



Contribution ID: 125

Type: Oral

Ganga/DIANE and SAGA for solving complex scientific problems

Thursday, April 15, 2010 9:00 AM (20 minutes)

An increasingly important requirement for solving complex scientific problems is the simultaneous use of qualitatively different resources, from simple batch nodes to supercomputers. We present a project where Ganga and DIANE user-level middleware tools are used in connection with SAGA, a high-level Grid API, to achieve an improved quality of service when using resources across multiple computing infrastructures, such as EGEE, NDGF, OSG and TeraGrid. We discuss these advances in the context of LQCD simulations in theoretical physics and other applications in the current portfolio.

Detailed analysis

This project make a complementary use of three software components, Ganga, DIANE and SAGA to support complex scientific applications. We extend the WMS/DIANE pilot-job approach to provide a uniform policy-based pilot-job submission mechanism across these multiple Grids, and thus provide production level interoperability for a range of applications.

Ganga is a computational task-management tool, which allows for the specification, submission, bookkeeping and post-processing of computational tasks on a wide set of distributed resources. Ganga is primarily used by Atlas and LHCb experiments at LHC at CERN, and also by a number of smaller user-communities in the context of the EGEE project.

DIANE is a late-binding scheduler and optimizer of the task processing, taking advantage of the pilot-job paradigm. DIANE has been used successfully on a large-scale to provide an improved Quality of Service experience for end-users and to improve the reliability and efficiency of task processing.

Conclusions and Future Work

The systems which we describe will become increasingly important for ROSCOE and INSPIRE applications.

Impact

SAGA is a high-level API that is currently an OGF technical specification; it is targeted to become a standard in 2010. Several production grids worldwide currently support SAGA or are deploying it currently (including TeraGrid and UK's NGS). SAGA provides the ability to develop applications and tools that can work on and interoperate across different distributed systems.

SAGA is currently being deployed on several production grid infrastructure – including the US-TeraGrid, UK-NGS, Japanese-NAREGI, and thus applications can seamlessly be extended to interoperate across these different infrastructure.

This project will ultimately take advantage of concurrent effort in developing SAGA as a uniform access layer to other European middleware such as gLite, ARC and UNICORE.

Keywords

application support tools, pilot jobs, SAGA, Ganga, DIANE

URL for further information

cern.ch/diane

Authors: Mr MOSCICKI, Jakub (CERN); LAMANNA, Massimo (CERN); Mr WEIDNER, Ole (LSU); Dr JHA, Shantenu (LSU)

Presenters: LAMANNA, Massimo (CERN); Dr JHA, Shantenu (LSU)

Session Classification: DCI programming

Track Classification: Programming environments