



EMI Execution Service Status

Bernd Schuller
Jülich Supercomputing Centre
On behalf of the EMI ES team

Outline

- Some background
- EMI Execution Service
 - Port types
 - Delegation
 - Processing model and state model
 - Data staging
 - Job description

Background: EMI

- EMI is a “joint venture” of European middleware providers
 - ARC, UNICORE, gLite, dCache, ...
- Goals
 - Support existing gLite, ARC, UNICORE, dCache
 - Harmonisation, evolution, integration
 - Less maintenance effort, reduced duplication, ...
 - New developments: clouds, virtualisation, messaging,

Background: EMI ES

- Common specification by key people from ARC, UNICORE, gLite
- Based on
 - PGI and “AGU” spec (v0.42)
- Status and plans:
 - First version (EMI Milestone) due Nov. 2010, consists of document and xsd/wSDL
 - Prototype implementations in 2011 by all the providers
 - Second version will be done → feedback loop
- <https://twiki.cern.ch/twiki/bin/view/EMI/WebHome>

Basic assumptions and scope

- Service of the “compute element” type
 - ARC CE, Cream CE, UNICORE atomic services
- Accepts and manages single jobs that run on some backend
 - Cluster managed by an LRMS
 - “Standalone” compute node
- Out of scope (at least for now)
 - Brokering / forwarding of jobs
 - Multi-jobs (e.g. JSDL parameter sweep)

“Modules” and port types

- Activity Factory module
 - Create : create activities
 - ResourceInfo : GLUE2 service info
 - Delegation : manage credentials
- Activity Management module
 - ActivityManagement
 - ActivityInfo : GLUE2 activity info
 - Delegation

Basic ideas

- Where possible, operations support sets of requests (“vector operations”)
- Operations are asynchronous
 - Avoid slow replies
- Example:
 - Submission of many jobs
 - Service immediately replies with vector of IDs
 - Validation stage processed async.

Delegation

- Allow to delegate trust to the service, avoiding GSI
 - Currently *limited* to proxies
 - SAML assertions as a future option
- InitDelegation
 - Client asks for a X509 CSR from the server
 - Server responds with CSR + delegation ID
- PutDelegation
 - Client puts the signed CSR + delegation ID to the server, allowing the server to create a proxy certificate

Data staging

- Three dirs, available at various stages in the processing
 - stage in / session dir / stage out
- Modes
 - Server pull/push
 - Client push/pull

State model

- Primary states + state attributes
 - Accepted
 - Preprocessing
 - Processing-accepting
 - Processing-queued
 - Processing-running
 - Postprocessing
 - Terminal

State attributes

- Indicate some transient condition, or convey additional information about the primary state
- Examples in “Preprocessing”
 - Client-stagein-possible
 - Provisioning
 - Client-paused
 - Server-stagein

Job description

- Model + XML rendering
- Elements
 - Job metadata (job name and such)
 - Resources
 - Executable + its attributes
 - Runtime environments
 - Data staging specification

Resources

- Similar to JSDL, but improved (?!)
 - Operating system
 - Value ranges
- VERY similar to JSDL: resource specification:
 - Nodes, processors, memory
- Under discussion
 - Scalable runtime request (i.e. linked to a benchmark value)

Executable + its attributes

- Executable
 - Input/output/error redirect
 - Arguments
 - Environment
-
- Executable can be provided by a Runtime environment (“abstract” application)

Runtime environments

- Provide the job's “ecosystem”
- Used by the ES to generate “concrete” job
- Provides environment variables, path settings, etc.
 - e.g. “module load deisa”
- May provide predefined executable
 - e.g. “Gromacs” → /opt/gromacs/bin/gromacs

Summary

- EMI ES effort will provide a joint web service interface to ARC, gLite, UNICORE compute services
- Understood and agreed (sort of)
 - Basic modules and port types
 - Processing model
 - Data staging
- Work in progress
 - Job description and its XML rendering
- Agreement needed within EMI
 - Delegation

EMI ES team

- Massimo Sgaravato, Eric Frizziero, Luigi Zangrando (gLite)
- Martin Skou Andersen, Aleksander Konstantinov, Balazs Konya, Oxana Smirnova (ARC)
- Shahbaz Memon, Shiraz Memon, Bernd Schuller (UNICORE)
- <https://twiki.cern.ch/twiki/bin/view/EMI/WebHome>



Thank you!

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