



EUROPEAN MIDDLEWARE INITIATIVE

# Compute Area Status and Plans

Marco Cecchi – INFN CNAF

Massimo Sgaravatto – INFN Padova

## Y1 Status

- DJRA1.1.1 was broken down into very detailed tasks and comprised ~30 objectives, grouped in 22 in DJRA1.1.2:
  - 16 green (100% complete)
  - 5 yellow
  - 1 red ( $\leq 10\%$  complete)
- The reasons for those yellows/red are mainly because of:
  - Time taken by the first EMI release
  - GLUE2.0/infoproviders 'affair' for CREAM
    - Nonetheless some interesting developments
  - Some other very specific objectives not completely fulfilled

## Y1 obstacles/1

- Time taken by the first EMI release
  - All compute area components made it to Kebnekaise
  - Great achievement!
  - But it required a very big deal of effort for several reasons
  - E.g. issues with system dependencies
    - Just as an example, in EMI-1, gLite WMS had to port mod\_fastcgi to mod\_fcgid
  - Testing and certification issues
    - Many dependencies on other services (in particular for gLite compute)
      - 4 RCs for EMI-1, but able to test our software only starting with RC3
    - RC changing while PT certifies its components --> require testing afresh each time
  - The preparation of the next releases should require less time as the release process converges :-)

## Y1 obstacles/2

- GLUE2.0/infoproviders ‘affair’ for CREAM/WMS
  - CREAM cannot be made responsible for other’s people work with the magic wand or, worse, only by an authoritative decision
  - A reasonable way to go is not clear even for Y2
    - Integration of ARC LRMS modules could be another possibility, discussed in these days
  - gLite-CLUSTER had to be provided in the meantime
    - Declared initially out of scope in EMI
    - This decision had impacts also on glue2 support
- Maybe many too low level objectives specified
  - Somewhat risky and even not required by the project...

## Some Y1 successes/1

- Y1 high-level objectives were carried out with a high degree of success, in particular:
  - Definition of the EMI-ES
    - Completed around the times of the 1<sup>st</sup> EMI AH
    - EMI-ES was kept abstracted from security (see Aleksandr talk)
      - This is good on one side
      - Some is complaining about lack of direct integration with Security
    - Now the implementation has to start

## Some Y1 successes/2

- Integration with Argus
  - Done for gLite (CREAM)
  - Prototypes for UNICORE and ARC
  - Work to be finalized
    - Issues with policy language
      - » UNICORE not able to implement all the policies that can be implemented with their existing system
    - Common XACML Authorization Profile still needs to be implemented
    - EMI-ES integration
- Support for GLUE 2.0
  - Done in ARC and UNICORE
  - Partially done in gLite
- Several other new developments to address user requirements

## Roadmap for Y2/1

- Implementation of the EMI-ES in servers and clients
  - Tight interaction with the client harmonization task in the latter case
  - Implementation started: to be provided for EMI-2
  - Implementations should reveal shortcomings
  - Supervise how each implementation deals with security
- Accounting Usage record
  - Agreement on compute accounting record and its support in computing elements
    - Should be provided for EMI-2 as well, quite important

## Roadmap for Y2/2

- Client harmonization
  - Common client APIs for EMI-ES
    - JAVA APIs
      - Starting point is UNICORE
    - C APIs (whence python bindings could be provided)
      - ARC and gLite (in ways to be defined yet)
      - Btw, gsoap licensing issue needs to be sorted out
  - EMI-ES CLI
    - Two options
      - Specific EMI-ES CLI
      - Integration of EMI-ES in existing clients
    - Anyway the PTB will be called for a final decision, soon
  - Align existing clients (e.g. options, arguments, output format, ...) ?
    - There is a specific objective on improving usability
      - Shift responsibility to users as much as we can :-)
        - » Ask them to find the top five most annoying errors, ask them how they would like a common output etc.



## Roadmap for Y2/3

- Support for batch systems
  - EMI mgmnt to provide this list ASAP
  - By PM20 each PT should grant full support
- Common parallel framework
  - The exercise of transparent user experience in running MPI jobs already achieved
    - For now it uses Runtime Environments of ARC, Execution Environments of UNICORE and MPI-start in gLite
  - MPI-start proposed as common framework
    - It is a set of helper scripts that run on the worker node
    - It does not impose any change in the CEs
    - Heard about strong opinions against this proposal ...
    - ... but actually no strong reactions here in the compute area session (if existing approaches used in ARC and UNICORE are kept)
  - Creation of a joint TF expected

## Roadmap for Y2/4

- Extend parallel computing capabilities
  - Support FPGAs and GPUs
  - Even more speculative, could result in trivial developments (parameters fwd) for the CE middleware POV
- Full support in EMI-2 for
  - Clients: SL5/32bits, Ubuntu x86\_64
  - Full distribution: SL6, Debian x86\_64
  - Most PTs say this won't be an issue
    - Uhm, let's see how it goes...
- Do something with clouds
  - This is very vague and risky
  - Compute area approach is merely: we'll stick with whatever infrastructure will come up
  - No fresh ideas about grid/cloud interactions yet
- Provide monitoring probes

## Roadmap for Y2/5

- Increase performance
  - Catch-all objective to include developments that PTs particularly care
    - ARC
      - Better scalability of job control directory
    - gLite Compute PT (formerly known as gLite Job Management PT)
      - High-availability for CREAM
      - Bulk job submission for CREAM client
        - » Already available for server
      - DAGMANless DAGs for WMS
        - » Better performance
        - » Significant reduction of lines of code
    - UNICORE
      - Introduction of parametric jobs

## Conclusions

- Overall good degree of achievements of Y1 high-level objectives
  - Keep up with the good work
- The agenda for Y2 is full
- Some tricky/delicate issues
  - Harmonization of clients/APIs
    - Also, EMI PEB should produce some reliable statement about possible gSOAP licensing issues
  - Common framework for parallel jobs
- DJRA1.1.2 deliverable (Compute Area Work Plan and Status Report) finalized
  - <https://twiki.cern.ch/twiki/bin/view/EMI/DeliverableDJRA112>



**Thank you!**

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