

EUROPEAN MIDDLEWARE INITIATIVE

DJRA1.4.2 – INFRASTRUCTURE AREA WORK PLAN AND STATUS REPORT

EU DELIVERABLE: D5.4.2

Document identifier:	EMI-DJRA1.4.2-1277583- Infrastructure_Area_Work_Plan_M12-v1.0.doc
Date:	30/04/2011
Activity:	JRA1
Lead Partner:	CERN
Document status:	Final
Document link:	http://cdsweb.cern.ch/record/1277583

Abstract:

This deliverable contains the detailed work plan of the Infrastructure Services technical area compliant with the overall EMI Technical Development Plan. The plan is released early in the project life and updated every year including a status report on the achievements of the past 12 months compared to the planned objectives. The status report at M03 will cover the state-of-the art while the work plan at M36 will provide recommendations for further work.

Copyright notice:

Copyright (c) Members of the EMI Collaboration. 2010-2011.

See <http://www.eu-emi.eu/about/Partners/> for details on the copyright holders.

EMI ("European Middleware Initiative") is a project partially funded by the European Commission. For more information on the project, its partners and contributors please see <http://www.eu-emi.eu>.

This document is released under the Open Access license. You are permitted to copy and distribute verbatim copies of this document containing this copyright notice, but modifying this document is not allowed. You are permitted to copy this document in whole or in part into other documents if you attach the following reference to the copied elements: "Copyright (C) 2010-2011. Members of the EMI Collaboration. <http://www.eu-emi.eu>".

The information contained in this document represents the views of EMI as of the date they are published. EMI does not guarantee that any information contained herein is error-free, or up to date.

EMI MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, BY PUBLISHING THIS DOCUMENT.

Delivery Slip

	Name	Partner / Activity	Date	Signature
From	Laurence Field	CERN/JRA1	18/05/2011	
Reviewed by	Jedrzej Rybicki. Riccardo Zappi	JUELICH/JRA1, INFN/SA1	25/05/2011	
Approved by	PEB		01/06/2011	

Document Log

Issue	Date	Comment	Author / Partner
1	05/04/2011	Table of Contents	Laurence Field/CERN
2	18/05/2011	v0.2 to reviewers	Laurence Field/CERN
3	24/05/2011	v0.3 to reviewers	Laurence Field/CERN
4	01/06/11	v1.0 PEB approved version	PEB

Document Change Record

Issue	Item	Reason for Change
1		
2		
3		

TABLE OF CONTENTS

1. INTRODUCTION	7
1.1. PURPOSE	7
1.2. DOCUMENT ORGANISATION.....	7
1.3. REFERENCES	7
1.4. DOCUMENT AMENDMENT PROCEDURE.....	8
1.5. TERMINOLOGY	8
2. EXECUTIVE SUMMARY	10
3. INFRASTRUCTURE AREA STATUS REPORT.....	11
3.1. AGREEMENT ON COMMON INFORMATION EXCHANGE METHODS	12
3.2. GLUE 2 ADOPTION.....	12
3.3. INFORMATION SYSTEM CLIENT HARMONIZATION	13
3.4. EMI REGISTRY	13
3.5. MESSAGING	13
3.5.1 <i>Information System Use</i>	13
3.5.2 <i>Accounting Use</i>	14
3.6. SERVICE MONITORING AND MANAGEMENT	14
3.7. VIRTUALIZATION AND CLOUDS	14
4. INFRASTRUCTURE AREA WORK PLAN.....	16
4.1. MESSAGING INVESTIGATION; ACCOUNTING	16
4.1.1 <i>Harmonization/convergence aspects</i>	16
4.1.2 <i>Affected/involved components</i>	16
4.1.3 <i>Subtasks</i>	16
4.1.4 <i>Risks</i>	17
4.1.5 <i>JRA1 KPIs</i>	17
4.2. MESSAGING INVESTIGATION; SERVICE MONITORING AND MANAGEMENT	17
4.2.1 <i>Harmonization/convergence aspects</i>	17
4.2.2 <i>Affected/involved components</i>	17
4.2.3 <i>Subtasks</i>	17
4.2.4 <i>Risks</i>	18
4.2.5 <i>JRA1 KPIs</i>	18
4.3. IMPLEMENT THE COMMON EMI REGISTRY	18
4.3.1 <i>Harmonization/convergence aspects</i>	18
4.3.2 <i>Affected/involved components</i>	18
4.3.3 <i>Subtasks</i>	18
4.3.4 <i>Risks</i>	18
4.3.5 <i>JRA1 KPIs</i>	19
4.4. GLUE 2 INFORMATION MODEL SUPPORT	19
4.4.1 <i>Harmonization/convergence aspects</i>	19
4.4.2 <i>Affected/involved components</i>	19
4.4.3 <i>Subtasks</i>	19
4.4.4 <i>Risks</i>	19
4.4.5 <i>JRA1 KPIs</i>	20
4.5. MESSAGING GUIDELINES	20
4.5.1 <i>Harmonization/convergence aspects</i>	20
4.5.2 <i>Affected/involved components</i>	20
4.5.3 <i>Subtasks</i>	20
4.5.4 <i>Risks</i>	20



EUROPEAN MIDDLEWARE INITIATIVE

4.5.5	JRA1 KPIs.....	20
4.6.	EMI CLOUD STRATEGY AND ARCHITECTURE	20
4.6.1	Harmonization/convergence aspects	20
4.6.2	Affected/involved components.....	20
4.6.3	Subtasks	20
4.6.4	Risks.....	21
4.6.5	JRA1 KPIs.....	21
4.7.	IMPLEMENT ACCOUNTING RECORD PUBLISHING VIA MESSAGING	21
4.7.1	Harmonization/convergence aspects	21
4.7.2	Affected/involved components.....	21
4.7.3	Subtasks	21
4.7.4	Risks.....	21
4.7.5	JRA1 key performance indicators addressed by the objective.....	21
4.8.	INFRASTRUCTURE AREA SIMPLIFICATION PLAN	21
4.8.1	Harmonization/convergence aspects	21
4.8.2	Affected/involved components.....	21
4.8.3	Subtasks	22
4.8.4	Risks.....	22
4.8.5	JRA1 KPIs.....	22
4.9.	SERVICE MANAGEMENT	22
4.9.1	Harmonization/convergence aspects	22
4.9.2	Affected/involved components.....	22
4.9.3	Subtasks	22
4.9.4	Risks.....	22
4.9.5	JRA1 KPIs.....	22
4.10.	IMPLEMENTATION OF THE INFRASTRUCTURE AREA SIMPLIFICATION PLAN.....	23
4.10.1	Harmonization/convergence aspects	23
4.10.2	Affected/involved components.....	23
4.10.3	Subtasks	23
4.10.4	Risks.....	23
4.10.5	JRA1 KPIs.....	23
4.11.	MONITORING INTERFACES	23
4.11.1	Harmonization/convergence aspects	23
4.11.2	Affected/involved components.....	23
4.11.3	Subtasks	23
4.11.4	Risks.....	23
4.11.5	JRA1 KPIs.....	23
4.12.	CLIENT TOOL USABILITY	24
4.12.1	Harmonization/convergence aspects	24
4.12.2	Affected/involved components.....	24
4.12.3	Subtasks	24
4.12.4	Risks.....	24
4.12.5	JRA1 KPIs.....	24
4.13.	OPERATING SYSTEM STANDARDS.....	24
4.13.1	Harmonization/convergence aspects	24
4.13.2	Affected/involved components.....	24
4.13.3	Subtasks	24
4.13.4	Risks.....	25
4.13.5	JRA1 KPIs.....	25



EUROPEAN MIDDLEWARE INITIATIVE

4.14. PLATFORM SUPPORT	25
4.14.1 Harmonization/convergence aspects	25
4.14.2 Affected/involved components	25
4.14.3 Subtasks	25
4.14.4 Risks	25
4.14.5 JRA1 KPIs	25
4.15. MONITORING PROBES	25
4.15.1 Harmonization/convergence aspects	25
4.15.2 Affected/involved components	26
4.15.3 Subtasks	26
4.15.4 Risks	26
4.15.5 JRA1 KPIs	26
4.16. EMI RESOURCE INFORMATION SERVICE	26
4.16.1 Harmonization/convergence aspects	26
4.16.2 Affected/involved components	26
4.16.3 Subtasks	26
4.16.4 Risks	27
4.16.5 JRA1 KPIs	27
5. CONCLUSIONS	28

1. INTRODUCTION

1.1. PURPOSE

This document reports on the progress with respect to the objectives for Infrastructure Area over the first year and presents the work plan for year two and three. The work plan is an extended technical description of the high-level objectives that are defined in the overall EMI Technical Development Plan.

1.2. DOCUMENT ORGANISATION

An executive summary of the document can be found in Section 2. Section 3 reports on the progress with respect to the objectives for Infrastructure Area over the first year and the work plan, covering year two and three, is presented in Section 4.

1.3. REFERENCES

R1	DNA1.3.1 - Technical Development Plan (M2) http://cdsweb.cern.ch/record/1277540
R2	Infrastructure Area Twiki https://twiki.cern.ch/twiki/bin/view/EMI/InfrastructureArea
R3	Infrastructure Area Indico http://indico.cern.ch/categoryDisplay.py?categId=3099
R4	EMI Description of Work (Public DoW) https://twiki.cern.ch/twiki/pub/EMI/EmiDocuments/EMI-Part_B_20100624-PUBLIC.pdf
R5	MJRA1.13 - Agreement on common information exchange methods http://cdsweb.cern.ch/record/1277582
R6	GLUE Specification v. 2.0 http://www.ogf.org/documents/GFD.147.pdf
R7	OGF http://www.ogf.org
R8	EMI Registry Design https://twiki.cern.ch/twiki/pub/EMI/EMIRegistry/EMIRegistryDesign-v0.2.doc
R9	EMI Messaging Guidelines https://twiki.cern.ch/twiki/bin/view/EMI/EMIMessagingGuidelines
R10	Designing the Next Generation Grid Information System http://www.cern.ch/lfield/CHEP2010.pdf
R11	OpenLDAP Replication http://www.openldap.org/doc/admin24/replication.html
R12	ActiveMQ http://activemq.apache.org
R13	EMI Cloud and Virtualization Task Force https://twiki.cern.ch/twiki/bin/view/EMI/EmiJra1T5TaskForceCloudandVirtualization
R14	EMI Messaging Requirements



EUROPEAN MIDDLEWARE INITIATIVE

	https://twiki.cern.ch/twiki/bin/view/EMI/EMIMessagingRequirements
R15	EMI Enterprise Messaging Solutions Survey https://twiki.cern.ch/twiki/pub/EMI/EMIMessaging/EMI-JRA-MESSAGING_SOLUTIONS_SURVEY-v1_0.FINAL.pdf
R16	EMI Service Management Task Force https://twiki.cern.ch/twiki/pub/EMI/EmiJra1T5TaskForceServiceManagement/TF_service_management.doc
R17	Nagios http://www.nagios.org
R18	Ganglia http://ganglia.sourceforge.net
R19	Usage Record Working Group http://www.gridforum.org/gf/group_info/view.php?group=ur-wg
R20	DJRA1.4.2 Infrastructure Area Work Plan and Status Report (M12) DJRA1.4.2 Infrastructure Area Work Plan and Status Report GANTT (M12) http://cdsweb.cern.ch/record/1277583
R21	DJRA1.4.1 Infrastructure Area Work Plan and Status Report (M3) http://cdsweb.cern.ch/record/1277581
R22	DNA1.3.2 Technical Development Plan (M11) http://cdsweb.cern.ch/record/1277543

1.4. DOCUMENT AMENDMENT PROCEDURE

This document can be amended by the authors further to any feedback from other teams or people. Minor changes, such as spelling corrections, content formatting or minor text re-organisation not affecting the content and meaning of the document can be applied by the authors without peer review. Other changes must be submitted to peer review and to the EMI PEB for approval.

When the document is modified for any reason, its version number shall be incremented accordingly. The document version number shall follow the standard EMI conventions for document versioning. The document shall be maintained in the CERN CDS repository and be made accessible through the OpenAIRE portal.

1.5. TERMINOLOGY

API	Application Programming Interface
ARC	Asynchronous Resource Connector
BDII	Berkeley Database Information Index
CHEP	Computing in High Energy Physics
DCI	Distributed Computing Infrastructure
DGAS	Distributed Grid Accounting System
DSR	Domain Service Registry
EMI	European Middleware Initiative
GLUE	Grid Laboratory Uniform Environment



EUROPEAN MIDDLEWARE INITIATIVE

GOC DB	Grid Operations Centre Database
GSR	Global Service Registry
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
OGF	Open Grid Forum
PEB	Project Executive Board
PTB	Project Technical Board
SAGE	Simple API for Grid Application
URL	Uniform Resource Locator
WAN	Wide Area Network
YAIM	YAIM Ain't an Installation Manager

2. EXECUTIVE SUMMARY

This Infrastructure Area work plan document gives the status of the Infrastructure Area work after the first year of the EMI project, concrete work plans for the second year and outlines plans for the third year. These plans exclude day-to-day maintenance and focus on major changes that are required to meet the strategic goals of the EMI project. Section 3 provides a report on the first year and the work plans for year two and three are given in Section 4.

The status report provides an update on the work done to achieve the objectives as described in DNA1.3.1 Technical Development Plan [R1] and DJRA1.4.1 Infrastructure Area Work Plan and Status Report (M3) [R21]. These status reports are organized by objectives, as given in DNA1.3.1, and the status at the time of writing is given. All of the objectives planned for the first year were met without exception. There were a number of major highlights that are worth mentioning here. A design document for the EMI Registry was produced and has been endorsed by the EMI Project Technical Board. The EMI Registry represents a major new component for EMI and as such good progress with this is encouraging. Information system infrastructure components now all support the GLUE 2 data model. GLUE 2 support is an important objective for EMI and the support GLUE 2 in such core components is critical for success. Messaging is another major new area for EMI and a great deal of work has been done in this area. Messaging guidelines are now available internally to EMI developers and a number of investigations have been carried out that evaluate the potential benefits of adopting the messaging system within EMI.

The work plans for the second and third year is organized by the objectives as described in DNA1.3.2 Technical Development Plan (M11) [R22]. These are the objectives from DNA1.3.1 that have been updated to reflect the experience of the past year. Furthermore, some new objectives have been added to address the requirements from EMI customers, notably EGI and WLCG. One of the main objectives for the second year is to provide an implementation of the EMI Registry. This is an aggressive timescale and although achievable, there are risks associated to the development of a new service, which will have to be carefully monitored during the course of the year to make sure the service can be released in time at least as a convincing prototype. In addition there will not enough time to obtain feedback from production deployment. Both risks can be mitigated by including an extra repository for EMI containing experimental software. This would enable new projects to be released outside of the main release cycle and for early adopters to evaluate them before their inclusion in the main EMI release. The Infrastructure Area implication plan is another major objective for the second year. It is important for this plan to be delivered early-on so that there is enough time left in the project to implement the suggestions. Related to this plan is the harmonization of the Resource Information Services via the provision of the EMI Resource Information Services (ERIS).

3. INFRASTRUCTURE AREA STATUS REPORT

This chapter reports on the progress during the first year with respect to the technical objectives as defined in DNA1.3.1 Technical Development Plan [R1] for the Infrastructure Area. Each subsection describes the status with respect to a specific objective. A few comments on the general progress in the Infrastructure Area are provided first.

In order to help manage the Infrastructure Area objectives, communicate these objectives with the Infrastructure Area participants and to ensure transparency, a Twiki page [R2] was created to provide information relevant to the Infrastructure Area. This page details for the Infrastructure Area:

- Outstanding Actions
- Objectives as defined in DNA1.3.1 Technical Development Plan
- Objectives as defined in DJRA1.4.1 Infrastructure Area Work Plan and Status Report
- EMI 1 Road Map for components
- The Product Teams, including the contact details
- Task Forces, including a link to the home page of the Task Force
- Details about communication

Three mechanisms are used for communication between Infrastructure Area participants. Most of the communication is conducted via the Infrastructure Area email list. Phone conferences take place when required; the agenda, along with the minutes, are recorded in Indico [R3]. Announcements for such phone conferences are communicated via the Infrastructure Area email list. Finally face-to-face meetings are arranged when required and when convenient. Examples of such were the workshops and parallel sessions that took place at the EMI All Hands Meeting in October 2010.

Many of the topics related to the Infrastructure Area, as defined in the EMI Description of Work (EMI DoW) [R4], were only vaguely described due their investigative nature. In order to solidify such investigative topics into concrete objectives, the following Task Forces were created.

- Cloud
- Virtualization
- Service Management
- Service Monitoring
- Accounting

As an outcome of the first EMI All Hands Meeting in October 2010, it was noticed that there was a great deal of synergy between the Cloud and Virtualization Task Forces. As a result, the PTB made a decision to combine these Task Forces. Conversely, even though Service Management and Monitoring are tasks described in the EMI DoW, separate Task Forces were created for Service Management and Service Monitoring due to the different maturity and adoption of related technologies within grid Infrastructures.

During the first year, the rationale for including the Logging and Bookkeeping Product Team in the Infrastructure Area was questioned. The main application of this component is to obtain, store, process and make available job status messages for the gLite middleware. As such, the first priority should be to achieve harmonization within the Compute Area, however, there are currently no such objectives defined in this respect. As a result, during the PTB meeting on 9th November 2010, a decision was made to move the Logging and Bookkeeping component from the Infrastructure Area to the Compute Area to help focus effort.

Logging and Bookkeeping as a generic infrastructure service in EMI only makes sense if there are generic middleware requirements that are fulfilled by this service. Within EMI it is envisaged that a

messaging infrastructure will be used for transporting messages and the Logging and Bookkeeping service is being adapted to make use of this infrastructure. The open question is what generic middleware requirements are required on top of the messaging infrastructure? As messaging infrastructure only transports messages, publishing, processing and archiving responsibilities are left to the clients. Certain functionality, for example trying to re-publish a message if the messaging infrastructure is unavailable or archiving messages, may be common to many clients. In such cases, a common component could be included in EMI to provide such functionality. Such requirements must be first been understood and only then can it be seen if a component from the Logging and Bookkeeping service could be used to provide the required functionality. In some cases, it may be more efficient to leverage the Product Teams experience to create a new component from scratch. As a result, a recommendation has been made to the PEB to rename this topic Event Publishing, Processing and Archiving, to better reflect the focus of this endeavor.

3.1. AGREEMENT ON COMMON INFORMATION EXCHANGE METHODS

The agreement on common information exchange methods was defined as a cross-area objective in DNA1.3.1 and solidified by milestone MS38 Agreement on common information exchange methods [R5]. The milestone outlines the steps that were taken in order to ensure that a common information exchange mechanism is available for the needs of accounting, information provisioning and messaging. This objective was achieved by addressing the two fundamental aspects of the information exchange: data models and data transport. In summary, EMI components should fully support the GLUE 2 [R6] data model from the Open Grid Forum (OGF) [R7]. For accounting, EMI components will support the publication of a compute usage record that leverages the preexisting work from the OGF and the EMI Data Area will define a storage usage record as no such record has been defined within the scope of OGF. In general, applications that make use of the messaging system are responsible for the data models they use. The messaging PT will, as part of their guidelines, provide information on how to best to construct messages. The EMI PTB will monitor emerging use cases and if it is envisaged that a common data model is required, it will initiate an activity which will ensure that a model is defined. In addition, relevant data models will be synchronized with other projects (e.g. EGI) and communities where appropriate. As for data transport, a number of objectives were included in the EMI work plan to achieve harmonization by removing duplicated solutions and to make use of the messaging infrastructure where appropriate.

3.2. GLUE 2 ADOPTION

The adoption of GLUE 2 is a major objective for the EMI 1 release and is referred to in the Description of Work, DNA1.3.1 and MS38. GLUE 2 is a standardized Grid information model that enables information system and Grid interoperability via the use of agreed attribute naming and semantics. Production infrastructures running today will need to seamlessly migrate to GLUE 2 so that day-to-day operations may continue without interruption. This results in the need for GLUE 2 to be added to EMI components in a backwards compatible way.

A service update for the top-level BDII was produced which provides a mechanism for the top-level BDII service to obtain GLUE 2 information from the site-level BDII service. In addition, a new mechanism was implemented for obtaining the list of LDAP URLs for the site-level BDII services directly from the GOC DB. For this update, the opportunity was taken to significantly improve the configuration through the inclusion of new configuration packages which simplifies the YAIM configuration. The result is that now both the top-level and site-level BDII included in the EMI 1 release fully support GLUE 2.

Support for publishing resource information in GLUE 2 by the ARC LDAP based information system are part of nordugrid-arc version 1.0.0 which is available in the EMI 1 release. This feature is not enabled by default, but can be turned on by adding `infosys_glue2_ldap="enabled"` in the `[infosys]` section of the `arc.conf` configuration file. In addition, the ability to aggregate this information as the site-level was also included.

The EMI 1 release also includes a version of the Common Information Provide that publishes GLUE 2 information about the computing and storage resources.

3.3. INFORMATION SYSTEM CLIENT HARMONIZATION

Information system client harmonization was discussed at various PTB meetings, especially during the discussion on whether or not to include the SAGA Service Discovery API in EMI. The decision made by the PTB to include the SAGA Service Discovery API means that four information system clients are available in the EMI 1 release. This will enable the project to obtain customer feedback on each product and the resulting information will be used during the section year to harmonize the information system clients for the EMI 2 release.

3.4. EMI REGISTRY

The participants of the EMI Registry Product Team brought together many years of production experience from contemporary DCIs and middleware initiatives. This combined experience helped to identify limitations in the existing approaches and will help pave the way to address the future service discovery requirements. A document has been produced [R8] that describes the proposed architecture of the EMI Registry.

This document also details the low-level specification and design of the individual components. An overall problem description is presented that outlines the system environment in which the EMI Registry will be operating and the main use cases it intends to address. Attention is also given to high-level non-functional requirements such as scalability and robustness. A detailed requirements section outlines the main functionality which the EMI Registry must address. The proposed architecture introduces two main components the Domain Service Registry (DSR) and the Global Service Registry (GSR). Services register and unregister with their local DSR, which is that authoritative information source for that domain. A higher-level DSR may aggregate information from lower-level DSRs. At the top-level the GSR, which is itself a globally distributed service, provides a global view of all services. For the information model, GLUE 2 is leveraged for the service description, which will enable the EMI Registry to be interoperable with other middleware that is GLUE 2 compliant.

The document was submitted to the PTB who accepted it with minor comments. This gives a green light for the continuation of the EMI Registry work and an implementation plan is in preparation.

3.5. MESSAGING

Guidelines have been given to the EMI Product Teams in the form of a *crash course on messaging* as part of the wider EMI Messaging Workshop that was part of the EMI 2010 All Hands Meeting which took place in Prague in November 2010. These guidelines have then been refined, extended, published on the web [R9] and announced to the EMI Messaging mailing list.

Messaging requirements from the other EMI Product Teams have collated and documented [R14]. A survey of the various options for an EMI messaging service has also been produced [R15]. Both of these documents have been published and disseminated within EMI.

3.5.1 Information System Use

During the summer 2010, an investigation into alternative transport mechanisms for the gLite information system was conducted. A full description of this investigation, along with the results, have been summarised in a paper [R10] that was presented at the CHEP 2010 conference. In this investigation, the use of syncrepl [R11] and ActiveMQ [R12] were both evaluated as replacement mechanisms for the transportation of changes between the site-level and top-level BDIs. It was found that the use of syncrepl over the WAN was not robust and would fail after a period of time. Further investigation is required to identify the cause of this failure. In contrast, the use of the ActiveMQ messaging infrastructure proved to be reliable for the duration of the test. Transporting only the changes using the messaging system reduced the amount of information transported by two orders of magnitude. The update time was also reduced by 56.4% for the LAN deployment scenario and 34.8%

for the WAN deployment scenario. The simplicity of syncrepl deployment and that it was successful while working over the LAN, suggests that it would be a good candidate for transporting changes from the resource-level to site-level. The use of messaging, while more complicated to set up, as it requires the deployment of a messaging infrastructure, is effective over the WAN and hence would be a good candidate for transporting changes from the site-level to top-level. However, a full cost-benefit analysis still needs to be considered.

3.5.2 Accounting Use

A new sensor module has been developed and certified for the DGAS client, which uses the ActiveMQ messaging system as mechanism to transport usage records between the computing services and the site-level accounting server. A new usage record producer, based upon OGF Usage Record [R19] standard as been developed and certified which also leverages the messaging infrastructure. The porting of the legacy transport mechanism for storage accounting to the ActiveMQ infrastructure is still in progress and not yet complete. However, the porting of the legacy transport mechanism to the Active MQ infrastructure for transporting usage records between sites and regional accounting databases has been completed.

APEL has used a private ActiveMQ broker network in production on EGI since mid 2010. Work started in year 1 to move this to the EGI production ActiveMQ broker network. The server side has been developed in EGI and the corresponding glite-apel client has been developed in EMI in python to use STOMP like most other production use of ActiveMQ.

A new authentication model has been used based on signing the messages with the host certificate of the client and encrypting them with public key of the central server. The decrypted messages can then be archived for auditing purposes after the usage records have been extracted.

The new server corresponding to this version of glite-apel can handle multiple message types, leaving it well placed to support schema evolution and usage records from other clients such as storage accounting records from SEs.

3.6. SERVICE MONITORING AND MANAGEMENT

Service monitoring requirements were obtained in conjunction with first EGI survey of site managers. The results from this survey [R16], which includes the state-of-the-art, showed that site managers are interested in dedicated management interfaces for Grid services. In addition, half of the site managers would like a messaging system to convey both management request from the administrator to the Grid service and monitoring information from the Grid service to administrators (and other interested parties).

The survey showed that there are many open-source monitoring tools, however most of the survey participants were using Nagios [R17] and Ganglia [R18]. As different monitoring systems are used, either an agnostic (regarding both used monitoring solution and middleware) format for describing Grid service monitor data is required or one monitoring solution should be chosen for everything related to Grid service monitoring.

3.7. VIRTUALIZATION AND CLOUDS

Initially there were two task focuses, one focusing on virtualization and the other focusing on Cloud computing. After the EMI all hands meeting the task forces merged as there was a great deal synergy between them. The resulting Cloud and Virtualization Task Force produced a document [R13] that contains a survey of existing cloud technologies and details the requirements, which also considers input from the EGI and other DCI projects. It also includes the high-level system components design which will form the basis for implementations and site deployments.

While EMI is advancing the Grid middleware functionalities to support the scientific computing requirements, it is essential to enable Grid services responding to the emergence of on-demand computing, dynamic runtime environments, and a deployment over different types of virtual



EUROPEAN MIDDLEWARE INITIATIVE

infrastructures. Consequently, this will generate new set of requirements for Grid providers maintaining middleware components within the scope of virtualization and Clouds, thus the task force document have drafted the set of use cases. The report also identified the candidate EMI services to be easily deployed over and/or functionally extended for dynamic Cloud computing environments.

The task force document describes the Cloud service and deployment models, with emphasis on grid related scenarios. As EMI is aiming at the harmonization of existing different middleware components which will be available and ready to use for wide range of scientific communities, it is important to define a unified system model - considering general grid middleware presence, which should define a basis of how EMI components can easily be enabled and interact within cloud-based infrastructures. Based on this research, the document then goes on to define a reference architecture for the Cloud model to be used in EMI.

4. INFRASTRUCTURE AREA WORK PLAN

This section gives the work plans of the Infrastructure Area in order to attempt to achieve the prescribed technical objectives from DNA1.3.2 Technical Development Plan (M12). The subsections below follow the technical objectives, giving a description of the objective, the benefits to EMI, sub tasks and risks.

Where appropriate, the JRA1 KPIs are given. These are:

- KJRA1.1: Number of adopted open standard interfaces.
- KJRA1.2: Number of interoperable interface usage.
- KJRA1.3: Number of reduced lines of code.
- KJRA1.4: Number of reduced released products.

The Infrastructure Area GANTT chart is available at [R20]. It illustrates the start and finish dates of the Infrastructure Area technical objectives and their sub-tasks. Dependencies between sub-tasks, where relevant, are also included.

4.1. MESSAGING INVESTIGATION; ACCOUNTING

This technical objective (DNA1.3.2 ref: I1) is to investigate and report on the possible use of the messaging in the accounting area. The investigation should consider the products from the messaging product team and understand what possible benefits could be realized by its integration into the accounting architecture. A report should be produced that describes the investigation, along with an implementation plan detailing how these benefits can be realized.

4.1.1 Harmonization/convergence aspects

Currently in EMI there are two accounting solutions; APEL and DGAS. The use of a common messaging infrastructure, along with defined usage records, would improve compatibility between the components. The messaging infrastructure would serve to decouple the accounting publishers and consumers and results in the components being interchangeable. In addition, the current mechanisms used to transport the accounting information can be decommissioning resulting in a reduction in maintenance cost for both the software and service operation.

4.1.2 Affected/involved components

APEL parsers, APEL Publisher, HLR Clients, HLR Sensors.

4.1.3 Subtasks

ID	Description	Responsible	Target Date
1	Identify a task leader for the investigation	Accounting Task Force	M12
2	Identify the possible usages of messaging with the existing accounting architecture	Task Leader	M13
3	Cost-benefit analysis of the possible usages	Task Leader	M13
4	Identification of the optimal usage, which should be implemented	Task Leader	M13
5	Definition of an implementation plan	Task Leader	M13
6	Produce a report that covers the investigation, cost-benefit analysis and	Task Leader	M14

implementation plan		
---------------------	--	--

4.1.4 Risks

There are three risks associated with this objective. The first risk is the availability of a task leader who has available effort and the relevant competencies for this task. The second risk is that the cost-benefit analysis reveals that the adoption of messaging provides no net benefit. The third risk is that it is not possible to adopt messaging in an identical way for both solutions. This would result in the harmonization aspect of this objective not being achieved as envisaged.

4.1.5 JRA1 KPIs

None, this is an investigation only.

4.2. MESSAGING INVESTIGATION; SERVICE MONITORING AND MANAGEMENT

This technical objective (DNA1.3.2 ref: I2) is to investigate and report on the possible use of the messaging for service monitoring and management. The investigation should consider the products from the messaging product team and understand how these could be used to fulfill service monitoring and management use cases. A report should be produced that describes the investigation, along with an implementation plan detailing how service monitoring and management via messaging can be realized.

4.2.1 Harmonization/convergence aspects

Within EMI there is currently no common method to neither monitor nor manage services. The messaging infrastructure could provide a common *data bus* that could be used to transport monitoring and management information. With the absence of an industry standard for service monitoring and management, EMI may be required to provide a proprietary solution for the benefit of its customers. This endeavor would harmonize the service monitoring and management methods for all EMI services.

4.2.2 Affected/involved components

This does not affect involve any components as it is an investigation.

4.2.3 Subtasks

ID	Description	Responsible	Target Date
1	Identify a task leader for the investigation	Service Monitoring and Management Task Force	M12
2	Gather the service monitoring and management requirements	Task Leader	M13
3	Investigate how messaging could be used to meet these requirement	Task Leader	M13
4	Provide a design of a possible solution based on messaging	Task Leader	M13
5	Production of a report that documents the requirements and describes the design of a possible solution.	Task Leader	M14

4.2.4 Risks

This is a high-risk objective. The first major risk is that the service management and monitoring requirements are not clearly defined. The second major risk is that this objective may be affected by factors that are outside the control of the project. For example, standards for data center management has been a hot topic for a number of years, however, it is not clear that the adoption rates of such standards have gained such critical mass to become ubiquitous. Providing a proprietary solution for EMI services only, may not be beneficial to customers who deploy EMI services. Such a solution may not easily integrate with the other service monitoring and management solutions that are currently used in customers data centers. With the absence of a clear standard, any work in this area may not deliver the intended benefits. However, the goal of this objective is to investigate the possible use of messaging rather than implement a service monitoring and management solution in all EMI services. As such this risk is contained for now but should be considered with respect to further work in this area.

4.2.5 JRA1 KPIs

None, this is an investigation only.

4.3. IMPLEMENT THE COMMON EMI REGISTRY

This technical objective (DNA1.3.2 ref: I6) is to provide an implementation of the EMI Registry in time for the EMI 2 release. A design document [R8] for the EMI Registry has already been produced and endorsed by the EMI Project Technical Board. The document serves as the basis for the implementation plan.

4.3.1 Harmonization/convergence aspects

The EMI Registry service will provide a common method for services to publish their existence and for consumers to discover available services. Such functionality is a core part of the Grid middleware and will provide a robust foundation on which to build. An adapter for the SAGA Service Discovery API will be produced to enable the EMI Registry to be queried using a standardized API.

4.3.2 Affected/involved components

EMI Registry, SAGA SD API.

4.3.3 Subtasks

ID	Description	Responsible	Target Date
1	Provide an implementation plan	EMI Registry PT Leader	M13
2	Provide the implementation by following the plan	EMI Registry PT Leader	M20
3	Integrate the EMI Registry into the EMI 2 release	EMI Registry PT Leader	M22

4.3.4 Risks

The major risk of this objective is that the software is not ready in time for the EMI 2 release. With limited effort available, it is not possible to assign extra effort to this task in order to meet the deadline. The result is that it is highly likely that this deadline will be missed. The ability to provide this software after the release date for EMI 2 should be considered to mitigate this risk.

Another major risk is the significant effort will be invested in an unproven design and approach. To mitigate this risk the new design will require prototyping to validate the design before investing in a production quality implementation. It may be necessary to revise the design and approach. The

progress of the implementation should therefore be closely followed to ensure that each stage is progressing as planned and revise the implementation plan when necessary. The implementation plan should consider such an approach and contain relevant milestones so that the progress can be measured and evaluated.

The final major risk is that the software is not fit for the purpose when first released. Previous experience has shown that it can take a couple of years to production harden new services. In order to mitigate the risk, the expectation of the new service should be managed to reflect this reality. In the first release it should be labeled as experimental or beta release. This will identify to customers that it still requires production hardening and assure them that issues found will be actively addressed.

The risk that the deadline EMI 2 will be missed and that the first release may not be fit for the purpose suggests that EMI should have another repository for such experimental software. This would ensure that customers do not install the software by accident and it is clear that the software may not initially be fit for the purpose.

4.3.5 JRA1 KPIs

This may affect KJRA 1.1, KJRA1.2, KJRA1.3, KJRA1.4.

4.4. GLUE 2 INFORMATION MODEL SUPPORT

This technical objective (DNA1.3.2 ref: I5) is to ensure that all information components fully supports the GLUE 2 information model. This includes the development of a validation tool for use in quality assurance.

4.4.1 Harmonization/convergence aspects

The adoption of GLUE 2 is a major harmonization objective for the EMI project. It will enable a common information exchange mechanism to be used between EMI components. The use of a common model will also enable the issue of duplicated functionality to be addressed.

4.4.2 Affected/involved components

ARIS, EGIIS, ARC Grid Monitor, ARC infoproviders, CEMon, BDII, Glue model, gLite service info providers, gLite site info provider, gstat-validation, lcg-info and lcg-infosites, SAGA-SD, SAGA-IN, CIP

4.4.3 Subtasks

ID	Description	Responsible	Target Date
1	Identification of components that are required to support GLUE 2	Area Leader	M13
2	Assign a change request for each component	Area Leader	M14
3	Track the change requests to ensure that they are addressed on time	JRA1 Leader	M20

4.4.4 Risks

This objective needs to be considered with respect to the overall harmonization of the information system components. As some components will be deprecated, it may not be necessary to add GLUE 2 support. A holistic approach is required in this area and the move to GLUE 2 is just one aspect.

There is a possibility that not all the change requests may be addressed by the deadline. In such cases and expected delivery date for when that change will be available will be given with an explanation for the delay.

4.4.5 JRA1 KPIs

This may improve KJRA 1.1, KJRA1.2, KJRA1.3 and KJRA1.4.

4.5. MESSAGING GUIDELINES

This technical objective (DNA1.3.2 ref: I7) is to provide guidelines for third parties to integrate messaging into their service or application based on the EMI experience.

4.5.1 Harmonization/convergence aspects

This will enable non-EMI services and other application to make use of EMI messaging.

4.5.2 Affected/involved components

This does not affect involve any components as it is a document.

4.5.3 Subtasks

ID	Description	Responsible	Target Date
1	Produce the guidelines	Messaging PT Leader	M22
2	Disseminate the guidelines	Messaging PT Leader	M24

4.5.4 Risks

The risks associated with the objective are negligible as draft guidelines are already available.

4.5.5 JRA1 KPIs

None, the end product is documentation.

4.6. EMI CLOUD STRATEGY AND ARCHITECTURE

This technical objective (DNA1.3.2 ref: I4) is to provide a Cloud strategy and architecture for EMI. It should cover the the state-of-the-art with respect to Cloud computing and virtualization technology. An analysis should be carried out to understand how EMI could benefit from such technology and possible architectural changes to EMI middleware stack should be suggested. This work is to be carried out by the Cloud computing task force.

4.6.1 Harmonization/convergence aspects

This gives insight into future strategies for EMI.

4.6.2 Affected/involved components

This does not affect involve any components as it is an investigation.

4.6.3 Subtasks

ID	Description	Responsible	Target Date
1	Describe the available Cloud computing and virtualization technologies and standards	Task Force Leader	M12
2	Investigate how EMI could benefit from the adoption of such technology	Task Force Leader	M14
3	Document the possible directions of EMI middleware stack to the virtualized environments	Task Force Leader	M16

4	Produce a report contain the above information	Task Force Leader	M18
---	--	-------------------	-----

4.6.4 Risks

The only risk for this objective is that the report may not be ready on time or lacks sufficient depth.

4.6.5 JRA1 KPIs

None, the end product is documentation.

4.7. IMPLEMENT ACCOUNTING RECORD PUBLISHING VIA MESSAGING

This technical objective (DNA1.3.2 ref: I9) is provide or adapt the accounting publishers for compute and Data Area services to use the common messaging system.

4.7.1 Harmonization/convergence aspects

The harmonization and convergence aspects of this objective were outlined in 4.1.1.

4.7.2 Affected/involved components

The affected/involved components were specified in 4.1.2.

4.7.3 Subtasks

ID	Description	Responsible	Target Date
1	Obtain Compute Usage Record specification	Accounting Task Force Leader	M13
2	Obtain Storage Usage Record specification	Accounting Task Force Leader	M14
3	Create the publisher according to the specifications	APEL and DGAS PT Leaders	M16
4	Integrate into the EMI release	APEL and DGAS PT Leaders	M28

4.7.4 Risks

The main risk for this objective is the assumption that the investigation showed a positive net benefit from this change. The other risk is the availability of agreed compute and storage usage records.

4.7.5 JRA1 key performance indicators addressed by the objective

This may improve KJRA 1.1, KJRA1.2, KJRA1.3 and KJRA1.4.

4.8. INFRASTRUCTURE AREA SIMPLIFICATION PLAN

This technical objective (DNA1.3.2 ref: I3) is provide a plan for substantial simplification and reduction in the number of Infrastructure Area CLIs, libraries, internal components and services.

4.8.1 Harmonization/convergence aspects

This plan contains simplifications and reductions via the removal of components where functionality is duplicated.

4.8.2 Affected/involved components

The list of affect/involved components will be revealed during the investigation.

4.8.3 Subtasks

ID	Description	Responsible	Target Date
1	Provide an architecture documenting the current situation	Infrastructure Area Leader	M14
2	Identify possible changes that could achieve the desired results	Infrastructure Area Leader	M15
3	Achieve agreement within EMI on which changes to pursue	Infrastructure Area Leader	M17
4	Produce a report showing the architecture, desired changes and implementation plan	Infrastructure Area Leader	M18

4.8.4 Risks

The main risk with this objective is achieving consensus within EMI on what changes to pursue, especially those that involve the deprecation of existing components. This risk can be mitigated by referring contentious issue to the Project Technical Director for arbitration.

4.8.5 JRA1 KPIs

None, the end product is documentation.

4.9. SERVICE MANAGEMENT

This technical objective (DNA1.3.2 ref: I8) is to investigate service remote management interface for compute, data, security and infrastructure services, including remote configuration change and service management, utilizing the messaging system.

4.9.1 Harmonization/convergence aspects

None, this is a new feature.

4.9.2 Affected/involved components

This does not affect/involve any components as it is an investigation.

4.9.3 Subtasks

ID	Description	Responsible	Target Date
1	Gather the requirements on service management	Task Force Leader	M20
2	Produce a state-of-the-art report on service management	Task Force Leader	M22
3	Suggest a direction for EMI with justifications for this decision	Task Force Leader	M24

4.9.4 Risks

The main risk for this objective is the requirements gathering task. It may not be possible to clearly gather the requirements from the EMI customers.

4.9.5 JRA1 KPIs

None, the end product is documentation.

4.10. IMPLEMENTATION OF THE INFRASTRUCTURE AREA SIMPLIFICATION PLAN

This technical objective (DNA1.3.2 ref: I10) is to implement the plan for simplification and reduction in the number of Infrastructure Area CLIs, libraries, internal components and services.

4.10.1 Harmonization/convergence aspects

The harmonization and convergence aspects of this objective were outlined in 4.8.1.

4.10.2 Affected/involved components

The affected/involved components will be identified in 4.8.2.

4.10.3 Subtasks

The definition of subtask will be defined in the implementation plan.

4.10.4 Risks

The main risk associated with this objective is the timely delivery a realistic implementation plan.

4.10.5 JRA1 KPIs

This may improve KJRA 1.1, KJRA1.2, KJRA1.3 and KJRA1.4.

4.11. MONITORING INTERFACES

This technical objective (DNA1.3.2 ref: X14) is to adapt or implement monitoring interfaces, sensors and providers for infrastructure services to allow the use of standard monitoring tools preferably based on the common EMI messaging system.

4.11.1 Harmonization/convergence aspects

This objective is part of an overall EMI objective to ensure that all EMI services use the same monitoring interface.

4.11.2 Affected/involved components

The list of affect/involved components will be revealed during the investigation.

4.11.3 Subtasks

ID	Description	Responsible	Target Date
1	Contribute to the definition of the monitoring interface	Task Force Leader	M16
2	Identify the list of components which need to implement this monitoring interface	Infrastructure Area Leader	M17
3	Add change requests for each component	Infrastructure Area Leader	M18
4	Monitor each change request to ensure timely delivery	JRA1 Leader	M30

4.11.4 Risks

The main risk associated with this objective is the definition of a monitoring interface. This interface may not be defined in time for the work to be carried out before the end of the project.

4.11.5 JRA1 KPIs

This may improve KJRA 1.1, KJRA1.2, KJRA1.3 and KJRA1.4.

4.12. CLIENT TOOL USABILITY

This technical objective (DNA1.3.2 ref: X8) is to improve usability of client tools based on customer feedback by ensuring a) better more informative, less contradictory error messages b) coherency of commands line parameters.

4.12.1 Harmonization/convergence aspects

This objective is part of an overall EMI objective to improve the usability of client tools.

4.12.2 Affected/involved components

The list of affect/involved components will be revealed during the investigation.

4.12.3 Subtasks

ID	Description	Responsible	Target Date
1	Obtain a client tool policy document for EMI	Infrastructure Area Leader	M13
2	Cross check the policy document with the client tools in the infrastructure to identify improvements	Infrastructure Area Leader	M16
3	Add change a request for each improvement	Infrastructure Area Leader	M16
4	Monitor each change request to ensure timely delivery	JRA1 Leader	M22

4.12.4 Risks

The main risk associated with this objective is the availability of a client tool policy document.

4.12.5 JRA1 KPIs

This may improve KJRA1.3, KJRA1.4.

4.13. OPERATING SYSTEM STANDARDS

This technical objective (DNA1.3.2 ref: X6) is to adhere to operating system standards for service operation and control regarding configuration, log and temporary file location and service start/status/stop.

4.13.1 Harmonization/convergence aspects

This objective is part of an overall EMI objective to improve the adherence of EMI services with the Operations System standard practices.

4.13.2 Affected/involved components

The list of affect/involved components will be revealed during the investigation.

4.13.3 Subtasks

ID	Description	Responsible	Target Date
1	Obtain OS adherence policy document for	Infrastructure Area	M13

	EMI	Leader	
2	Cross check the policy document with services in the Infrastructure Area to identify improvements	Infrastructure Area Leader	M14
3	Add change a request for each improvement	Infrastructure Area Leader	M14
4	Monitor each change request to ensure timely delivery	JRA1 Leader	M18

4.13.4 Risks

The main risk associated with this objective is the availability of an OS adherence policy document.

4.13.5 JRA1 KPIs

This does not address any KPIs.

4.14. PLATFORM SUPPORT

This technical objective (DNA1.3.2 ref: X7) is to port, release and support EMI components on identified platforms (full distribution on SL6 and Debian 6, UI on SL5/32 and latest Ubuntu).

4.14.1 Harmonization/convergence aspects

None.

4.14.2 Affected/involved components

The list of affect/involved components will be revealed during the investigation.

4.14.3 Subtasks

ID	Description	Responsible	Target Date
1	Identify which components must be available on which platforms	Infrastructure Area Leader	M14
2	Ensure that each component is available on that platform for the release date.	JRA1 Leader	M22

4.14.4 Risks

The risks associated with this objective are difficult to ascertain. The problems to be overcome will only be known once an attempt is made to port the software to a specific platform.

4.14.5 JRA1 KPIs

This does not address any KPIs.

4.15. MONITORING PROBES

This technical objective (DNA1.3.2 ref: X4) is to provide and support monitoring probes for EMI services (e.g. Nagios).

4.15.1 Harmonization/convergence aspects

This objective is part of an overall EMI objective to ensure that all EMI services have monitoring probes.

4.15.2 Affected/involved components

The list of affect/involved components will be revealed during the investigation.

4.15.3 Subtasks

ID	Description	Responsible	Target Date
1	Identify a list of components where monitoring probes are required	Infrastructure Area Leader	M14
2	Add change a request for the delivery of each probe	Infrastructure Area Leader	M14
3	Monitor each change request to ensure timely delivery	JRA1 Leader	M18

4.15.4 Risks

The main risk associated with this objective is the agreement on the format of the monitoring probe.

4.15.5 JRA1 KPIs

This does not address any KPIs.

4.16. EMI RESOURCE INFORMATION SERVICE

This technical objective (DNA1.3.2 ref: *additional development items*) is to agree upon and implement a standardized/common mechanism for obtaining service and resource information from Grid services (“local information”). The goal of this objective is to provide a single EMI Resource Information Service (ERIS). Currently within EMI there exist three resource-level information services, which results in a considerable duplication of effort. While the resource information service is not a complex service (it is a single service that enables remote access to information about a single resource), the development and support effort for the require resource information providers is non-trivial. The provision of a single service that also has a well specified interface for information providers would be of a huge benefit EMI and system integrators in general. As all the existing resource-level information services are GLUE 2 compliant, the can be used to help infrastructures migrate to GLUE 2. Therefore, the ERIS only needs to focus on the GLUE 2 information model.

4.16.1 Harmonization/convergence aspects

This objective removes the current duplication with respect to the resource-level information services within EMI. In addition, and more importantly, it will simply the development and support of information providers, which is required during system integration efforts. The removal of the need to support legacy information models will simplify the code base and ensure convergence on the GLUE 2 information model.

4.16.2 Affected/involved components

This objective affects the resource-level BDII, the ARC Resource Information Service (ARIS) and the Common Information Provider (CIP).

4.16.3 Subtasks

ID	Description	Responsible	Target Date
1	An agreement within EMI on the common approach	Infrastructure Area Leader	M14
2	An implementation plan should be produced	PT Leader	M15

3	Integrate in to the EMI 2 release	PT Leader	M24
---	-----------------------------------	-----------	-----

4.16.4 Risks

There are two main risks associated with this objective. The first is that the agreement within EMI is not achieved according within the desired time frame. The second is that it is not possible to provide a working version of the service before the code freeze for EMI 2.

4.16.5 JRA1 KPIs

This may improve KJRA 1.1, KJRA1.2, KJRA1.3 and KJRA1.4.

5. CONCLUSIONS

This document has described the status of the work performed by the EMI Infrastructure Area in the first year of the project according to the objectives as outlined in DJRA1.4.1. All of the first year objectives were met without exception. The concrete work plan for the second year and outline work plan for the third year were provided, based on objectives defined in DNA1.3.2.