

Computational steering on the Grid

OGF20 7 May 2007

Panel Discussion

Notes: Martin Turner

Write-up: Helen Wright, Martin Turner

Work by the SAGA development group in asking for use cases has identified a steering demand but only in a proportion of submissions received. There is possibly a chicken-and-egg problem, in that more users and potential users would create more demand, but that left the question of how to get the initial process started. A show of hands revealed most of the audience considered themselves developers rather than users.

There is a problem with the current infrastructure for providing large-scale scientific code support and issues were raised as to how this would be funded. A recent initiative by EPSRC to make HPC applications from existing codes was a welcome boost though naturally this came with limitations due to finite funding. There was general agreement that (in the UK at least) research funding pressures militated against developments in computational science: should users wishing to pursue science have to become application writers; how is this burden on them viewed by outside evaluators; and, how should it be resourced?

From a computing centre point of view there is a lot of diversity in the requests users make and the time might be right to try standardising some aspects, such as which firewall port numbers to connect via.

There was definite need within APIs for standardised hooks to middleware, to avoid having to re-invent the wheel for different computing contexts. Analysis of detailed use cases could clarify this need and there was potential for these use cases then to become demonstrators of any API that was developed. Different versions of SAGA code (termed adapters) might be needed to fulfil this task. The question remained as to whether standardisation at a number of different levels was required, since different types of user might need differing degrees of complexity in an API. Another question that would have to be resolved was how computational steering would integrate with existing structures such as gLite/Globus and, again, different SAGA adapters could resolve this problem. It was recognised that helping early adopters of an API would accelerate the growth of the user base. The successful example of the MPI API was mentioned, in respect of understanding why, and at what point, its adoption took off.