Collaboratory for Research on Electronic Work
School of Information, University of Michigan

TeraGrid Evaluation Report, Part 2: Findings from the TeraGrid User Survey

Ann Zimmerman

Magia Krause

Katherine A. Lawrence

Thomas A. Finholt

July 2008

Collaboratory for Research on Electronic Work
School of Information
University of Michigan
1075 Beal Avenue
Ann Arbor, MI 48109-2112
http://www.crew.umich.edu

Table of Contents

Executive Summary

TeraGrid is a national infrastructure that integrates multiple resources at distributed resource provider facilities. In 2006, the National Science Foundation awarded a grant to the University of Michigan’s School of Information to conduct an external evaluation of TeraGrid. One of the main objectives of the evaluation study was to assess TeraGrid's progress in meeting the needs of its users. This report describes the results from the *TeraGrid User Survey*, a major activity in support of this objective.

Purpose and Methods

The purposes of the TeraGrid User Survey were to gain insight into the characteristics of those who use TeraGrid and to understand similarities and differences in the needs, motivations, and commitment of different types of TeraGrid users based on factors such as their experience with supercomputers, frequency of TeraGrid use, stage of career, field of research, gender, and age. The major constructs we sought to measure in the survey were informed by literature on technology acceptance and use, the influence of personal characteristics on use or intention to use technology, and the affect of social and organizational factors on computer use. We adapted items used in prior studies to increase the reliability and validity of the survey and the repeatability of the study.

The survey sample was constructed using data from the TeraGrid central database. Our population included all users who were active between 1 October 2005 and 30 September 2006 and principal investigators associated with active projects. We stratified our population along two criteria: 1) the largest allocation associated with a user (i.e., DAC, MRAC, and LRAC), and 2) the field of science associated with projects. We selected a total of 595 individuals, representing a random, stratified sample proportional to the distribution of users by field and allocation category. The 48-item questionnaire was administered online. We received a response rate of 52% for a final sample of 311.

Results

This survey's findings help to characterize TeraGrid users and their patterns of usage. In addition, they identify relevant relationships between usage patterns and users' satisfaction with TeraGrid which should help future implementation and budgeting decisions. The population of TeraGrid users is highly educated and most have at least several years of supercomputing experience. Those that use TeraGrid more frequently are also greater users of TeraGrid support, more strongly identified as TeraGrid users, perceive themselves as more experienced, and are more positive about TeraGrid’s usefulness, ease of use, and the facilitating conditions for using TeraGrid. These associations suggest that TeraGrid can improve its users’ experience by scaffolding those who are less frequent users. In sum, the population of TeraGrid users is generally satisfied with TeraGrid’s services and support, but there is room for improvement, particularly in support of those who—due to allocation limitations, unfamiliarity, or perceived barriers—use TeraGrid less frequently. By creating an experience for these infrequent users that more closely matches the experience of frequent—and satisfied—users, TeraGrid can further improve the perceived quality of its offerings.

1. Introduction

TeraGrid is a national infrastructure that integrates multiple resources at distributed resource provider facilities.[1] Following a 5-year construction phase, TeraGrid became operational in late 2004. At the time of the survey, TeraGrid's resources included more than 150 teraflops of computing capability and greater than 15 petabytes of online and archival data storage.[2] High-performance networks provide rapid access and retrieval to data. TeraGrid supports a variety of use cases ranging from exploiting a single TeraGrid resource to combining resources across sites.

In late spring 2006, the National Science Foundation (NSF) awarded a grant to the University of Michigan's School of Information (UM-SI) to conduct an external evaluation of TeraGrid. The primary goals of the evaluation were a) to give NSF leaders and policy makers general data to help them in making strategic decisions about future directions for cyberinfrastructure; and b) to provide specific information to TeraGrid managers to increase the likelihood of TeraGrid success. One of the main objectives of the UM-SI evaluation study was to assess the needs of TeraGrid users in order to assist NSF and TeraGrid in measuring progress toward meeting those needs and to provide information for planning purposes. Part 1 of this report describes the full range of methods that were employed toward this particular goal, including a user workshop, interviews, and a survey of current TeraGrid users. The latter activity is the subject of this report.

The purposes of the TeraGrid User Survey were to gain insight into the characteristics of those who use TeraGrid and to better comprehend their needs. At one level, we were interested to gain a picture of TeraGrid users in terms of attributes such as experience with supercomputers, frequency of TeraGrid use, stage of career, field of research, gender, and age. Beyond this, the aim was to understand similarities and differences in the needs, motivations, and commitment of different types of TeraGrid users based on their characteristics and other factors. In addition, to the findings from the survey, this report presents the conceptual frameworks that guided the development of survey constructs. This document also describes construction of the survey sample, design and administration of the questionnaire, and methods of data analysis.

[1] At the time of the survey, there were nine resource providers: Indiana University, National Center for Atmospheric Research (NCAR), National Center for Supercomputing Applications (NCSA), Oak Ridge National Laboratory (ORNL), Pittsburgh Supercomputing Center (PSC), Purdue University, San Diego Supercomputer Center (SDSC), Texas Advanced Computing Center (TACC), and University of Chicago/Argonne National Laboratory (UC/ANL).

[2] At the time of this report, the resources had grown to 750 teraflops of computing capability and more than 30 petabytes of online and archival data storage.

2. Background: Technology Adoption

The scientific users of high-performance computing (HPC), which we define as individuals from disciplines that share the need for HPC resources and services, have received little attention from scholars. This is the case despite the steady increase in the use of HPC by academic researchers following NSF's establishment of three supercomputer centers in 1985 (Graham, Snir, & Patterson, 2005, p. 13) and the fact that simulations made possible by HPC are now considered a 'third way' of doing science (Buetow, 2005; Rogers, 1998).[3] Topics that *have* been studied include the history and evolution of supercomputers (Elzen & MacKenzie, 1994; Schneck, 1990), the cultures and practices of disciplines that rely on HPC such as physics (Galison, 1997), and the role of simulations and models in science (Humphreys, 1990; Sismondo, 1999). While prior research in these areas provided useful background, it offered few insights to guide the design of a questionnaire to be administered to TeraGrid users. For example, what factors influence a user's level of satisfaction with TeraGrid, affect patterns of use, or help predict the use of grid computing? Although these questions have not been studied within the context of HPC use, there is a rich and varied body of research that has investigated technology adoption. This includes models of technology acceptance and use, the influence of personal characteristics on use or intention to use technology, and the affect of social and organizational factors on computer use. In the sub-sections below, we review the literature that informed the major constructs we sought to measure in the TeraGrid User Survey.[4]

2.1 Technology Acceptance and Use

The field of information systems (IS) has developed a number of theoretical models in an attempt to explain user acceptance and usage of information technology. The Unified Theory of Acceptance and Use of Technology (UTAUT) compared eight of the most prominent models, integrated elements from each into a unified model, and then validated the model (Venkatesh et al., 2003). According to UTAUT, four constructs play a significant role as direct determinants of user acceptance and usage of information technology:

- *Performance expectancy* is the degree to which an individual believes that using the system will help him or her to attain gains in job performance.
- *Effort expectancy* is the degree of ease associated with use of the system.
- *Social influence* is the degree to which an individual perceives that important others believe he or she should use the system.
- *Facilitating conditions* is the degree to which an individual believes that an organizational and technical infrastructure exists to support use of the system.

[3] Supercomputers were certainly in use before this time. However, according to Rogers (1998), prior to the mid-1980's access to supercomputers by academic researchers, mainly physicists and chemists, was limited primarily to those who had grants or contracts from the Department of Energy (DOE).

[4] *Constructs* are the elements of information that the survey seeks to measure (Groves et al., 2004, p. 41). These can be complex ideas such as consumer optimism or more straightforward such as consumption of bottled water.

UTAUT posits three direct determinants to explain the use of or intention to use information technology (performance expectancy, effort expectancy, and social influence) and two direct determinants of usage behavior (intention and facilitating conditions). Performance expectancy, or what other models refer to as *usefulness*, is the strongest predictor of acceptance or intention to use a particular technology in UTAUT and many other models. Factors such as computer self-efficacy and attitudes toward technology (defined as a user's overall affective reaction to using a system) were not found to be significant. We used the UTAUT model as a conceptual frame to analyze TeraGrid use and the intention to use grid computing. In addition, we adapted questionnaire items developed in prior surveys to measure the four constructs listed above.

The context of TeraGrid also offered an opportunity to extend the UTAUT framework. Information systems models have been developed through the study of relatively simple, individual-oriented information technologies such as word processing software (Davis, Bagozzi, & Warshaw, 1989; Venkatesh et al., 2003), whereas TeraGrid use involves multiple steps, specialized knowledge, and the ability to demonstrate efficient use of the system. Table 1 lists the hypotheses we formed based on UTAUT, data we collected through interviews (see Part 1 of this report), and other literature. The rationale for some of the hypotheses shows the exploratory nature of the study; it was difficult to form firm hypotheses at this stage of our knowledge of TeraGrid users specifically and HPC users generally.

H1: There will be a positive relationship between perceived usefulness of TeraGrid and simulation/modeling approach. Rationale: Simulation requires compute power in order to 1) extend the realization of complex natural phenomena so they can be understood scientifically; 2) test systems that are costly to design or to instrument, or 3) replace experiments that are hazardous, illegal, or forbidden.
H2: There will be a positive relationship between perceived usefulness and frequency of use, and frequency of use will be related to allocation level. Rationale: Larger allocations have more service units available for use. On the other hand, most projects with large allocations have multiple users, so use may be spread out and not all users may be frequent users.
H3: There will be a positive relationship between perceived usefulness and personal innovativeness. Rationale: See text stating that TeraGrid users might be considered early adopters of HPC.
H4: There will not be a significant relationship between discipline and perceived usefulness. Rationale: For those who use TeraGrid, it is necessary to answer research questions of interest. This will mitigate differences that might otherwise be expected to exist between disciplines.
H5: There will be a positive relationship between the use of self-developed codes and perceived ease of use. Rationale: An HPC expert stated that those who do not write their own codes but rely on commercial software are forced to wait when something new comes along in terms of supercomputer architecture (Anonymous, 2000). He also implied that those who write code have more intimate knowledge of supercomputers and how to use them. On the other hand, Graham and colleagues (2005) noted that commercial codes and some community codes are very large, so porting them can be difficult.

H6: There will be a positive relationship between perceived ease of use and facilitating conditions such that those who find TeraGrid easy to use perceive that they have people, documentation, and guidance to assist them in using TeraGrid.
H7: There will be a positive relationship between experience with supercomputing and perceived ease of use.

Table 1: Hypotheses regarding technology acceptance

2.2 Personal Innovativeness

Other studies have analyzed the influence of personal characteristics on the adoption of or intention to use a particular technology. As noted above, individual characteristics have not generally been found to significantly affect acceptance of a technology. Personal innovativeness, however, is one characteristic that has been shown to have some influence on the use of new technologies (Lu, Yao, & Yu, 2005). Innovative individuals tend to take more risks and may be more confident in their ability to handle a new technology. For these reasons, they may be more positively disposed to use a new technology without clear perceptions regarding its usefulness or ease of use. We reasoned that at least some part of the population of TeraGrid users could be considered early users of a new technology. At the time we administered the TeraGrid User Survey, more than ninety percent of TeraGrid users had allocations at one or more of three major supercomputing centers: the National Center for Supercomputing Applications (NCSA), San Diego Supercomputing Center (SDSC), and/or Pittsburgh Supercomputing Center (PSC). Further, a large percentage of TeraGrid allocations are made to a small percentage of investigators. Many of these so-called "hero users" had been using HPC for 15 years or more. Thus, TeraGrid users might be characterized by a high-degree of personal innovativeness based on their use of advanced computing resources. However, we also learned in our interviews that the use of TeraGrid and/or HPC was a necessary tool for the conduct of science in particular areas such as large-scale molecular dynamics and quantum chronodynamics. Thus, for these users and members of their research teams, including postdocs and graduate students, there is not really a choice about whether to use HPC. However, they may have options in terms of *where* to compute, which is also something we sought to measure.

H8: There will be a positive relationship between the use of grid computing and perceived usefulness of TeraGrid. Rationale: TeraGrid was designed to enable this usage mode.
H9: There will be a positive relationship between perceived usefulness of grid computing and personal innovativeness. Rationale: Grid computing is difficult to do, so those who use TeraGrid in this way are more likely to be innovative.

Table 2: Hypotheses regarding personal innovativeness and grid computing

2.3 Social and Organizational Contexts

Finally, social and organizational contexts are potentially important elements in technology adoption. Individual characteristics such as personal innovativeness may influence technology acceptance or intention to use a particular technology, but the culture, organization, and work practices of research fields are other possible sources of influence. Previous studies have compared scientific disciplines along dimensions such as competitiveness (Knorr-Cetina, 1999), need for access to scarce resources such as specialized instruments, and the nature of the work, including the physical scale of the research, agreement on research questions and methods, and the need for help (see Birnholtz, 2007 for a review). Birnholtz (2007) investigated the influence of social factors and the nature of work on the likelihood of a researcher to collaborate at a particular point in time. He found that differences in the nature of the work in different fields explained more about a researcher's propensity to collaborate, although social factors had subtle effects. Others studies have shown significant differences in use of the Internet and computer-mediated communication (CMC) by scientific field (see Walsh & Roselle, 1999 for a review). CMC and Internet use have also been shown to have a positive relationship with collaboration and scientific productivity (e.g., Hesse et al., 1993; Kaminer & Braunstein, 1998; Walsh et al., 2000).

It was difficult to form hypotheses regarding relationships between social and organizational factors and the various disciplines represented by the population of TeraGrid. First, while TeraGrid is a networked computer system, it does not facilitate communication among users. Instead, the network serves to tie multiple resources together, so users can submit jobs to one or more of them and/or transfer data from one place to another. Second, we anticipated from the interview data that domain was unlikely to be a significant factor to explain differences between users since the common need for access to advanced computational resources seemed likely to outweigh such differences. Further, in fields such as social science, where we might expect to see differences, there were not enough users to reliably detect them if they existed. Although we had few hypotheses about these constructs, they have been important in past studies, and we reasoned that baseline data on the social and organizational contexts of the TeraGrid user population would be useful to recognize changes that might occur over time as new communities utilize TeraGrid. We used items developed by Birnholtz (2007) to measure collaboration propensity, which integrated nature of work aspects, and items he adapted from Walsh and Hong (2003) to measure scientific competition.

2.4 Other Hypotheses

Based on interview data, we hypothesized that allocation size and frequency of use would be related and more frequent users would identify themselves more strongly as TeraGrid users.

H10: There will be a positive relationship between frequency of use and allocation level such that those with larger allocations will use TeraGrid more frequently.
H11: There will be a positive relationship between frequency of use and identification as a TeraGrid user such that those who use TeraGrid more often will identify themselves more strongly as a TeraGrid user.

Table 3: Other hypotheses

3. Methods and Data Analyses

In this section we discuss the methods employed in the TeraGrid User Survey. This includes a description of the sampling scheme, questionnaire design and content, and administration of the survey.

3.1 Sample Selection and Composition

The survey sample was constructed using data from the TeraGrid central database (TGCD). In order to use NSF high-performance computing resources, including TeraGrid resources, prospective users must prepare and submit a proposal. Typically, when proposals are accepted and projects are granted an allocation, they are assigned to one of three award categories: development allocations (DACs), medium resource allocations (MRACs), and large resource allocations (LRACs). The awards differ based on the number of service units allotted, ranging from 30,000 for DACs, between 30,000 and 200,000 service units for MRACs, and over 200,000 service units for LRACs. Service units are generally defined as "equivalent to either one CPU-hour, or one wall-clock-hour on one CPU, of the system of interest" although exact definitions vary based on resource platform according to the NSF Cyberinfrastructure Resource Allocations Policy document.[5] Data on DAC, MRAC, and LRAC TeraGrid projects, along with information on users associated with those projects, including their names, postal addresses, and email addresses, are stored in the TGCD.

We constructed a sampling frame for our target population based on the data available to us from the TeraGrid central database. We defined our population of users in two ways. First, we included all users who were active between 1 October 2005 and 30 September 2006. *Active users* were defined as those who had consumed at least one service unit during the selected timeframe. Second, we included all Principal Investigators (PIs) associated with projects that were active during the specified time period even if they themselves had not consumed any service units. We chose to do this because we felt their opinions about TeraGrid were valuable even without direct, hands-on experience. We limited our population to users of the past year in order to measure recent use of TeraGrid because survey methodologists have found that memory fades with time, thus decreasing the likelihood of response accuracy (Tourangeau et al., 2000). The following were excluded from the survey: TeraGrid staff, Science Gateway Community Users, and users selected to pilot test the survey.

We stratified our population along two criteria: 1) the largest allocation associated with a user (e.g., DAC, MRAC, and LRAC), and 2) the field of science associated with projects as taken from the TeraGrid central database. We chose these strata in order to analyze any significant differences among the categories. We oversampled the fields of engineering and geoscience to help ensure that if response rates were low we would have data to make significant conclusions about these areas. We also included all users from social and behavioral

[5] See http://www.cipartnership.org/Allocations/allocationspolicy.html#_Toc116808729

sciences.[6] Since some users simultaneously have awards of more than one type, we decided to associate users with their largest allocation award in order to avoid duplications in the sampling frame. Associating users with their largest allocation appropriately identifies users who exclusively have smaller allocations (DACs, MRACs) within the timeframe of our target population. We selected a total of 595 individuals, representing a random, stratified sample proportional to the distribution of users by field and allocation category and including the oversampled areas mentioned above.

We received 311 valid surveys, which represents a response rate of 52%. Response rates were similar across all strata, ranging from 42% to 71% by field and between 52% and 54% by allocation category.

3.2 Survey Design and Content

The survey was designed to meet several goals. The primary purpose, as stated previously, was to "get a picture" of the TeraGrid user population according to various attributes and to understand similarities and differences in the needs, motivations, and commitment of different types of TeraGrid users based on factors such as their experience with supercomputers, frequency of TeraGrid use, stage of career, field of research, gender, and age. Another goal of the survey was to provide information of particular interest to TeraGrid not included in the items we developed. In this section we describe the overall survey design and the construction of the items and questions. The final questionnaire, which is available in Appendix C, contained 48 items.

There were three overriding considerations in the design of the survey. First, we wanted to limit the time and mental effort required to complete the survey as previous research has shown that they affect response rate. We considered these factors throughout the construction of the survey, and we piloted the survey with expert and non-expert users to insure we had met these goals. Second, we needed to create a survey that would be relevant to the broad array of individuals who make up the TeraGrid user community. For example, specific questions about the allocation proposal process would have been difficult to ask because only some TeraGrid users participate in this process. Third, wherever possible and appropriate, we used or adapted items from previous surveys as this approach improves the reliability and validity of survey items and the repeatability of a study. The items used to measure personal innovativeness (Questions 4-5), collectivist orientation (Q6-7), scientific competition (Q8-10), collaboration propensity (Q11-15), intention to use grid computing (Q28-30), and technology acceptance and use (Q23-27 and Q31-33) were drawn from prior studies as described earlier in this report. Several items were contributed by TeraGrid managers based on information they wished to collect from their users (Q3, 21, & 35-38). We also used two questions (Q19 & 24) that appeared in a user survey developed jointly by NCSA and SDSC. Survey items also originated from the qualitative data we gathered through interviews. We wished to test nascent hypotheses about the affect of research approach (Q2) and the types of codes used (Q19) on the needs of TeraGrid users. Finally, we selected and adapted items

[6] Oversampling occurs when certain groups are sampled with higher probabilities than others. This provides enough cases to complete analysis of subgroups of the population.

used in prior surveys to measure constructs such as stage of career, frequency of use, and other demographic information.

When we had a draft of the survey, we pre-tested it with naïve users and subject experts to insure that our survey items and questions were relevant and would be understood and easily answered by our respondents. These suggestions were considered and many were incorporated into the final version of the survey.

3.3 Survey Administration

The survey was administered using SurveyMonkey.com and was available from December 5, 2006 to January 7, 2007. The invitation to take the survey was sent by mail in early December 2006. Each envelope contained a cover letter from us that explained the nature of the study (Appendix A), an endorsement letter from the Director of TeraGrid (Appendix B), and a $2 bill as a cash pre-incentive (Birnholtz et al., 2003; Church, 1973).

Our cover letter also informed individuals that they would receive an e-mail message within the next week that contained a direct link to the survey, and it provided a URL and a unique identification number for those who wished to take the survey immediately. The unique identification code enabled us to track who responded to the survey. Keeping track of respondents also allowed us to send reminders only to individuals who had not yet completed the survey.

The relative newness of TeraGrid presented some challenges, which we attempted to address in the survey administration. Namely, we anticipated that some respondents might be confused as to why they were identified as a TeraGrid user and included in our survey sample. Our interviews with TeraGrid personnel and with users (see Part 1 of this report) indicated that many individuals with TeraGrid accounts utilize resources and services at one or two TeraGrid sites, much as they did before the existence of TeraGrid. With the exception of Science Gateway developers, many of the users we interviewed had heard of TeraGrid, but the details of TeraGrid were unclear to them.[7] In addition, when Resource Providers added resources to TeraGrid they often attempted to make the transition transparent to users. For instance, the number of TeraGrid users increased significantly in the first and second quarters of 2006. This jump was due to NCSA and SDSC officially making all their resources available to TeraGrid on April 1, 2006. Pre-TeraGrid allocations on these resources were simply transferred; no action was required by users, and they did not receive new logins or passwords. Thus, users continued to work as they always had, largely unaware of the change that had taken place. We took two steps to inform potential respondents about why they were selected to complete the survey. First, the cover letter from TeraGrid listed all the Resource Provider sites. Second, the letter from us stated that individuals were chosen to receive the survey because they used resources at one or more of these sites; this information was repeated in the introductory text to the survey (see Appendix C).

[7] Several responses to the open-ended questions (Q40-41) reinforced what we learned in the interviews. For example, one respondent wrote: “I’ll try to use TeraGrid to find out other barriers.” Another said: "I did not have any access to TeraGrid.”

3.4 Data Analyses

With the exception of the two open-ended questions (Q40-41) all analyses were conducted using the statistical package SPSS®. We calculated descriptive statistics for all variables and conducted tests of association between variables based on the appropriate method in specific instances. These tests are described in the results section. The responses to the open-ended questions were coded according to major categories that emerged from an analysis of the data.

4. Results

This section begins with a review of data on the general attributes of the 311 respondents from the survey sample. The remaining results are grouped according to some of the survey's key topic areas. A review of the descriptive statistics found that the distribution of the responses was not range restricted. The questions did vary significantly from multivariate normality for skewness ($Z= -2.41$, $p<.02$) but not for kurtosis ($Z=0.67$, $p=.50$). A visual examination of the shape of the distribution showed that many of the variables had a slight bias towards higher scores, but in most cases they did not look very different from the normal distribution curve.

4.1 General Attributes of Respondents

Of the responses we received, 82% were from males and 15% were from females.[8] The age of respondents ranged from 20 to 85, and approximately half of those who provided their year of birth were younger than 35 (mean age=37). Most respondents had a PhD degree or equivalent (70%) and were affiliated with a research university (88%). Half of all respondents received their highest degree after 2000 and another third received it in the 1990s. Faculty comprised 52% of respondents, and the remainder was made up of students, postdocs, and research staff. We did not find significant relationships between these general attributes and major constructs related to technology adoption. Frequencies of TeraGrid use and allocation size, as discussed later in this section, were more useful in distinguishing users from each other.

[8] Except for field of science, percentages have been rounded to whole numbers. In the case of gender, a small number of respondents chose not to answer this question.

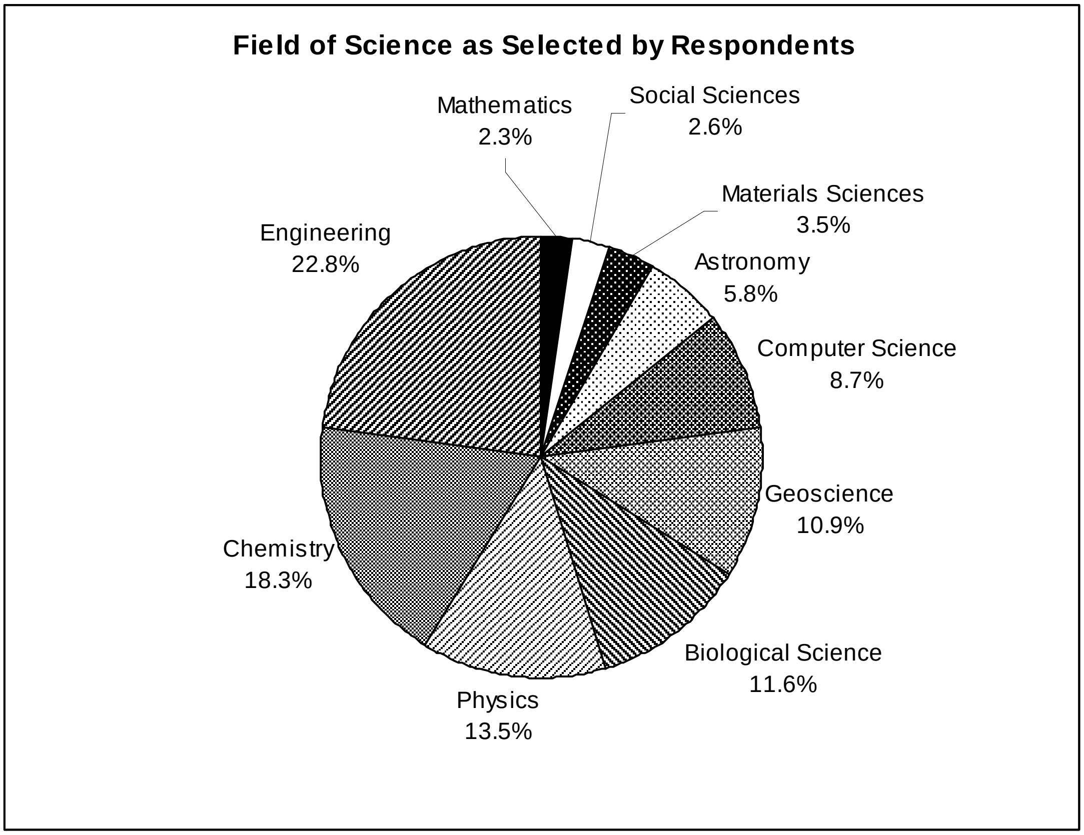

Figure 1

Figure 1 shows the field of science that respondents indicated most closely represented their research. Most respondents described their research approach as simulation/modeling (71%). The remainder was split between theoretical/analytical (16%) and experimental/observational (12%) methods. Hypothesis 1, which stated that there would be a positive relationship between perceived usefulness of TeraGrid and simulation/modeling approach, was not supported.

> **H1**: There will be a positive relationship between perceived usefulness of TeraGrid and simulation/modeling approach. (Not supported)

Surprisingly, only 77% of respondents indicated that they use local workstations. This may be because some interpreted "local workstation" as a dumb terminal allowing access to TeraGrid. Respondents use local resources extensively. Local clusters with 64 or fewer processors are used by 62% while 36% use local clusters of 56 to 128 processors.

We analyzed the list of respondents based on the information we had regarding the largest allocation associated with a user and found that 56% of respondents had DAC allocations, 23% had MRAC allocations, and 21% had LRAC allocations.

4.2 Nature of Research

Results indicate that more than half the respondents face a competitive scientific environment as 60% agreed or strongly agreed that the competition for prizes or widespread recognition in their field is intense. In spite of the competitive nature of their fields, respondents overwhelming perceived collaboration as important. About three-quarters responded that it was necessary in their field and more than 90% agreed or strongly agreed that it is useful in solving problems of interest to them and for accessing people with expertise helpful to them.

4.3 Supercomputer Use

The items in the section on supercomputer use asked respondents about their experience with supercomputers, the importance of supercomputers to their research, access to supercomputers, and the resources they need to accomplish their research. Respondents generally have several years of experience with supercomputers; 61% have been using supercomputers for 3 or more years, and 44% have been using supercomputers for five or more years. However, 14% indicated they had less than one year of experience. When asked to describe their experience with supercomputers relative to others in their field, respondents appeared to be modest, with 41% reporting that they were not at all experienced or just somewhat experienced (the lower two categories), while 31% described themselves as very or extremely experienced (the upper two categories). While the perceived level of experience is difficult to generalize across disciplines because use of HPC various within and across domains (see Part 1 of this report), it is interesting to note that this variable was significantly associated with allocation level (for Gamma coefficient, $p<.001$). Respondents with DAC allocations were more likely to place themselves in the lower two categories, whereas those with MRACs and LRACs tended to place themselves in the upper three categories.

As we expected based on interviews, there was overwhelming agreement among respondents that supercomputers are necessary to answer research questions of interest; 53% strongly agreed with this statement and 38% agreed. Nearly half of them (45%) use codes developed by themselves or their group and augmented with third-party routines or libraries. Thirty-nine percent use third-party software, some of which is augmented with their own routines or libraries. Fields of science demonstrate significant differences in the software codes that they use. Chi-square tests of association are statistically significant ($p<.001$), and the differences are primarily as follows: Biologists and chemists are more likely to use third-party codes or third-party codes augmented with some of their own routines/libraries. Computer scientists and astronomers favor codes developed entirely by themselves or their group. Geoscientists use codes developed by their group and augmented with third-party software.

We offered respondents a list of supercomputer resources and asked which ones they need for their research. Computer systems, not surprisingly, were used by almost all respondents. The next most prevalent resources used were persistent online storage, user services support, and visualization software (see Figure 2). Nearly half the respondents (49%) have access to supercomputer resources through their institution's supercomputer facility. Access to other options is limited. Department of Energy resources were accessible to 16% and 14% make use of state or regional supercomputer facilities. The remaining choices were each selected by 5% or less of respondents.

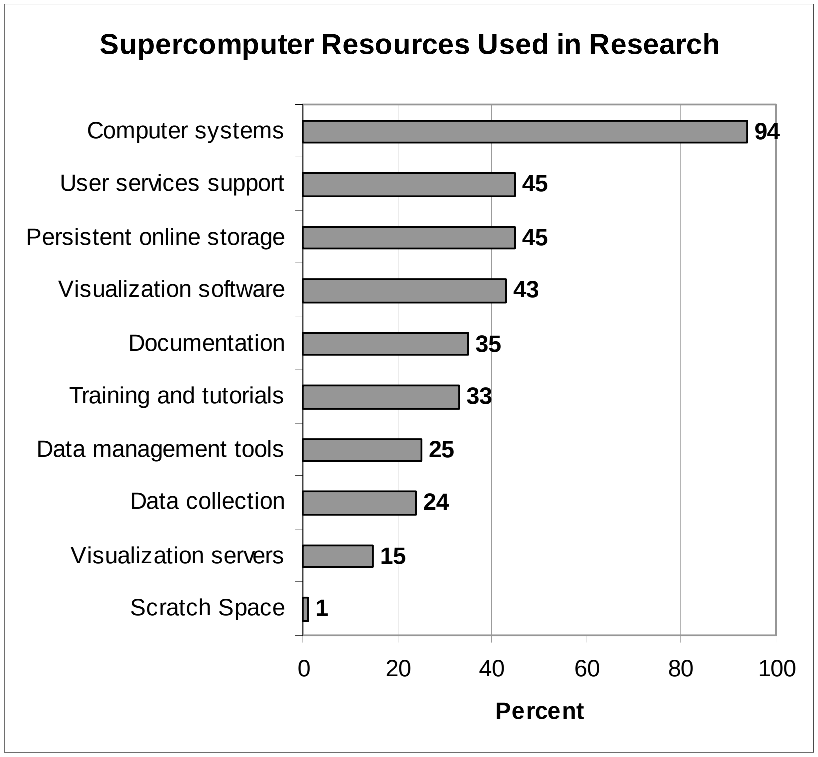

Figure 2

4.4 Grid Computing

We were surprised to find that 45% of those surveyed said they currently use grid computing capabilities as most of our interviewees did not compute in this way, and statistics from TeraGrid do not show significant use of this mode. Ten percent indicated that they did not know if they use grid computing. Of those who stated that they use grid computing, 84% use TeraGrid resources for that purpose. The same percentage expects to continue to use grid computing and almost as many find grid computing useful in their research. However, only 45% agree or strongly agree that it is easy to become skillful at grid computing, and 37% are neutral on this issue. Respondents' use of grid computing is at all stages; 27% of those using grid computing have used it in production runs, 21% have experimented with test runs, and another 21% have investigated the capabilities offered by grid services or software. Because we could not be certain how respondents perceived the items in this section of the survey, we chose not to analyze the data beyond the generation of descriptive statistics. We are, therefore, unable to state whether the following two hypotheses are supported.

H8: There will be a positive relationship between the use of grid computing and perceived usefulness of TeraGrid.

H9: There will be a positive relationship between perceived usefulness of grid computing and personal innovativeness.

Of those who do not yet use grid computing, 56% believe that grid computing would be useful in their research, and 26% expected to experiment with grid computing in the six months following the survey. One-third of respondents believe they would find it easy to become skillful at grid computing. Again, this was a somewhat surprising finding because while some interviewees noted it could be helpful to them, most were not interested in pursuing grid computing for various reasons (see Part 1 of this report).

4.5 Use of TeraGrid

The section on use of TeraGrid consisted of three parts. In the first part (Q31-33), we sought to measure factors shown in other studies to affect technology adoption: 1) usefulness of TeraGrid to respondents' research, 2) ease of using TeraGrid, 3) facilitating conditions supporting their use, and 4) the degree to which others influence their use. We constructed and tested the statistical validity of scales for each of these constructs. We used Cronbach's alpha (Cronbach, 1951) for reliability estimates. A reliability estimate of .7 or better is considered to be good for early stages of research, and an estimate of .8 is advised for basic research (Nunnally, 1978).

The scale for usefulness of TeraGrid was composed of four questions that fit well together. The reliability of the scale was excellent (.95) and the items were highly correlated. The mean response was 4.0 (scale of 1 to 5 where 5 equals strongly agree) with a standard deviation (SD) of .81. The 4-item scale for ease of use had a very good reliability coefficient of .82, with a mean of 3.5 (SD=.72) (i.e., between neutral and agree). The 3-item scale for facilitating conditions had an adequate reliability coefficient of .73 and a mean of 3.4 (SD=.72). The two items measuring the degree to which others influence the use of TeraGrid were not well correlated (r=.43) and therefore did not show good reliability (alpha=.60). An exploratory factor analysis (EFA) of the items verified that the items for the three main scales (usefulness, ease of use, and facilitating conditions) did not load on unintended factors (i.e., demonstrated discriminant validity). The maximum likelihood factor analysis, using varimax rotation, showed good factor structure with eigenvalues greater than 1 for all 3 factors, which accounted for 74% of the variance. The factor structure does suggest, however, that ease of use correlates with the other two scales.

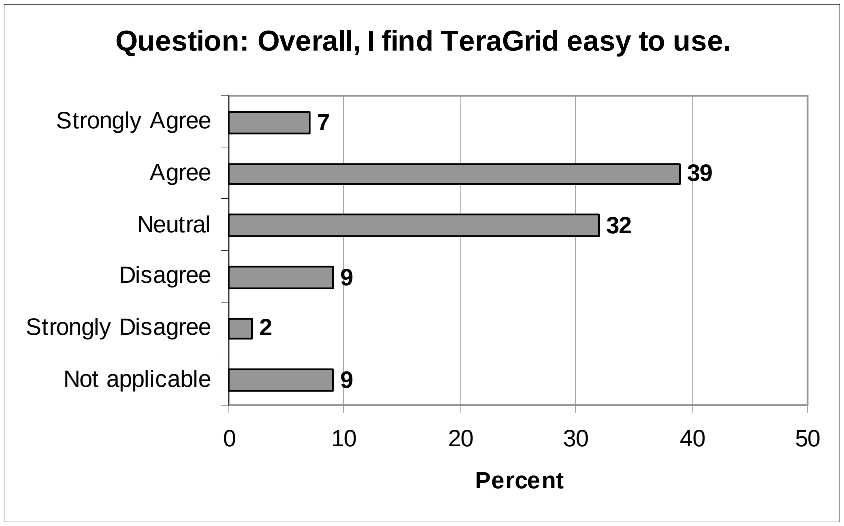

Figure 3

These results indicate that respondents have a bias in favor of the usefulness of TeraGrid and are favorable to a somewhat lesser degree in terms of the ease of use and facilitating conditions for their use of TeraGrid. Looking at individual questions in each of these scales, nearly three-quarters of respondents (73%) perceive TeraGrid as generally useful in their research, and a similar percentage (69%) agreed or strongly agreed that TeraGrid increases their research productivity. TeraGrid was not as widely perceived as being easy to use (see Figure 3). Further, the hypothesis regarding ease of use and self-developed codes was not supported nor was ease of use and experience with supercomputing.

H5: There will be a positive relationship between the use of self-developed codes and perceived ease of use. (Not supported)

H7: There will be a positive relationship between experience with supercomputing and perceived ease of use. (Not supported)

Almost half (46%) of the respondents agreed or strongly agreed that TeraGrid is easy to use, but items for the ease of use and facilitating conditions scales contained high percentages of neutral responses compared to other sections of the survey. These three scales are also correlated. Ease of use is significantly correlated with both usefulness ($r=.57$, $p<.001$) and facilitating conditions ($r=.52$, $p<.001$). To a lesser extent, usefulness and facilitating conditions are correlated ($r=.39$, $p<.001$). Our hypothesis regarding usefulness and field was supported:

H4: There will not be a significant relationship between discipline and perceived usefulness of TeraGrid. (Supported)

However, contrary to expectations, willingness to try information technology is not correlated with usefulness. Therefore, the hypothesis below was not supported.

H3: There will be a positive relationship between perceived usefulness and personal innovativeness. (Not supported)

The relationship between experience with supercomputing and personal innovativeness was significant, but the correlation was very low.

The second portion of questions on TeraGrid use asked respondents about their frequency of using TeraGrid, their use of and satisfaction with TeraGrid support, and the degree to which they identified themselves as TeraGrid users (Q34-36 & Q39).[9] Results show that survey respondents are frequent users of TeraGrid. More than half (53%) indicated they used it daily or weekly and 22% used it monthly. Frequency of use is also significantly associated with allocation size (p=.011). For example, respondents with DAC allocations (who represented 56% of those who responded to the survey) represented more of the quarterly or less frequent use (between 60% and 70% of those reporting those levels of use). MRACs are slightly more heavily weighted toward monthly use and LRACs more toward daily or weekly use. Thus, the following hypothesis was supported:

H11: There will be a positive relationship between frequency of use and allocation level such that those with larger allocations will use TeraGrid more frequently. (Supported)

More than 70% of those surveyed had contacted TeraGrid support at least once in the past year (see Figure 4). Respondents were largely positive about the support they received; 57% indicated that were satisfied or extremely satisfied (see Figure 5), with an average (mean) evaluation of 3.8 out of 5 (SD=0.8). Of the 16% who were neutral about their satisfaction with TeraGrid support, more than half of those were people who have used support services fewer than twice in the last year.

[9] We also asked whether respondents were affiliated with a Science Gateway or ASTA project. Since very few respondents indicated that they were (7% for Gateways and 2% for ASTA), we do not discuss these results.

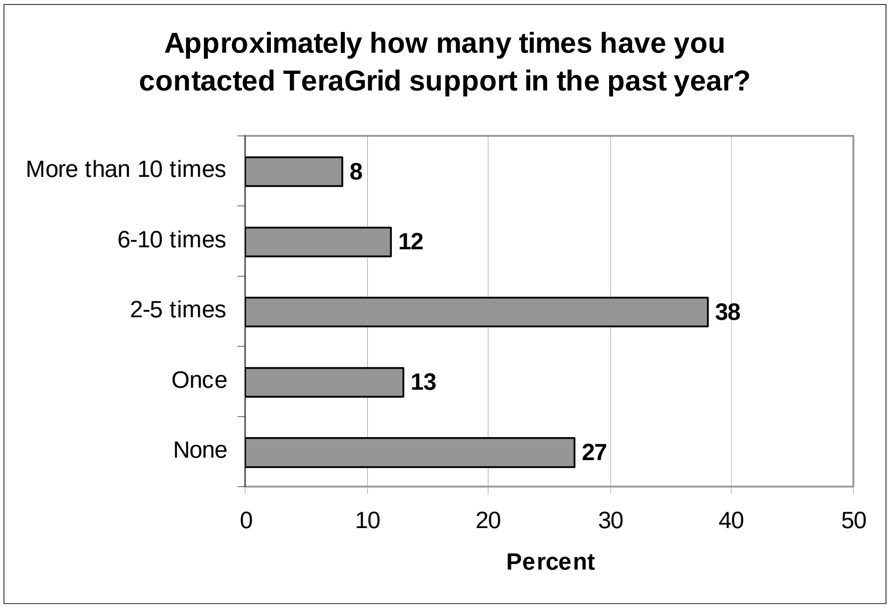

Figure 4

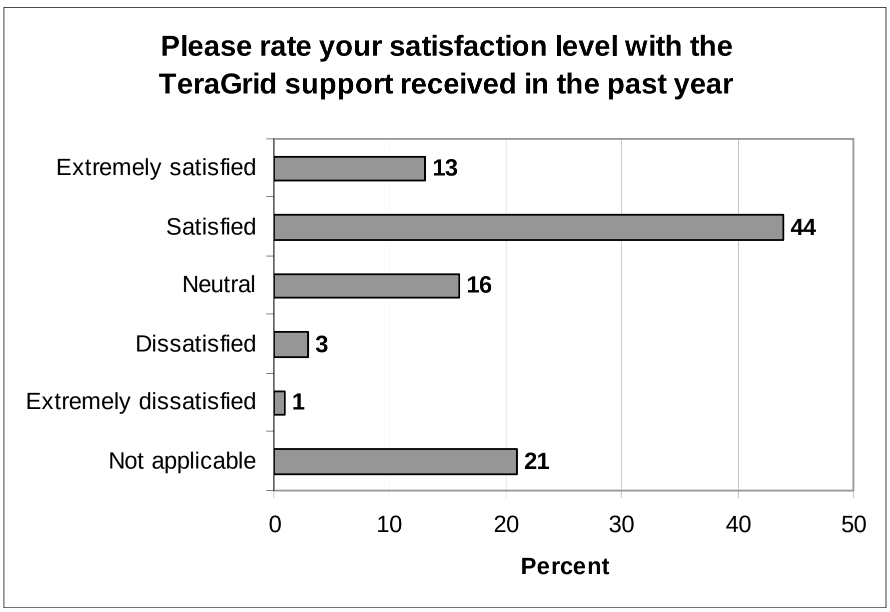

Figure 5

The frequency with which respondents used TeraGrid and TeraGrid's support services are associated with the various measures of satisfaction. Not surprisingly, there is a strong association between the frequency of using TeraGrid and the frequency of contacting support ($p<.001$) such that the more frequently a respondent uses TeraGrid, the more likely he or she also contacts support more frequently. These more frequent users are also more satisfied with the support provided by TeraGrid ($p<.001$). In fact, those who contacted support six or more times are much more likely to report that they are extremely satisfied with this service whereas those who have little contact with support are more likely to report neutral. Finally, the frequency of using TeraGrid is positively associated with perceived usefulness, ease of use, and facilitating conditions. The statistically significant difference in perceptions of usefulness and ease of use, however, is between those who never or only quarterly use TeraGrid and those who use it on a monthly, weekly, or daily basis. Likewise, those who perceive facilitating conditions to be higher are those who use TeraGrid, even if only on a quarterly basis. Based on these findings, there was support for the following hypothesis:

> **H6**: There will be a positive relationship between frequency of use and perceived usefulness, ease of use, and facilitating conditions. (Partially supported)

The degree to which respondents are satisfied with TeraGrid support is significantly correlated with their evaluations of usefulness ($r=.43$, $p<.001$), ease of use ($r=.52$, $p<.001$), and facilitating conditions ($r=.51$, $p<.001$). However, a *t*-test comparing group means based on frequency of contacting TeraGrid support indicates that the significant differences are between those who have not contacted TeraGrid support at all and those who have used it frequently (six or more times). Likewise, those who perceive themselves as having less experience relative to others in their field also contact TeraGrid support less frequently. Specifically, those who have contacted TeraGrid support five or fewer times in the last year rate themselves significantly lower in experience when compared to those who have contacted TeraGrid support more than ten times. (It is important to remember that this is *perceived* level of experience calibrated to others in the same field, not an absolute measure.)

Slightly more than half of those surveyed (52%) identified themselves as TeraGrid users (see Figure 6). Like earlier items in this section of the survey, this one elicited a substantial number of neutral responses (30%). The frequency of use and the degree to which individuals identify themselves as TeraGrid are significantly positively associated ($p<.001$) such that people who use TeraGrid more often identify themselves as TeraGrid users. For example, those who use TeraGrid only quarterly more commonly responded as neutral, disagreeing, or strongly disagreeing that they identify themselves as TeraGrid users. Likewise, allocation size is significantly associated with identification as a TeraGrid user ($p=.002$). There was also a strong association between identification as a TeraGrid user and the mean responses on the usefulness and ease of use scales ($p<.001$). Thus, the following two hypotheses were supported.

> **H2**: There will be a positive relationship between perceived usefulness and frequency of use, and frequency of use will be related to allocation level. (Supported)

> **H10**: There will be a positive relationship between frequency of use and identification as a TeraGrid user such that those who use TeraGrid more often will identify themselves more strongly as a TeraGrid user. (Supported)

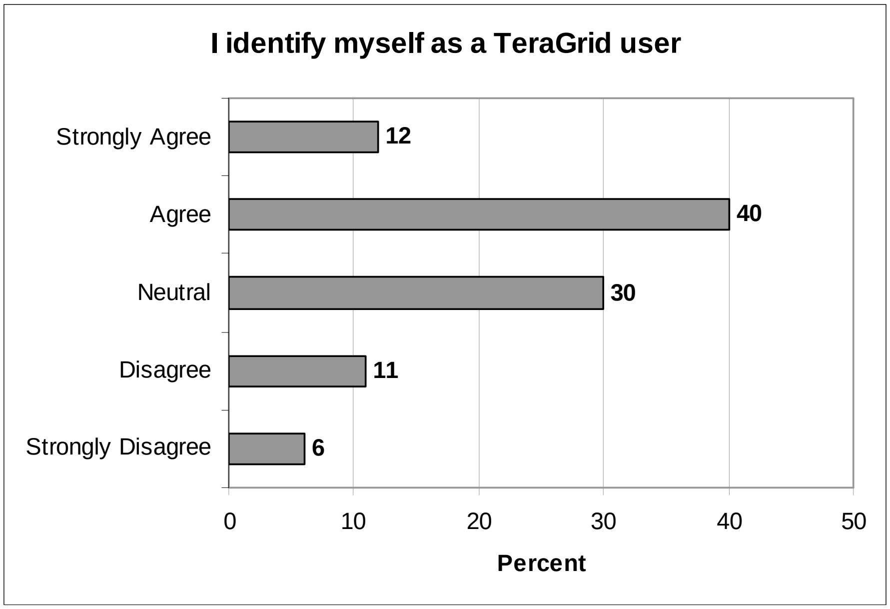

Figure 6

Two open-ended questions completed the section on TeraGrid use. These questions asked respondents to identify the two most significant barriers to their use of TeraGrid and the two things that would make TeraGrid more useful to them. We coded these responses (identifying larger categories into which the responses fit) and categorized them according to fourteen types of barriers and the same number of categories of improvements. These barriers and improvements are displayed in Figures 7 and 8. (Note that Figure 7 is listed in order of decreasing percentages and categories in Figure 8 are arranged in the same order as Figure 7.)

The top barriers to the use of TeraGrid are in the area of job submission, scheduling, and the turnaround of jobs. Most of the comments complained about long queue wait times. The next greatest concerns are documentation, support, and training. Included in this category are comments about the steep curve in learning to use TeraGrid; a lack of training opportunities, including tutorials, as well as time for learning; selecting the appropriate resource(s) to use; a lack of up-to-date, easy to find, and/or user-friendly documentation on topics such as installed software, libraries, and compilers, and "how-to" information; and complexity of the overall system and the common software. Applications software was also identified as a barrier, but this issue was mentioned only slightly more frequently than the bulk of remaining issues. These responses were consistent with the data we collected in interviews and at the user workshops. See Appendix D for representative responses in each barrier category and respondent suggestions to the most commonly perceived barriers.

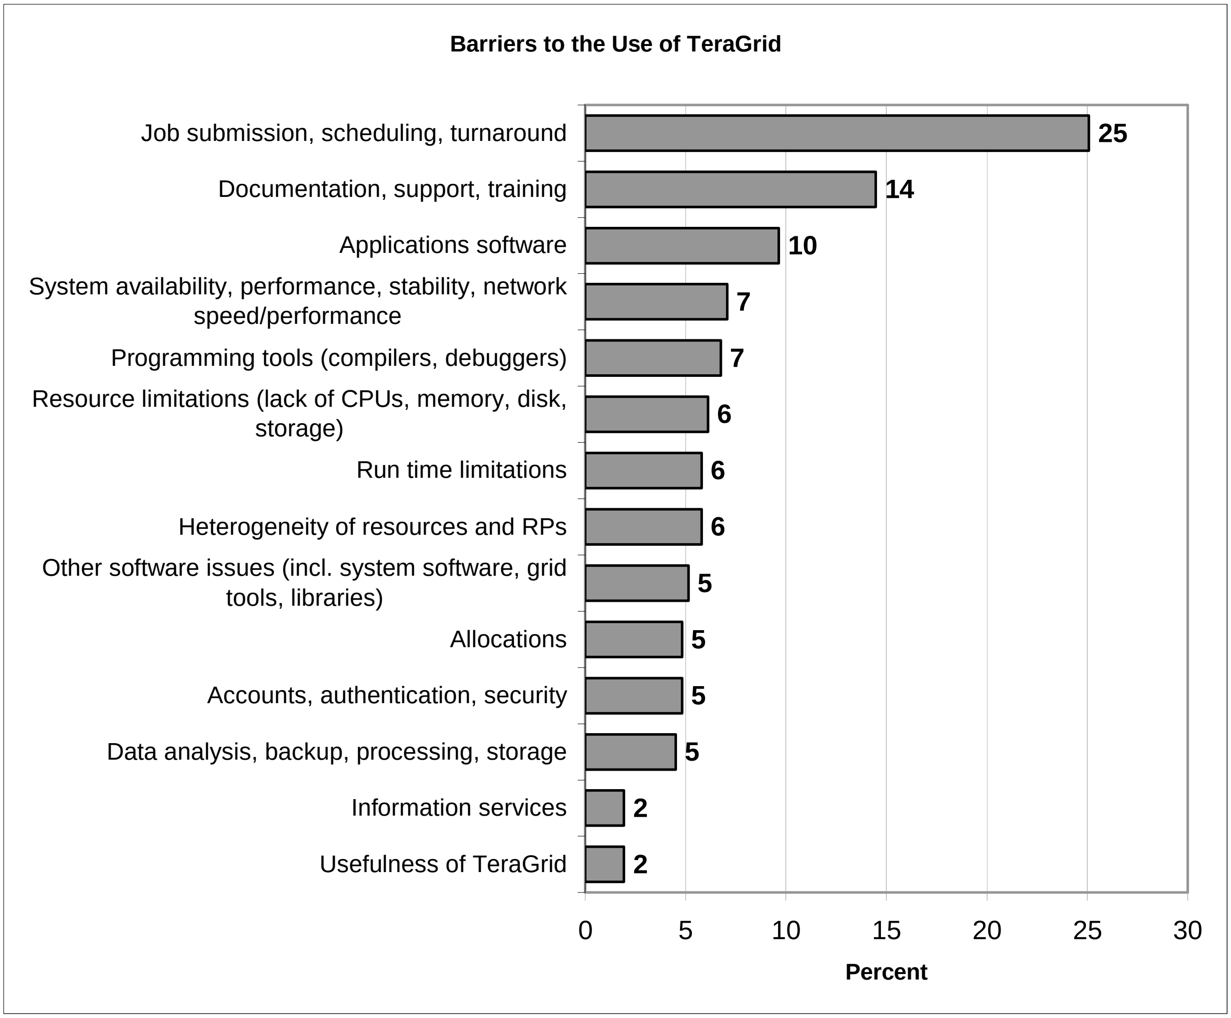

Figure 7[10]

Job submission, scheduling, and turnaround time and documentation, support and training were also the top areas noted as being in need of improvement, but the remaining suggestions were not prioritized in the same way as the barriers. The third most common suggestion was to improve resource limitations (such as lack of CPUs, memory, disk capacity and storage).

[10] Percentages here are calculated relative to an N of 311; however, 33% of respondents did not answer this question.

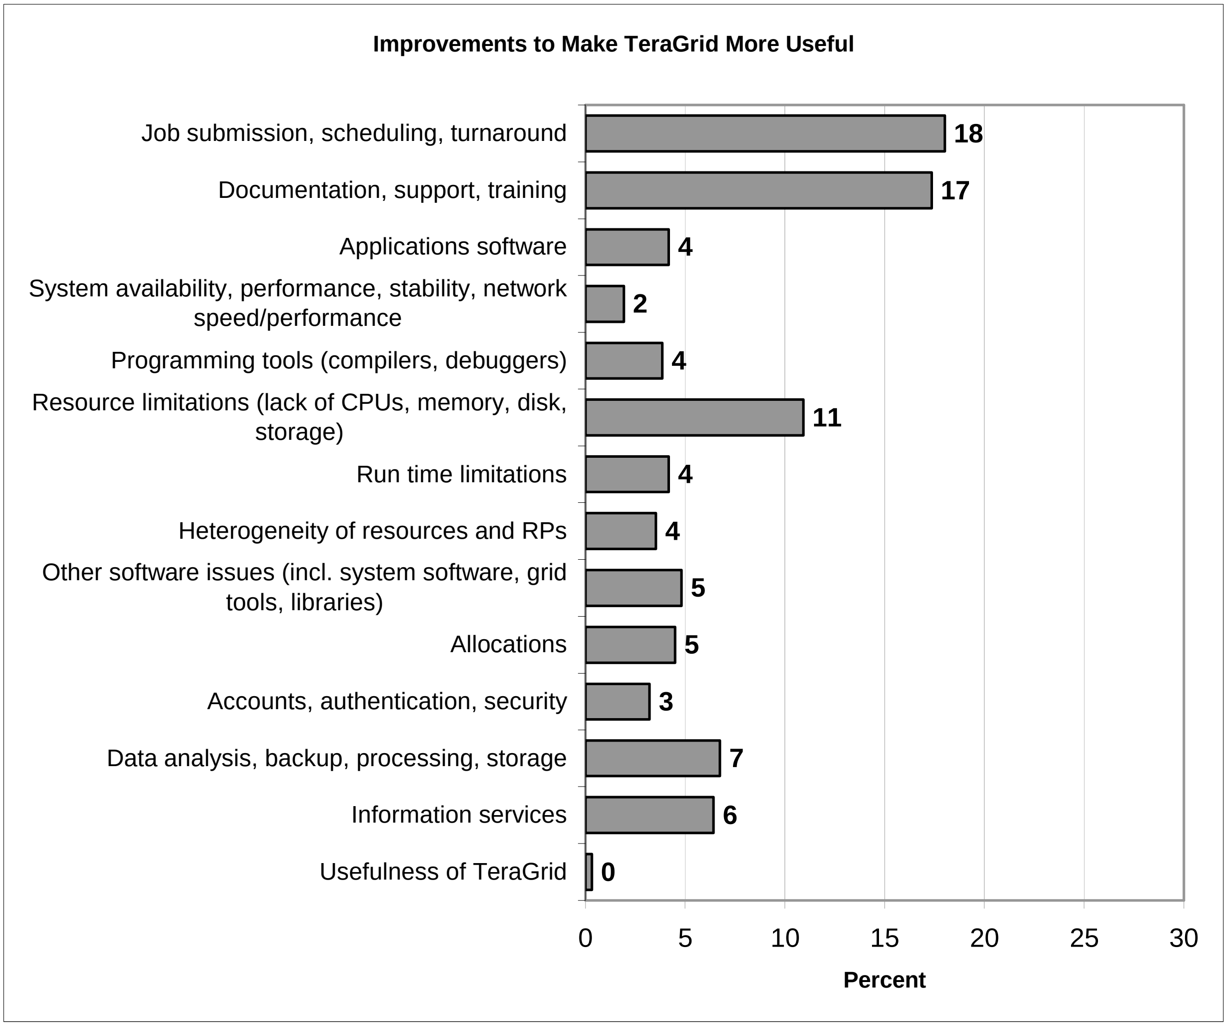

Figure 8[11]

The responses to open-ended questions indicated that some individuals were responding to use of TeraGrid as a grid facility, whereas comments from others indicated they did not consider themselves to be TeraGrid users. In terms of the latter, respondents noted that they did not perceive a need for TeraGrid in their current research, or they did not understand what it could do for them, especially in terms of increasing their productivity.

[11] Note that the sequence of categories in Figure 8 matches that in Figure 7. Percentages here are calculated relative to an N of 311; however, 40% of respondents did not answer this question.

5. Discussion and Practical Implications

This survey's findings help to characterize TeraGrid users and factors that affect their patterns of use. In addition, they identify relevant relationships between usage patterns and users' satisfaction. This information should help TeraGrid design or revise programs to better meet the needs of various types of users.

5.1 The TeraGrid User Population

Survey responses suggest that the TeraGrid population is predominantly male and well educated. About half have received their highest degrees in the last 10 years, and most are affiliated with research universities. Graduate students, postdocs, and research staff comprise about half of those who completed the survey, and the other half is faculty. Since TeraGrid collects limited demographic data beyond institutional affiliation, it is not possible for us to know if these findings are reflective of the TeraGrid population. Based on what we know about the gender makeup of the fields that are the predominant users of TeraGrid and the environment in which TeraGrid is used, it is not surprising that most users are male and highly educated. Although our hypothesis regarding a positive relationship between self-developed codes and ease of use was not supported, there were strong associations between discipline and the nature of the software codes used. The implications of this are that those who use third-party codes have a lower barrier to entry, but in the future they may have greater issues with portability because they cannot update or improve their codes on their own. Conversely, those in computer science or astronomy have a higher barrier to entry, but they may have more flexibility at a future point in time. Others have observed, though, that the size of community codes can make them difficult to port (Graham et al., 2005). Regardless, the types of codes used will affect the appropriate strategy for helping users adapt codes to more capable resources, for example.

Almost two-thirds of those surveyed have been using supercomputers in their research for three or more years. However, a small proportion have less than one year of experience, which may have implications for documentation, training, and/or support, especially if there is a regular stream of individuals with little experience becoming new TeraGrid users each year. The majority of respondents also find supercomputers to be essential to their research. The percentage of respondents who indicated that they use grid computing capabilities is surprising based on TeraGrid data on usage modes and on our interview data; both which show that grid computing as we defined it is not so common. Based on results from the interview portion of this study, respondents likely have different interpretations of what constitutes grid computing. Alternately, those who responded to the survey may be more likely than others in the sample to use grid computing, although it is not possible to measure this based on the data.

5.2 Nature of Research

Although many respondents indicated that their fields are highly competitive, an even greater number strongly perceive collaboration as important or necessary. These results may point to opportunities for new TeraGrid services or enhancements to existing ones. For example, our interviews indicated that those whose primary research method is simulation are interested in collaborations with experimentalists. Perhaps there is a role for TeraGrid to play in bringing users together and in helping users to support each other as the TeraGrid user population grows.

5.3 User Support and Satisfaction

Responses to our questions on TeraGrid use revealed patterns that may help guide future plans for user support services. TeraGrid is widely perceived as useful. Respondents are generally positive, but less so, about the ease of use and availability of facilitating conditions for their use of TeraGrid; slightly fewer than half indicated that TeraGrid is easy to use. These findings may reflect a lack of clarity about what TeraGrid is due to its newness and to users' propensity to identify themselves with the particular resource provider sites that they use rather than with the entire TeraGrid infrastructure. These factors may have made it difficult for some respondents to answer questions about their use of TeraGrid. In terms of ease of use, it is also the case that supercomputers in general are not easy to use. Nevertheless, respondents do use TeraGrid on a regular basis, but more frequent use is associated with larger allocation size as one would expect. As Davis and his colleagues (1989) noted, "Users may be willing to tolerate a difficult interface in order to access functionality that is very important, while no amount of ease of usefulness will be able to compensate for a system that doesn't do a useful task." Still, the findings from all data collected in the study (survey, interviews, user workshop) show that even frequent users would benefit from a system that is easier to use.

Users also make use of TeraGrid support services, and overall this support is viewed positively. Those who have used TeraGrid support once or not at all are more likely to be neutral about its quality. Frequent users of TeraGrid also tend to use support services more frequently, and these more frequent users tend to be the most satisfied with support. Frequency of use and high satisfaction with TeraGrid support is also positively associated with perceptions of usefulness, ease of use, and facilitating conditions for using TeraGrid. These relationships suggest several practical implications for supporting users. For one, a user's first contact with support may be crucial. Second, for those who use TeraGrid infrequently (e.g., quarterly), TeraGrid might create a list of the five or ten things that most people forget or provide customized "startup packages" to make it easier for these users to get started or to remind them how to use the system. The helpdesk might also maintain a user database so that they are more familiar with the issues facing those who call for support, based on these general usage patterns.[12] For instance, they could usefully know the user's allocation, research field, past issues, and codes used.[13]

Curiously, those who perceive themselves to have less experience relative to others in their field also contact TeraGrid support less frequently. Because we cannot identify causality, this

[12] TeraGrid helpdesk staff have tools available to query the central database to find out what projects users are associated with and where they have accounts.
[13] Knowing the general type of codes used would also help TeraGrid find appropriate people to test new grid software and hardware.

raises several questions: Do those who consider themselves to have more experience feel that they have learned more as a result of talking to TeraGrid support? Are these users more confident about contacting TeraGrid support whereas less experienced users feel intimidated about doing so? Or is it simply that the self-identified “less experienced” are those with smaller allocations and therefore have had less of an occasion to contact support?

The open-ended responses identifying barriers and improvements reinforce some issues raised elsewhere in the survey and point to useful avenues for further data collection. Concerns with job submission, scheduling, and turnaround could be usefully addressed by those developing and enhancing the TeraGrid resources and shared infrastructure. Issues with documentation, support, and training echo the responses from those who are less frequent users of TeraGrid and should be further explored to identify whether these perceptions are a cause of or a consequence of infrequent use. (Although allocations put a hard boundary on frequency of use, variability in the size of job submissions means that frequency of use is not solely determined by allocation size and therefore might be influenced by improvements to the system.) One surprising result is the difference in the order of barriers and improvements by frequency of mention. One possible explanation is that users may believe that certain improvements can be more readily addressed by the TeraGrid staff. For example, even though applications software is the third most popular barrier, it is not a priority for improvement; however, it may be that applications software is primarily developed outside the scope of the TeraGrid resource providers. Conversely, resource limitations—seen as a less pressing barrier—are more within the scope of what TeraGrid (or its funding sources) could improve. Likewise, data analysis, backup, processing, and storage along with information services are not top areas of improvement, but they appear to be much more important than their relative standing among the barriers would suggest, perhaps because these are also system limitations more easily controlled by the TeraGrid resource providers.

The positive associations between frequency of use and so many key variables related to satisfaction with TeraGrid suggest that if TeraGrid helps infrequent users feel more like frequent users, they will likely improve the overall satisfaction of the population of users. Because allocations are a limited resource, giving people greater access to resources is not a viable option. However, the provision of additional documentation and support systems targeted at the new or occasional users could bridge these differences. Differences may also be attributable to inadequate knowledge about TeraGrid and what it provides them. Simply increasing awareness of the larger system and the role of individual resource providers might help users tap into the broader knowledge and support base that is available to them already.

In sum, the population of TeraGrid users is generally satisfied with TeraGrid’s services and support, but there is room for improvement, particularly in support of those who—due to allocation limitations, unfamiliarity, or perceived barriers—use TeraGrid less frequently. By creating an experience for these infrequent users that more closely matches the experience of frequent—and satisfied—users, TeraGrid can further improve the perceived quality of its offerings.

6. Limitations

This study was necessarily exploratory due to the absence of prior work on which to base firm hypotheses regarding factors that affect the needs of TeraGrid users, particularly since we expected that differences between users might be hard to detect because of their common need for TeraGrid and other high-performance computing resources. The relationship between frequency of use and multi-scale items measuring usefulness, ease of use and facilitating conditions as well as other variables such as allocation size, were useful, however, in teasing apart the needs of users. Scales consisting of only two items typically do not have good reliability. This may have contributed to the problems with the scales for personal innovativeness and social influence.

We anticipated that many respondents would be confused about why they had been identified as TeraGrid users. Although we attempted to mitigate this problem, there is evidence, particularly from responses to open-ended questions and to the questions on grid computing that users have different perceptions about what TeraGrid *is*, and this is likely to have affected their responses. The degree to which this limits the survey's findings depends on how TeraGrid perceives itself in terms of its vision, mission, and goals (see Part I of this report).

Finally, the lack of a suitable control population hindered our ability to compare TeraGrid users with users of other HPC facilities. This limitation is also a future research opportunity.

7. Future Work

To our knowledge, this is the first survey to study the characteristics and needs of users of high-performance computing. As such, it provides baseline data on one such population and helps to explain factors that affect their use and satisfaction. It would be useful to compare users of TeraGrid with users of other grids in the U.S. and elsewhere such as the Open Science Grid, the United Kingdom's National Grid Service, or the Distributed European Infrastructure for Supercomputing Applications. It would also be informative to compare users of TeraGrid with those who use other U.S. facilities of such as DOE or NASA. Investigations such as these would generate information on similarities and differences in user characteristics and behavior and needs across multiple types of environments that would be helpful, for example, in international collaborations, and in developing common strategies for user support, training, etc.

Second, as new communities begin to make use of TeraGrid's resources and services, it will be important to track changes in the characteristics, behavior, and needs of users. Periodic surveys would help TeraGrid to adjust plans and approaches as required to serve new users. A challenge for future work will be to find ways to assess the needs of those who will access TeraGrid resources through science gateways, and thus, who may not be aware that they are using TeraGrid.

Acknowledgments

We thank David Hart, San Diego Supercomputer Center, for answering our many questions regarding the TeraGrid central database and for providing the data on which the survey samples were drawn. Comments on the survey were received from several TeraGrid users, from colleagues and students at the UM School of Information, from survey experts at other institutions, and from TeraGrid personnel.

This report is based upon work supported by the National Science Foundation under Grant No. OCI 0603525 to the University of Michigan. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

Anonymous. (2000). The ongoing evolution of scientific computing: Talking with Sidney Karin. *Ubiquity* 1(18).
Retrieved March 30, 2008, from http://www.acm.org/ubiquity/interviews/s_karin_1.html

Birnholtz, J. P. (2007). When do researchers collaborate? Toward a model of collaboration propensity. *Journal of the American Society for Information Science and* Technology 58(14), 2226-2239.

Birnholtz, J.P., Horn, D.B., Finholt, T.A., & Bae, S.J. (2003). The effects of cash, electronic, and paper gift certificates as respondent incentives for a web-based survey of a technologically sophisticated sample. School of Information, University of Michigan. Retrieved March 30, 2008, from http://www.crew.umich.edu/publications/03-02.pdf.

Buetow, K. H. (2005). Cyberinfrastructure: Empowering a "Third Way" in biomedical research. *Science*, 308(5723), 821-824.

Church, A. (1993). Estimating the effect of incentives on mail survey response rates: A meta-analysis. *Public Opinion Quarterly* 57(1), 62-79.

Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika, 16*, 297-334.

Davis, F. D., Bagozzi, R. P., and Warshaw, P. R. (1989). Use acceptance of computer technology: A comparison of two theoretical perspectives. *Management Science* 35(8), 982-1002.

Elzen, B., & MacKenzie, D. (1994). The social limits of speed: The development and use of supercomputers. *IEEE Annals of the History of Supercomputing* 16(1), 46-61.

Galison, P. (1997). *Image and logic: A material culture of microphysics*. Chicago: University of Chicago Press.

Graham, S. L., Snir, M., & Patterson, C. A. (Eds.). (2005). *Getting up to speed: The future of supercomputing*. National Academies Press, Washington, DC.

Groves, R. M., Fowler, F. J. Jr., Coupers, M. P., Lepkowski, J. M., Singer, E., and Tourangeau, R. (2004). *Survey methodology*. Hoboken, NJ: Wiley.

Hesse, B., Sproull, L., Kiesler, S., & Walsh, J. (1993). Returns to science: Computer networks in oceanography. *Communications of the ACM*, 36(8), 90-101.

Humphreys, P. (1991). Computer simulations. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1990, vol. 2, 497-506.

Kaminer, N., & Braunstein, Y. M. (1998). Bibliometric analysis of the impact of Internet use on scholarly productivity. *Journal of the American Society for Information Science*, 49(8), 720-730.

Lawrence, K. A., & Zimmerman, A. (2007). *TeraGrid Planning Process Report: August 2007 User Workshops*. Collaboratory for Research on Electronic Work, School of Information, University of Michigan, Ann Arbor, MI. Retrieved March 30, 2008, from http://www.teragridfuture.org/system/files/TeraGrid+User+Workshops+Final+Report.pdf

Lu, J., Yeao, J. E., & Yu, C-S. (2005) Personal innovativeness, social influences, and adoption of wireless Internet services via mobile technology. *Strategic Information Systems* 14, 245-268.

Nunnally, J. C. (1978). *Psychometric theory* (2nd ed.). New York: McGraw-Hill.

Rogers, J. D. (1998). Internetworking and the politics of science: NSFNET in Internet history. *Information Society* 14(3), 213-228.

Schneck, P. B. 91990). Supercomputers. *Annual Review of Computer Science* 4, 13-36.

Sismondo, S. (1999). Models, simulations, and their objects. *Science in Context* 12(2), 247-260.

Tourangeau, R., Rips, L.J., & Rasinski, K. (2000). *The psychology of survey response*. Cambridge: Cambridge University Press

Venkatesh, V., et al. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly* 27(3), 425-478.

Walsh, J. P., Kucker, S., Maloney, N. G., & Gabbay, S. (2000). Connecting minds: Computer-mediated communication and scientific work. *Journal of the American Society of Information Science and Technology* 51(14), 1295-1305.

Walsh, J. P., & Roselle, A. (1999). Computer networks and the virtual college. *Science, Technology and Industry Review* 24, 49-78.

Appendix A: Cover Letter from UM-SI Research Team

November 29, 2006

ID Number
Name
Address
City State

Dear Colleague:

You have been selected to participate in a study of TeraGrid users funded by the National Science Foundation and being conducted by the University of Michigan (UM). TeraGrid is a distributed infrastructure that integrates high-performance resources across nine resource provider facilities. These resource providers have made some or all of their computing systems available on the TeraGrid. *You were identified as a TeraGrid user because you use one or more of the resources available on the TeraGrid*.

This study is part of a larger effort to better understand how to support the computing needs of scientific and engineering research communities. You may not directly benefit from this study; however, the results of this effort will help drive the development of future tools and other technologies to support research in your field and others. It is not possible for us to understand the relevant factors without responses from individuals engaged in a range of activities, which means that *your response is very important to us.*

As a token of our appreciation for your efforts, please accept the enclosed cash gift. In the next week, you will receive an email message from us that contains a link to the survey. Or, you can take the survey right now by going to **www.teragrid-survey.org** and entering the ***Survey ID#.*** If you choose to participate, you can also sign up to receive a summary of the results via e-mail. **If possible, please complete the online survey by December 20, 2006.** *It is very important that you do not pass the survey onto another individual since respondents have been scientifically selected.*

Your participation in completing the survey is voluntary. You may skip questions, and you are free to withdraw at any point. Your responses will be used for research purposes only and will be kept in secure locations at the UM. Only primary members of the research team at the UM will have access to these data. The information you provide in the survey on the website will be kept confidential. Furthermore, all personal information will be presented only in an aggregate form in reports and publications. Individual responses will not be identifiable. If you have any questions regarding your rights as a participant in this research, please contact: Institutional Review Board, 540 E. Liberty Street, Suite 202, Ann Arbor, MI 48104-2210, Tel: 734-936-0933, email: irbhsbs@umich.edu.

Thank you in advance for taking time to complete this important survey. If you have additional questions or concerns, please contact us via e-mail at teragrid-survey@umich.edu or by calling 734-764-1865.

Sincerely,

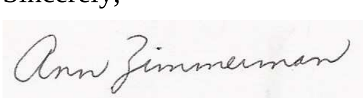

Ann Zimmerman, PhD
Research Investigator

Thomas A. Finholt, PhD
Director, Collaboratory for Research on Electronic Work

Appendix B: Letter from TeraGrid Director

November 27, 2006

Dear TeraGrid User:

I want to bring to your attention a study being conducted by Ann Zimmerman and Tom Finholt, researchers at the University of Michigan's School of Information. They are analyzing how high performance computing in general—and TeraGrid more specifically—can best support your research needs. This survey will be administered to a group of approximately 1,000 scientists and engineers across the United States. I strongly encourage you to take 15-20 minutes to complete this survey.

This study is supported by the National Science Foundation through a grant to the University of Michigan's School of Information and is part of a larger study to better understand researchers' requirements for high-performance computing. The results of this survey will directly influence these efforts and will have important implications for users of TeraGrid as well.

As the Director of TeraGrid, I am excited about the potential of this study to provide information that will enable scientists and engineers to conduct research in new ways. As TeraGrid continues to grow and evolve, we want to make sure that the services and resources that we are developing and providing are best suited to meet the needs of the larger research community. To learn more about the TeraGrid project, visit *teragrid.org*.

Again, I hope that you will take some time to participate in this important study.

Sincerely yours,

Charlie Catlett
Director, TeraGrid, Grid Infrastructure Group
Senior Fellow Computation Institute
University of Chicago/Argonne National Laboratory

TeraGrid is coordinated through the Grid Infrastructure Group (GIG) at the University of Chicago, working in partnership with nine Resource Provider sites: Indiana University, Oak Ridge National Laboratory, National Center for Supercomputing Applications, Pittsburgh Supercomputing Center, Purdue University, San Diego Supercomputer Center, Texas Advanced Computing Center, University of Chicago/Argonne National Laboratory, and the National Center for Atmospheric Research

Appendix C: Survey Instrument

Welcome

You have been selected to participate in a study of TeraGrid users funded by the National Science Foundation and being conducted by the University of Michigan (UM). This survey is part of a larger effort to better understand how to support the computing needs of scientific and engineering research communities. We greatly appreciate your taking 15-20 minutes to share your opinions with us.

TeraGrid is a distributed infrastructure that integrates high performance resources across nine resource provider facilities. *You were identified as a TeraGrid user because you use one or more of the resources available on the TeraGrid*.

Your participation in completing this survey is voluntary. You may skip questions, and you are free to withdraw at any point. Your responses will be used for research purposes only and will be kept in secure locations at the UM. Only primary members of the UM research team will have access to these data. Furthermore, any personal information will be presented only in an aggregate form in reports and publications. Individual responses will not be identifiable. If you have any questions regarding your rights as a participant in this research, please contact:

Institutional Review Board
540 East Liberty Street, Suite 202
Ann Arbor, MI 48104-2210
Tel: 734-936-0933
E-mail: *irbhsbs@umich.edu*

Field of Research

1. Please select the area of science that *most closely* represents your research

□ Astronomy
□ Biological Science
□ Chemistry
□ Computer Science
□ Engineering
□ Geoscience
□ Mathematics
□ Materials Science
□ Physics
□ Social Science
□ Other (please specify)

2. Please select the category that *best* describes your research approach.

□ Theoretical/Analytical
□ Experimental/Observational
□ Simulation/Modeling
□ Other (please specify)

General Technology Use

3. Which of the following computer resources do you use? Please check all that apply.

□ local workstation
□ local cluster (less than 64 processors)
□ local cluster (65-128 processors)

Please indicate the extent to which you agree with the following statements *with regard to the use of information technology to support your research.* For the purpose of this survey, information technology is defined as the use of computers to process, store, retrieve, and transmit information.

4. Among my peers, I am usually the first to try out new information technology.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

5. In general, I am hesitant to experiment with new information technology.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

Nature of Research Field

Please indicate the extent to which you agree with the following statements.

6. Researchers in my field typically work alone.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

7. In my field, most major research advancements are made by individuals working alone.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

8. The competition for prizes or widespread recognition in my field is intense.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

9. In addition to my collaborators, I feel safe in discussing my current work with other persons doing similar work.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

10. I am concerned that the results of my current research might be "scooped" by other researchers working on similar problems.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

Research Collaboration

Please indicate the extent to which you agree with the following statements.

11. Collaboration is necessary in my field.
□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

12. Collaboration is useful in solving research problems that are of interest to me.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

13. Collaboration allows me to access people with expertise that are helpful to me.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

14. Collaboration allows me to access resources (e.g., computers, instruments, data) that I could not otherwise use.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

15. Other researchers in my field who do collaborative work are successful in their research careers.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

Supercomputer Use

This section asks about your use of supercomputers. For the purpose of this survey, supercomputers refer to those computing systems (hardware, software, and applications) that, at a given point in time, provide close to the best achievable sustainable performance on demanding computational problems. In answering the questions in this section, please consider all supercomputers you have used and not only those available on the TeraGrid.

16. When did you first begin to use supercomputers in your research?

□ Less than 1 year ago
□ 1-2 years ago
□ 3-4 years ago
□ 5+ years ago
□ Not applicable

17. In comparison with others in your field, how would you rate your experience in using supercomputers to achieve desired outcomes in your research?

□ Not experienced at all
□ Somewhat experienced
□ Experienced
□ Very Experienced
□ Extremely Experienced
□ Don't know

18. Please indicate the extent to which you agree with the following statement.

Supercomputers are necessary to answer research questions of interest to me.
□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

19. Please select the response that best describes the codes that are *most important* to the achievement of your research goals.

□ Third-party software (e.g., commercial software, community codes)
□ Third-party software, augmented with routines or libraries written by you or your group
□ Codes developed by you or your group, augmented with third-party routines or libraries
□ Codes developed entirely by you or your group
□ Not applicable
□ Don't know

20. Besides TeraGrid, do you currently have access to supercomputer resources through any of the following? Please check *all* that apply.

□ Department of Defense (DOD)
□ Department of Energy (DOE)
□ National Aeronautics and Space Administration (NASA)
□ National Center for Atmospheric Research (NCAR)
□ Open Science Grid (OSG)
□ State or regional supercomputer center
□ Your institution's supercomputer system or facility
□ Not applicable
□ Don't know
□ Other (please specify)

21. Which of the following supercomputer resources do you need for your research? Please check *all* that apply.

□ Computer systems
□ Data collections
□ Data management tools
□ Documentation
□ Persistent online storage
□ Training and tutorials
□ User services support
□ Visualization software
□ Visualization servers
□ Other (please specify)

Grid Computing

This section asks about your use of grid computing. For the purpose of this survey, *grid computing is defined as a hardware and software infrastructure that enables users to apply the resources of many computers to a single problem*.

22. Do you currently use grid computing capabilities?

□ Yes (continue with questions 23-27)
□ No (continue with questions 28-30)
□ Don't know (continue with questions 28-30)

Grid Computing Continued

For the purpose of this survey, *grid computing is defined as a hardware and software infrastructure that enables users to apply the resources of many computers to a single problem.*

23. Have you used TeraGrid resources in your grid computing work?

□ Yes
□ No
□ Don't know

24. How would you describe your use of grid services or software (e.g., Globus Toolkit, Condor, GridShell) in the past year?

□ Have used grid tools in production runs
□ Have started to experiment with grid software in test runs
□ Have heard about and investigated capabilities offered by grid services or software
□ Don’t know
□ Not applicable
□ Other (please specify)

Please indicate the extent to which you agree with the following statements with regard to grid computing.

25. I expect my usage of grid computing to continue in the future.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

26. Overall, I find grid computing useful in my research.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

27. It is easy for me to become skillful at grid computing.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

Grid Computing Continued

For the purpose of this survey, *grid computing is defined as a hardware and software infrastructure that enables users to apply the resources of many computers to a single problem.*

28. I would find grid computing useful in my research.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree
□ Don’t know

29. I would find it easy to become skillful at grid computing.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree
□ Don't know

30. During the next 6 months, I plan to experiment with grid computing in my research.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree
□ Don't know

Use of TeraGrid

This section seeks your general opinions about your use of TeraGrid. In the sections that follow this one, you will be asked for more specific information about your use of TeraGrid. For the purpose of this survey, *TeraGrid use is defined as the utilization of any TeraGrid compute resources and non-compute resources (e.g. data storage and management, servers, networks, visualization, and tools and software).*

31. Please indicate the extent to which you agree with the following statements with regard to TeraGrid.

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree	Not applicable
Use of TeraGrid enables me to accomplish research tasks more quickly.						
Use of TeraGrid increases my research productivity.						
Use of TeraGrid is important to help me achieve my career goals.						
Overall, I find TeraGrid useful in my research.						

32. Please indicate the extent to which you agree with the following statements with regard to TeraGrid.

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree	Not applicable
It is easy for me to become skillful at using TeraGrid.						
It is easy for me to get help from TeraGrid support when I need it.						
I find it difficult to get TeraGrid to do what I want it to do.						
Overall, I find TeraGrid easy to use.						

33. Please indicate the extent to which you agree with the following statements with regard to TeraGrid.

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree	Not applicable
A specific person (or group) is available to assist me in using TeraGrid.						
Up-to-date documentation on the software that I need to accomplish my work on TeraGrid is available to me.						
Guidance is available to me in the selection of TeraGrid resources.						
Other people I work with think I should use TeraGrid.						
Researchers in my field who use TeraGrid have more prestige than those who do not use TeraGrid.						

34. On average, how frequently have you used TeraGrid in the past year?

□ Daily
□ Weekly
□ Monthly
□ Quarterly
□ Twice
□ Once
□ Never

35. Approximately how many times have you contacted TeraGrid support in the past year?

□ None
□ Once
□ 2-5 times
□ 6-10 times
□ More than 10 times

36. Please rate your satisfaction level with the TeraGrid support received in the past year.

□Extremely Satisfied
□ Satisfied
□ Neutral
□ Dissatisfied
□ Extremely Dissatisfied
□ Not applicable

37. Are you affiliated with a Science Gateway project that has a TeraGrid allocation?

□ Yes
□ No
□ Don’t know

38. Are you currently or have you previously been affiliated with a TeraGrid ASTA (Advanced Support for TeraGrid Applications) project?

□ Yes
□ No
□ Don't know

Please indicate the extent to which you agree with the following statement.

39. I identify myself as a TeraGrid user.

□ Strongly Agree
□ Agree
□ Neutral
□ Disagree
□ Strongly Disagree

40. What are the two ***most significant*** barriers you have encountered in your use of TeraGrid? (Each line is limited to 175 characters.

41. What would make TeraGrid more useful to you? Please list the two things that are ***most important*** to you (e.g. services, functions, tools, policies, etc.) (Each line is limited to 175 characters.)

Demographic Information

42. Which of the following *best* describes your current professional status?

□ Assistant Professor
□ Associate Professor
□ Professor
□ Junior Research Scientist
□ Senior Research Scientist
□ Research Programmer
□ Research Assistant
□ Postdoc
□ Graduate Student
□ Undergraduate Student
□ Other (please specify)

43. What kind of institution are you affiliated with?

□ Research university (PhD granting institution)
□ Teaching university or college
□ Government agency
□ Nonprofit organization
□ Commercial business or service provider
□ Other (please specify)

44. What is the name of the institution you are affiliated with?

45. What is the highest academic degree you have obtained?

□ Bachelor's
□ Master's
□ PhD or equivalent
□ Other (please specify)

46. In what year did you obtain your highest degree?

47. What is your gender?

□ Male
□ Female

48. In what year were you born?

Appendix D: Responses to Open-Ended Questions

The survey included two open-ended questions. One asked respondents to the name the most significant barriers they had encountered to their use of TeraGrid (Q40), and the other requested suggestions for things that would make TeraGrid more useful to them (Q41). Responses were analyzed and grouped according to fourteen types of barriers and the same number of categories of improvements. A sample of responses in each coded category to the question regarding barriers is given below. This is followed by things that respondents noted would make TeraGrid more useful to them as regards the top 3 barriers to TeraGrid: job turnaround time; documentation, support, and training, and applications software. Most of the suggested improvements in each category were opposing statements to the given barriers. For example, a common suggestion to improve lengthy waits in the queue was to find ways to shorten queue times. We do not list this type of response below. Instead, we focus on more substantive improvements made by respondents.

Job submission, scheduling, and turnaround

Barriers

Long queue times!!!!!!!!!!!!!!!!!!!!!!

Queue is slow/clogged especially close to major holidays

After I submit my job I have to wait too much time for my job to run. For example, my job takes just a day, but I have to wait more than one week after submitting my job

Queue backups: This not only hinders production but is also extremely detrimental to model development.

Since October 2006, queue waits for 96-hour jobs have been 14-21 days!

Jobs stay queued up for a long time (weeks at times) mainly because TOOOOO many processors are perpetually engaged in executing priority jobs

I need to run many small, relatively short simulations (i.e. 8-16 processors per job), but there doesn't seem to be a queuing system on NCSA appropriate for this.

Long queue times/not enough processors

Queuing system will halt sometimes

Lack of a real test queue

Too many single processor/serial jobs slowing throughput on the system

Priority given to larger jobs versus small ones

Lack of support for medium-sized (~64-processor) jobs

Limitation in number of jobs that can be submitted to queue (Machines are set up for large parallel jobs rather than many independent runs, which I need.)

Queue penalty for requesting long run times

The fact that I have to be in the queue for a long time and after that send my work several times to get it all done. In my cluster I run the whole work once; in TeraGrid I run the job in pieces.

Suggestions

Knowing how the batch system works

Automated co-scheduling

Some way to submit jobs to a single queue for execution on multiple systems (of same architecture).

A tool that identifies the shortest queue in my available TeraGrid resources

Jobs terminated abruptly due to system crashes should have highest priority when resubmitted.

Documentation, support, training

Barriers

Had no idea what I was doing

I am extremely busy and have little time to learn necessary details.

Enough dedicated time (I teach at a PUI) to access and learn how to use TeraGrid

Finding the time to make good use of the resources

Getting started−different environments on different systems

Understanding the complexity of TeraGrid

Learning commands

Getting my graduate students to learn the ropes

New students working on my research team must train themselves on how to use the system

Going to help@teragrid.org vs. help at the local computing center

Cannot ask questions face-to-face

Lack of real-time technical support

Lack of homogeneity in user support with some very good and some not so good advice

It's hard to find the appropriate cluster I want to use

Figuring out what resources were available

The lack of on-line documentation

Lack of up-to-date online information regarding installed software, libraries, and compilation

Poor online help

Simpler web site needed

Navigating through the web pages to figure out how to set up a passkey to access IU systems

Learning how batch jobs function on the clusters

Learning curve associated with Grid middleware (Globus, etc.)

Apparent complexity of system software has made me reluctant to experiment with Grid computing

It's hard to figure out how to send jobs to the grid to run anywhere and not just at a particular site.

Support for third-party products (i.e., Intel compiler)

Parallel computing knowledge

My own programming abilities

Suggestions

An introduction to TeraGrid (how to use it, how to use it efficiently, etc.)

Training seminar/courses/online tutorials on using the system

Tutorials on how to use the software tools installed on each site (i.e., step-by-step instructions and working example code

Better documentation on how to use the third party software installed on TeraGrid resources. For example, location of scripts to run these programs if they exist, scratch space location.

More up-to-date documentation on Gridshell

More up-to-date web pages describing new machines, incorporating user experience

If there were an interface where I could submit job exactly as I do locally, but to a port that is 'TERAGRID' or something straightforward

Example code

More case studies on scientific software

Public folder with examples (can also be from users of TeraGrid)

Applications software

Barriers

Initial software set-up

Porting software to specific TeraGrid machines

Front end vectorization of our code

Getting software working correctly

Parallelization of in-house codes

Modifying our home-grown simulation tools to work successfully in a grid environment

Difficult to use CHARMM on TeraGrid

The difficulty to have custom versions of NAMD running

The need to rewrite code for checkpointing

Commercial software (for computational simulations) is not supporting the kind of operating systems of some supercomputers.

Availability of commercial software

Availability of compiled 'standard' software

Lack of updates as to new software packages

Trying to figure out how to run third-party software on TeraGrid

Poor third-party software management, poor maintenance of third-party software

Lack of easy to use parallel MATLAB interface

Parallelization of third-party codes

Getting an open-source application to run successfully

Software scalability

Suggestions

A service to help adapt my existing scripts for use on TeraGrid

Support and provide graphical interface tools for running community codes

Short job scripts for different software (NSchem, GROMACS, NAMD, Gaussian)

System availability, performance, and stability and network speed and performance

System stability issues

BlueGene hardware and stability issues

Platform instability

Intermittent machine reliability

Cluster downtime

gpfs reliability on the IA-64 cluster

Unavailability of certain machines due to maintenance all the time

Interruption to work due to computer crashes

Machine crashes have resulted in pushing jobs all the way down in queues.

If a job is not finished, everything is lost.

Network speed is not high enough.

Remote network connection speed

Insufficient access to TG networks and supporting hardware to drive high-volume data transfers

Programming tools (compilers, debuggers)

Compiler differences on machine.

Not the best compilers in the more appropriate nodes

The compiler and linker guidance for some important software, e.g. NAMD, CHARMM are not available at some TERAGRID slots.

I often compile with g77 my codes. TeraGrid, I think, doesn't support it.
Manuals on compilers

Some differences between the compilers and the internal representation of numbers in TeraGrid respect to other computers used by me that implied different results in computations.

Very long compile times on login nodes

Remote debugging of our parallel codes is hard.

The difficulties in debugging my code on the system

Resource limitations (lack of CPUs, memory, disk, storage)

More CPUs

Memory limitations of each processor

Lack of accessible nodes with large enough shared memory

The 1 GB memory barrier does not allow me to debug large programs quickly

My jobs require hundreds of GB of memory, hundreds of processors, and terabytes of disk space

The difficulty in running a large problem with a limited home directory

The availability of more high-SMP machines would be useful

Storage is very limited

Small home disk quota

Limited hard disk space

Run time limitations

Time limit

The 18 hours time limit

Each job only can run 12 hours but the initialization needs 1 hour for the large dataset.

Time limit per job was/is set to 18 hours. This reduced the range of sensible ab-initio calculations tremendously.

The walltime limit in the script is not quite the real walltime limit of the run.

After the job is running, I cannot extend the time limit from 48 to 100 hours

Getting ability to run continuously for extended time− weeks

The unavailability or difficulty in getting access to long-time (say 5 days) jobs

18 hour limit for each submission, when longer (50+ hrs) simulation is needed

Modifying software to accommodate runtime limits

Heterogeneity of resources and RPs

Different operating systems and architectures

The heterogeneous environments of the various centers. For example: different queues, different policies, different submission software (e.g., PBS, mpirun)

Each system has a different procedure for submitting jobs and for storing data in long-term storage. It would be easier to learn and use if all were the same.

Managing my knowledge of all the queuing systems, system environments unique to each machine

Learning unique aspects of each system

The servers do not have a unified user interface. I have to learn how to use it each time I switch to a new server.

Non-uniform rate structures

Different login procedures at different sites

Different library locations/versions on each machine

Lack of conformity between platforms event to the extent of accounts

Inconsistent development environment

Maintenance of different archival areas

It is unclear how data is shared (if at all) between TeraGrid sites

Difficulty in making sense of available resources. (TeraGrid is loosely organized; making it hard to become skilled on more than one available machine.)

The software stack available: I'm using a large package and to install it on different machines is too much effort.

Constantly changing APIs and incompatible versions causing extra work just on implementing my scripts and preventing me from spending that time on research

Other software issues (including system software, grid tools, libraries)

It is hard to find the location of the library that I need. (ex. parallel HDF5)

Unsupported libraries or tools

gridftp protocols are pretty opaque and have syntax that is too arcane

Tools for inter-site work and transfers are overly cryptic.

Lack of stable, robust and standard grid computing software

Allocations

Bureaucratic overhead

Limited allocation

The complicated and slow allocation process

Applying for compute time is a time-consuming process.

Allocation process is slow and demanding.

Writing proposals -- lots and lots of work to actually write these

Others' hesitation in applying for allocations

My grant allocation was 1/3 less than what I proposed.

Having proposals reviewed once for science, then a second time, by non-experts, for computing allocation.

CPU allocation limited or denied due to lack of benchmarking demonstration

It's not easy to get the Medium Resource Allocations after I used up the Development Allocations.

Allocation of more hours for researchers who need them would improve efficient use of the TeraGrid resources.

Access to TeraGrid should be open to everyone who needs it.

Accounts, authentication, security

Multiple login names

Complex login management procedures; difficult to reset password once the original expires

Username differences

Don't like the portal for adding new users

Understanding how to obtain passwords, log in, and configure accounts

Getting username and password through mail

Passwords that were mailed to me were incorrect

Getting account

Setting accounts

Availability and consistency of account information

Inability to easily track allocation depletion per day on a per RP basis. A historical report should be easy to acquire

Passwords/ssh keys/certificates

Getting certificate

Firewall policies make it difficult to connect TeraGrid resources to non-TeraGrid resources

Security is important, but it is the biggest issue for user to access resources easily. In fact, it is not so important in research areas (personal opinion).

Grid security: ssh and scp work so much better, when you run jobs for months you can't be renewing certificates all the time. Have the people who design this software ever done a big calculation?

Data management, movement, analysis, and storage

Data management

Large amount of time it takes to write data to files, save/get them from storage, etc.

Slow file transfers to/from sites and hard-to-use file transfer capabilities

Inability to transparently move data between TeraGrid sites

Narrow bandwidth for data transfers

Manually transferring data consumes a lot of time

Tansfer of large datasets across multiple sites

Difficulty to visualize results on TeraGrid machines

Limited nodes for post-processing, visualization

Lack of data analysis applications available

Managing workflow on all the machines simulataneously (especially data processing and storage)

Data storage policies: How much storage can I use and for how long?

Persistent storage

Information services (Most of the responses in this category simply said "services.")

Services

Accurate updating on the TeraGrid User Portal

A more detailed and transparent queue status information system

Not so many emails but maybe just a webpage, which gives current stage of different computers (down, up, under maintenance, etc.)

Universal monitoring system

If there are any users occupying most of the slots, I hope the information will be available to other general users

Usefulness of TeraGrid

Not clear need for it in current research

Lack of knowledge of possibilities

I do not see how it would enhance my research productivity or that of my group.

I'm not sure that it would really benefit my research. What I need is a fast computer with lots of memory. It's not clear that my productivity would improve by using distributed resources.

Fear of switching to something different than traditional supercomputing resources

Inertia (things work now - why change?)

If one has 100 processors, but 100 people want to use it, there is no advantage to using TeraGrid.

It is not clear to me why there is a need to have grid resources (with respect to traditional supercomputers).

I use NCSA machines, but know little about TeraGrid. I know that the NCSA machines are part of TeraGrid. That is about it.