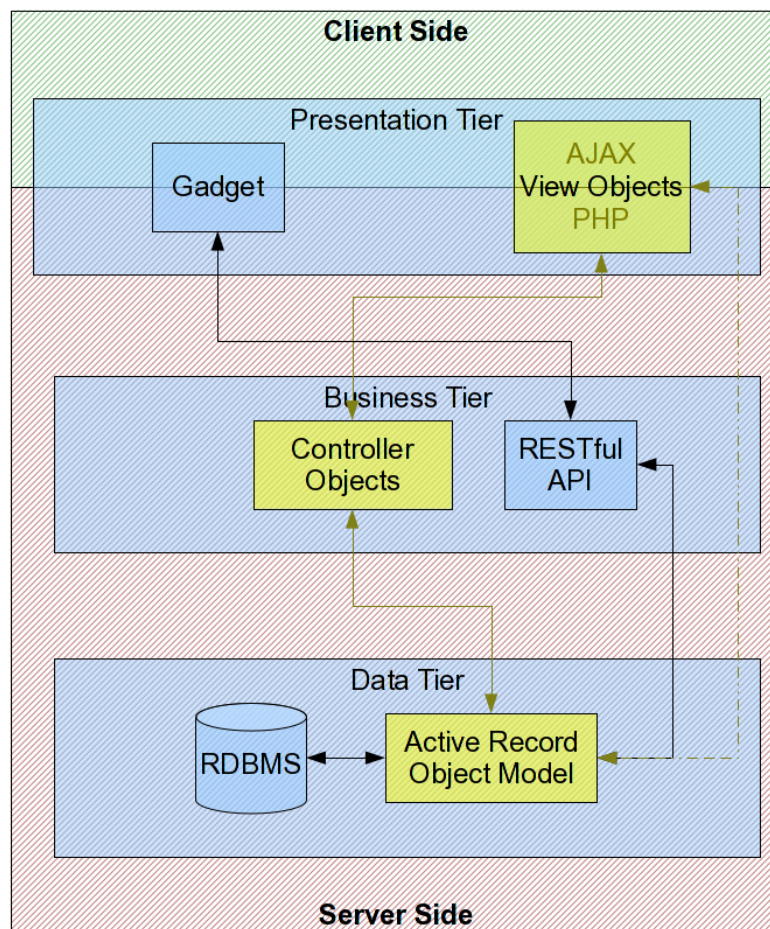


Image 1: Existing architecture. The server-side MVC design can be seen in yellow.

Such a re-factoring will require almost completely rewriting the presentation tier code, translating to major workload. Yet, the advantage is obvious, since time spent will be earned back in the future, by centrally coding, maintaining, and debugging presentation logic for all client types.

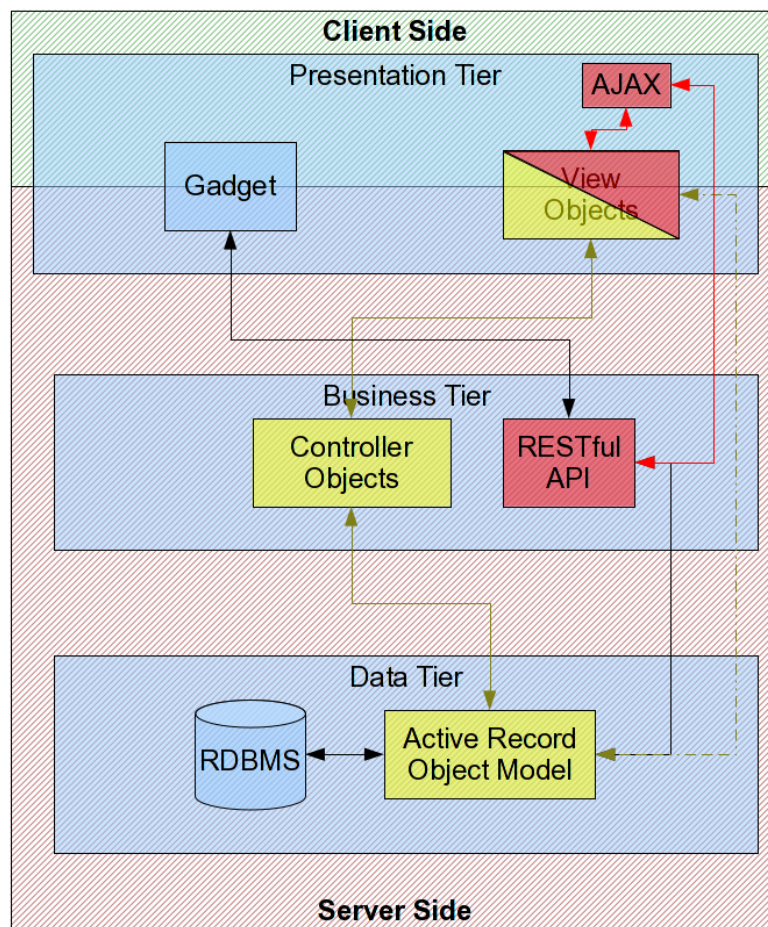


Image 2: Proposed architecture. The server-side MVC design can be seen in yellow, whereas the new client-side MVC can be seen in red.

3.2 Information Retrieval

3.2.1 S3: Re-implement searching/filtering mechanism

In its current state, data searching and filtering through the portal is limited to one entry type at a time (e.g. applications XOR people XOR VO's, etc.), and provides for no nesting nor field-wide search for a given entity (e.g. search for the string "test" in either name or description). The reasons for this are due to the fact that at the time the filtering mechanism was designed, there weren't many entity types, and also to the fact that data transformation in the data tier was limited, given the RDBMS of choice, if one makes the sensible decision of constricting oneself to viable queries with regards to computational time.

The migration to PostgreSQL is expected to alleviate computational time for such complex queries, and allow for the design of an extended filtering mechanism, which will meet the criteria described above.

3.2.2 S4: Applications Entry Tagging Mechanism

As information stored in the system rapidly grows larger in volume, finding what one seeks becomes a non-trivial task. Even though such a task may greatly benefit from a powerful filtering mechanism, giving the ability to the user to tag existing information with custom metadata can help fill in the gaps and provide for better searching, especially since such a feature can be carried out by people, for other people who search for the same type of information to benefit from. We plan on providing a set of predefined tags attributed automatically by the system to certain entry types (e.g. in order to distinguish between native entries, and entries imported by other similar systems), as well as free-string tags which can be attributed by the users on demand.

3.3 Quality of Information

3.3.1 S5: User Comment/Ranking System

Quality of information can be a sensible subject in large data stores. One of the actions planned in order to help ensure content quality is the provision of quality related metrics, such as comments and rankings. These metrics should be provided by authenticated end users through the course of time, and should be visible by everyone. Nevertheless, UCST members (managers) may need to have stronger rights upon them (e.g. adding special flags such as “hidden” or “to be improved”).

3.3.2 S6: Entry Problem Reporting System

There have been several cases during the past six months, where a certain application entry raised disputes concerning the “ownership” of the referenced application, or the use of data, such as logos or names, which belonged to other projects. Moreover, given the fact that it is fairly easy to obtain an EGI SSO, it is not beyond the realm of possibilities that we be faced with malevolent entries, of the “spam” type. Therefore, there have been plans to provide end users with the ability to “report” a certain application entry, providing a short text about the type of infringement, so that appropriate measures may be taken on by the UCST, or the submitter of the said entry, who would also be notified via e-mail.

3.3.3 S7: Broken Link Detection Notification System

One of the requests made by the UCST pertaining to QoI, is the detection of application entries with broken links in their URLs section. Such entries should be corrected by the submitter's team (who should be notified via e-mail), or be considered for deletion by the UCST, in case of no response withing a sensible time frame. To this purpose, a web spider will be developed and deployed, which will periodically scan URLs and perform the appropriate actions.

3.3.4 S8: Application Revocation Mechanism

As a catch-all case for application entries that present some kind of problem not identified by any of the above cases, a mechanism for temporal removal of said entries from class of presentable information should exist. This means that in such cases, said entries should not appear in index pages, nor search results, not should they be accounted for by statistics and export functionalities, etc.). Therefore, multiple modifications should be made in all code that relates to fetching application entries in the data tier.

3.3.5 S9: Detect and promote or remove non-finished apps based on status

During the EGEE era, application entries stored in the AppDB featured a status property, which represented their development status as far as gridification was concerned. However, since the EGI era, it has been deemed that only completed applications (i.e. fully gridified and in production) should exist in the database. Therefore, the status field is planned to be removed, and a list of application entries that are not marked as completed should be compiled and forwarded to the UCST, in order to contact the respected development teams, and either promote their status, or remove them from the database.

3.3.6 S10: Application/tool name uniqueness validation mechanism

A recurring issue which has been raised during the past term is related to the naming of application entries. There have been several cases where users did not check for the existence of an entry for a particular application, and thus re-added it. Moreover, there have been other cases where users created entries referring to new applications, which nevertheless used the same name with an existing yet different application, or with a fork thereof. While the later cases are “legitimate” for all matters, they do seem to create turmoil and disorientation. Therefore, in order to avoid future problems of this nature, it has been decided that name uniqueness should be imposed. Existing entries sharing the same name are few, and a comprehensive list should be compiled and forwarded to the UCST in order to separated them. Once uniqueness is imposed, entries that would otherwise use the same name, could instead be deployed with a diacritic suffix. Users registering new entries will be thus prompted in such cases.

3.4 Dissemination

3.4.1 S11: Notification Services (email/RSS)

Apart from information retrieval on-site, delivery of selected information off-side is equally important. Therefore, subscription services are to be developed, providing mailing lists and RSS feeds to the end users, based on criteria of their choice. Furthermore, separate privileged lists should exist in order to notify the UCST or NGI representatives about certain events such as new entry additions, role verification requests, etc.. Such privileged lists could also be used to inform the UCST about event concerning action falling under § (Quality of Information)

3.5 Cross-browser Compatibility

3.5.1 S12: Investigation and estimation of Internet Explorer 9 (IE9) compatibility

Microsoft Internet Explorer 8.0 compatibility is still a work-in-progress, mainly due to the fact that the constant addition of new features during the past term greatly impaired out ability to thoroughly test. The recent release of MSIE 9.0 means that effort should now be focused on this new version, which might prove to be more compatible with fewer modifications to the codebase. Moreover, compatibility of mobile browsers, such as the native Android browser, Mozilla Fennec 4.0, etc. should also be considered, time permitting, especially under the context of § pertaining to architecture.

3.6 EGI User Support Platform

3.6.1 S13: Provide write-access through API

One prospective feature that has been much discussed, both internally (NA3) and otherwise, is the provision of write-access through the web-API. This has been requested for cases where

- ⤴ another site wants to maintain its own database, yet integrate with the AppDB, by presenting our entries, and synchronizing its entries with the AppDB
- ⤴ another site wants to provide a customized view of the AppDB contents, without maintaining a separate database of its own, yet procure its users with the ability to create/modify entries on-site (i.e. without redirecting to the AppDB)
- ⤴ another site wants to permanently migrate/hand over its data over the AppDB.
- ⤴ variations and combinations of one or more of the above

All of the above cases obviously require some kind of write-access to the AppDB data tier, and even though alternatives to a write-enabled API have been investigated (e.g. migration tools, redirection workflows, peer pull-pull access, etc) -- each with their own pros and cons -- the verdict weights towards the former. This has been made much clear during the EGI User Forum 2011, held in Vilnius, Lithuania, where questions made by the users after the AppDB presentation mainly regarded this issue. Moreover, discussions made afterwards, with representatives from initiatives such as GISELA [R 12] and IGI [R 13] for instance, also favoured the write-enabled API.

Our team, therefore, will try to allocate effort into extending the existing RESTful API prototype with write-access, time permitting, according to Table 1. If that is not possible, then this work-item will be transferred to the next term (Y2B). Important Issues that are to be considered while redesigning the prototype include

- ⤴ security implementation: authorization may be based on existing SSO design, including non-personal (i.e. site-wide) accounts, or extended with certificated features. Delegation of credentials should also be addressed, when operations are non-personal.
- ⤴ synchronization: the case of interacting sites, which maintain their own database as well, will require a synchronization scheme/protocol, in order to avoid collisions and loss of data. This issue would most probably be address through a technical document proposing the scheme, for prospective partners to agree on.

3.6.2 S14: Develop new/extend existing gadget with write access to the AppDB

As with the case of read-only access, write-access could also be provided via gadgets -- as an out-of-the-box solution -- provided that the API is extended first. This could be done by either extending the existing gadget, or by providing a different dedicated gadget.

Providing this kind of service via gadgets, along with other possible services, will be investigated by our team after the completion of §. Depending on input from other NA3 partners, regarding their respective services, data exchange schemes between gadgets might be considered, towards forming a common framework for the EGI User Support Platform.



4 CONCLUSION

The AppDB subtask has used the second six-month term of the project to:

- ⤴ Provide a read-only RESTful web-API
- ⤴ Integrate with the Operations Portal to provide VO data
- ⤴ Provide a web gadget for other sites to integrate with the AppDB
- ⤴ Implement various improvements requested by users

The plan for the next six-month term, includes:

- ⤴ Major enhancements, initially focused on architecture
- ⤴ New requested features, primarily focused on quality and delivery of stored information

It should be noted that end of Y1 marks the end of a Rapid Application Development (RAD) phase, where the design of architectural aspects took place iteratively, in parallel to and in conjunction with code writing, and requirement gathering. This process has worked well up to now, forming a well established codebase and service "image". However, the project has grown rather large and complex for development to continue successfully in this manner, for fear of loss and duplication of effort, as well as future shoddy architectural decisions. A faithful implementation of the proposed Scrum/Agile methodology, paired with priority-based queuing of new requirements into the running work-plan can help ensure better results during the terms to come.

5 REFERENCES

R 1	http://appdb.egi.eu
R 2	https://na4rs.marie.hellasgrid.gr/
R 3	https://documents.egi.eu/document/276
R 4	http://operations-portal.egi.eu/
R 5	Requirements collection document
R 6	http://grid.in2p3.fr/lavoisier/
R 7	https://wiki.egi.eu/wiki/GOCDDB/PI/Technical_Documentation
R 8	https://wiki.egi.eu/wiki/TNA3.4_Technical_Services
R 9	http://appdb-pi.egi.eu
R 10	https://wiki.egi.eu/wiki/AppDB_Gadget_Editor#Success_Stories
R 11	http://appdb.egi.eu/gadgets/editor
R 12	http://www.gisela-grid.eu/
R 13	http://www.italiangrid.org/communities/applications/porting/igi_application_database