

Nov 01, 06 21:44

response-refs-naw.txt

Page 1/4

Comments from NAW 30 Oct 2006

=====

I've put the refs comments between [...]

ref PP/E003060/1

main critique is one of defining specific s/w development detail.

[1. In section 2.7, it says that the system has been scoped to support the professional astronomy community in the UK, but the case gives no idea of how this has been done, and goes on in the same paragraph to say "we expect that the computing power to support this will eventually be provided via departmental servers". I think that their strategy in this regard needs clarifying.]

The process has been one of requirements capture and input from the community via the Science Advisory Group, interaction with and assessment of future UK mission needs.

Eventual usage of departmental servers is probably essential when a large fraction of data analysis is actioned via the AstroGrid VO system. The AG s/w infrastructure is constructed to allow for inclusion of per institute resources, which they can configure to support their own communities. This will ensure that running AstroGrid means having to resource a whole new layer of CPU/ Storage. Rather the AstroGrid s/w infrastructure interfaces to compute h/w in place (either NGS, Campusgrids or institutional compute).

Politically it is already the case that departments are finding access to funds for new compute easier to come by if they make the case for shared compute, Witness campus grids being funded via SRIF funds (e.g. Cambridge and elsewhere) where departmental groups put in joint bids to provide shared compute resources. Thus the Astrogrid model fits well into this and is technically able to support such a model.

[2. Related to this, is the issue of making use of other grid capabilities. NGS is mentioned in section 5.13, but again no very clear picture emerges of their plans to allow astronomers use NGS, or of how this fits into their planned deployment. No mention at all is made of GridPP. Is this entirely irrelevant to their plans? Whilst one can only guesstimate the rate at which usage of the VO will ramp up, I would like to see some plan for what they are assuming, how the CPU power available will ramp up to match this estimate, and where this power will come from.]

Work with NGS is progress to allow authenticated access to applications configured on the NGS compute facility. Pete Oliver from NGS is a member of the AGDAG.

An example piece of work in this direction is the VOTECH test running AstroNeural (check with John Taylor) on EPCC compute, with the job actioned via the AG workbench.

Also, with the upcoming move to Taverna the ease of submission to NGS will improve as Taverna is OMII-UK supported infrastructure, with a commitment to be NGS compatible.

In terms of GridPP, they are moving to make their CPU accessible via the NGS (via the EGEE work), thus once AG can submit to NGS it will be able to submit to GridPP compute. .

re the CPU question - I think the answer to this is that the provision of CPU is demand led. In other words, we provide a base CPU resource through our AstroGrid to support initial user workflows

Nov 01, 06 21:44

response-refs-naw.txt

Page 2/4

and applications. As time progresses we integrate CPUs via those of the Data Centres, and then mainstream institute compute. I think we don't need a specific plot of CPU usage here as we are not actually bidding for large amounts of extra CPU in AG3. However, we must be able to give them an explanation of how institutes will be able to plug in CPUs - in a controlled manner, and how they can regulate who gets what use of what CPU (so the story will revolve around our community model). Keith can put together a suitable explanation.

[3. On the software applications side, I found a similar lack of detail in how the project intends to progress, though I am pleased to see that AstroGrid has inherited some of the most talented of the applications programmers who used to work for Starlink.]

[(a) Much prominence is given to the role of "Science Calls" with regard to applications development, but it is unclear to me whether these calls are intended to drive the development of new applications as such, or just of pipelines which roll together existing software.]

I think here we have to point out that AG3 is not really about the creation of new applications as such. Rather AG3 will provide the APIs to which new applications can be developed (hitting the 3 E's - economic, through cheaper code development based on access to powerful VO APIs; efficient, through the production of higher capability code for the same amount of resource utilising the fact that the VO provides lots of extra capability 'for free' e.g. see the tools utilising astro-runtime now; and effective, new tools will be hitting the data that the astronomers want).

In terms of the science calls - that is more about helping aid the prioritisation of access to legacy apps - which data to provide access to first and so on. This recognises that we'll be resource limited in terms of applications.

[(b) Whilst Science Calls are a fine idea, they do not amount to a strategy for science tools development. There are hints in the proposal that there may be more to their approach than just these calls, but we are not really told what. I think the panel should ask.]

Indeed - a perceptive comment. As noted above AG3 doesn't fund tools development - except in a very limited manner. For wider tools development AG could work with teams who propose the development of new apps and partner them in bids to PPARC. Or PPARC could organise apps calls - and stipulate that apps interface to AstroGrid. There are a number of models. I think we should be more rather than less in this, as we actually have a good vehicle for delivery. I have seen many projects which are considerably less effective.

[c. Portsmouth are specifically mentioned in connection with data mining expertise, but it is entirely unclear to me just what their role is within the project.]

One for Andy and Bob Nichol

[4. Some other things I would have liked to hear about are their plans for the Portal, which is likely to be crucial in encouraging new users to take to the system, their plans for documentation and online help (very thin at present), and their approach to robustness, given the importance of software being fail-safe for large batch jobs.]

Our portal to the VO is via the workbench. The legacy browser based

Nov 01, 06 21:44

response-refs-naw.txt

Page 3/4

portal has been discontinued as being inappropriate for the VO.

The website is undergoing a significant overhaul with a focus on increased usability in support of the science end user. We should point to www2.astrogrid.org. Likewise, significant effort is now being put into improving the on-line help materials i think <http://www2.astrogrid.org/science/science-examples-stars/iphas-catalogue-images/> is a good example.

Robustness - covered in the PPRP response

[5. Finally, I do not really understand the distinction they try to draw (in sect. 6.3) between Working Allowance, and Contingency. It seems to me that the two are one and the same, and they simply developed an interesting and useful approach to estimating a reasonable sum of money.]

We follow the definitions outlined by PPARC see the PPRP handbook

section 4.4.2 Working Allowance - held within the project and justified in Appendix D of the AG3 proposal. The working allowance is included within the project bid resource estimate:

'This is used to cope with the known uncertainties in the project. It can be calculated in a number of standard ways (e.g. the formulaic risk assignment method used by DoE). The Working Allowance should be estimated by the project team and scrutinized by the Projects Peer Review Panel (PPRP).'

section 4.4.3 Project Contingency - held centrally by PPARC, not included in the project costings:

'This is used to cope with the things that go wrong that cannot reasonably be predicted at the beginning of the project. The intention is that there should be reasonable certainty (better than 95%) that the project can be completed within the total estimate including contingency.'

Referee number 37:

[My main concern is the confusion in the document between the provision of national infrastructure and the continuation of what was essentially a research project. How much of the proposed middleware is in a sufficiently mature state to be an effective research tool. The maturity of the middleware will determine whether there is a universal takeup of VO tools.]

For this we need to point to the PPRP response to question 3 - where are we at the end of AG2. Basically we will have a mature set of s/w fit for deployment with the caveats that we are a year or so behind original estimates - thus the slightly enhanced resource bid for the first year (coupled with the deployment pulse in the first two years).

Thus AG3 is very much about provision of a service, and that all VO engineering is in support of this service - supporting the service in a sustainable manner as the internet and the VO evolve.

[The proposal under-estimates the requirement for training and help-desk support particularly in the early years of VO take up. Shift resources from A3 to A1]

I could see that there might be opportunities to offer enhanced user

Nov 01, 06 21:44

response-refs-naw.txt

Page 4/4

support at least initially. Our management model allows for re-deployment of effort if required.

I do not think that in the AG3 period we'll be in the position to support AG3 VO engineering or technical research via FP7 funding. However, we will aim to explore this route (as we have responded to opportunities for supplementary funding in AG2 with FP6 programmes, e.g. VOTECH Euro-VO DCA).