
Argonne Leadership Computing Facility 2010 User Survey Results

Methodology:

Two versions of the Annual User Survey were created and circulated in 2010, the INCITE/ALCC Survey and the Discretionary Survey. Discretionary users were not invited to participate in user calls last year and were not given access to a Catalyst for support, so the Discretionary Survey did not contain questions about these two support mechanisms.

- Version #1: Used for Discretionary users – 15 questions
- Version #2: Used for INCITE and ALCC users – 16 questions

Both survey instruments were multiple-choice style, with comment boxes for open-ended responses where appropriate. Upon conclusion of the survey, respondents were invited to enter into a random drawing for one of three new iPod Nanos.

SurveyMonkey, an online survey tool, was utilized for both surveys. A link to the appropriate survey was emailed to each eligible user on December 20, 2010. Reminder emails to non-responders were sent on two occasions before the survey was officially closed on January 14, 2011.

Response Rate by Category (# responses/# sent):

- INCITE: 119/393 (30.3%)
- ALCC: 10/34 (29.4%)
- Discretionary: 120/420 (28.6%)
- All: 249/847 (29.4%)

Category Determination:

Some users have multiple projects in different categories (i.e., INCITE, ALCC, Discretionary). In these cases, the user's response was determined by his/her highest category. For example, if the user was on an INCITE project, the user's response was categorized as INCITE; if not, and he/she was on an ALCC project, the response was categorized as ALCC. Otherwise, the response was categorized as Discretionary.

Follow-up Contact Option:

Via the survey tool, respondents were given the opportunity to request follow-up from the ALCF regarding their responses to the survey. All those requesting a follow-up were contacted by their Catalysts (for INCITE/ALCC) or the User Services group (for Discretionary). Each was provided an individual response to any issues mentioned in their survey response.

PART I – INCITE USER SURVEY RESULTS

1. The Catalyst Program provides projects with a dedicated ALCF staff person (a performance engineer or computational scientist) to maximize your use of ALCF resources. ALCF Catalysts are Charles Bacon, Ramesh Balakrishnan, Graham Fletcher, Kumar Kalyan, Ray Loy, Vitali Morozov, James Osborn, Scott Parker, Katherine Riley, Nick Romero, Tim Williams, Jeff Hammond, and Marta Garcia.

On average, how frequently this past year did you communicate with your Catalyst?

Total Responses	110
Daily	2
Weekly	9
Monthly	60
Never	39

2. Please complete this statement from the responses provided below. The amount of communication and support from my Catalyst in 2010 was:

Total Responses	109
Just right	77
Not enough	4
N/A	28

3. Ideally, how often would you like to communicate with your Catalyst?

Total Responses	32
As needed	26
Daily	0
Weekly	0
Monthly	3
Other	3

Respondent comments: 7

- I have had just 1 or 2 incidents when I have needed assistance in the past 3 years.
- Some support I required was missing.

- We do our large production runs at ORNL, but I do use the BG/P at the ALCF for scaling test, ensuring platform-independence, and while participating in workshops organized by the ALCF. As such, I don't need regular contact with my catalyst.
- Depends on whether there is a problem or not!
- Depends on my schedule -- when I develop new code, I may have detailed questions. At other times I need no help.
- I joined the project late in the year. So never had real opportunities to collaborate with my catalyst. In the next year, I hope to have a closer collaboration with my catalyst.
- We are trying to move to threads and believe that the catalyst might have helpful expertise for this. Also for code optimization.

4. Please tell us about any additional services you would like your Catalyst to provide.

Respondent comments: 22

- Quick answers to simple questions?
- It would be great if more tutorials regarding the compilers and compiling environment are available. The IBM compiler on Intrepid sometimes cannot properly compile codes that are available on other Linux/Unix platforms with standard gcc/icc + MPI environment.
- How about having the Catalyst forward to the users any relevant news and information about ALCF and its systems. Something that is not generally announced to all users but can affect a project in particular. Things such as improvements in the HPSS interface, compilers, libraries. New optimization tricks, new tools, etc.
- We are pleased with the support we get. We look forward to additional help with GAMESS, GPAW, and POLYRATE.
- How to optimize # of processors used in the machine to get the best possible performance not only in terms of speed but also in terms of work done (i.e. including the wait time in queue).
- 1. Provide counselling about how to optimize my own code. 2. Give some suggestions about improve efficiency of my own code. 3. Also hope them could teach/guide us to solve problem step by step.
- I'm happy.
- He was pretty awesome so I can't really think of any additional services.
- Tim has been exceptionally helpful, generally going the extra mile. Our success would not have been possible without all his help. I can't think of any service which Tim does not already handle.
- Tim W. did a good job. He got us the information we needed.

- Better/more information about libraries (ANL-developed and otherwise) installed on ALCF machines.
- Thought he did a great job and was always responsive when I had a question or issue.
- Nothing in my mind. The wiki is a good spot to get most info I want.
- While I have interacted minimally with my catalyst, the few interactions have been great and have been the result of some seemingly unsolvable problem that got resolved. So I'm very happy with the way it works now.
- For example: Available solutions for particular supercomputing problems that has been already solved or correctly oriented, i.e best IO solutions, etc... In case that the problem is a new one, some orientation of the available resources at Argonne that may the user is not aware.
- It would be helpful if ALCF had resources to officially support contacts outside weekday working hours, for when there are system problems.
- All of the catalysts I have interacted with have been very responsive and helpful, whatever my requests. I cannot think of any additional services at the current time.
- Perhaps a designated hour or two a month where I could be sure to contact my catalyst, with quick response.
- Assistance about future projects linked to INCITE. Information on similar projects and technologies which we should use too.
- I often have problems with my Cryptocard. It keeps getting deactivated. My catalyst was helpful in sorting that out.
- Giving us advance warning and opportunities for testing our performance tools software prior to IBM kernel upgrades.
- I think he's been terrific. He's been there when I needed him and that's what we need.

5. Based on your experiences this year with your Catalyst, please provide a rating for each of the following statements.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
Having access to my Catalyst benefited my project.	4.3	4	5	0.6	0.8	79
My Catalyst is able to assist me with issues I bring to his/her attention.	4.4	5	5	0.5	0.7	77
My Catalyst is prompt and professional in our	4.5	5	5	0.5	0.7	77

dealings.						
Total Responses						233

Respondent comments: 15

- Our Catalyst, Tim Williams, is very knowledgeable in our field of plasma physics and that helps a lot when discussing possible optimizations and the constraints that are imposed by the physics. Tim is great. Very professional.
- Ray Loy is my Catalyst. He is very helpful with my project!
- The few times I had a question for my catalyst, he could not answer it; I find it more useful to contact directly somebody with expertise in the area(s) where I need help.
- We are very pleased so far with the assistance we received, especially from Graham Fletcher. We expect to tool up for greater usage soon, after several of my group visit Argonne in a week or so.
- So far, I haven't contacted my Catalyst, thus, I cannot make any comments.
- I would not be the right person in this project to be interacting with our Catalyst, so my answers should in no way be seen as a complaint or criticism.
- Even when we got grumpy that the system wasn't doing what we wanted, he stuck by us and helped us understand how to get it working again.
- I've technically contacted him exactly once over the last year. This was mostly account-setup related. My work on ALCF machines have not required his assistance so far, but I can anticipate scenarios where I would require his expertise on the software stack and support for scientific kernels on the machines.
- I interacted with Ray Loy, and he was very helpful, nice job Ray!
- The interactions have been few, but always really useful. The usefulness of the catalyst is measured in this quality, not the number of hours/e-mails exchanged.
- Tim has been great!
- I am a member of a few projects with time on Intrepid. My comments are based primarily on my interactions with Kumar Kalyan and with Ray Loy.
- I joined the project late in the year. So never had real opportunities to collaborate with my catalyst. In the next year, I hope to have a closer collaboration with my catalyst.
- For the purposes of this survey, I will use Charles Bacon as my Catalyst even though I, too, am a Catalyst.
- My project is a bit weird since it's an OS project. But in those areas where I needed some contact and support he's been great.

6. Did you participate in one or more of the monthly ALCF User Calls in 2010?

Total Responses	103
-----------------	-----

Yes	17
No	86

Respondent comments: 67

“I did not participate in the calls because _____”

- 12 cited scheduling issues, too busy, no time.
- 11 said they didn't know about the User Calls.
- 14 said they had no need or the calls were not necessary for them.
- 3 said they were not using the ALCF or were using it infrequently.

Other specific respondent comments:

- external user
- I am a relatively small user doing mostly code development. Other in the same allocation DO participate in the calls.
- I will participate in the future. It will be helpful.
- All my questions were promptly answered by support/catalyst. As I had no outstanding issues at the time of the calls, I did not participate in them.
- 1. We are still tooling up. 2. Dr. Tishchenko, my co-PI, has been the main user so far. 3. We received other useful communications on a 1:1 basis.
- Too lazy to try, but I do hope to join next time!
- I have been 100% satisfied with the throuput I've been getting on Intrepid, and thus did not feel the need to participate in the user calls.
- Not relevant at this phase in the project.
- No obstacles required it... e-mail updates are frequent and technical issues are handled within our project.
- I didn't have the time to spare. Moreover, I didn't begin using the machine in earnest until late in the year.
- I didn't keep track of when they were.
- I've not had the need and I am still uncertain of any constructive role from which I can contribute by my participation in those calls.
- We already had a lot of experience with BlueGene systems. The assistance we needed was very specific and was obtained through direct interactions with our Catalyst and ALCF specialists.
- At too low of a level.
- Unaware of it. Also, probably not useful (for us), as we are tool developers. We talk to admins and developers.
- I use ALCF only sporadically, and I do not expect to get useful information from these calls. I prefer information dissemination over the web, or via email if time sensitive.
- All my production runs went smoothly and we have a good understanding of the scaling of our code. We had a lot of contact with our catalyst ~2 years ago

when we were encountering new problems. Other people from our project may have participated in user calls this year, but not me.

- Direct communication with my catalysts and other contacts within the ALCF have been sufficient to address my questions and problems.
- I joined the project late in the year.
- I didn't feel the need to. My catalyst is very accessible.
- We were able to resolve system- or code-related issues (if any) directly with our Catalyst.
- Just signed up.
- Not sure what their benefit is.
- I have opened my account too recently.
- I am not a PI on the project.
- As tools providers, our support needs are different from typical users.
- I get sufficient information about resources from my Catalyst.
- It's not really relevant to my needs. We are not doing applications, we are much lower level.

7. Did you attend an ALCF-sponsored workshop in 2010?

Total Responses	110
Yes	28
No	82

8. Please rate the effectiveness of the workshops you attended in addressing the following topics.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
Introducing me to ALCF staff and services	4.3	4	4	0.5	0.7	27
Getting my project up and running	4.2	4	4	0.5	0.7	23
Providing relevant/helpful training	4.2	4	4	0.3	0.6	27
Providing access to experts	4.4	4	4	0.4	0.6	28
Helping resolve performance issues	4.2	4	4	0.5	0.7	24
Total Responses						129

Please share any additional comments about workshops you attended, or topics for future workshops you wish to see.

Respondent comments: 2

- I am commenting on the BGQ early science startup
- I helped provide training about performance issues. We didn't have any to resolve ourselves.

9. ALCF provides user support via the web, through email, and by phone through our service desk. In regard to user support you have received, please rate the following topics.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The ALCF staff provides accurate, complete assistance and/or answers to my questions.	4.5	5	5	0.3	0.6	104
The ALCF staff is courteous and professional.	4.7	5	5	0.2	0.5	104
Assistance from the ALCF staff is prompt.	4.6	5	5	0.3	0.6	103
Online supporting documentation is helpful.	4.0	4	4	0.7	0.9	100
ALCF support is available when I need it.	4.3	4	4	0.4	0.7	105
Total Responses						516

Respondent comments: 17

- I would consolidate the documentation about the systems. Some of it can be found on the main web pages and other in the wiki FAQ. That is sometimes confusing.
- Web pages are very helpful. More technical issues are answered by Ray Loy very promptly.
- I would feel better if the Accounts web site were signed with an SSL certificate that my browser recognized out of the box (in particular, Safari).
- The performance of the front end nodes is rather low when compiling large applications with xlc.
- Some level of weekend support -- even operations level with some on-call capability -- would be greatly beneficial. we had several occurrences when a weekend slot was unproductive because systems problems could not be

resolved.

- To effectively operate a machine like Intrepid support MUST be available 24 hours a day, 7 days a week in some form because the size and complexity of the machine ensures that some glitch or malfunction will happen on a regular basis. Support during off hours to field these glitches is essential, especially since that's when reservations are scheduled. The computing facilities at LLNL is an example of a good support model.
- The biggest support issue was that ALCF wanted us to start our full system runs after ALCF staff had gone home for the day. We wasted several time slots early in the project because no one from ALCF was there to see that there was broken hardware. When we were allowed to start a job during working hours, ALCF staff quickly found the broken hardware.
- Very recently I had questions about the Double Hummer floating point unit in the BlueGene/P. I found little information on the web, and only the IBM xlc compiler documentation said anything detailed enough to be useful for programming. I wish there was more documentation, but I assume this is mostly out of ALCF's control.
- The website is good but could be better. For example, the NERSC website is more complete and easier to navigate.
- Support outside of weekdays happens thankfully some but cannot be counted on and reportedly is not funded by DOE or ALCF. With an enormous system as ALCF provides, inevitable system problems would be covered 24-7, but this is not an ideal world.
- While I cannot remember any specific instances at the moment, my impression is that it can be difficult to find the information that I need in online documentation. This may be due to the nature of my questions, but I usually have to contact ALCF staff directly to get my questions answered.
- It is not so easy to search through the online documentation.
- Only needed this help once, through phone, for password reset. However, help was quite effective.
- "ALCF support is available when I need it" is agree rather than strongly agree because I'm in a very different timezone most of the time (Europe) so my mornings, I don't get any support at all... I don't really mind that though - that was clear at the outset that this would be the case. There is a lot of information on the ALCF wiki - indeed googling on the internet at large often takes me right there. But I find it is sometimes misleadingly outdated.
- Robert and Charles are especially helpful.
- Documentation about concerns we encounter supporting performance tools is slim. Often we have to resort to reading kernel or system code.
- You should all get a raise. :-) I love working with the ALCF.

10. Please rate each of the following resources you accessed through the ALCF this year.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The performance tools/debuggers I needed were available.	4	4	4	0.6	0.8	88
The libraries I needed were available.	4.1	4	4	0.5	0.7	95
Total Responses						183

Please list any other tools that would be helpful.

Respondent comments: 12

- NCL and NCO on intrepid front end, to quickly check output without having to log into eureka.
- I haven't used the performance tools yet.
- I did not need a debugger much on Intrepid during the past year.
- Debugging on Intrepid is reduced to 'printf'. Interactive jobs would be useful (on all platforms).
- Intrepid desperately needs the ability to start an xterm. LLNL has this with SLURM, and we run interactive sessions and much more easily debug BG/P jobs from the command line. Cobalt's support for interactive sessions is limited and restricts users to gdb and totalview. Either cobalt needs an upgrade, or some other method needs to be put in place for debugging on Intrepid. I've even heard this complaint from application developers at Argonne, who aren't happy with the state of debugging in Cobalt.
- But sometimes difficult to know which libraries were available (e.g. found out through Ray that /soft/apps/current is the right place to look, didn't see that documented anywhere).
- I could have used an installation of HDF5 1.8.5-patch1 with both C++ and Fortran bindings for Blue Gene/P. (I built it myself by now, which wasn't easy because of the cross compiler.)
- I know how to cross-compile libraries for the compute nodes so it doesn't really matter, but this year I have needed to install: libxml, gsl, lua. It would save time if these were pre-installed.

- I will need Trilinos in the very near future. I have determined that a version has been installed on LLNL BG/P systems, and have requested that my Sandia collaborators determine whether this can be used for a port to Intrepid. Any help in getting Trilinos installed on Intrepid would be appreciated.
- I work on performance tools (HCPToolkit). We want some of our tools to run on the BG/P back-end nodes. We found that in many cases it is impossible -- without significant hacks and contortions -- to use the back-end version of GCC (mpicxx, mpicc) with autoconf/automake/libtool to create a statically linked application.
- Profiling with TAU is great. Debugging on BGP is a huge pain in the ass. Using GDB or Totalview is too difficult to bother so I just use printf.
- nano editor.

11. Please rate the following statements using the scale provided.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The disk and tape storage provided was sufficient.	4.2	4	4	0.6	0.7	97
The policy of giving higher priority to capability jobs (>32K cores) is reasonable.	3.9	4	4	1.0	1.0	99
I was able to run capability jobs (>32K cores).	4	4	5	1.0	1.0	72
I am satisfied with the scheduling of my jobs.	3.9	4	4	0.7	0.8	101
I was able to utilize Eureka for visualization and data analysis.	3.7	4	3	0.8	0.9	46
Total Responses						415

Respondent comments: 10

- Of course I would like faster turn around.
- Scheduling is sometimes mysterious - long jobs would sit in the que for a week, and it seems I could get better throughput by splitting the job up into N shorter submissions.
- I haven't run with >32K cores. I haven't used Eureka.
- We had problem initially getting yorick built and running on eureka, but i think steve langer has sorted it out. we haven't tried VisIt on the data set yet,

but will do so soon.

- The constant need to ensure that large system software builds and large datasets reside on the larger lustre compute-node file system is somewhat annoying. A home partition of 2GB (?) is fairly restrictive for someone who has to work with a wide range of scientific software systems. It forces me to develop and compile on the lustre file system itself, the act of which is the cause of/results in slow file system performance. That becomes an unproductive way to work. I would very much prefer to store all my software systems in the faster home file system, compile to a binary there which can then be copied into the lustre system from which compute work can be carried out. Comment on 32K jobs - my largest run has so far been on 24K, hence the N.A.. I do, however, expect to run on far larger configurations in the future.
- It would be good to be able to submit more concurrent jobs on Eureka.
- I have not used Eureka, but this is my fault -- the machine seems to be very useful.
- From my user viewpoint, it is a mystery as to why my jobs are scheduled when they are scheduled. I may just not know the correct commands to determine estimated start time (and have been too lazy to ask my catalyst).
- Currently I'm only working on surveyor. I believe >32k cores won't be a concern for me for quite some time. :)
- Our project is certainly a "capability" project, but the jobs themselves are not capability jobs by your definition. Our optimum job size is 4K to 8K. But we run many thousands of them in a project, requiring tens of millions of SU's. There are no other resources that can handle a project of that size whether at a University or DOE. Such a lack of alternative resources is, I believe, the real justification for a leadership class facility. I would lower the bar in this definition. We can always package our jobs so they require 32K or more cores, but run independently on 8 internally assigned partitions. This wastes resources because the independent jobs don't finish simultaneously. So we don't do it. But why encourage users to play such games?

12. What other resources could we provide that would be helpful to you now or in the future?

Respondent Comments: 19

- A BG/Q system! ;-)
- More/Faster Frontend nodes on the BG/P for large compile jobs/post processing.
- Keep doing the workshops -- though I didn't attend in 2010, I have participated in the past, and might do so again in the future, depending on the topics.
- Help with GPAW, GAMESS, and POLYRATE.

- I got all I need so far from ALCF!
- Is there a queue status option that could give an estimated job start time?
- Early access to Blue Gene/Q hardware.
- Debuggers for the production machines (also requires interactive sessions for production machines).
- A scheduler/resource manager with a better interactive mode. I'd love to see SLURM installed, though I don't know what the likelihood of that is.
- It would be nice if each submitted job were automatically performance profiled, and if a high-level report was available on the web. This could include e.g. global MPI latency/bandwidth, floating point instructions, memory bandwidth, cache misses etc. The report should be so high-level and brief that it is understandable to a "casual" user.
- Continue to update the ALCF wiki. It is an excellent resource.
- Wish we can have a real-time GUI web portal to see job scheduling and maintenance plans, etc., altogether.
- Testbed for medium sized jobs with quicker turnaround.
- Some tutorials about how to use Paraview in Eureka for postprocessing. Not only how to start Paraview at Eureka, but real example of project that are willing to prepare some slices about how they did it.
- IDL (Interactive Data Language) on Eureka for data analysis.
- Other machines! Capacities to transfer files faster and with safety
- POWER6 login nodes to compile faster.
- access to surveyor for small-scale testing and testing before new OS deployment.
- We're going to need a way to put our own file server on Intrepid's IO node network – that's an issue that will hold us back from testing when the time comes.

13. Are you able to achieve the science goals for your project?

Total Responses	107
Yes	100
No	7

Please tell us why you are unable to achieve your project's science goals.

Respondent comments: 8

- Not so much yet, but this is not actually a computational resource issue as much as it is an issue of lack of sustained external funding for the work. The method is very promising and aside from I/O issue scales extremely well.
- Our code does not perform well on the BlueGene/P architecture. (We have

access to many TeraGrid systems that we use for science, but it would be nice if we could use BG/P as well.) We are currently analysing why our code does not perform well.

- I work to support the scientists in their mission but do not have specific science goals myself. This question is not relevant to my work.
- the INCITE "steady use" demand caused us to rush some experiments at the early part of the year and left us with no time for experiments at the end of the year.
- In progress.
- Just starting. We think and hope we will succeed.
- Went over on allocation.
- We've not been able to do Intrepid runs. The User Agreement is structured in a way that makes it impossible for IBM staff to sign it. Would be nice to work this out.

14. Overall, how would you rate your experience with the Argonne Leadership Computing Facility in 2010?

Total Responses	110
Excellent	64
Above Average	41
Average	5
Below Average	0
Poor	0

Average Rating: 4.5

Respondent comments: 7

- Everybody at ALCF is very friendly and professional.
- Fast supercomputer, reasonable waiting time, good management/support!
- Flexibility in terms of allowing installation of the Coraid and Plan 9 server machine is appreciated.
- Much better than ORNL, slightly better than NERSC, and on a par with LANL, and SNL.
- I used ALCF only sporadically, so I cannot answer this question.
- I don't like the comparison - I don't work with enough different computer centers for "average" to make much sense to me. I am redefining "Above Average" to be "Good".
- ALCF is the greatest supercomputing facility on planet earth.

15. Would you like someone to contact you regarding your responses to this survey?

Total Responses	101
Yes	10
No	91

PART II – ALCC USER SURVEY RESULTS

1. The Catalyst Program provides projects with a dedicated ALCF staff person (a performance engineer or computational scientist) to maximize your use of ALCF resources. ALCF Catalysts are Charles Bacon, Ramesh Balakrishnan, Graham Fletcher, Kumar Kalyan, Ray Loy, Vitali Morozov, James Osborn, Scott Parker, Katherine Riley, Nick Romero, Tim Williams, Jeff Hammond, and Marta Garcia.

On average, how frequently this past year did you communicate with your Catalyst?

Total Responses	9
Daily	0
Weekly	0
Monthly	4
Never	5

2. Please complete this statement from the responses provided below. The amount of communication and support from my Catalyst in 2010 was:

Total Responses	9
Just right	4
Not enough	1
N/A	4

3. Ideally, how often would you like to communicate with your Catalyst?

Total Responses	5
As needed	5
Daily	0
Weekly	0
Monthly	0
Other	0

Respondent comments: 0

4. Please tell us about any additional services you would like your Catalyst to provide.

Respondent comments: 3

- Services currently provided are satisfactory.
- There were troubles with the initialization of my account. It would have been great if my catalyst could have helped get that fixed rather than waiting until a report was needed.
- Examples that are up-to-date.

5. Based on your experiences this year with your Catalyst, please provide a rating for each of the following statements.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
Having access to my Catalyst benefited my project.	4.8	5	5	0.3	0.5	4
My Catalyst is able to assist me with issues I bring to his/her attention.	4.8	5	5	0.3	0.5	4
My Catalyst is prompt and professional in our dealings.	4.0	4	5	1.5	1.2	5
Total Responses						13

Respondent comments: 2

- Did not need to interact with my catalyst in the past.
- The catalyst was professional, but this person was unable to help me. We had to go to their supervisor. It wasn't their fault, they just weren't able to authorize things, I think.

6. Did you participate in one or more of the monthly ALCF User Calls in 2010?

Total Responses	9
Yes	4
No	5

Respondent comments: 2

- Some other group member did participate.
- My issues were of an individual nature, when the came up.

7. Did you attend an ALCF-sponsored workshop in 2010?

Total Responses	9
Yes	1
No	8

8. Please rate the effectiveness of the workshops you attended in addressing the following topics.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
Introducing me to ALCF staff and services	4.0	4	n/a	n/a	n/a	1
Getting my project up and running	5.0	5	n/a	n/a	n/a	1
Providing relevant/helpful training	5.0	5	n/a	n/a	n/a	1
Providing access to experts	5.0	5	n/a	n/a	n/a	1
Helping resolve performance issues	4.0	4	n/a	n/a	n/a	1
Total Responses						5

Please share any additional comments about workshops you attended, or topics for future workshops you wish to see.

Respondent comment:

- Perhaps the BGQ kickoff meeting was not a 'workshop'

9. ALCF provides user support via the web, through email, and by phone through our service desk. In regard to user support you have received, please rate the following topics.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The ALCF staff provides accurate, complete assistance and/or answers to my questions.	3.7	4	4	0.9	1.0	7
The ALCF staff is courteous and professional.	3.9	4	4	0.5	0.7	7
Assistance from the ALCF staff is prompt.	3.7	4	4	1.2	1.1	7
Online supporting documentation is helpful.	3.8	4	4	1.1	1.0	8
ALCF support is available when I need it.	3.9	4	4	0.5	0.7	7
Total Responses						36

Respondent comments: 2

- Online documentation is very helpful and a resource I refer to frequently. The ALCF Wiki, in particular, serves as an example that other computing facilities should adopt. Assistance from support@alcf.anl.gov is the primary resource for logistical details and diagnosing hardware/software errors.
- Much of the information I have seen on the wiki is out-of-date.

10. Please rate each of the following resources you accessed through the ALCF this year.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The performance tools/debuggers I needed were available.	4.1	4	4	0.4	0.6	8
The libraries I needed were available.	4.0	4	4	1.1	1.1	8
Total Responses						16

Please list any other tools that would be helpful.

Respondent comments: 2

- On the ALCF Wiki, I would like to have the compiler options and flags used to make each software package explicitly listed on the Applications & Libraries page for Intrepid and other machines. This is currently done for some unsupported packages, but not for most supported packages.
- I would like Trilinos available.

11. Please rate the following statements using the scale provided.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The disk and tape storage provided was sufficient.	4.0	4	4	0.3	0.5	8
The policy of giving higher priority to capability jobs (>32K cores) is reasonable.	3.5	3	3	0.6	0.8	8
I was able to run capability jobs (>32K cores)	4.0	4	4	0.8	0.9	6
I am satisfied with the scheduling of my jobs.	3.5	3.5	4	0.9	0.9	8
I was able to utilize Eureka for visualization and data analysis.	3.0	3	n/a	n/a	n/a	1
Total Responses						31

Respondent comment:

- I am satisfied with the scheduling and queuing policies. The relatively quick turn-around time and diverse options provides a good balance.

12. What other resources could we provide that would be helpful to you now or in the future?

Respondent comment: 1

- None.

13. Are you able to achieve the science goals for your project?

Total Responses	8
Yes	8
No	0

Please tell us why you are unable to achieve your project's science goals.

Respondent comment: 1

- I was very late getting access to the system. But my project will catch up. It was simply more complicated.

14. Overall, how would you rate your experience with the Argonne Leadership Computing Facility in 2010?

Total Responses	8
Excellent	2
Above Average	4
Average	2
Below Average	0
Poor	0

Average Rating: 4.0

Respondent comments: 0

15. Would you like someone to contact you regarding your responses to this survey?

Total Responses	8
Yes	0
No	8

PART III – DISCRETIONARY USER SURVEY RESULTS

1. How many discretionary allocations has this project had?

Total Responses	112
One	58
Two	12
Three	3
Other	3
Unsure	36

2. What ALCF system(s) is this discretionary project currently running on?

Total Responses	141
Surveyor	64
Intrepid	67
Unsure	10

3. At what stage is this discretionary project?

Total Responses	135
Getting code to compile	22
Debugging	23
Scaling to one rack	22
Scaling beyond one rack	48
Other (please specify below)	20

Respondent comments: 27

- Received INCITE award to continue.
- Scaled to 131K cores.
- Preparation for major production runs.
- Working on code physics testing and implementation. Also parallel I/O is an issue.
- Scaling to the full Intrepid system.
- Research and development on runtime system mechanisms and

implementation, enabling optimization of code running atop the improved runtime.

- Modernization of current scalable MPI code to Hybrid MPI-OpenMP.
- I don't know what others are doing. I'm just porting some software.
- Further development.
- It would be good if your survey noted the allocation you're asking about. Some people have more than one allocation.
- Successfully accomplished.
- I am typically running on 512, a reasonable compromise between efficiency and productivity. The project is the stage of physics validation and benchmarking against other codes, and should this benchmark converge, lead to acceptance in the Accelerator Community.
- The NRCM version of WRF is running in production. A second discretionary allocation for wind energy simulation (project heron) is in the early stages, but since we're using WRF we expect this will also scaled to multiple racks.
- I believe our time regarding this project has run out. We were scaling beyond one rack while we were running on Intrepid.
- System software research.
- Goal of project is to maintain and improve implementation of X10 programming language for BG/P systems. We've successfully scaled some programs to (and beyond) one rack, but other programs need significant work and the language implementation itself is evolving.
- Simple flow problems were scaled very successfully to 4096 cores. Problems involving more complex physics still require more debugging to run on Intrepid.
- We used the allocation for calculations with commercial code (NAMD).
- We had issues in getting good scaling beyond one rack. So we moved the project to other resources, which are more suited for the demands of the project.
- Using surveyor largely for testing purposes, we have a BG/P ourselves and at times it's helpful to have another BG/P on which to run for comparisons.
- Implementation of new computational algorithms for better scaling and speed.
- We have completed our work on Intrepid.
- Used commercial code (NAMD).
- We scale fine up to 32 Racks with no I/O. And we are now working on I/O performance for large core counts.
- Done – Scaling.
- In production and nearing the end of the project.
- Not sure.

4. What is NEXT for this discretionary project?

Total Responses	100
Apply for time on Intrepid	26
Continue scaling on Intrepid	39
Apply for an INCITE award	24
Apply for an ALCC award	11

Respondent comments: 27

- ESP and INCITE running.
- Unsure as the project funding for this work has fizzled for the time being.
- Work remains to get things working w/Zepto on Surveyor.
- Using machines at LLNL.
- Further scaling work is needed. So far, have only run up to 4 racks
- The developed software will be integrated in the stack of software tools employed for INCITE project.
- The mpqc4bgp project led to a BG/Q ESP project. Use of this code for existing
- INCITE projects has been discussed but not acted upon.
- NWChem is going to be used by multiple existing INCITE projects.
- The purpose of PARTS is to maintain parallel runtime systems upon which applications are built. Half of what the project tries to accomplish is maintaining the status quo (i.e. stuff continues to work) while the other half is getting said stuff to perform better at scale (e.g. ARMCI).
- Ask Jeff Hammond...
- Scaling on Magellan, and using gridftps on Intrepid.
- Should it be necessary, running with more different parameter, depending on the physics merit.
- The project concerns developing a unified build system for a fusion simulation framework that works from laptop to LCF, and then using that framework for science applications. To date, much of the work has focused on getting builds for surveyor and intrepid to successfully complete and to integrate ALCF builds results into the integrate build and test system for FACETS. This includes understanding and incorporating the ALCF specific modifications required for config/make/install toolchains to also work on ALCF systems. In addition, FACETS can also run regular scaling tests as part of the regression testing (i.e. test not only that code commits do not break the code but that they also do not cripple performance). Finally, FACETS uses some of the platforms where it can successful deploy builds to run physical simulation campaigns for fusion relevant science studies.
- Not sure.

- Continued research.
- We will not re-apply.
- Not using ALCF resources as of now for this project.
- Continue to have access to Surveyor for testing.
- Work has been completed.
- Will likely not re-apply.
- We have applied for an INCITE award in June 2010. The reviews specified that we should continue requesting discretionary allocations to continue the scaling research. Our project is also working on Jaguar and the feedback from work on both systems is very valuable to establish robust scaling beyond 100,000 cores.
- Will probably try for an INCITE award again this year.
- We have just applied for and ALCC award and hope to continue the work we have done already.
- TeraGrid.
- Not sure.
- Not yet decided whether we apply for an award.

5. Please tell us about any performance issues you are currently experiencing with your application.

Respondent comments: 40

- The code is finite element-based Lattice Boltzmann method for multiphase flows-- aside from algorithmic problems related to simulation of turbulent flows, we were having some issues with I/O and need help implementing a more efficient I/O method (currently all processors write to same file one at a time--obviously this is terrible for scaling). Also, there were some issues with node memory usage--FELBM is a little heavy--though we were not able to investigate these completely.
- I am currently working with applications that need I/O.
- At the debugging stage there are no performance issues.
- Small numbers (<100) of RAS notices about exceeding the DMA receive FIFO, which don't appear to be impacting overall performance.
- PVFS is unstable sometimes.
- Data sets have limited scalability (climate models with fixed grids).
- Efficiency is 66% on 16K cores.
- Scaling of application at over 100,000 cores.
- The eigensolver doesn't scale. We have a solution for this - Elemental - and started implementing it literally this week.
- PNNL software is terrible. Both NWChem and GA/ARMCI are awful spaghetti code. Unfortunately, I have only found time to rewrite one of them.
- Too early to say.

- Please try to improve on your time-allocation reviews. We have renewed the time allocation on our project two times and every time it has taken almost a month.
- Lately there have been some surprising differences in execution times between Surveyor and Intrepid that do not appear to be related to I/O times.
- Our application has two major parts. One of them is a purely serial process that can be fully decomposed across multiple processors. The other part requires parallel computation to solve, but the actual problem being solved is not intrinsically large so it does not decompose well when the number of processors exceeds the essential granularity of the problem. We are working on identifying and exploiting other parallelisms in our problem.
- a.) Despite producing results, there is room for improvements. I am using the VORPAL "commercial" (i.e., via SBIR grants) code, beta versions. Currently we still have one infinite loop to debug, in one case, and a file corruption (or data corruption) in another case. Perhaps the latter problem is due to instabilities in the High Performance Parallel file system we are using. In production mode, on 512 nodes, running a maximum of 12 hour continuously, it is not at all uncommon to have to wait 2 to 3 days to get in the queue. O.K., I have patience, but since I am old fart, this might not be the most convincing argument to justify my fat salary...
- We are unable to use the SIMD instructions on the Intrepid processors because our app (WRF) is single precision.
- Need to reduce MPIall to all communication in electrostatic calculations (multiple FFTs per simulation step) in reactive molecular dynamics simulations.
- Not many really. They are currently pretty standard issues of the combinatorics of compiler vendors, version numbers, OS's and various libraries and implementations. In general, no big problems, but much maintenance to keep things working well in such a changing environment. FACETS has not yet pushed extreme scaling limits to explore and ease hard scaling bottlenecks at tens or hundreds of thousands of processors.
- We do not currently have time allocated.
- The application implements a remote method invocation (RMI) communication system using prefired MPI receive messages. The current issue we're facing is that using settings from previous systems we're seeing that we fire too many messages and run out of memory quickly. Additionally, finding the supported set of timers and other monitors to measure memory usage has been difficult.
- So far so good.
- A lot of TLB misses and poor single-node locality issues. At a coarse-grain the application has excellent scalability as it is embarrassingly parallel.
- Our I/O method slows the code's performance. Parallel I/O is under

investigation.

- The code used does not scale well beyond one rack. As ALCF is meant for higher end computing, we were suggested to look for resources more suited for our demands.
- Scaling of MPI-OpenMP code is not perfect.
- File I/O.
- We are having difficulties to scale i/o for large core counts.
- Code runs very slowly on individual cores under Zeptos (mostly but perhaps not all) due to clockspeed. This makes it hard to fit jobs into walltimes. Probably need to add parallelism to make BG/P appropriate for app.
- Currently we see dramatic degradation of performance beyond one rack compared to jaguar where our application scales over 70% on 80,000 cores. We are not sure of the cause of this; it is still being investigated.
- Seemed to crash on 1 rack due to memory issues.
- I hope to get to soon get to the point where I can where I can have performance issues. :-)
- Back end compilers can be frustrating for building codes with many dependencies.
- Using lammps which does not scale beyond 16 processors.
- Not quite good strong scaling when the input problem size is small; some IO issues.
- I/O; Node balancing.
- Our code has been already ported to Cray XT* platforms and demonstrated to scale up to 74,000 cores in hybrid mode of operation (MPI+OpenMP). However, we have encountered number of problems during the porting stage and also performance profiling on Blue Gene platforms. We are hopeful that we will be able to isolate these issues and test the code on comparable scale.
- Some specific settings with VASP do not run for more than 256 cpus, which may be a VASP bug.
- Compile time for some C++ code is 10x slower then other machines.

6. What year do you anticipate submitting an INCITE proposal for this project?

Total Responses	85
2013	29
2014	0
2015	0
Undecided	56

7. Did you attend an ALCF-sponsored workshop in 2010?

Total Responses	97
Yes	18
No	79

8. Please rate the effectiveness of the workshops you attended in addressing the following topics.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
Introducing me to ALCF staff and services	4.1	4	4	1.1	1.0	17
Getting my project up and running	4.1	4	5	1.2	1.1	17
Providing relevant/helpful training	4.1	4	5	1.2	1.1	17
Providing access to experts	4.2	5	5	1.4	1.2	18
Helping resolve performance issues	4.1	4	5	1.4	1.2	15
Total Responses						84

9. What general topics would you like to see covered at workshops for discretionary projects? (Select all that apply.)

Total Responses	257
Performance profiling, code optimization	62
Debugging	39
Scaling	48
I/O Performance	49
Data analysis/Visualization	26
Hands-on Time	30
Other	3

Respondent comments: 5

- Time distribution.
- Training for advanced users.
- Overall tips to run effectively on a system which is over-committed...
- No clear preference. All [listed choices] look valuable. FACETS specific interests might lie in the ability to understand the deployed software infrastructure (OS, compilers, MPI implementations, PETSC, ...) that affect FACETS build and performance.
- I think best use of time for us during a workshop is spending time with experts on interpreting the performance profiling results we acquired. So much effort is spent to show how to collect the data but small amount of time is spent on how to interpret the data and resolve bottlenecks pointed with the data collected as this will have the highest impact. We are computational scientists not performance profiling experts. So having the opportunity to interact with people having this expertise in a workshop dedicated for this purpose will be great. As a prerequisite all participants could be expected to pass the stage of porting and already collected performance profiling data.

10. ALCF provides user support online, via email, and by phone through our service desk. In regard to user support you have received, please rate the following topics.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The ALCF staff provides accurate, complete assistance and/or answers to my questions.	4.4	5	5	0.5	0.7	92
The ALCF staff is courteous and professional.	4.6	5	5	0.4	0.6	93
Assistance from the ALCF staff is prompt.	4.5	5	5	0.6	0.8	93
Online supporting documentation is helpful.	3.9	4	4	0.7	0.8	92
ALCF support is available when I need it.	4.2	4	4	0.6	0.8	90
Total Responses						460

Respondent comments: 10

- All of our interactions with ALCF staff has been very positive and constructive up to now.
- The ALCF staff forgot to send me the Cryptocard token. I got it few months

after I applied the account.

- It's my fault that I work a lot nights, weekends and holidays so staff is not always around when I am.
- Online materials are poorly organized. There was also some confusion regarding an extension of the time to use our allocation. Communication could be much better.
- The status of the BG/P machines should be available online. This can help to avoid the maintenance.
- I've had problems finding online documentation for surveyor. Perhaps an annual email containing links to the online documentation would solve this issue.
- I have found nothing to complain about from ALCF's operations and support teams -- and that is rare coming from a normally crank HPC user.
- Everyone in User Services is fantastic.
- The Wiki is awesome.
- one of the best.

11. Please rate each of the following resources you accessed through the ALCF this year.

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The performance tools/debuggers I needed were available.	3.8	4	4	0.6	0.8	74
The libraries I needed were available.	3.9	4	4	0.7	0.8	85
Total Responses						159

Please list other tools that would be helpful.

Respondent comments: 12

- Trilinos.
- The ability to retroactively change the core dump setting for a job (i.e., if I see that a job is probably in an infinite loop, have some way of immediately making it dump core).
- libxml2 boost module.
- hdf4.
- Numpy and python that runs on compute nodes without fiddling and is optimized. I was able to compile these myself with some effort.
- debuggers are available but it is hard to use them remotely.
- Graphical text editor such as gedit or gvim.
- Attempting to analyze the performance problem I mentioned earlier is

painful because the set of performance timers supported on the platform is limited, and it wasn't easy to find the documentation detailing what is supported.

- XL is not current (cannot compile all C++ code).
- The way we have to use GDB on BGP is hideous.
- I hate using Totalview.
- using Totalview remotely is problematic due to connection speed.

12. Please rate the following statements using the scale provided:

	Mean	Median	Mode	Variance	Std. Dev.	Response Count
The disk and tape storage provided was sufficient.	4.2	4	4	0.5	0.7	78
The policy of giving higher priority to capability jobs (>32K cores) is reasonable.	3.9	4	4	0.7	0.9	81
I was able to run capability jobs (>32K cores).	3.7	4	4	1.1	1.1	47
I am satisfied with the scheduling of my jobs.	3.9	4	4	0.7	0.9	83
I was able to utilize Eureka for visualization and data analysis.	3.4	3	3	0.8	0.9	25
Total Responses						314

Respondent comments: 6

- Have some >32k jobs in the queue right now...
- I was developing code for a tight deadline, and getting reasonable job turnaround times (within hours for 8k nodes) required special priority increases from the system administrators. I was able to get one 32k-node job run, but it failed due to a bug, and I was told it was unlikely I'd get to run another within a few days, so turnaround times at that scale are long. I have been able to get 15-minute 4k and 8k jobs run overnight in the last few days, though.
- During certain weeks it seemed that small (1/2, 1, 2 rack jobs) were started inordinately slowly, taking perhaps 3 days. This seemed to happen right after the scoring method changed.
- We have not yet explored capability jobs yet.
- 2 to 3 days waiting to get some answer is tediously slow!

- For debugging programs, it would be nice to get quick turnaround on less than 64 nodes, which seems to be the minimum allocation on Surveyor. Unless there some way that I don't know about to have something like 30 different tests all share the same allocation? Well, it's only a temporary need for this stage of the development, nothing really important unless others say the same thing.

13. What other resources could we provide that would be helpful to you now or in the future?

Respondent comments: 25

- storage space.
- GPUs.
- I desperately need MPI dynamic process management to work. :-)
- Larger BG system.
- More powerful nodes with 4 or 8 Gb memory.
- GPU clusters.
- More powerful processors.
- Different sparcification of the resources?
- Small BG/P that available all the time.
- GPU-based clusters.
- More support for fast file transfers between ALCF sites not in the globus network.
- Access to ALCF training materials (including videos of lectures) for those who cannot attend the training at the particular scheduled time.
- A queuing solution for jobs that do not scale beyond one rack.
- BG/Q.
- Access to GPU-based clusters.
- Unlimited allocations for all with no wait time . . .and a pony. :-)
- Online help.
- Powerful login nodes for compilation: it is taking much longer than other leadership-scale machines.
- More advice on codes to use for particular objectives.
- More training workshop and tutorial.
- Heterogeneous system.
- Lessons learned type of documentation that is compiled based on the experiences of users or staff while porting or performance tuning their applications. This may be quite application dependent but I am sure there will be bits and pieces that will be applicable to wider audience and save them the time to figure it out by themselves.
- Long running times (a few days) for jobs using around 1000 CPUS.
- New versions (4.5 or 4.6) of GCC, Better documentation for DCMF (hosted at

ANL, but ANL is probably not responsible for it). Some kind of faster turnaround for short (<1h) jobs even at larger scale (rather than all of them going into prod-capability and competing with the 12- or 24-hour jobs).

- Just a general comment that fits in this box. This Survey talks about a "Discretionary project" and then talks about ALCC and INCITE. None of this was defined. They are only meaningful to project PIs.

14. Are you able to achieve the science goals for your project?

Total Responses	92
Yes	78
No	14

If no, why not?

Respondent comments: 14

- In preparation phase.
- Project was not finished. Will continue next year.
- Limitations on my own time to devote to the project were the main problem - nothing related to ALCF is "at fault".
- I'm merely contributing software that should help others achieve science goals, so the question doesn't really apply to me, or not yet anyway.
- I was about to give you a Norman answer, but the reason why VORPAL results are in doubt have little to do with Computer Science, so I say yes for now.
- Require more CPU time.
- We thought we had been granted additional time to use our allocation and finish debugging our program on Intrepid, but our allocation expired before we were finished.
- Still working on it.
- We are still trying to get proper scaling beyond one rack.
- Need to complete successful scaling to >10 racks.
- I just have begun to use the machine.
- Difficulty in scaling up the run the lammeps. Still trying.
- Not yet as we are still working on scaling of the code to approach where we are on Cray platforms.
- I just started, but haven't been able to run yet because a SNAFU on the Project allocation.

15. Overall, how would you rate your experience with the Argonne Leadership Computing Facility?

Total Responses	96
Excellent	44
Above Average	42
Average	7
Below Average	3
Poor	0

Rating Average: 4.3

Respondent comments: 4

- When I contact the support staff, they are excellent and responsive and knowledgeable. HOWEVER, the Automation of support seems consistently broken. (1) I went through a process to get logged into Surveyor to find that my project had only intrepid allocation. Time on the development platform should be automatic -- why bog down the big machine with porting tests? (2) My cryptoCard for Intrepid did not come until I made contact two additional times (3 weeks later). My colleague had the same problem. (3) I was given a project name that had no allocation. All of this was resolved promptly and politely, but none of it should have been necessary. From Dec 9 to Jan 14 to be able to type "qsub".
- The IBM compiler seems to be much less forgiving than the Intel compiler we generally use. Our code is still under development, but simulations which run successfully on our clusters do not run on Intrepid without a lot of additional work.
- The maintenance of Surveyor and Intrepid should be in different time, so that there is a machine available to users. Otherwise have a small BG/P for debugging purpose which can be available all the time. Spend one workday for maintenance is not efficient.

- By average I mean neutral.

16. Would you like someone to contact you regarding your responses to this survey?

Total Responses	89
Yes	12
No	77

PART IV – Conclusion

Overall, the satisfaction with ALCF is high. The following table summarizes the individual sections and aggregate scores.

		INCITE	ALCC	Discretionary	All
Number Surveyed		393	34	420	847
Number of Respondents		119 (30.3%)	10 (29.4%)	120 (28.6%)	249 (29.4%)
Overall Satisfaction	Mean	4.5	4.0	4.3	4.4
	Variance	0.3	0.6	0.6	0.5
	Std Dev	0.6	0.8	0.7	0.7
User Support	Mean	4.5	4.0	4.3	4.4
	Variance	0.6	1.2	0.7	0.7
	Std Dev	0.8	1.1	0.8	0.8
Problem Resolution	Mean	4.5	3.9	4.4	4.4
	Variance	0.4	0.8	0.6	0.5
	Std Dev	0.6	0.9	0.8	0.7
All Questions	Mean	4.3	4.0	4.2	4.2
	Variance	0.7	0.8	0.9	0.7
	Std Dev	0.8	0.9	0.9	0.9