Report on Engage Survey

TR-10-05
February 2010

RENCI Technical Report Series
http://www.renci.org/techreports

Report on Engage Survey

James Howison and Jim Herbsleb
Carnegie Mellon Computer Science
Scientific Software Ecosystems Research Project

February 2010

Contents

Chapter 1

Introduction

The VOSS SciSoft research project at CMU*, conducted a survey of the Engage VOs contacts. The survey was prepared by James Howison, with input from Jim Herbsleb and John McGee, from Engage. The individually identifiable results are confidential to CMU (this was done to ensure participants were comfortable speaking honestly on the survey) and this report thus avoids identifying individual responses. The report presents an overall summary of the respondents answers to the questions, using quotes where they are not personally identifiable.

1.1 Survey Logistics

The survey was conducted using the SurveyMonkey tool, and invitation emails were sent out by the internal SurveyMonkey system. The first contact was on 12 Dec 2009 and was a "heads up" sent from John McGee's email address, notifying potential respondents of the survey. This was closely followed with the initial invitation. A further invitation email was sent Jan 4, and another Jan 13. On Jan 27 personalized emails were addressed to individual non-responders.

1.2 Sampling and Response Rate

The sampling frame was a list provided by Engage, containing Names, Email addresses and basic information about potential respondents organizations. This was produced from Engage VO's CRM software.

There were 114 contacts in the original sampling frame. There were 113 participants successfully imported into SurveyMonkey and who received the first message. There were 4 bounced email addresses and a quick scan did not reveal additional contact

*This work was supported by NSF grant #0943168. The opinions expressed within are those of the authors and not necessarily those of the NSF

details. There were 6 active opt outs. In total there were 44 responses and 68 from whom nothing was heard. Of those 44 responses 9 indicated that they had not had contact with Engage (answered "No" to the initial eligibility question), leaving 34 responses to analyze. This gives an eligible sampling pool of $(113 - 4 - 9 = 100)$ with 6 opt-outs (94); of the 68 non-respondents we do not know their eligbility, but assuming a similar rate there were 14 ineligible $(68 \times (9/44) \approx 14))$, for a total eligible frame of 80 $(94 - 14)$ from which 34 responses were obtained, a response-rate of 42.5% (34/80). Table 1.1 shows when those responses were received.

Table 1.1: Response Timetable

Date	Message Type	Responses
21 Dec 2009	Broadcast	11
4 Jan 2010	Broadcast	4
13 Jan 2010	Broadcast	5
27 Jan 2010	Personalized	14
Overall		34

1.3 Analysis

Results were downloaded from SurveyMonkey into an Excel file. Some manual editing of responses was required for analysis, for example where respondents indicated in the comment field that "Neutral" responses should count as "Not Applicable". Graphics and statistics were produced with *ggplot2*, an *R* package (Wickham, 2009; R Development Core Team, 2004).

Chapter 2

Results

2.1 Respondent Demographics and Experience

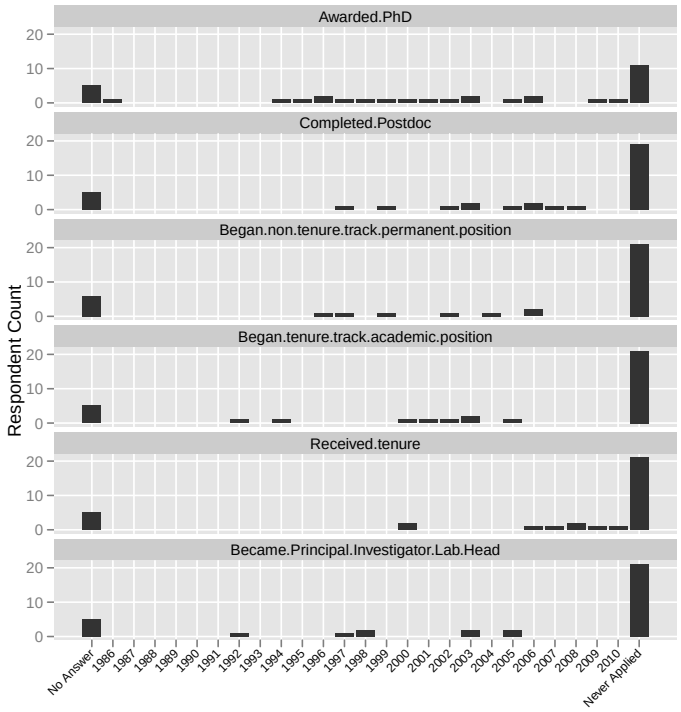

Figure 2.1: Demographics. Year at which condition first applied to respondent.

Figure 2.1 shows a summary of respondent by career stages. Clearly there is a good spread, with some bias towards current PhD students. There were 8 people who identified themselves as lab heads.

Computational Infrastructure Experience

This question shows the experience range of the answer group, see Figure 2.2. The majority of the answers are clustered towards the right (ie most recently), with increasing recency down the page. The respondents split into two groups in terms of performing scientifically useful computations, those that began in the 80s or 90s and an equivalent number that began more recently. Only a few have never performed scientifically useful calculations. Understandably local clusters have been in use longer, beginning in the late 90s and most participants have used these. Grid computation, of course, has only been used since 2001, with most participants only using it in the past 2 years (lining up with Engage's period of existence).

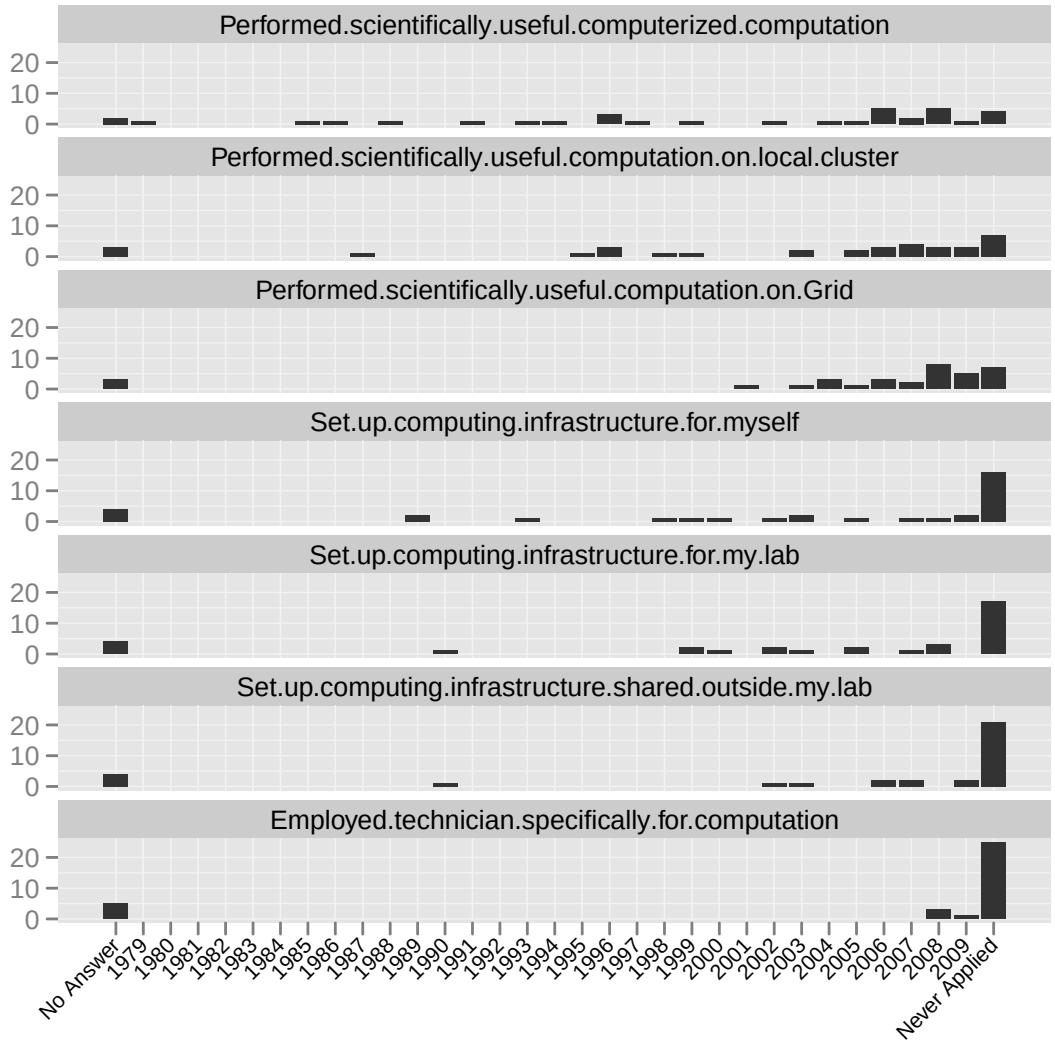

Figure 2.2: Experience. Year at which condition first applied to respondent.

Experience with Engage

The initial question of the survey asked for a summary of experience with Engage. Table 2.1 summarizes these results, allowing individual respondents to be counted in multiple categories, but not all respondents provided an answer in this space.

Count	Summary
4	Attended seminar and learned of OSG
4	Only Discussed possibilities
18	Help to run jobs
8	Received help adapting software
2	Received help in setting up OSG infrastructure

Table 2.1: Summary of experience with Engage responses

A subsequent question asked this in more detail, seeking the most advanced stage of interaction with Engage. Figure 2.3 shows these results.

The majority of participants had run jobs, with many used in scientific publications. Text answers included summary of jobs and time run (e.g. "a 10,000 iteration job which took about 10hrs on a 4 core compute node, was successfully made to run on a full scale of 100,000 iterations in 4 hours on Engage resource pool") but were most counts of scientific papers result from the work (e.g "In the last 2 years I've obtained computational results that have spawned 4 papers published in scientific journals")

Notably the final graph shows that the vast majority of respondents have not employed a technician for computation. This may be affected by the question, which should probably have asked about their lab for this question, since the way it was asked may simply reflect the age/experience of the participants

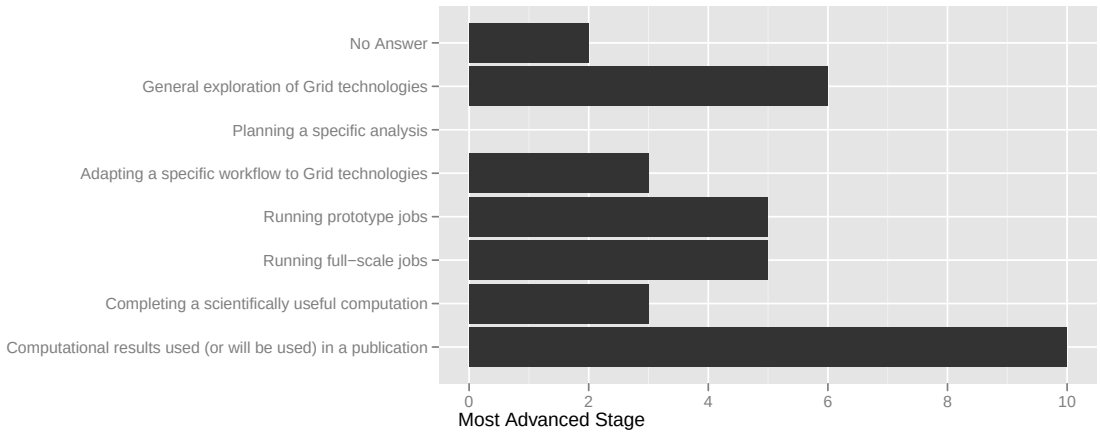

Figure 2.3: Most advanced stage of Grid/Engage use

2.2 Resources and Bottlenecks

Resources used in combination or as an alternative to OSG Grid

This question asked respondents to indicate whether, in their computational scientific work they used a set of resources (from laptops to cloud computing, such as Amazon EC2). The possible answers were "Neither", "Combination", "Alternative" or "Both" (indicating that they used these in different ways in different projects). The results are shown in Figure 2.4.

The results show very little use of Cloud computing (just 4 respondents), that non-OSG Grid resources are used, both in combination and as an alternative (15 respondents). Unsurprisingly Local Clusters are broadly used (both in combination and as an alternative and in 5 cases Both). Personal Computing resources (Servers, Desktops and Laptops) are still heavily relied on. Given their substantially smaller processing capacities this might indicate the persistence of small scale analyses, or that researchers are willing to wait to use these more controllable and easier to program for environments.

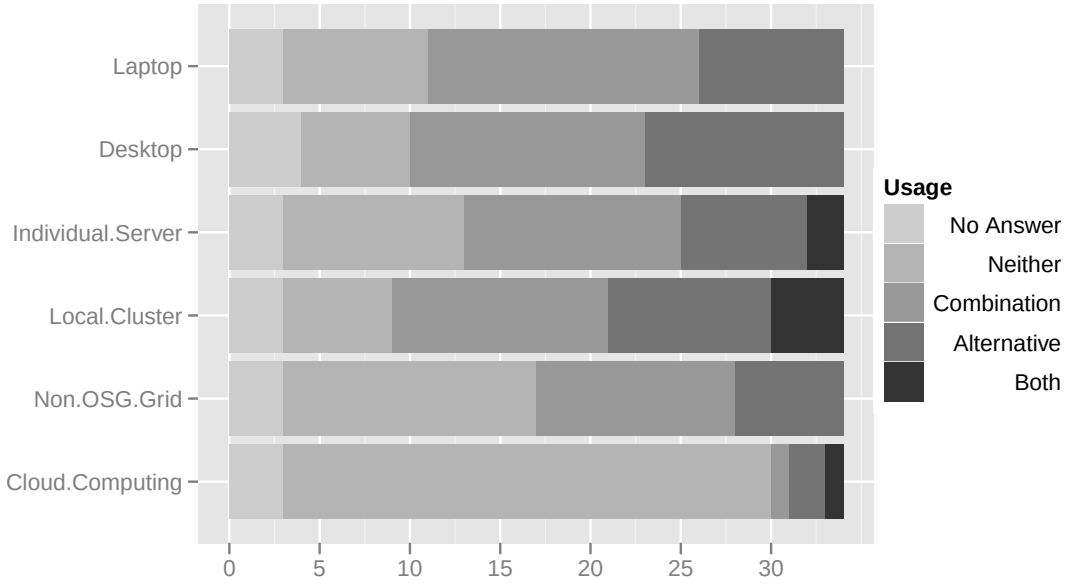

Figure 2.4: Technologies used in combination or as alternative to Grid computing

Research Bottlenecks

We asked two questions regarding the bottlenecks which impact people's ability to conduct research. We asked about a number of factors, from cycles to bandwidth to ability to communicate scientific needs to technical people. The possible answers were No Impact, Slows, Severely Slows and Blocks my research. The results are shown in Figure 2.5, the high number of "No Impact" (and "No Answer") responses indicates that generally people are able to conduct their research unimpeded. The factors indicating the most impact (reading the darker colors from the right) were: Adapting software, Post-processing, Cycles and bandwidth.

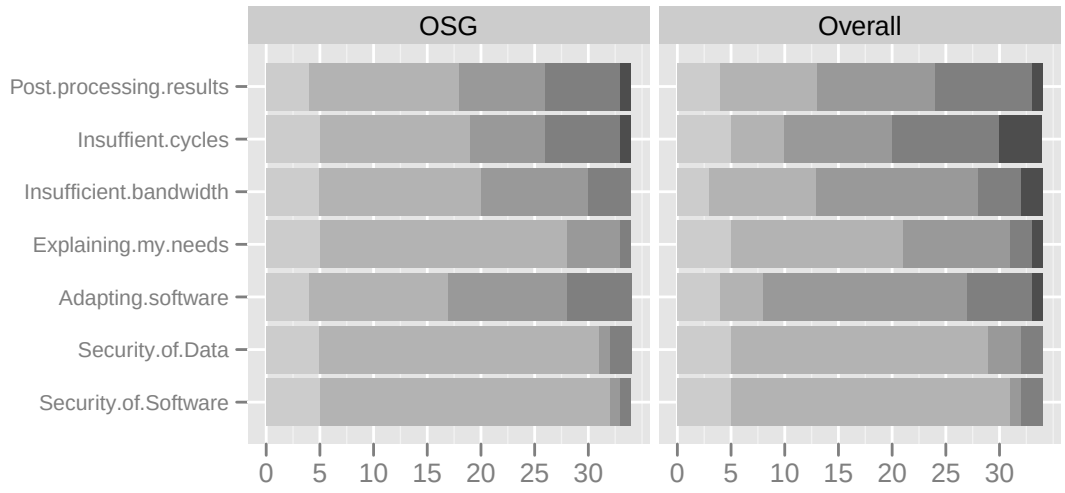

Figure 2.5: Counts of Research bottlenecks (see weighted figure for legend)

Figure 2.6 zooms in on those impediments that did exist. The scale has been weighted to emphasize the ordering of slows, severely slowing and blocking. The result of this shows that Post-processing and cycles had the highest overall impact (since they actually blocked research). Nonetheless adapting software and bandwidth continue to show impact.

Notably concerns about security and privacy of data and software were not seen to have substantial impacts on the ability to conduct research (although later one respondent did indicate ability to ensure that one is complying with HIPPA regulations as important.).

The same question was asked about their research in general (grid or non-grid); shown on the right-hand side of Figure 2.5 and 2.6. Overall respondents indicated that these factors had a larger impact on their research overall than for OSG work. The ordering is broadly comparable, with security and privacy of data and software having little

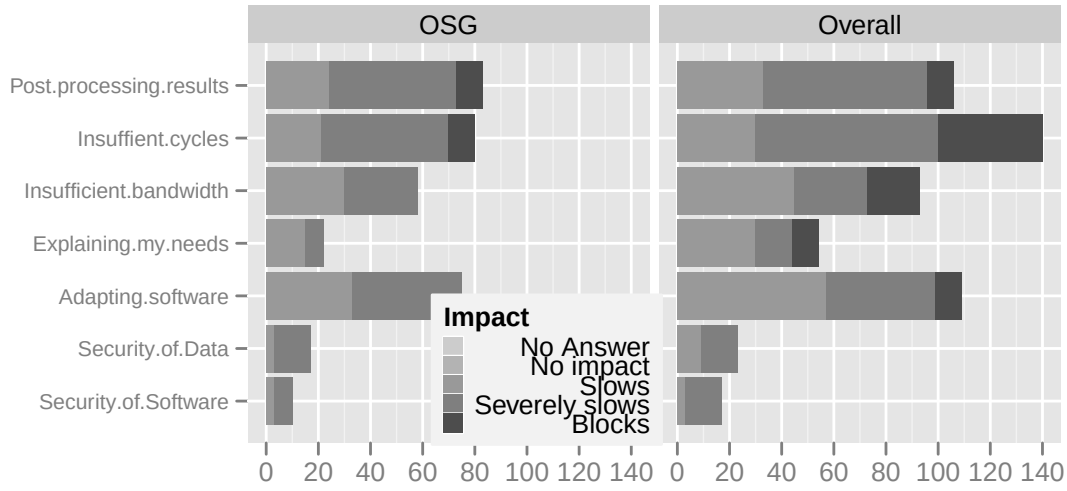

Figure 2.6: Weighted Research Bottlenecks (Block=10, Severe=7, Slows=4, No Impact or No Answer=0)

impact. Notably adapting software had a slightly higher overall impact; while this is an issue for OSG it is a strong issue overall. In the weighted comparison it is clear that cycles are a strong issue overall, blocking and severely slowing research, but less so with OSG (which makes sense since that is the focus of OSG's infrastructural effort). Overall explaining needs to technical people had a greater impact than it did for OSG, which speaks to the quality of Engage's assistance in this population.

Chapter 3

Satisfaction with Engage

3.1 Overall Satisfaction

Overall the survey respondents indicated great satisfaction with Engage across all the aspects of their work which we surveyed. A clear caveat in this type of survey, however, is that there is a risk that those who were dissatisfied with Engage would simply have refused to take the survey. On the other hand, they may have seen this as an opportunity to vent, which did not happen.

Question	n	Min	Max	$\bar{x}$	$\widetilde{x}$	s
Responsiveness	29	0	5	3.7	5	1.8
Professionalness	29	0	5	3.8	5	1.8
Knowledge.of.grid.technologies	29	0	5	4.0	5	1.6
Ability.to.understand.my.scientific.questions	29	0	5	3.2	3	1.8
Ability.to.understand.my.computational.needs	28	0	5	3.5	5	1.9
Technical.skills	29	0	5	4.1	5	1.6
Ability.to.help.solve.my.problem	29	0	5	3.7	5	1.9

Table 3.1: Summary statistics for Engage Satisfaction Questions

Figure 3.1 and Table 3.1 show the distribution of answers. There were no Dissatisfied or Extremely Dissatisfied answers. Using a zero-centered weighting scheme of -5,-3,0,3,5 the overall average response was 3.708 with a median of 5, showing the skew towards extreme satisfaction. All questions had a median of 5, other than "Ability to understand my scientific questions," which also had the lowest mean.

The open-ended responses provides additional information.

One respondent identified "Ability to help solve my problems" as his lowest rating, commenting, "Engage has provided canned solutions for problems, often in a way that does not help me understand what they have done or how to expand/extend their

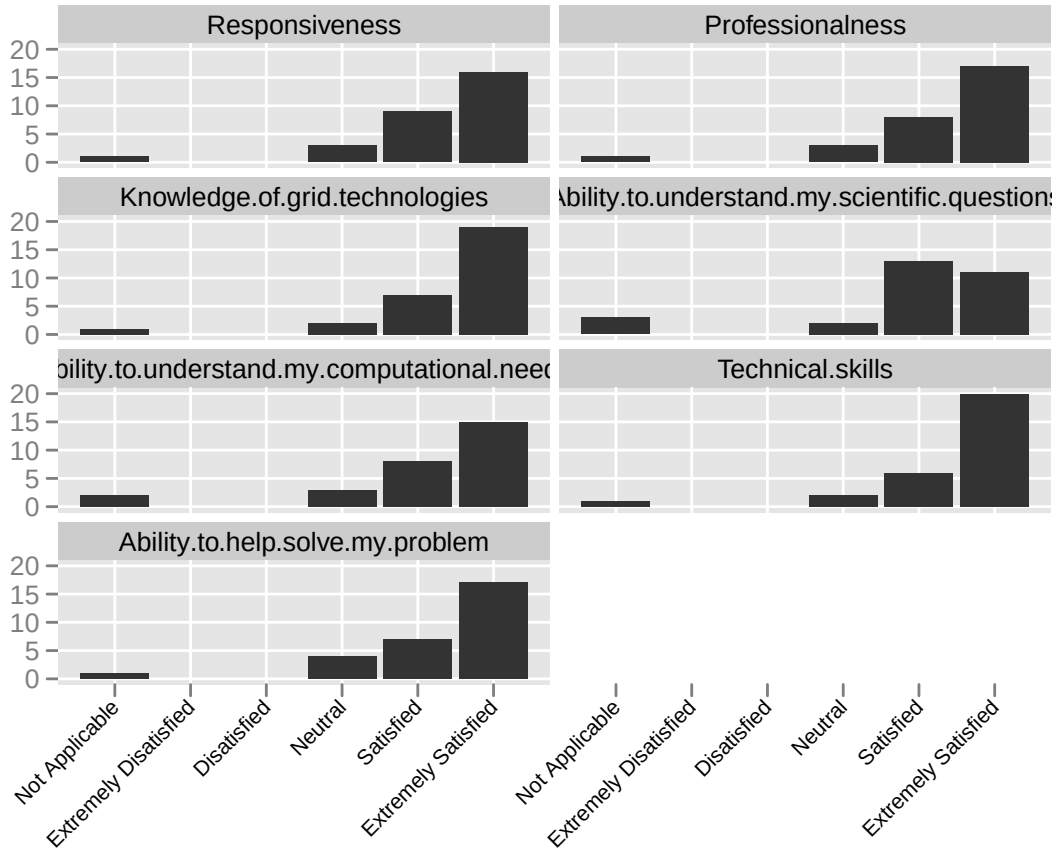

Figure 3.1: Satisfaction with Engage

efforts to the next stage of my research or computational scale." (This respondent had only run prototype jobs). This should be balanced against multiple comments along the lines of "Staff/RENCI group were very helpful" (sometimes citing individual staff members). One answer stated that they actively recommended Engage to others.

The only other mildly negative comment highlighted some coordination delay resulting from RENCI staff being on furloughs.

3.2 Most important thing that should be changed?

This question was "Please use this space to describe the most important thing you would like to see changed, based on your experience with Engage?" There was a high response rate to this question, with answers, falling into five clusters, ordered by how the number of responses mentioning the issue.

More resources on Grid

The first cluster spoke of a desire for more resources and capabilities on the Grid. These included personal space allocations (quota increases), faster speeds and the ability to submit OSG jobs from their workstation, avoiding "Engage VO job submission host". Two respondents identified a desire to use "large-scale" MPI applications. One respondent requested a mechanism for dealing with private/HIPPA sensitive data (this respondent was one of the few to identify privacy of data as a research blocker, saying that it "Severely Slows" their research).

Improve ease of certificates

The second cluster spoke of a need to improve the ease of signup and renewal. Four respondents identified this, saying that current certificate processes are "a bit complicated" and "confusing" and they need to be "simplified".

More system state information

The third cluster spoke of a desire to see more system state information. This was identified by two respondents. One asked to "bring back OSG_APP and OSG_DATA and OSG_TMP information" and to add other environment variables from the OSG stack, such as GLOBUS_LOCATION. The other asked generally for a more informative system load tool.

More visibility of Engage's internal processes

The fourth cluster sought more visibility of Engage's internal processes. One respondent remarked that "the server gets rebooted for various reasons" and asked for more announcements of such activities. The second asked for "the software management practices of engage" to "become more available for local use". This is not entirely clear to us but may refer to wanting more visibility on status of software adaptations.

Engage staff and researcher relationship

The fifth cluster addressed issues of the relationship between Engage technical staff and researchers. One respondent asked for efforts to make interfacing with the Grid easier for users, so that less assistance was needed. The other respondent in this cluster identified a need for more knowledge transfer. They described contacting a successful Engage assisted user for more information on working with the Grid and found that user to be unable to give much explanation, referring to the Engage staff. They wrote that in this case the Engage staff seemed like a "research assistant in managing the computations."

3.3 Aspects of Engage found most useful

The final question asked was "What was the aspect of Engage that you found most useful?" Again the response broke into clusters: People, CPU-time and Learning.

People

Respondents consistently complimented the Engage people (9 responses), citing their helpfulness, their "knowledge of the tricks and quirks of using OSG," their ability to script interaction with OSG resources. The model of having single point of contact was mentioned as well, including two citations of particular Engage staff members.

CPU-time

Access to CPU-time was mentioned in 13 responses. These included simple and straightforward comments like "computing power" as well as more detailed comments. One respondent cited the fact that Engage had access to more OSG resources than other VOs they were associated with. Another mentioned that OSG time was excellent for when their local cluster was busy. Finally one respondent mentioned that the access to cycles was very "black box" and "pleasant to deal with."

Learning

The third cluster of responses (5) focused on learning from Engage. They cited with appreciation the Engage information sessions and one simply wrote "empowerment". Two respondents (in longer responses) wrote that Engage and knowing that the resource was available was helping them think about future research plans in new ways, "has had a tremendous impact on how we think about and plan our future computational research."

Two respondents did not cite any aspects of Engage in this section saying either they'd forgotten because it was too long ago, or that they'd never "engaged with engage."

3.4 A respondent reflects

One respondent wrote a long piece reflecting on why they hadn't used Engage and OSG "to their potential." They write that, at the time of engagement, "code AND scientific maturity were inadequate at the time," relating to the fact that their code was still maturing and they discovered many scientific issues as they progressed ("adequate understanding of model behavior to know where to spend the compute cycles").

It appears their code developed in a way that makes it hard to move to the Grid now that it is reaching maturity and they reflected that perhaps education on how to modify MPI code in general terms would have been useful "so we would understand better how that process would interact with our future design choices, rather than modifying the code." They conclude their reflection by suggesting that the right time to Engage on specific projects is once the model is working locally, showing that the code and science had matured.

To summarize, this response suggests that early general education about infrastructural resources helps guide actions during the maturing of scientific inquiry, and that spending man hours adapting software should be reserved for scaling relatively mature projects. This respondent was "Extremely Satisfied" with all aspects other than the two "Ability to understand" science/my computation needs questions, where they indicated "Neutral." At numerous places through the survey this respondent reflects that they have not had sufficient manpower to engage as fully as they'd wish. They also only cite "Adapting Software" as "Slowing" their research and primarily work with Local Clusters as an alternative to the Grid.

Chapter 4

Feedback on Survey overall

The final box was a request for feedback on the survey. Here respondents noted a small number of aspects of the survey which could have been improved, mostly pertaining to targeting of questions (asking for Not Applicable responses). Future surveys should have a stronger ability to take split paths, for example reserving questions only for people who had run jobs.

Additional feedback mentioned that they expected to be asked about how well Engage actually empowered participants:

> *This survey didn't seem to capture questions about whether people who had worked with Engage had felt personally empowered in their grid computing. Did they go forward to setup their own VO? Did they introduce other colleagues to OSG? What portion of their overall grid usage was through Engage, and at what point did they dis-Engage to act as independent grid users? How much did they understand about what Engage was doing for them? How visible was the process of establishing a computational workflow to them? How much did they participate in this process? Were they able to understand and analyze the results they were given? Were errors and problems visible to them, and if so, were they able to filter these out or work to resolve/overcome the implications of grid-induced errors/faults (or did they get lost in the deluge of data from grid-scale computing, or simply delegate all responsiblity for this challenges to the Engage team)?*

These questions seem useful for future assessments of the Engage program.

Bibliography

R Development Core Team (2004). R: A language and environment for statistical computing. Technical report, R Foundation for Statistical Computing, Vienna, Austria.

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer, 1 edition.