



Contribution ID: 53

Type: **Poster**

Job management in gLite

Monday, April 12, 2010 5:42 PM (3 minutes)

The gLite WMS has been designed and implemented to provide a dependable, robust and reliable service for efficient and transparent distribution and management of end-user requests to high-end resources shared across a production quality Grid.

The WMS comes with a fully-fledged set of added-value features that hide to end users the complexity of such a heterogeneous and ever growing infrastructure and enable, thanks to a flexible, service oriented and general architecture, applications coming from largely different domains

Detailed analysis

The WMS is responsible to translate users' requirements and preferences into concrete operations, interactions and decisions, in order to bring the execution of a request for computation, storage and the like (also known as 'job') to a successful completion. This is done transparently, while acting on behalf of the user.

Several types of jobs are supported: simple, intra-cluster MPI, interactive, collections, parametric and work-flows in the form of directed acyclic graphs.

The Grid is a complex system and errors can occur at various stages throughout the so called submission chain. The WMS has the ability to automatically recover from infrastructure failures by implementing resilient strategies which include resubmission and retry policies. Additional benefits concern sandbox management - with support for multiple transfer protocols, compression and remote access - data-driven match-making, automatic credential renewal, service discovery and optimisations for collections such as bulk-submission and matchmaking.

Job tracking information in terms of relevant events, milestones and overall status can be retrieved and used by the WMS via the so called Logging & Bookkeeping service.

Conclusions and Future Work

After all these years operating in the EGEE infrastructure, the latest WMS releases have reached unprecedented stability and a performance which can smoothly accommodate for the current needs. By the end of EGEE-III, the WMS will have extended its support to more architectures and platforms.

Nevertheless, a new and challenging era is coming which will require the whole gLite stack to deal with other middleware distributions and an expanded use base. Consequently, the WMS will have to be deeply involved in managing different computing paradigms, standards, services and emerging technologies.

Impact

Managing a grid job, from submission to completion, typically involves coordinating and interacting with a number of different services: computing elements, storage elements, information systems, data catalogues, authorization, policy and accounting frameworks, credential renewal. In this respect, the WMS, especially by virtue of his central, mediating role, has to deal with a wide variety of people, services, protocols and interfaces. Interoperability with other Grids must also be taken into account in this scenario.

On the user's side, the WMS exposes a Web Service based interface in accordance to the WS-I profile, which defines a set of Web Services specifications to promote interoperability. Access to the WMS is also granted

by a dedicated User Interface and APIs which are available in C/C++, Java and Python bindings.
Furthermore, the WMS fully endorses the Job Submission Description Language, an emerging standard which aims at facilitating interoperability in heterogeneous environments, through the use of an XML based job description language that is free of platform and language bindings.
On the resource's side, both legacy and OGSA/BES based interfaces are supported.

Keywords

Job Submission and Management, Resource brokering, Interoperability, Grid Computing, Metascheduling

URL for further information

<http://web.infn.it/gLiteWMS>

Author: Dr CECCHI, Marco (INFN)

Presenter: Dr CECCHI, Marco (INFN)

Session Classification: Poster session

Track Classification: Scientific results obtained using distributed computing technologies