



EUROPEAN MIDDLEWARE INITIATIVE

A-REX / ARC CE Status and Plans

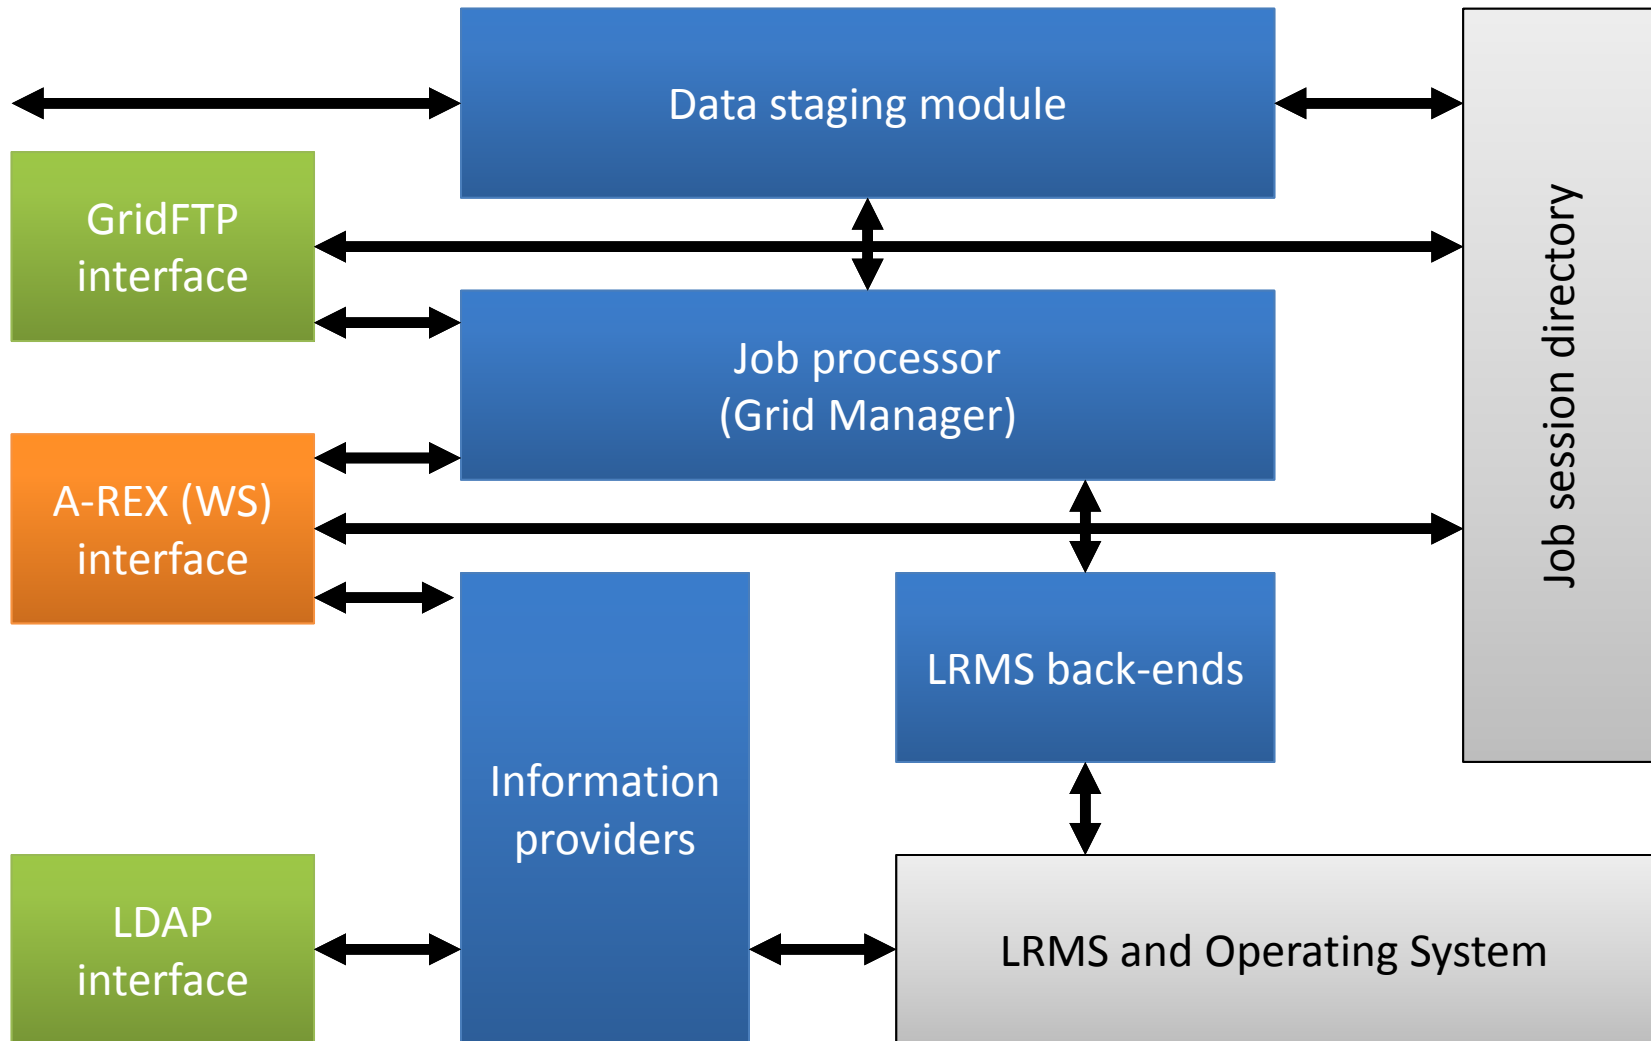
Oxana Smirnova

Lund University

EMI All-hands meeting

30 May 2011, Lund, Sweden

ARC computing service in a nutshell



What's new in EMI-1 - summary

- Completely re-factored job processor: accommodates for A-REX, among others
 - Is now in production
 - Old Grid Manager reached End-of-Life
- Load balancing (“ARC-CE HA”)
- Brand new data staging framework
 - Turned off by default in order not to affect production sites
 - Essentially, an internal file transfer service
- Support for Glue2 (LDAP)
 - XML rendering defined
- Some others (not in the scheme):
 - Janitor dependency on Redland RDF libraries removed
 - ARGUS plugin
 - crl-update replaced with fetch-crl

Interfaces still used in production

- GridFTP interface
 - Job submission and control
 - Access to session directory
 - Fine-grained authorization using legacy ARC groups
 - Can be transparently phased out
- LDAP interface
 - NorduGrid, Glue1.3 (and now GLUE2) information schemas
 - Information about service
 - Information about jobs
 - Coarse-grained information about authorization – VOs and gridmaps
 - Difficult to phase out as WLCG explicitly relies on it

New WS Interface

- A-REX
 - Job submission and control
 - Access to session directory
 - GLUE2 XML information schema (pre-standard)
 - Information about services
 - Information about jobs
 - Fine-grained authorization
 - gridmaps
 - ARC authorization policies
 - GACL for per-job access control
 - Incomplete XACML support
 - Legacy groups – work in progress

Objectives - 2011

- Definition of a common SAML profile all over the middleware stacks: **in progress**
- Publish coherent GLUE2.0 version information: **done in LDAP, to do in XML**
- Agreement on UR: **in progress**
- Nagios probes: **WLCG ones exist, few more in the pipeline**
- Adhere to OS standards for service operation and control: **done**
- Integration of the compute area services with the ARGUS authorization framework: **ARGUS plugin exists, XACML profile yet to be implemented**
- Interactive access to CE: **access to job session directory is in place**
- Common parallel execution framework: **shown to work via RunTime Environment**
- Implementation of the EMI-ES interface: **in progress**
- Interface to batch systems: **framework finalized, all relevant batch systems interfaced**

Objectives - 2012

- Implement common data library (relevant for ARC-CE): **todo**
- Implement UR: **todo**
- Support on SL6, Debian 6 and latest Ubuntu: **Debian 6, RH6, Ubuntu 11.04 are supported already**
- Introduce minimal DOS protection via configurable resource limits: **all conceivable limits are in place already**
- Replace GSI with TLS/SSL: **done in A-REX; common delegation - todo**
- Semi-automated configuration of service backends: **todo**
- Implement the EMI cloud strategy within compute area: **todo**
- Complete the rewrite of components utilising the new AuthN libraries: **todo**
- Increase performance of EMI services: **in progress**

Objectives - 2013

- Multi-core jobs on new architectures (FPGAs, GPGPUs etc): **todo**
- Investigate instrumentation tools: **todo**
- Implement extensions for new architectures: **todo**

Conclusion

- Good list of well-defined tasks
 - Some are not well defined yet, expect them to be improved
- Good progress
- Large picture is still blurry
 - Suppose a site wants to replace old LCG-CE with an ARC-CE – are we ready?
 - Where are we w.r.t. clients?
 - Where is Grid?