

Feedback on INCITE Program by Industrial Awardees

14 December 2009

Introduction

The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, established in 2003, offers over a billion hours of supercomputing time to universities, other research institutes, and industry. The INCITE program has sought to enable increased usage in industry of the some of the most powerful supercomputers and to encourage collaboration between researchers and the scientific and technical staff at DOE's computing facilities.

In late 2008, three participants of the INCITE program were interviewed (as part of a broader Council on Competitiveness case study process¹) to better understand their experiences and ascertain the effectiveness and appeal of the operational process. Participants interviewed included DreamWorks, P&G, and Boeing. This report is structured into three main areas: 1) Why use INCITE? 2) What Works Well? and 3) What Needs Improvement?

Why Use INCITE?

Based on the interviews, it was clear that use of the INCITE program is driven by three main reasons: encouragement by key personnel, ability to address very large scale problems, and the ability to validate code without impacting production machines or budget.

Key Personnel Encouraged Submissions of INCITE Proposals

All three sites had contacts, either directly with the DOE labs or with a member of the Council on Competitiveness, who encouraged submitting the proposal.

*We heard about it through Ed Leonard, who has worked with CoC for several years.
– DreamWorks*

I was informed of the program by my director [Tom Lange], so [it was] someone inside P&G. I think Tom Lange had heard from Suzy [Tichenor, of the CoC] about the program. – P&G

It came through Tom Zacharia [of Oak Ridge National Labs] and our involvement with the Council. – Boeing

¹ Supported under DOE award DE-FG02-06ER25795.

INCITE Provides Resources to Push the Limits of Scalability

Modeling more features of a product, adding another input variable, or processing a current job at breakneck speeds is a continual goal. Yet computer resources have been a gating factor. All participants viewed the INCITE program as an opportunity to access resources to achieve this dream and run a larger simulation or test the scalability of an existing program.

Ed and I also sort of developed the idea for what we would do with such a resource, which was to explore on a grander scale a problem that we had become very interested in, in the last few years. So the specific problem is to figure out how to get a scalable and yet full production quality rendering system so that the artists can still work at their best but get much lower turnaround time. – DreamWorks

One [benefit] is that we're able to look at the larger cases. So, we can look at problems that we couldn't look at before. It allows us to produce the validation data sets that we need to give ourselves the confidence that we've got a process that we can turn over to our design engineers. – Boeing

I went all the way up to 32,000 [processors], just for a benchmark though. That was just basically to say, "If I were able to use this many processors, could I produce data that much faster?" My benchmarking said yes, I could. The program scales that well for the large systems that I was simulating. – P&G

Validating Code without Impact on Production Machines or Budget

Testing new code or testing the scalability of upgraded code inside these companies often requires sacrificing valuable and needed cycles on in-house production systems. INCITE allowed participants to access larger, known HPC systems without impacting their internal production systems or purchasing a new HPC system.

An important aspect was that it would provide us access to a larger HPC system than we have. ... If I had to do that on our own production HPC system, ... I would've been competing with sort of urgent production demands for resources. The other important part was that we could work to validate our ideas on systems without affecting our production, which is our bread and butter basically. – DreamWorks

I think a major value to us is that we were running on a machine that was known to scale well, having solved other people's problems, so that if we had a problem... we would be reasonably confident that the problem was actually something in our code and not something about the way the machine was configured because we did it ourselves. So I think it was important to have sort of a known scalable piece of hardware so that we could eliminate a whole class of problems... – DreamWorks

Try before you buy. And it used to be with NASA, when NASA started up their CFD lab at NASA Ames, back in the '70s, we used to do all this kind of work down there. But then, in the '90s, NASA pretty much got out of that game. – Boeing

What Works Well?

The feedback was very positive among the three participants. Access to the HPC systems helped them to validate and expand their work. The staffs at Argonne and Oak Ridge were responsive and helpful. And it provided positive visibility of HPC capabilities and usage to management within the participant's organizations.

Staff and Responsiveness was "Awesome"

Praise for the staff and their support at Argonne and Oak Ridge was overwhelmingly positive.

The program in particular with Argonne worked very well because the number of awarded projects was rather small, especially in the first year. So it was very easy to get the time of the administrators or the application specialists, and they were very responsive to our needs. – P&G

... what [Argonne] has really improved on this last year, [is that now] they have biweekly or monthly calls with the users where it's basically an open forum teleconference, so that we can call in and complain or request information or help them change their policies if we're not getting things done the way we want them. – P&G

I had really a very good experience with the staff there. I felt we were taken care of very well by the staff there. – DreamWorks

Positive HPC Visibility to Management

By utilizing the INCITE program, visibility of the participants projects and usage of HPC technology increased at their companies.

Internally for me it's great because this is something on the radar of my management and certainly something that they've rewarded me for, taking advantage of public resources to advance our understanding. – P&G

You know, from an executive-level viewpoint, I think it's absolutely a fabulous program. – Boeing

If you go back to January, February of 2007, there was the INCITE press release with Ray Orbach and others, our boss, Mike Garrett, was present for that, for the award and for the press release. – Boeing

It was an enormous catalyst on our side also because this project was sort of high profile; executives were interested in it; everybody wanted to know how it was going, and that in itself kind of allowed a certain amount of latitude that it might not have otherwise had. – DreamWorks

What Needs Improvement?

The participants provided useful feedback on areas of the INCITE program which they felt could be improved. Three common themes amongst the participants were the application process, scheduling, and security issues. Additionally, one participant shared his challenges in using the HPC systems for interactive jobs. Finally, the lack of collaboration or lack of discussion suggests that one of the key goals of the INCITE program might not have been fully achieved.

Application Process Can Be Arduous

The application process frustrated two respondents. DreamWorks was concerned with intellectual property rights and the amount of time the legal process took, resulting in less time to actually work. In contrast, P&G was running open-source code, not concerned about the IP issues but concerned with choosing the correct resource. P&G also was concerned the length of the initial application process was becoming too tedious to attract industry users.

*We had some IP contract issues that took our armies of lawyers quite a bit of time to work through. I gather some of the other commercial applicants had similar issues, so that took awhile. We lost a lot of time in the first year to that. So we technically had two years, but I would say that we probably had closer to a year and a quarter to actually work, if that makes sense.
– DreamWorks*

[The contract process] was very smooth. I mean especially because we weren't asking for proprietary work to be done on the system, so I was aware of the work on others' contracts; and P&G does have I think a master agreement with Argonne, so we could have proceeded very easily down that path, but we chose initially to stay away from that. – P&G

When applying to the INCITE program, I was not very clear which computer I should be asking for time on. ... I think there is a workshop on how to write your INCITE proposal. Perhaps as part of that, there should be a service to advise people on what they should be asking for in terms of the resources. – P&G

I did realize that the application procedure changed from the first year to the second year. The second year, when I looked at the full application, it seemed like it was a pretty tedious process for industry. Academics are used to these long proposal processes; but in industry, we're used to writing a one-page memo and asking for something, and we either get it or we don't. The application process for some of your industry partners might click them off, but the renewal process was very straightforward and painless. There's a big discrepancy between a new application and a renewal application. – P&G

Scheduling Affected Desire to Run Very Large Jobs

Scheduling very large jobs can be a daunting task. The availability of system resources for large jobs was difficult to predict. Both DreamWorks and P&G revealed they ran smaller than anticipated jobs because of perceived scheduling issues.

Staff at both Argonne and Oak Ridge offered solutions. Argonne established a reservation system while Oak Ridge offered to shut everyone else out of the system to allow DreamWorks to test the scalability to 10,000 nodes. However, DreamWorks did not feel comfortable interfering with the work of all the other users on the system and only scaled to 600 nodes. Both DreamWorks and P&G stated that they were able to validate the scalability of their application without reaching their maximum node/cores initially desired.

One of the main users on the system is running small jobs for very short amounts of time, and he has done this basically to work around the fact that he was not able to get larger jobs on the system to run longer just because of the demand on the system when he was trying to use it. ... He doesn't know far enough in advance to schedule the resources. Now I made him aware that we can get reservation blocks from Argonne, and that has worked very well for me. – P&G

*The typical job was not very large because in practice, I didn't really have a lot of patience waiting for the scheduler to give me the nodes I wanted, so typically I probably was running on the order of a few dozen to maybe a hundred nodes; and the largest I ran, I think, was about 600. ... But on the flipside, again to the lab's credit, they seemed quite willing actually to shut everybody else out of the machine for some period of time and let me run at whatever scale I wanted.
– DreamWorks*

Protection of Code and Data

For sites running proprietary code or data, protection of this info was extremely important. The issues ranged from ownership of the software to isolation of the private data to the transfer of private data between the remote site and the lab. Both Boeing and DreamWorks were able to achieve a comfortable protection level by using high encrypted connection levels and firewalls. P&G in contrast was conducting an open science project and had no concerns with security of the data.

In order to do the visualization, the data transfer is such that you can't do it from Seattle to Oak Ridge. You have to move that solution back to Seattle. And therefore, it takes Oak Ridge out. For example, [visualization] also ran into some problem with the bandwidth, and also with the security. ... We do not want to advertise our data elsewhere, everywhere else around the world. And we have to go through a secure shell. We have to go through a whole bunch of firewalls. So, security is also a problem, yes. – Boeing

*We needed to retain ownership of the software that we were porting to these systems exclusively.
– DreamWorks*

We had considered going to the point of sending people, maybe with hard drives full of our production data. We never really reached that point. We actually managed to get almost everything we wanted by my logging in remotely using highly encrypted SSH connection.
– DreamWorks

I thought [the data protection] was great. I really felt like the people I was working with bothered to truly understand what my concerns were ... acting very much in my interest and proposing solutions that would address my needs. So that was a very positive aspect, I think, of the whole thing. – DreamWorks

I am clear on the level of data protection and I am comfortable with that level of data protection. This is an open science project, so I don't expect confidentiality. – P&G

Technical Challenges for New HPC Users

Unlike Boeing and P&G, working with Cray supercomputers was new to DreamWorks. The participant had not worked on a Cray and was unfamiliar with its compilers and its operating system. Not only did technical limitations affect him, but he also felt scheduling was not supportive of interactive jobs. Though the challenges were greater than for the other participants, he was able to achieve his objective and was extremely satisfied with the help and responsiveness of the program.

There's no mechanism to connect to a running process on that machine through a socket, and what I had hoped to do was basically set up the client for our rendering tool on my machine here in Southern California and remotely run – basically manage a rendering session where the work was being done on the machine at Oak Ridge and essentially pipe back the images that were being produced at whatever break-max speeds I could've seen to my display locally.
– DreamWorks

I would never have imagined that I could've essentially asked for a feature in an OS on a machine like this and actually had people say, Oh, yeah, we want that too, and Yeah, we should do that. Then basically within a year or a year-and-a-half, Cray had it. It's kind of amazing. I really didn't expect that. – DreamWorks

I think I underestimated the challenges of the differences between the two operating systems, but I'm not sure that it would've been any easier if I had picked any other supercomputer to run on.
– DreamWorks

[I need] to further explore the scheduling challenges posed by computing, interactive, and batch jobs. – DreamWorks

I think it'd be also interesting to explore what could be done also related to the IP issues but in a way that would allow a more sort of intellectual collaboration. – DreamWorks

Collaboration with Scientists and Technical staff Was minimal

As stated in INCITE's 2010 Call for Proposals, "A key aspect of the program is to connect leaders of the projects with scientific and technical staff at the computing facilities." This goal did not seem to be fully achieved for these industry users. Based on interviews with the participants it was not a point that was emphasized, suggesting collaboration with the scientists did not occur. DreamWorks did regretfully express their initial hope of intellectual collaboration did not come to fruition. Due to the intellectual property issues, DreamWorks realized it was not going to happen.

We had developed something [together], but I had done all the coding. But he knew what I did because we talked it through. I don't really think we would've been able to expect that he would somehow act like he didn't know this if say Pixar came to do a contract with Oak Ridge, but I think that we must have had a much more sort of low-level concern that if they've seen anything, they own it. – DreamWorks

We're getting access to experts in the sense that they're helping us set up our systems, and we could've used them; but as a practical matter, we couldn't. – DreamWorks

*In effect, ... I didn't work with [him] on any of the ideas that I was really deploying.
– DreamWorks*

INCITE – A Success for All

In the end, all three participants were satisfied with the access, the support, and their results. All three would consider reapplying.

The answer to all of those things [apply again or recommend] is absolutely. I would apply again. I mean I would apply with perhaps more open eyes this time, but I would, I think it was a huge success for us. – DreamWorks

I hope that the INCITE program remains intact and continues for a long time to come. – P&G

*The question is, who can not only discover it first, but who can get it into their products and services first and capture the market? That's what's critical. INCITE is an enabler to doing that.
– Boeing*

Appendix A: Additional Quotes and Issues

Access to New Hardware

And then, the third one is – and this is a big one – is that it gives us very early-on access to the new emerging hardware. – Boeing

Running Large Jobs

Difficult to run a large scale job because so many users running small jobs, have to wait a long time for scheduling – since operating in interactive mode and no idea given to when time slot available to get 10,000 nodes – but on the flipside again to the lab's credit, they seemed quite willing actually to shut everybody else out of the machine for some period of time and let me run at whatever scale I wanted. Because I felt, I did have to admit feeling a little bit out of place essentially kicking off everybody else in the world, who was trying to get something done on this machine just so I could test something but... (National labs may have to encourage use).

Use of Training:

We are mainly working remotely. We have participated in some of the user group meetings or training sessions that were given at Argonne, but we're very comfortable working directly from our offices here in P&G. – P&G

They offered but I don't ... It wasn't worth, no. I mean they had some sort of training on how to use the Cray itself and actually I was going to go and for some reason didn't. I can't remember why. That probably would've been somewhat helpful, but I was able to figure out how to use it, using the online documentation so yeah; no, not a ... – DreamWorks

Proprietary Codes

Proprietary codes have a hard time scaling and, as well, there are license issues. – P&G

Computer Time Allocations Accounting Report Desired

One of the things that I thought should have been improved is at Argonne, they don't do very good accounting of how much of the resource you've been using; and it seems that they're not really enforcing the limitations in terms of the allocated time, or at least they're not, they haven't shared basically the amount of the allocation that you've used to date. – P&G

- *Stats on use of computers and which are better for my apps.*
- *Want to know how much used, how much left.*
- *Told to work as normal but still felt uncomfortable pushing it to limit.*

If we made any mistake, it's that we overestimated how much time we needed [to validate their code] – DreamWorks