



Overview of the EMI development objectives

Balázs Kónya (Lund University)

EMI Technical Director

EMI All Hands Prague, 23rd November 2010

Outline

- Development principles
- EMI software stack
- Technical objectives
- Timeline



Development principles

- Support components in production
 - Re-active & pro-active maintenance
 - Phase-out plans
- Consolidate and streamline
 - Harmonize and integrate components
 - Remove duplications
 - Merge and re-use components
- Evolve middleware
 - Extend functionality and operational capability
 - Harden components
 - Cautiously introduce new components
 - preferably to be based on existing code or off-the-shelf 3rd party solutions

EMI portfolio

- 98 components
 - 2 planned, 1 alpha, 3 beta, 6 ready, all the rest in-production
 - 37 services; others are client, library or internal
 - lots of internal components (especially within security)
 - 11 components are already tagged as "phase out", several others as "investigate phase out"
- All components originate from the 3+1 middleware stack
 - So far only two new "planned" EMI components listed: EMI Service Registry and EMI Messaging layer
- The components belong to one of the four areas (compute, data, security, infrastructure) and taken care of one of the 29 Product Teams

Component table



COLLABORATION PROGRAMS

Doc. Identifier: EMI_DNA1.3.1_draft.docx

Date: 21/11/2010

	Component	PT	Area	Type	Status	Maintenance and Development plan			
						Phase out	Maintenance	Harmonization	Evolution
17.	<u>LocalLDAP</u>	ARC Information System	I	service	in-prod	investigate	pro-active	integrate	no
18.	EGIIS		I	service	in-prod	investigate	support	integrate	no
19.	ARC Grid Monitor		I	client	in-prod	no	support	integrate	yes
20.	<u>ARC infoproviders</u>		I	internal	in-prod	no	support	integrate	yes
21.	<u>update-crls</u>	ARC Security Utils	S	internal	in-prod	yes	support	no	no
22.	<u>nordugridmap</u>		S	internal	in-prod	investigate	support	merge	yes
23.	<u>arcproxy</u>		S	client	in-prod	no	support	integrate, merge	yes
24.	HED	C Container	I	service	ready	no	pro-active	no	yes
25.	HED security		S	internal	ready	no	pro-active	integrate, merge	yes
26.	HED LIDI		I	internal	ready	no	pro-active	integrate	yes

Technical objectives

- Define directions and priorities for EMI software stack evolution
 - Product Team developments must follow the objectives
 - Development work not addressing the objectives is not allowed
- Form the EMI technical plan (DNA1.3.1)
 - Originate from the DoW, Maintained by the Project Technical Board and Tech Director
 - Structured around 4+1 areas
 - An objective corresponds to a new feature, consolidation task or a preparational task
 - High level formulation
- Dow vs. DNA1.3.1 vs. area workplans vs. PT dev plans
 - Dow: "vision"
 - Tech objectives of DNA1.3.1: concretization of the vision, directions
 - Area workplans: clarification of tech content for each objective
 - PT development plans: specification of the development tasks

Tech objectives table



COLLABORATION PROGRAMS

Doc. Identifier: EMI_DNA1.3.1_draft.docx

Date: 21/11/2010

	Cross area Technical Objectives	Due	Components
9.	The legacy <u>Globus</u> security infrastructure (GSI) will be replaced with a common security solution based on TLS/SSL still including the delegation capability.	M24	DPM, <u>dCache</u> , <u>StoRM</u> , <u>gfal</u> , WMS, CREAM, A-REX, HED, <u>libarcclient</u> , <u>libarcdata2</u>
10.	Adapt or implement monitoring interfaces, sensors, providers for compute, data, security and infrastructure services to allow the use of standard monitoring tools preferably based on the common EMI messaging system.	M24	all EMI services
11.	Investigate service instrumentation interface for compute, data, security and infrastructure services, including remote configuration change and service management, utilizing the messaging system.	M24	all EMI services
12.	Complete migration to the new <u>AuthN</u> libraries.	M36	<u>Dcache</u> , DPM, <u>StoRM</u> , LFC, FTS, A-REX, HED, WMS, CREAM, CEMON, UNICORE services, U. gateway, U. sec. <u>libs</u> , ARGUS, VOMS, Hydra, SLCS, STS, <u>Trustmanager</u> , and all the corresponding clients

Compute area technical objectives (1/2)

1. Glue 2.0 support in job management services and client tools.
2. Implementation of the agreed common job submission and management methods in all the CEs and compute clients.
3. Provide limited interactive access for at least one EMI Computing element.
4. Support for the agreed compute accounting record (UR).
5. Consolidation and harmonization of compute area clients/APIs.

Compute area technical objectives (2/2)

6. Extend job definition language, resource information (GLUE model) and job management service capabilities so that EMI compute clients are able to request access to virtualized resource managers and appliances.
7. Successful computational usage of emerging computing models i.e. clouds with EMI components (scaling out to clouds).
8. Provision of a common MPI execution framework, a “backend” across the different computing services to allow users to execute parallel applications in a uniform way.
9. Extend the parallel computing capabilities to better address multi-core jobs on all emerging architectures resources, multi-node execution on interconnected clusters; and special scenarios like advanced topologies, FPGAs, GPGPUs

Data area technical objectives (1/2)

1. All storage elements publishing initial GLUE 2.0 storage information.
2. Using https instead of httpg for the SRM protocol as a prototype implementation in one storage element and client (library).
3. All storage elements offering support for the http(s) protocol.
4. All storage elements offering at least a prototype-level support for the "file://" access protocol.
5. File Catalogue Access from UNICORE
6. One storage client is capable consuming GLUE 2.0 information published by storage elements.
7. All storage elements publishing full set of GLUE 2.0 storage information and EMI clients are capable consuming that.
8. Storage elements offering support for the WebDav protocol.

Data area technical objectives (2/2)

9. Using https instead of httpg for the SRM protocol as a production implementation in all the storage elements and clients.
10. Overall consolidation of data area by adopting a consistent interpretation of SRM.
11. Providing a common set of data access libraries at least between gLite and ARC.
12. Solve the synchronization problem of the storage elements and the file catalogue.
13. Integration of SRM-based access into UNICORE storage management
14. Completed migration to the common set of data access libraries.
15. Add support for storage space usage accounting on the SE/FTS side, including the refinement, definition and adoption (if/when applicable) of relevant standards.

Security area technical objectives (1/2)

1. Agreement on a minimal common set of security attributes to be used in policies.
2. Simplified management of security credentials by reducing the complexities of handling certificates and integrating different security mechanisms like Shibboleth and Kerberos across the EMI stack that allows users to use their own authentication system to access a ``Grid''.
3. Provide common authentication libraries supporting X.509 and optionally SAML.
4. Consolidation and reduction in the number of security CLIs so that the users don't have to face the very different clients and utilities.

Security area technical objectives (2/2)

5. Agreement and full support for a common single X.509 and SAML based Attribute Authority Service integrated with all EMI components.
6. Substantial simplification and reduction in the number of security area libraries, internal components and services.
7. Provide a transparent solution for encrypted storage utilizing ordinary EMI SEs.

Infrastructure area objectives (1/2)

1. Provide early internal guidelines for integrating messaging into potential EMI target components.
2. Design a common EMI service registry that is required in order to discover all the service endpoints of the different middleware components.
3. Investigate possible use cases for a common standard messaging system in the accounting area.
4. Investigate possible use cases for a common standard messaging system for the service monitoring and management.
5. Investigate possible use cases for a common standard messaging system for the information services and L&B.

Infrastructure area objectives (2/2)

6. Implement the common EMI Registry.
7. Fully utilize and support the GLUE2 information model.
8. Provide guidelines for 3rd parties to integrate messaging into their service/application based on the EMI experience.
9. Explore the modifications necessary in the EMI services to take advantages of the elasticity of the clouds resource management model while provisioning grid services within virtual machines (“grid in a cloud” scenario).
10. Implement or adapt the accounting record publishers of compute and data area services to use the common messaging system.
11. Consolidation and reduction in the number of information system discovery APIs and CLIs.

Cross area objectives (1/2)

1. Define the Information Flow architecture describing messaging and non-messaging based information exchange of the EMI components (e.g. service registry, information system, accounting, monitoring, and instrumentation). A common information exchange between the EMI components is preferable.
2. Investigate possible use cases for a common standard messaging system in the computing area. UMD capability
3. Investigate possible use cases for a common standard messaging system in the data area.
4. Evaluate integration scenarios with off-the-shelf computing cloud systems to be able to execute grid jobs on those (scaling out to clouds).
5. An EMI-blessed delegation solution for at least the computing area.
6. Definition and implementation of initial support for the common SAML profile all over the middleware stacks.

Cross area objectives (2/2)

7. Integration of the compute area services with the ARGUS authorization framework.
8. Initial integration of the storage elements with the ARGUS authorization framework.
9. The legacy Globus security infrastructure (GSI) will be replaced with a common security solution based on TLS/SSL still including the delegation capability.
10. Adapt or implement monitoring interfaces, sensors, providers for compute, data, security and infrastructure services to allow the use of standard monitoring tools preferably based on the common EMI messaging system.
11. Investigate service instrumentation interface for compute, data, security and infrastructure services, including remote configuration change and service management, utilizing the messaging system.
12. Complete migration to the new AuthN libraries.

Roadmap: timeline

- First phase, EMI-I
 - important technical agreements
 - consolidation plans for compute and data
 - design and early prototypes
 - additional new capabilities for production ready components.
 - EMI-1 release due April 2011.
- Second phase, EMI-2
 - consolidation plans for the remaining two areas
 - Some design and prototypes
 - Most intensive development phase resulting production ready features
 - EMI-2 due April 2012
- Third phase, EMI-Final
 - completing the consolidation plans
 - bringing the prototypes to production level
 - EMI-3 (or Final) release due April 2013

EMI phase 1 (2011 April)

- Agreements:
 - Execution Service interface
 - Accounting records (compute, storage)
 - Common security attributes
 - Common SAML profiles
 - Messaging use cases
 - Replacement of legacy GSI
 - EMI delegation
 - AAI "strategy"
- Consolidation plans
 - Compute area clients & APIs
 - Data access libraries (EMI_datalib)
 - Common authentication library (EMI_authlib)
- Design or early Prototypes
 - File catalogue and SE synchronization
 - EMI Service Registry
- Production ready (on top of EMI-0)
 - GLUE2 support in compute area

EMI phase 2 (2012 April)

- Consolidation plans
 - Security area components
 - Information system components
- Design or early Prototypes
 - EMI_authlib
 - EMI_datalib
 - Messaging-based service instrumentation

EMI phase 2 (2012 April)

- Production ready (on top of EMI-1)
 - CEs and clients with EMI interface
 - CEs with EMI accounting record
 - Consolidated compute area CLIs and APIs
 - GLUE2 support in data area
 - Glue2 support in infra area
 - All SEs supporting "file://", https, WebDav
 - Adoption of a fully consistent SRM implementation within data area
 - SEs and LFC synchronization
 - Transparent integration with AAI
 - VOMS as the single common EMI Attribute Authority Service
 - ARGUS and compute area services integration
 - EMI Service Registry
 - Messaging-based accounting publishers for compute and data area
 - Messaging-based monitoring sensors for all EMI services

EMI phase 3 (2013 April)

- Production ready (on top of EMI-2)
 - CEs capable scaling out to Clouds
 - CEs with consolidated MPI support
 - Consolidated data components including migration to EMI_datalib
 - Complete support for storage space accounting
 - Consolidated information system components
 - Consolidated security area components including migration to EMI_authlib
 - ARGUS and data area services integration
 - "Cloud-friendly" EMI services

References

- DNA1.3.1 Technical Development plan
twiki.cern.ch/twiki/bin/view/EMI/DeliverableDNA131
- Technical area workplans
 - DJRA1.1.2 - Compute area work plan and status report
 - DJRA1.2.1 - Data area work plan and status report
 - DJRA1.3.1 - Security area work plan and status report
 - DJRA1.4.1 - Infrastructure area work plan and status report
- Component dev. plans for EMI-1
- DSA1.2 Software Release plan



Thank you

EMI is partially funded by the European Commission under Grant Agreement INFSO-RI-261611