Appendix B: Analysis of Web Survey Results

A quantitative analysis of the Web survey is valuable to supplement the more qualitative conclusions presented in Sections 2 and 3.

User and Application Profiles

Comparisons with the Hayes Report, when available, are shown on the right. Where appropriate, each survey result is followed by a brief analysis and conclusions.

5. Principal categorization as a researcher:	Number of Responses	Response Ratio
University (regular or research) faculty	321	45%
University staff	82	12%
Post-Doctoral	79	11%
Graduate Student	99	14%
Undergraduate Student	9	1%
Federal laboratory	66	9%
Private sector	34	5%
Other	21	3%
Total	711	100%

6. Field of Research matched to NSF Divisions or Directorates: *(check all that apply)*	Number of Responses	Response Ratio	Hayes Report (1995)
Physics	190	26%	14%
Chemistry	119	16%	11%
Astronomy	65	9%	7%
Materials Research	64	9%	5%
Mathematical Sciences	67	9%	3%
Earth and Atmospheric Sciences	194	27%	Not included
CISE - Computer Science	136	19%	20%
ENG - Engineering	118	16%	19%
BIO - Biological Sciences	117	16%	6%
GEO - Geosciences	75	10%	4%
SBE - Social, Behavioral, and Economic Sciences	8	1%	1%
Other *(please specify)*:	51	7%	9%

7.	Do you presently use, or anticipate using in the future, digital libraries or federated data repositories?	Number of Responses	Response Ratio
I use them now		175	34%
No, but plan to do so in the future		117	23%
No, and have no plans to do so in the future		101	19%
Unsure or no opinion		127	24%
	Total	520	100%

11.	If your code(s) is/are platform-specific, which of the following platforms do you use?	Number of Responses	Response Ratio
Compaq Alpha/Linux		58	17%
Compaq Alpha/Tru64		75	22%
Cray/UNICOS		118	34%
IBM/AIX		111	32%
Intel/Linux		175	51%
Intel/MS Windows		61	18%
Intel/Solaris		19	6%
Hewlett-Packard/HP-UX		36	10%
PowerPC/Mac OS		29	8%
SGI/IRIX		159	46%
SGI/Linux		38	11%
Sun/Solaris		133	39%
Other *(please specify)*:		37	11%
	Total	718	100%

10.	Is/are your code(s) portable or platform-specific (e.g., must run on an Intel Linux system -- see next question for a sample list of machines)? If some are and some are not (e.g., principal application is portable but the visualization code is not), please so state in the comment line.	Number of Responses	Response Ratio
Portable		511	72%
Platform-specific		113	16%
Mixed *(please comment)*:		145	20%

12.	Is/are your code(s) parallelized? If some are and some are not (e.g., principal application is parallelized but the visualization code is not), please so state in the comment line.	Number of Responses	Response Ratio
Yes		382	54%
No		160	23%
Mixed *(please comment)*:		196	28%

13.	If your application is parallelized, which of the following tools are used? *(select all that apply)*	Number of Responses	Response Ratio
MPI		443	81%
OpenMP		136	25%
Shared memory parallelism via platform-specific directives		113	21%
Parallelism via language constructs (e.g., High-Performance Fortran)		66	12%
Automatic (i.e., compiler-based) parallelization		89	16%
Other *please specify*:		71	13%

14.	Which of the following is/are the primary factor(s) limiting the execution speed of your application(s)? *(Select all that apply.)*	Number of Responses	Response Ratio
CPU speed		506	71%
Cache size		202	28%
Memory bandwidth		290	41%
Number of processors		317	45%
Internal communication network bandwidth (e.g., message passing)		257	36%
External network bandwidth (e.g., Internet2)		51	7%
I/O bandwidth		149	21%
Data archive bandwidth		71	10%
Other		31	4%
Don't know		65	9%

15.	Which of the following is/are the primary factor(s) limiting the size or duration of your runs? *(Select all that apply.)*	Number of Responses	Response Ratio
Cache size		80	12%
Memory size per processor		227	33%
Total memory size		219	32%
Number of processors		300	43%
Disk space availability		143	21%
Queue configuration (e.g., short jobs that are delayed substantially because of large memory requirements)		166	24%
Other		70	10%
Don't know		76	11%

16.	Do you generally perform the analysis of output generated by your application on a machine different from that on which the application runs?	Number of Responses	Response Ratio
Yes		531	74%
No		184	26%
	Total	715	100%

17.	Do you generally perform, or have a need to perform, the analysis of output generated by your application as the application is running?	Number of Responses	Response Ratio
Yes		312	44%
No		405	56%
	Total	717	100%

18.	Please express the resources used for analysis/mining/visualization of output generated by your application as a percentage of the resources required for the application itself.	Number of Responses	Response Ratio
0-25%		471	66%
25-50%		116	16%
50-75%		38	5%
75-100%		17	2%
100-150%		8	1%
150-200%		3	0%
Over 200%		9	1%
Don't know		47	7%
	Total	709	100%

A total of 677 respondents specified the size of their data or output files that are analyzed/mined/visualized. Multi-gigabyte data sets was quite common and 100+ gigabyte sizes were rare. However, two respondents noted that their data sets range in size from between 1-10 terabytes.

2. Have you (or your collaborators or students) used any PACI Center during the past year?	Number of Responses	Response Ratio	Hayes Report (1995)
Yes	389	71%	75%
No	159	29%	25%
Total	548	100%	

3. Which of the following facilities have you used during the past year (check all that apply)?	Number of Responses	Response Ratio	Hayes Report (1995)
NCSA - National Center for Supercomputing Applications	175	34%	31%
PSC - Pittsburgh Supercomputing Center	117	23%	10%
SDSC - San Diego Supercomputer Center	187	36%	26%
NCAR - National Center for Atmospheric Research	115	22%	22%
Other Federal (NASA, DoE...) Supercomputer Center(s)	118	23%	
State or regional Supercomputing Center	36	7%	
My University's Supercomputing System or Facility	139	27%	
My own, or my department's, supercomputing facility	191	37%	
VIEW Other (*please specify*):	65	13%	

4. To what degree has the use of national supercomputing centers (NSF or otherwise) influenced your research? In addition to hardware, storage, and software, be sure to consider issues such as consulting and scientific collaboration.	Number of Responses	Response Ratio	Hayes Report (1995)
No Influence	70	13%	8%
Moderate Influence	151	27%	15%
Significant Influence	126	23%	29%
Essential Influence	203	37%	48%
Total	550	100%	

The above statistics indicate an increasingly bi-modal structure in the use of high-performance computing resources: large supercomputing centers and departmental or research-group facilities. Furthermore, the response to question 4 indicates a substantial increase in the impact or importance of national supercomputing centers on research. This and other information suggests that users desire increased investments in high-end computing as well as local facilities to facilitate their usage. Increasing emphasis is being placed on the use of the latter, especially in light of affordable technology (e.g., 160 gigabytes of disk space available for $300, which means that affordable terabyte storage capability already is available on the desktop).

6.	How would you rate the PACI Centers in providing the resources you need (this includes cycles, tools, consulting, etc)?	Number of Responses	Response Ratio
Poor 1.		24	5%
Fair 2.		47	10%
Good 3.		142	31%
Very Good 4.		163	36%
Excellent 5.		75	17%
	Total	451	100%

Comments regarding the above noted that PACI staff were helpful, with problems centering around an inadequate number of cycles, long queues, poor turn around, overcrowding, the need for more memory, and the increasing difficulty of actually using parallel machines (compared to the autoparallelizing compilers on, for example, Crays) and obtaining performance that represents a reasonable fraction of theoretical machine peak.

8.	What specific limitations, if any, have impeded your progress at the PACI Centers?	Number of Responses	Response Ratio
Resources oversubscribed in general		186	52%
Poorly-designed job queueing		76	21%
Inadequate consulting support		44	12%
Poor management and administration		23	6%
Cumbersome policies		57	16%
VIEW Other *(please elaborate)*:		142	40%

9.	Are the strategies for resource allocation at the PACI Centers appropriate, fair, and effective?	Number of Responses	Response Ratio
Yes		236	49%
No		30	6%
Unsure or no opinion		216	45%
	Total	482	100%

With regard to allocations, PACI needs more cycles, users noted difficulty in obtaining sufficient time as well as large number of processors (e.g., 64-128 processors for several days). A preponderance of respondents also noted the lack of multi-year grants of time as a major limitation to research grants that cover multiple years.

11.	Does an appropriate balance exist at the PACI Centers with regard to hardware, software, tools, and personnel?	Number of Responses	Response Ratio
Yes		194	40%
No		44	9%
Unsure or no opinion		243	51%
	Total	481	100%

Most of those providing written responses to Question 11 noted that too much emphasis is placed on raw hardware performance compared to tools and especially personnel. Specifically, users commented that many tools never proceed beyond the experimental stage to full deployment, and that although existing personnel are excellent, they are spread far too thin compared to the sophistication of the hardware and software environments they are tasked with supporting. Overwhelmingly, users support significant increases in support personnel. Further, users note that investments in new directions (e.g., Grid technologies) appear to be slow in yielding tangible benefits to the broader community.

13.	If you are new or relatively new to high performance computing, are you satisfied with the mechanisms available at the PACI Centers to bring new users on board?	Number of Responses	Response Ratio
Yes		172	44%
No		31	8%
Unsure or no opinion		188	48%
	Total	391	100%

Only 84 written comments were received out of the 391 individuals who responded to Question 13. Most already had been working with high performance computing, and several noted that the switch from vector-based machines to other paradigms had been difficult. For those who clearly were new to high performance computing, the experience in using the PACI centers was judged to be positive.

15.	Do you feel that practices/infrastructures/modalities used in other countries could be effective if adopted by the US?	Number of Responses	Response Ratio
Yes		49	10%
No		45	10%
Unsure or no opinion		378	80%
	Total	472	100%

For those users who expressed familiarity with infrastructures in other countries, the overwhelming sentiment expressed in written responses to Question 15 is that the United States is leading the way in the provision of hardware, network connectivity, supporting software, and collaboration tools. Further, the PACI infrastructure clearly is unique (other countries are adopting it) and thus has a significant impact on science and engineering progress. Several respondents expressed regret that the US does not have access to Japanese supercomputers

Disciplinary Impacts

There exists universal sentiment in the community that significant discovery has been enabled by the PACI centers, and that many, even more significant discoveries will be possible in the future. A good portion of these are anticipated to occur at the intersection of disciplines as well as in the context of societal implications, and made possible by Grid and related capabilities. Multidisciplinary teams will continue to proliferate, and efforts must be made to support them. Likewise, respondents noted that the proliferation of powerful, affordable desktop and departmental or research-group computers have had a dramatic impact on the ability to perform exploratory research, analyze data, and extend research in new directions. A sample of individual responses to a question regarding disciplinary impacts to date is given below. In many cases, respondents noted that the impacts were too numerous to mention.

- Simulation of newly discovered planets.
- Discovery and explanation of critical phenomena in gravitational collapse.
- First molecular dynamics simulations to show the sequence of actions of a polymerase and its role in the DNA repair process.
- Complete simulations of solid rocket booster firings.
- Demonstration of the practical predictability of individual thunderstorms and their related weather.
- Simulation of the large-scale structure of the universe.
- Fully coupled climate model simulations showing the agreement between observed and predicted increases in the global mean temperature.
- Largest simulations to date of Einstein's equations of general relativity.
- Simulation of fully three-dimensional coupled flow and heat transfer in a turbine engine.

A sample of individual responses to a question regarding anticipated discoveries for the future is given below.

- Mining of massive data sets across disciplines (e.g., weather and population).
- Upscaling fine resolution ecosystem models to broader scales.

• Determining the nature of dark energy.
• The generation of gravitational waveforms from numerical simulations of colliding black holes and neutron stars.
• Simulation of the complete life cycle of a tornado.
• A complete representation of the coupled magnetosphere-ionosphere-thermosphere system.
• Prediction of code performance on high-end computers.
• Use of the Grid to solve massive large practical problems.
• Understanding of protein and DNA folding and unfolding.
• First principle predictions of structures for protein domains.
• Large modeling of interactions among regions of the human brain.
• Tools for data analysis and knowledge extraction.
• Fully 3-D imaging of small-scale structures deep within Earth interior.
• Major advances in the treatment of subgrid scale turbulent processes in large eddy simulations.
• Extremely long-term and more realistic simulations of the entire coupled climate system.
• New insights into chaotic systems and properties of fluid turbulence.
• Complete simulations of the Earth-Sun system.
• Greatly improved tropical cyclone predictions.

A long-standing question regarding the provision of resources for the community is whether centers should serve primarily one or multiple disciplines. As shown below in Question 19, approximately 1/4 of all respondents believe that a PACI center or program organized around a specific discipline would be of greater value to them as a user compared to the present multidisciplinary organization of the centers. However, those who responded in the affirmative also noted that, despite possible advantages, the necessary division of resources to create such centers would lead to an overall reduction of quality and capability. Further, such centers would tend to maintain historical boundaries between traditional disciplines, which is incongruent with the future of science and engineering research and education. In the context of the current PACI framework, however, respondents expressed a clear desire for greater depth of consulting expertise within specific disciplines. They further noted that PACI centers should be able to dedicate significant resources to large disciplinary projects. defined periods of time.

19.	Would a PACI Centers Program organized around disciplines (i.e.; a center for Physics, another for Chemistry, etc.) be of greater value to you as a user?	Number of Responses	Response Ratio
Yes		118	24%
No		206	42%
Unsure or no opinion		171	35%
	Total	495	100%

Considerable emphasis has been given by the PACI program to facilitating interactions among disciplines. Question 21 shows that users have widely differing views regarding the effectiveness and even

the appropriateness of such a role. Based upon written responses, most users view the establishment of interdisciplinary collaborations as the responsibility of individual scientists, and many don't identify the PACI centers as the first point of reference for linking with other disciplines. The greatest value of the centers as a "melting pot" of disciplines appears to be the linking of domain scientists with computer scientists.

21.	Have the PACI Centers been effective in facilitating your interactions with researchers from other disciplines?	Number of Responses	Response Ratio
	Yes	123	26%
	No	155	33%
	Unsure or no opinion	195	41%
	Total	473	100%

Anticipated Use of Emerging Capabilities

This portion of the survey sought information regarding future use of emerging technologies such as the Grid, federated data depositories, and digital libraries. It also sought input about special needs, such as real time and on-demand availability of resources.

2.	Do you plan to run your application(s) across long-haul networks in a distributed fashion using the Computational Grid?	Number of Responses	Response Ratio
	Yes	116	22%
	No	258	49%
	Unsure or no opinion	153	29%
	Total	527	100%

The response to Question 2 above indicates, the written responses confirm, that the community as a whole is not aware of the Grid or concepts related to it and distributed Web services. Many noted that they lack expertise to modify their codes for execution across distributed resources, and that poor performance on existing parallel platforms might suggest equally poor performance across the grid. Further, numerous respondents appeared skeptical of the practicability of distributed methodologies, at least within the present management and facilities environments.

4.	Do you presently use collaboratory or knowledge networking (e.g., the Access Grid)?	Number of Responses	Response Ratio
Yes		61	12%
No		461	88%
	Total	522	100%

Likewise for Question 4, most respondents expressed lack of understanding about collaboratories and knowledge networks, though many noted that they presently are using, or soon plan to be using, the Access Grid to facilitate remote collaboration. Several commented that the Access Grid needs to become more reliable and cost effective to be practicable for community-wide use. Interestingly, when asked of their requirements for networked collaborations, most respondents said they had none.

7.	Do you presently use, or anticipate using in the future, digital libraries or federated data repositories?	Number of Responses	Response Ratio
I use them now		175	34%
No, but plan to do so in the future		117	23%
No, and have no plans to do so in the future		101	19%
Unsure or no opinion		127	24%
	Total	520	100%

9.	If you use digital libraries or federated data repositories, what are the primary impediments, if any?	Number of Responses	Response Ratio
Slow network speeds		118	55%
Ineffective data cataloging and/or difficulty locating required data		107	50%
Disparate formats among data types		112	52%
Inadequate documentation		80	37%
Inadequate or missing information on quality control		56	26%
Inadequate analysis tools		48	22%
Inadequate local storage for data analysis		34	16%
VIEW Other *(please elaborate)*:		40	19%

Questions 7 and 9, and their associated written responses, indicate a potentially significant increase in the need for digital libraries and federated data repositories. Several noted the lack of easy accessibility to historical data holdings, and the difficulty of dealing with multiple formats and data characteristics. However, most who provided written responses feel that data repositories are among the most important and challenging aspects of high performance computing and should receive considerable attention in the future.

10.	Do you conduct research that requires running codes and analyzing their output in real time (e.g., weather prediction software)?	Number of Responses	Response Ratio
Yes		124	24%
No		403	76%
	Total	527	100%

12.	Do you conduct research that requires management or control of remote devices or instruments?	Number of Responses	Response Ratio
Yes		52	10%
No		470	90%
	Total	522	100%

14.	Do you conduct research that requires real time data acquisition and/or cataloging?	Number of Responses	Response Ratio
Yes		113	22%
No		408	78%
	Total	521	100%

16.	Is network quality of service an important issue in your research?	Number of Responses	Response Ratio
Yes		327	63%
No		97	19%
Unsure or no opinion		93	18%
	Total	517	100%

Questions 10, 12, 14 and 16 above dealt with timeliness and related quality of service issues. A remarkable 24% of respondents noted that they conduct research that requires real time analysis of results, i.e., analysis that must be conducted as soon as the results are available, with the topic areas ranging from weather prediction (dominant response) to visualization and nano-materials research. A similar response was found for real time data acquisition and cataloging, while a smaller percentage of respondents noted the need for remotely controlling instruments. Changes in these percentages are difficult to anticipate, though the written comments suggest that several who do not require such capabilities now most likely will within the next 5 years.

A clear majority of respondents noted that network quality of service is important in their research, mostly in the context of speed, reliability, and security, generally in that order. We were surprised that nearly one fifth of those responding had no opinion or were unsure.

18.	Would your research benefit from advanced visualization technology (e.g., immersive or virtual reality systems)?	Number of Responses	Response Ratio
Yes		269	52%
No		121	23%
Unsure or no opinion		128	25%
	Total	518	100%

Slightly more than half of those responding to Question 18 indicated the need for advanced visualization technology. Many noted that existing systems are not suited to their needs, are too slow, and are too expensive and not practical (e.g., cave, power wall). Several indicated the need to visualize in dimensions greater than 4, and that visualization tools lag significantly both hardware and scientific application codes. In that context, it was noted that advanced visualization technologies are slow to move from the prototype phase (e.g., demonstrations) to practical implementation for use by the broader community. Finally, most of those providing written responses noted that visualization is a key component of their research methodology.

Other Future Needs

20.	Do you expect that your area of research will require access to supercomputing resources in the future?	Number of Responses	Response Ratio
Yes		479	91%
No		17	3%
Unsure or no opinion		28	5%
	Total	524	100%

The need for continued access to high end resources is underscored by the remarkable response to Question 20 above (93% responded in the affirmative in the Hayes report). Numerous written responses contained the phrases “stretching the limits,” “supercomputing is essential to my research,” and “I see no end to the need.” Many also noted that their most significant discoveries have been facilitated by use of the most power computing resources available (e.g., weather prediction, turbulence research, materials research, chemistry, bioinfomatics).

22. During the next several years, what fraction of your computing needs do you think could met by:

Percentage indicates total respondent ratio and parenthesis indicate actual number.	1 0-20%	2 21-40%	3 41-60%	4 61-80%	5 81-100%	6 Do not know
1. Desk top workstations	40% (210)	24% (125)	14% (73)	9% (48)	9% (45)	2% (10)
2. Mid Range Systems (Department/University owned with a market value from $100K to $2M)	25% (133)	28% (147)	20% (105)	10% (54)	6% (33)	3% (14)
3. Highest Performance Systems (National)	25% (133)	18% (94)	16% (85)	15% (80)	18% (94)	4% (22)

23. For the "Highest Performance Systems", how do you rank your projected needs in the following areas:

Percentage indicates total respondent ratio and parenthesis indicate actual number.	1 No Opinion	2 Unimportant	3 Important	4 Very Important	5 Critical
1. Processing speed	3% (15)	1% (7)	17% (86)	32% (167)	46% (238)
2. Large Memory size	3% (13)	7% (34)	21% (110)	33% (172)	35% (180)
3. Large Mass Storage availability	3% (18)	12% (64)	26% (135)	30% (155)	26% (132)
4. Networking Bandwidth for remote access	5% (24)	15% (76)	37% (191)	28% (142)	13% (67)
5. Large I/O Bandwidth (to local peripherals)	9% (45)	20% (102)	34% (173)	20% (104)	14% (73)
6. Network Bandwidth for inter-node communication	9% (45)	11% (56)	18% (94)	25% (128)	34% (174)
7. System availability and reliability	2% (8)	2% (12)	18% (91)	34% (177)	43% (221)
8. System security	5% (27)	23% (121)	35% (180)	19% (96)	15% (78)
9. Education and training	8% (42)	23% (118)	42% (215)	17% (88)	9% (44)
10. User support tools	6% (31)	16% (80)	41% (212)	24% (126)	11% (58)
11. Optimization	7% (34)	10% (54)	34% (173)	31% (162)	16% (82)
12. Consulting	6% (29)	17% (88)	42% (216)	22% (112)	12% (62)
13. Dedicated resources	13% (65)	25% (130)	33% (168)	17% (88)	10% (52)
14. Visualization	6% (33)	22% (111)	36% (185)	22% (114)	12% (63)
15. Data services	14% (73)	27% (138)	35% (182)	14% (70)	7% (38)
16. Code migration	13% (65)	24% (123)	34% (176)	18% (95)	8% (43)

The responses to Questions 22 and 23 above largely mimic those presented in the Hayes report, though with a general shift at the present time toward dependence upon personal workstations and departmental systems. Anticipated use of high performance national systems was found to be nearly identical to that in the Hayes report, and the percentages in Question 23 tended to shift slightly overall toward greater importance, with new categories (e.g., dedicated resources) clearly viewed as important.

24.	Rate below your view of the importance of national centers in providing the listed services (located at the centers) in relation to your research or educational projects:				
Percentage indicates total respondent ratio and parenthesis indicate actual number.	1 No Opinion	2 Unimportant	3 Important	4 Very Important	5 Critical
1. Diversity of computer architectures	11% (58)	26% (133)	31% (159)	21% (109)	10% (51)
2. Visualization facilities and software (including Virtual Environments)	10% (50)	28% (146)	34% (176)	20% (104)	6% (33)
3. Information processing	16% (80)	29% (148)	33% (170)	15% (79)	5% (24)
4. Consulting services	8% (39)	17% (87)	43% (221)	21% (107)	11% (58)
5. Third party or commercial software	11% (55)	31% (160)	34% (177)	15% (77)	7% (37)
6. Expertise in your research topic/Research Team Formation	11% (55)	37% (189)	28% (145)	17% (85)	6% (32)
7. Assistance in Code Porting and Optimization	10% (51)	18% (95)	37% (189)	21% (109)	13% (65)
8. Training	11% (55)	19% (98)	39% (203)	21% (108)	8% (41)
9. Communication software	15% (78)	22% (112)	38% (197)	16% (83)	5% (27)
10. Data Services (massive databases)	12% (64)	25% (130)	26% (132)	19% (100)	15% (75)
11. Repository for Data and Data Dictionaries	12% (64)	24% (123)	27% (138)	18% (94)	16% (80)
12. Internet software development	15% (79)	36% (183)	29% (150)	13% (66)	5% (24)

The responses to Question 24 also largely mimic those presented in the Hayes report, though again with a general shift at the present time toward greater importance of all items.

25.	What specific modalities do you feel are most appropriate for the future (check all that apply)?	Number of Responses	Response Ratio
Large essentially autonomous national centers		156	31%
Large coordinated national centers		290	58%
Interconnected collaborative regional centers		203	41%
Local centers		192	39%
Discipline-specific facilities		120	24%
Data-only facilities		47	9%
VIEW Other *please specify*:		29	6%

Finally, the responses to Question 25 indicate an overall sentiment toward non-discipline specific, interconnected collaborative centers and alliances.

Open-Ended Comments

Suggestions for meeting the needs of high-performance computing:

- Smoother transition of scales from local to regional to national
- Better coordination between NSF, DOE and NIH
- Need more powerful local/regional machines
- NSF grants should be linked to CPU allocations
- Integrate climate modeling with observations
- "PACI should have the middle ground between pure production and computer science research centers enabling collaborative research in high performance computing"
- Enhance network speed/bandwidth for data archives
- More high bandwidth links to facilities beyond national centers
- Concern/worry about the lag of academic computing versus DOE resources
- "PACI needs ambitions plans to ensure that the next generation of students are trained on state-of the art machines"

General Comments on the PACI program

- "PACI is very important"
- "Need to maintain the diversity of PACI centers"
- "Flexible powerful national centers are very useful"
- "PACI provides high end machines for the high end user"
- Current machines and strategy are not supporting high end science and engineering applications
- "Need to streamline allocations"
- "Terrible turnaround"
- "Turnaround too slow"
- "PACI needs capability, not just capacity"
- "Improve Queues"
- "Queues too long, machines overloaded"
- "Yearly grant applications a burden"
- PACI needs to maintain data archives, historical data and make them available"
- PACI needs to integrate "digital libraries data collections and persistent archives so as not to lose knowledge"
- "need data access accounts"
- "PACI needs to support a diversity of high-end users"
- "PACI needs many processors and multi-terabyte memory"
- "Too few users of 1,000+ processors"
- "Need large numbers of nodes with reasonable latency"
- "PACI should encourage usage of large processor sets"