

Advantages of adopting late-binding techniques through standardised interfaces for workflow managers.

Monday, 16 September 2013 09:00 (8h 30m)

Description of Work

The GWpilot features make it suitable to accomplish the requirements of workflow applications, based on provenance and precedence constraints, but not aware about pilot jobs. GWpilot is a new developed pilot-job framework based on GridWay that offers common possibilities already implemented in other pilot systems to overcome remote queues, to correctly fit tasks to pilots or to discard bad resources. It also discards a passive role in order to effectively coordinate the pilot-task matchmaking with advanced scheduling techniques such as pre-allocation, reservation, data-location awareness, etc. The number of pilot jobs and the hosting resources are progressively adjusted by GWpilot regarding the requirement of tasks queued and without intervention from the user. However stacking a workflow manager as Kepler over GWpilot is not a straightforward work. Although the installation of GWpilot is easy, users do not like to fight with middleware releases. Additionally simulations could take too hours to maintain active a Kepler instance on a server. Thus, local connectors must be discarded. It is then necessary to deploy a service on GWpilot that manages remote task submissions and to implement Kepler actors that can interface with it. The protocol selected must be accepted by OGF and EGI in order to unify the access criteria and promote the future extensibility of the system. For these reasons the new OGSA-BES interface for GridWay has been selected and new Kepler actors have been implemented to bind the GridWay BES interface for submitting jobs, checking their statuses and storing IDs in a simple way, as well as other necessary tools to simplify the creation of JSDL files or the automatic delegation and renovation of user credentials.

Moreover, Kepler will not schedule to GWpilot bag of tasks because they will be independently submitted, allowing to benefit from the resource appropriation made by pilots and to improve the performance. In addition, the scheduling of data pre-allocation on remote storage elements is not a direct responsibility of the pilot system, and although some task assigned to pilots could do this work, it is Kepler who submits these tasks. Additionally many Kepler users could use the framework at the same time. These circumstances imply to assure the reliability of OGSA-BES interface and GWpilot to levels never previously tested. Related with this, the OGSA-BES service also was modified to not perform additional procedures that reduce its performance, as staging input and input files for example.

Relevant URL (if any)

<http://kepler-project.org/> ; <http://www.gridway.org/> ;
<http://www.ciemat.es/portal.do?IDR=343&TR=C>

Printable Summary

Visual workflow systems are fundamental tools to implement and maintain complex scientific simulations, especially when researchers from different areas of knowledge are involved in their development. Diverse workflow managers allow using HPC platforms, databases or repositories in a single framework, which is highly valuable when we deal with complex large scaled problems. Additionally they can allow the access to grid resources, but directly interfacing with some specific middleware, as the Serpens suite for Kepler does. However this approach means to be tied to middleware implementations and, also, to support the performance slowdown due to waiting queues and failing jobs derived from unsuitable or untrusted resources.

A way for improving the grid execution performance is the use of pilot-job techniques, which somehow reserve the available resources by means of submitting regular grid jobs (the pilots) that lately allocate the application tasks for doing the real calculation. Huge collaborations, like the LHC ones, usually rely on ad-hoc pilot systems that fit their specific necessities, but a researcher working on his own, or out of these collaborations, does not count on a general purpose application. Few frameworks really offer the pilot-job

advantages to conventional users (DIANE or DIRAC for example). Nevertheless, all of them lack some features such as easy-installing, user-sharing or the absence of standardised interfaces to remotely access these systems. Other systems as Pegasus or P-Grade make use of Glide-in's to face this problem, with the drawback of the local dependence on Condor or the control loss of the pilot provisioning guidance. These issues limit their deployment for projects that establish the workflow managers as the support of their research.

A new approach that overcomes those limitations is presented in this work. It represents a unified, simplified, efficient, flexible, standardised and portable new mechanism to remotely manage pilot jobs and the tasks that compound a Kepler workflow. Kepler takes advantage of the GridWay modularity to build a new multi-level scheduling framework based on the performance and reliability provided with GWpilot; the portability and interoperability achieved by the OGSA-BES actor and interface. At the same time, the compatibility with legacy middleware remains by means of Kepler and GridWay drivers.

Thus, this work, already available to users, simplifies the creation, maintenance and execution of workflows since: only requires a scarce knowledge of the Kepler actors; offers an underneath management of user tasks and pilot jobs; and, allows a portable simulation that can be checkpointed and carried out on laptops thanks to the remote job delegation to GridWay and the job accounting stored in the Kepler local database. Moreover, task requirements will automatically be translated to pilot provisioning, allowing the creation of complex workflows. Additionally, the backward compatibility will facilitate the transition to the standardised OGSA infrastructure model that EGI promotes.

Primary authors: RUBIO-MONTERO, Antonio Juan (CIEMAT); PLOCIENNIK, Marcin (ICBP); ISMAEL, Marín Carrión (UCM); ZOK, Tomasz (ICBP)

Co-authors: PALAK, Bartek (ICBP); Dr HUEDO, Eduardo (Universidad Complutense de Madrid); CASTEJON, Francisco (CIEMAT); RODRIGUEZ-PASCUAL, Manuel Aurelio (CIEMAT); OWSIAK, Michal (ICBP); MAYO-GARCIA, Rafael (CIEMAT)

Presenters: RUBIO-MONTERO, Antonio Juan (CIEMAT); RODRIGUEZ-PASCUAL, Manuel Aurelio (CIEMAT); MAYO-GARCIA, Rafael (CIEMAT)

Session Classification: Posters display