High Performance Computing Facility

Operational Assessment Report

for the

Argonne Leadership Computing Facility (ALCF)

Pete Beckman, ALCF Director
August 6, 2010

Executive Summary

The High Performance Computing & Networking Facility Operational Assessment (HPCNOA) is an Office of Science (SC) programmatic management tool used to assess the status and quality of operation of scientific user facilities. The HPCNOA ("OA" hereafter) is used for evaluating the Facilities' plans for providing high-performance computing and network resources as well as support to the scientific user base. Relevant information from the OA is used to respond to the succeeding year's annual Operational Analysis data call for major Information Technology operations from the Office of Management and Budget (OMB) and the Office of the Chief Information Officer (OCIO).

The guidance for the FY '10 Operational Assessment indicated that facilities would provide:

- Responses to recommendations from the 2009 OA review,
- Performance data against the previous year's baseline plan,
- Performance results and projections for the next year.

The reporting period for financial data is fiscal year based, and for the current review, FY10 data is reviewed. All other performance metrics are for actual data from January through July, and for projected data from August through December for the review year.

There were six recommendations from the previous year's Operational Assessment Review (OAR). The ALCF has chosen to address some of the comments as well. The recommendations and the selected comments are addressed in the first section, titled "Responses to 2009 OA Recommendations".

For performance against the baseline plan, the guidance directed the Facilities to respond to seven different metric areas:

1. Customer Results
2. Business Results
3. Strategic Results
4. Financial Performance
5. Innovation
6. Risk Management
7. Cyber Security

Each of these metric areas is addressed, in order, in the *FY 2010 section* following the recommendation section. With the exception of *Financial Performance*, each section follows the format of: charge, an overview, and then each metric within the metric area is covered. For each metric, the description (from the approved metrics document for FY2010) is provided first, followed by this year's result, and, where appropriate, the prediction for the remainder of CY2010. If necessary, additional details (such as examples) are provided. The *Financial Performance* section follows the template provided by the program office.

Responses to 2009 OA Recommendations

Baseline Area 1: Customer Results

1) Comment and/or Recommendation: Publish survey

Comment and/or Recommendation

Consider publishing the survey results and ALCF responses on the ALCF website. This helps users understand that their input has been received and the center has taken steps to explain or improve the environment.

Response

The December 2009 survey results and ALCF response is available on the ALCF wiki at https://wiki.alcf.anl.gov/index.php/User_Satisfaction_Surveys. The ALCF will continue to publish future survey results and responses on the ALCF wiki. INCITE and ALCC users will be notified when the survey results are published.

2) Comment: Ticket categorization

Comment

ALCF receives an average of 1,546 new tickets per month, compared with NERSC's 500. The Operational Assessment report could categorize these tickets by topical area – account support, programming questions and the like. It should also give more specific example of user support.

Response

Due to a programming error in the analysis script, the statistics reported in the 2009 Operational Assessment Report (OAR) were incorrect. We share the request tracking system with other Argonne divisions and the system includes queues for non-ALCF resources. The analysis script counted other, non-ALCF related tickets when calculating the metric results. The script was corrected in September 2009. The corrected 2009 OAR data is shown below.

During the period August 1, 2008, to July 31, 2009, there were 3,438 tickets opened by users, an average of 287 tickets per month. 72% of the tickets were addressed within three days. The 2009 metric was 66%.

The ALCF has begun the process of converting to categorization of tickets. Some customized programming on the problem ticket system was required. The first category, account support, was put into place in March 2010. ALCF is in the process of determining the remaining set of categories and plans to survey NERSC, OLCF and NCSA for their categories and utilize this to inform the category selection. ALCF expects to have the process completed by December 2010.

We have provided some specific examples of user support in the *Baseline Area 1: Customer Results* subsection of the *FY2010 Metrics* section.

3) Comment: Discretionary projects in annual survey

Comment

The discretionary project users should be included in the annual user survey.

Response

At DOE's direction, ALCF concentrates its support resources on INCITE projects. ALCF has many more discretionary users than INCITE users, and including discretionary users in the survey would skew the results toward discretionary projects. In 2010, the ALCF will be adding ALCC projects to the INCITE projects in the top tier support category. With DOE's concurrence, the ALCF will continue to include only INCITE and ALCC projects in the annual user satisfaction survey. The ALCF will continue to use informal surveys to gather feedback from all users.

4) Comment: More information about user surveys

Comment

The Operational Assessment Report should include more information about the user surveys so that reviewers can have a more complete understanding of the survey results. This could include response rates, detailed survey results (e.g., made available online), and summaries of workshop surveys.

Response

Full details on the survey are provided in the user results published on the ALCF wiki (https://wiki.alcf.anl.gov/index.php/User_Satisfaction_Surveys). Requested details for the December 2009 user survey are also included the *Baseline Area 1: Customer Results* subsection of the *FY2010 Metrics* section.

Baseline Area 2: Business Results

1) Recommendation: Low utilization

Recommendation

The ALCF should investigate (if necessary) and report on the reasons for the utilization rate of 65.3%, whereas utilization of 80% may be expected. Is the low utilization number due to lack of allocation, lack of jobs, scheduling inefficiencies, an artifact of how it was measured, or some other factor?

Response

ALCF investigated the low utilization reported in the 2009 OAR. There were two main reasons for the relatively low utilization; 1) a bug in the backfill algorithm and 2) job mix combined with the partitioned nature of the Blue Gene. A bug fix for the backfill algorithm was deployed on Intrepid in early 2009. This change was reflected in the increased utilization during the last half of the 2009 OAR reporting period, but it was unable to offset the first half low utilization. In spring 2010, we addressed the job mix/partition-nature issue with the implementation of a series of scheduling policies. These scheduling policies were intended to improve capability utilization by favoring jobs 8 racks and larger, and restricting small long jobs to a subset (40%) of the system. This change had the added benefit of dramatically increasing utilization because it basically reserved a large portion of the system for large jobs. As a result, utilization was significantly better this year (81.1%) than last year.

2) Comment: Calculation of availability

Comment

The Operational Assessment Report should provide more complete information on how availability numbers are calculated. For example, the availability chart should state over what period the numbers were averaged: 30 day rolling average? Monthly averages?

Response

Full details on the method used to calculate availability have been provided in the 2010 OAR.

3) Comment: Utilization statistics

Comment

The Operational Assessment Report should provide more utilization statistics, e.g., time used by INCITE, Discretionary, or any other categories; and scientific time by science discipline.

Response

Charts showing utilization by science domain and utilization by allocation category (INCITE and Discretionary) have been provided in the 2010 OAR. The ALCF will continue to provide this information in all future reports.

4) Recommendation: Evaluate Cobalt vs. LoadLeveler

Comment

The Operational Assessment Report should include an evaluation of continuing to use Cobalt versus IBM's LoadLeveler.

Response

ALCF performed a comparison of Cobalt and LoadLeveler and evaluated whether to change resource managers. While it is true that Cobalt is not yet as mature as resource managers that have been around for decades, such as LoadLeveler, Cobalt provides some significant advantages. The decision was made to remain with Cobalt for the reasons listed below.

1. LoadLeveler does not allow for the specification of an alternate software stack (OS) with the submission of a job. This is required for several of the INCITE projects, as well as a number of the system software discretionary projects.
2. Cobalt is open source and also provides clean, modular API for developing and deploying new resource management and scheduling algorithms. The access to source code, and the clean API is what allowed use to make some of the modifications that improved our

utilization from less than 65% last year to more than 80% this year. As a result of modifications to Cobalt, we are now able to run capability jobs through the scheduler with reasonable wait times and largely without the use of advanced reservations, something that would have been difficult, if not impossible, to do without the ability to modify the source.

3. Cobalt provides a simulator that simulates not only the resource manager, but also the Blue Gene control system. It includes the ability to inject errors. The simulator can be run with historical data to test and validate proposed resource manager/scheduling modifications. This allows research projects from MCS (Argonne Mathematics and Computer Science division) develop, debug and test the impact of new scheduling or resource manager algorithms and policies without having to run on the actual system. The simulator was also used recently to develop a new harness for a regression test suite. The ability to simulate the resource management environment of the production system allowed the developer to try new methods of fully loading the system and gathering data from the control system without impacting the production system. Once the new harness was fully tested on the simulator, the developer was able to bring the new system up on Surveyor during a maintenance outage with only a few small tweaks. Another example of the benefit of the open source, modular API of Cobalt: we are currently in the process of developing a method of doing backfill based on historical trends for users rather than their requested wall times. This improvement in backfill performance resulted in improved response times and utilization; the paper was published at IPDPS, and won a best paper award.

5) Comment: Expansion factor

Comment

The Operational Assessment Report should describe and possibly improve the methodology used to calculate the expansion factor. For example, if wait times are calculated since job submission and users are allowed to submit hundreds of jobs at one time, this will inflate the factor. Some schedulers limit the number of jobs that a user can submit directly into the "run" queue and distinguish between jobs that are not yet eligible versus those eligible to run. The transition between these two states can be used to measure a more realistic wait time.

Response

Currently, ALCF calculates the expansion factor as:

timestamp of the job when it was submitted minus timestamp of the job when it was started

As pointed out by the reviewer, this causes our expansion factor measure to be artificially high (e.g. worst case). In the past year, we attempted to leverage statistical sampling capabilities in the resource manager. In the process of analyzing the data, we discovered a significant sampling error. Without the ability to utilize the statistical sampling, the data in the 2010 OAR remains calculated the same way as past years and is again artificially high. We have just completed work to record all state transitions and are in the process of validating in preparation for deployment. In addition, the scheduling policy currently only allows 20 queued jobs and 5 jobs running per user. We expect to be able to provide more realistic wait time data at next year's OAR.

6) Comment: User error category

Comment

Do the jobs marked as "user error" include ones that exceed their time limit, and if so, are there hung jobs in the mix?

Response

Yes, the "user error" category includes jobs that exceed their time limit. As a result, hung jobs are included in that category.

7) Comment: Consistent format and labels for metrics.

Comment

Make sure metrics are called out clearly and are consistent with the document.

Response

In the 2010 OAR, we have used the same format through the FY 2010 Metrics section, with a metric note in each subsection. In the case where multiple metrics can apply (e.g. availability), we have marked the metric that applies and why it applies.

8) Comment and/or Recommendation: Increase metrics.

Comment and/or Recommendation

Consider increasing metrics for future years/systems.

Response

Metrics have been increased over past years, both for availability and for capability. ALCF current metrics are standard for leadership class facilities and were developed with input from the ALCF program manager.

Baseline Area 3: Strategic Results

1) Recommendation: Publication reporting

Recommendation

The ALCF should track and report refereed publications and give more precise numbers of publications, e.g. "exactly 38 reported by users."

Response

The ALCF has improved its process for tracking user publications. Included in this year's report is a discussion about the new process, as well as the precise numbers of publications tracked. This information can be found in the Baseline Area 3: Strategic Results subsection of the FY 2010 Metrics section. We will continue to track and provide precise numbers in future OARs.

Baseline Area 4: Financial Performance

1) Recommendation: Complete picture of budgeting process

Recommendation

The financial section of the Operational Assessment Report should include a more complete picture of the budgeting process to demonstrate that the budget is well managed and has low risk of significant cost variances. This should include a comparison of planned costs to a "constraint" in the budget with highlights for what major categories the funds will be spent towards. Some may refer to this as a baseline but it may also be the authorized funds DOE gave/will give to the program. Showing the cost projections compared to the constraint and how the program plans to deal with projected costs over and under is needed to determine whether the program has the financial resources to complete to planned objectives.

Response

ALCF follows a process that was developed for annual budget formulation and monthly monitoring. Details about the process are provided in the *ALCF Budget Formulation and Monitoring Details* subsection of the *Baseline Area 4: Financial Performance* section. In addition, full details, as directed by the program office, including a comparison of budget versus actuals and cost projects, are provided in the *Baseline Area 4: Financial Performance section.*

2) Comment: Carry-forward

Comment

Is the committed cost budget authority that is carry-forwarded expected to pay for next year's contractual obligations or are these contractual obligations for FY09 that vendors did not deliver?

Response

Carry-forward is usually planned to pay for year-end commitments (often contractual obligations for the current year that have not yet been delivered) and continuity of operations for the next fiscal year.

3) Comment: Impact of not meeting staffing plan

Comment

What is the contingency plan if the ALCF does not meet its aggressive goal for hiring staff? What will happen to the budget? Depending on the overhead application this could be between $3.6M and $7.2M in total costs.

Response
The ALCF realizes we have an aggressive staffing plan. We will use contract and matrix staff to mitigate the risk, where possible. Because of the delay in hiring, it is causing a cost variance.

4) Comment: Power costs

Comment
Why are the power costs relatively constant in FY10? Are there guarantees that power rates will not increase and no expectation of additional equipment in FY10?

Response
The move to the new TCS building, which is more energy efficient than the previous building, is expected to offset any escalation in power rates in FY10. No new machines are planned for FY10, so machine power consumption should remain about the same.

Baseline Area 5: Innovation

1) Comment: Interactions: MCS or ALCF?

Comment
The discussion of the innovations as "interactions" made it unclear to what extent these activities were staffed and paid for by the ALCF, or whether some of them were performed by MCS staff on separately funded projects. Is there a Computer Science group within the ALCF that is responsible for these collaborative activities?

Response
ALCF does not have a Computer Science research group. We are fortunate that MCS does have a Computer Science research group that enjoys addressing real life issues that are related to their research focus. In response to this comment, ALCF has clearly delineated the responsible party for the contributions outlined in the *Best Practices and Innovations* subsection of the *Baseline Area 5: Innovations* section of this year's report. Many of the contributions are not from a single team, but are the result of a cooperative effort involving the ALCF Performance Engineering group, the ALCF Advanced Integration Group, and the MCS System Software teams.

Baseline Area 6: Risk Management

No comments or recommendations

Baseline Area 7: Cyber Security

No comments or recommendations

FY 2010 Metrics

Baseline Area 1: Customer Results

Charge Question 1: *Are the process for supporting the customers, resolving problems, and communicating with key stakeholders effective?*

FY '10 Operational Assessment Guidance: *Customer results cover measures of customer satisfaction with the services provided by the Argonne Leadership Computing Facility (ALCF), the resolution of problems encountered by the users, and overall support to the user base. For each of the following metrics, report results and provide projections using a methodology developed with the concurrence of the Federal Program Manager: Customer Satisfaction, Problem Resolution, and User Support.*

ALCF User Services and Outreach Overview

All active users of the ALCF belong to an approved active project. Membership in a project is made up of a principal investigator (PI), co-PIs, and project members. All INCITE, ASCR Leadership Computing Challenge (ALCC), and some discretionary projects, are given an allocation on ALCF production resources Intrepid (computer resource) and Eureka (visualization resource). INCITE and ALCC projects may receive time on the test and development compute resource, Surveyor, under special circumstances (e.g. project work is disruptive for other production users, typically this is a computer science system software project). Discretionary projects typically start on Surveyor as they work to scale up their applications. For CY2010, 60% of the computing time on Intrepid is allocated to INCITE projects, 30% to ALCC projects, and 10% to Director's Discretionary projects. ALCF provides all users with:

- Services Desk support
- Issue tracking and resolution
- Detailed web-based documentation
- Workshops and tutorials

In addition, INCITE and ALCC projects receive an increased level of support in the following areas:

- Customized introductions to ALCF resources
- Dedicated support from catalysts and performance engineers
- Specialized workshops and tutorials
- Assistance with preparing quarterly reports

Figure 1 shows the number of ALCC and INCITE projects over the lifetime of the ALCF. The spike in January 2010 is due to both 2009 INCITE and 2010 INCITE projects being active during the period of January 1 through January 14.

Figure 2 shows the number of Discretionary projects over the same time period. The number of discretionary projects fluctuates over time, because discretionary projects can begin at any time, and many are active for only a few months.

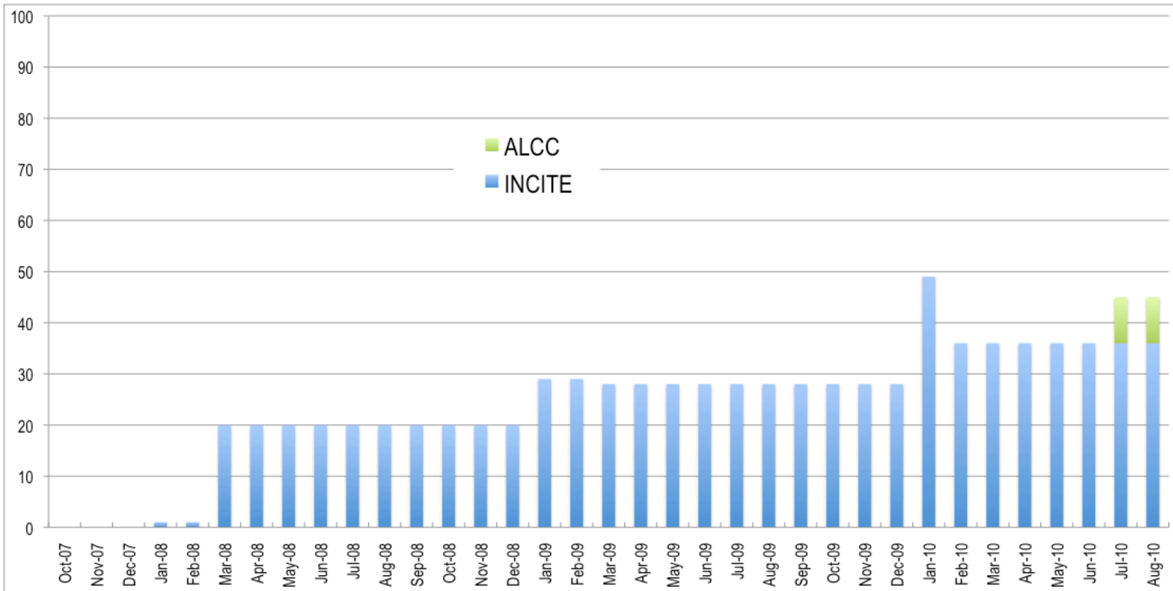

Figure 1: Number of active ALCC and INCITE projects by month.

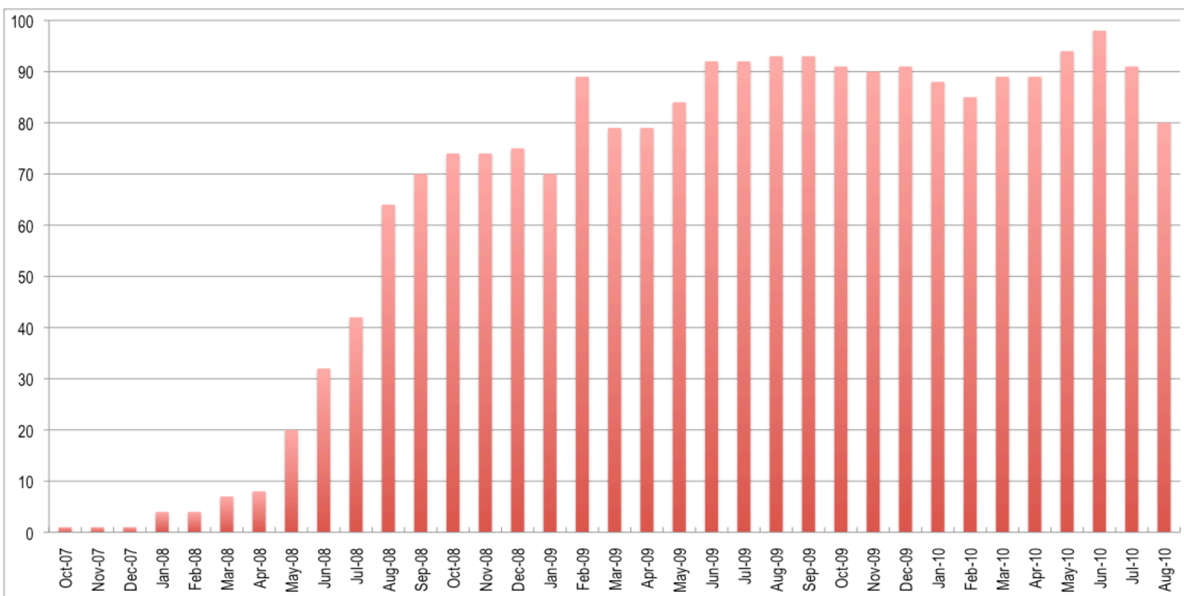

Figure 2: Number of active discretionary projects by month.

Metric 1.1 – Customer Satisfaction

The ALCF will conduct a user satisfaction survey in December of each year. The survey will be used as an annual tool to help gauge user satisfaction and to help identify areas for continued improvement. For the formal user satisfaction surveys going forward, the overall ALCF score will be at a level of 3 or higher (out of 5.0).

Metric 1.1 - Results

ALCF conducts a formal annual user satisfaction survey in December of each year. This survey is used to gauge the overall satisfaction of the INCITE, and, beginning in 2010, ALCC PIs. The survey gathers feedback on their experience with the ALCF support as well as their level of satisfaction with the ALCF resources. The 2008 annual survey was designed and reviewed with the help of a qualified external reviewer for coverage, clarity, and applicability to the user base. The 2009 survey was updated slightly. As part of the annual survey, users are asked to rate topical areas (or questions) on the survey using a Likert scale and are given the opportunity to respond to open-ended questions in a more qualitative manner. Surveys were sent to INCITE PIs in December 2009, with responses due in January 2010. PIs from 71% of the INCITE projects responded to the survey. Of those responding, 85% were part of projects renewed for 2009; the rest were new 2009 projects.

Survey results:

- Overall survey score: 4.6/5.0
- Overall satisfaction with the ALCF: 4.7/5.0
- ALCF staff provides accurate, complete assistance: 4.67/5.0
- Satisfaction with job scheduling and turnaround: 4.78/5.0

Metric 1.1 - Additional Details

Survey Open-Ended Questions

- The catalyst team was mentioned collectively (and many members, individually) numerous times for their very significant contributions during 2009.
- A few of the areas noted as weaknesses or offering opportunities for improvement include:
    1. Online documentation split between web and wiki
        - Action: All online documentation is being moved to the wiki.
    2. Long account creation time for foreign nationals
        - Action: Procedures have been developed and implemented to streamline the approval process for foreign national access. All required information is collected via on-line forms and support staff track and expedite requests. The process remains limited by the Argonne processes as mandated by DOE.
    3. Lack of weekend/evening support
        - Action: DOE has not funded ALCF for weekend/evening support. After hours support is provided on a best-effort basis.
    4. Lack of SLURM (interactive sessions)
        - Action: The addition of interactive capability to Cobalt is in progress, in the meantime, a limited workaround has been provided.
    5. Lack of MRNet
        - Action: We are currently investigating the request.
    6. Lack of ability to run MPI programs on the front-end
        - Action: Running MPI programs on the login nodes adversely affects

the other users, and is not the intended use of the login nodes. Users requesting to run MPI programs on the login nodes are asked to use eureka, which is designed specifically for post-processing, data analytics, and visualization. Eureka can launch single node jobs very efficiently.

7. Long scheduled downtime every week
    - Action: Scheduled downtimes have been reduced to 8 hours every other week, with exceptions for high-priority maintenance.

Results for the 2009 annual survey was posted to the ALCF wiki (https://wiki.alcf.anl.gov/index.php/User_Satisfaction_Surveys). Going forward, all survey results will be posted to the wiki.

In addition to the annual INCITE/ALCC user satisfaction survey, the ALCF uses other means of feedback (both formal and informal) to better understand how all ALCF users view our services and capabilities. Surveys are taken at each of the ALCF workshops and tutorials to assess the value delivered and to identify areas for improvement. For example, the 2010 Leap to Petascale workshop consisted almost entirely of hands-on time, since feedback from previous and potential attendees indicated this would be the most valuable format. The feedback data is disseminated to both the ALCF staff and the DOE program manager.

Metric 1.2 - Problem Resolution

ALCF problem resolution is measured on problem reports sent by the users to the ALCF problem tracking system. For FY10, 73% of ALCF user problems shall be addressed within three working days, either by resolving the problem or informing the user how the problem will be resolved. The percentage addressed within three working days will reach 80% in FY11.

Metric 1.2 - Results

During the period August 1, 2009 to July 31, 2010, users opened 3,214 tickets. 79% of the tickets were addressed in within three working days. During the period January 1, 2010 to July 31, 2010, users opened 1,976 tickets. 79% of these tickets were addressed within three working days.

Metric 1.2 - Projections

Though the number of tickets fluctuates from month to month, we assume the average rate will be similar for the remainder of the year. We project that users in CY2010 will open a total of 3,250 tickets and that ALCF will respond to 80% of them within three working days.

Metric 1.2 - Additional Details

This year's 79% of tickets addressed within three working days measurement was a significant improvement over the 72% from last year. We believe we are on track to reach the FY11 target goal of 80%.

ALCF uses a request tracking system based on the Request Tracker (RT) open source software. All support request emails are tracked, and their progress and resolution logged. The ALCF services desk is available via phone as well. Most calls are connected to existing tickets, and are tracked as part of the original ticket. A small number of calls are for new issues. Although ALCF carefully tracks problem tickets and their resolution, the primary focus is on solving the user's problem, not rapid ticket closure. The ALCF user base is a highly sophisticated group, and ALCF strives to work in collaboration with each user. The ALCF service desk responds to a wide variety of issues, from a simple password reset to complex hardware errors. The following are some representative issues and their resolution.

Examples of specific user support

As recommended in the review of the ALCF 2009 OAR, below are some specific examples of support provided to users. We have chosen a representative sample of the ALCF problem tickets.

New Account Request (alcf-support #34808)

An email request was received from someone who had heard about the ALCF from one of our collaborators and wanted to apply for an ALCF account. Support staff sent an email to the requester providing information about types of accounts (INCITE, Director's Discretionary) and their restrictions. Links to the account and project request pages were also provided, along with a suggestion to call the ALCF Helpdesk number if he had further questions or needed additional assistance. The ticket was closed.

Reset the ALCF Web Password (alcf-support #47827)

An email request was received from a user to reset their Accounts web page password. The only password provided for ALCF users is for their Account web page, which typically isn't accessed very frequently. Because of this, users frequently forget their password. [Passwords are not allowed for resource access (we require cryptocards for production resources, or cryptocard or ssh keys for research resources)]. The user in this ticket was advised to call the services desk number, as identity verification to reset passwords is not allowed using email. He called and a support staff verified his identity and provided a temporary Account page password. The support staff verified the user changed the password to a new one while on the call. The user reported that he was able to login with his new password and the ticket was closed.

Account Renewal Request that required a 593 Renewal (alcf-support #33667)

A user requested that his account be renewed (ALCF accounts expire after a period of time determined by the account type). Account renewals require the approval of the user's account sponsor. The ALCF support staff handling the ticket received approval from sponsor and completed the next step in the user's account renewal process. This user was a foreign national and the approval for his 593 form (Foreign Visit & Assignment Request) had expired. Site access for foreign nationals who already have site access approval requires the renewal of a Visit/Assignment Request form (593) upon expiration.

The user's information needs to be current, so the support staff contacted him and asked him to log into his Account page and confirm, or update where needed, his information was current. Once the staff had validated the information was current, the staff initiated the 593 renewal process, notifying the user that the process had been initiated. The user was also notified that the administrative security specialist would contact him in the case that additional information was required. Once the support staff received notification of Argonne Security clearance, the account was renewed with new 593 and account renewal expiration dates and the ticket was closed.

Cryptocard out of Sync/Locked (alcf-support #50193)

A user filed a problem ticket reporting that he was unable to login to Intrepid. The ALCF support staff contacted the user by phone, verified his identity, and gathered data from the user to determine if it was one of the following issues:

- Token was locked out
- Token response wasn't entered in UPPERCASE
- PIN was not prepended to the Cryptocard response
- Token was out-of-sync

From the data offered by the user, the support staff was able to determine the user's cryptocard was out-of-sync. The user was given the URL to the re-sync instructions on the website. The support staff stayed on the phone while the user walked through the instructions. Once the user had completed the instructions, he was able to successfully login to Intrepid and the ticket was closed.

Resource manager misconfiguration after a driver upgrade (alcf-support #50199)

A user filed a problem ticket reporting that he was having trouble submitting jobs on Surveyor. He was getting a "Failed to contact system component" error message. Support staff routed the ticket to the operations staff member on call, and notified the user that the ticket had been escalated to the operations team. The operations staff realized that permissions on a server directory the resource manager accesses had become set to incorrect values and corrected the problem. The operations staff notified the user that the problem had been corrected and, upon receiving a response back from the user that they were able to submit jobs again, closed the ticket.

Boot Failure (alcf-support #50098)

A user filed a problem ticket reporting that jobs were failing with a "Failed to boot the partition (failure #37)" error. He provided the job ID number, and the support staff routed the ticket to the operations staff member on call, and notified the user his ticket had been escalated to the operations team for further diagnosis. The operations staff investigated the issue, inspecting the relevant hardware, and determined that a part replacement was required. The part was replaced, the operations staff verified that jobs were able to boot on the partition and notified the user. The ticket was closed.

Syntax Error in Driver File (alcf-support #49359)

A user filed a problem ticket reporting that there was a syntax error in the mpixlf90 wrapper to the BG/P driver. Support staff routed the ticket to the operations staff member on call, and notified the user that the ticket had been escalated to the operations team. The operations staff investigated and determined that there was indeed a syntax error in the wrapper. The error was corrected and the user notified. The staff verified the problem was resolved and closed the ticket.

Network Link Error (alcf-support #48479)

A user filed a problem ticket reporting that his jobs on Surveyor were failing with the error "BE_MPI (ERROR): block error text 'E10000network link error.'" Support staff routed the ticket to the operations staff member on call, and notified the user that the ticket had been escalated to the operations team. The operations staff notified support staff a failing optical transceiver that was replaced that morning was believed to be the cause of the error the user reported. Support staff notified the user, verified the user's jobs were now successful and closed the ticket.

Metric 1.3 – Workshops, Tutorials, User Teleconferences and Application Support

ALCF will track its workshops, tutorials, monthly user teleconferences and application support provided to users and will provide quarterly reports to DOE.

Metric 1.3 - Results

Workshops and Tutorials

ALCF conducts workshops and tutorials to train new users, help existing projects improve, and introduce new techniques to experienced users. ALCF uses feedback from surveys throughout the year to establish interest and determine specific agenda items. ALCF workshops and tutorials are well attended, with attendees giving overall positive feedback. Detailed reports are provided to DOE after each workshop and in the *Customer Results* quarterly reports.

ALCF conducted a total of six (6) workshops and/or tutorials in the past year, five of which occurred between January 1, 2010 and July 31, 2010:

- TAU Workshop (Sept. 22-23, 2009)
- INCITE Getting Started Workshop (Jan. 27-29, 2010)
- Welcome to Magellan Day (Mar. 23, 2010)
- INCITE Proposal Writing Seminar/Webinar (May 17, 2010)
- Leap to Petascale Workshop (May 18-20, 2010)
- Large-Scale System Monitoring Workshop (May 24-26, 2010)

The Magellan workshops are reported in the OAR as the ALCF User Services and Outreach group provide logistical support for these workshops.

User Teleconferences

In CY2010, ALCF held six monthly user teleconferences, skipping only January, when it conflicted with the INCITE Getting Started Workshop. Agendas, attendance, and results are reported to DOE each quarter in the *Customer Results* report.

ALCF conducts monthly user calls with INCITE users. These calls are scheduled on the last Thursday of each month. Support staff, including catalysts, performance engineers, user services, system administrators, and storage engineers, are available on the call to answer any questions.

In the early days of the Intrepid resource, the calls were well attended, but as the system became more stable, attendance dropped. ALCF responded by opening the calls to all INCITE users, not just PIs, and advertising specific topics in advance. Attendance is still low, but the calls provide an important mechanism for users who have unmet needs to discuss them with a cross-section of ALCF staff.

In July 2010, users involved with the recently added ALCC projects were invited to the monthly user teleconference. Since ALCC projects receive the same level of support as INCITE projects, they will be included in the user teleconferences.

Application Support

ALCF records key support efforts and reports it to DOE on a quarterly basis in the *Customer Results* report.

The ALCF staff provides customized support to individual projects. INCITE projects, in particular, receive a high level of service, with each assigned to an ALCF staff member familiar with both the computational techniques and the domain science of the project. The following is a list of some significant achievements during the reporting period. Examples of the support are provided below, after the *Metric 1.3 Projections*.

Metric 1.3 - Projections

Workshops and Tutorials

ALCF plans to hold three more workshops during CY2010:

- Performance Tools and Scaling
- Early Science Program Kickoff
- Welcome to Magellan II

User Teleconferences

ALCF plans to hold four more user teleconferences, skipping the December one due to Argonne National Laboratory's holiday shutdown.

Application Support

ALCF will continue strong collaboration with users and will build on the current success.

Metric 1.3 - Additional Details

Examples of Applications Support

GPAW

ALCF staff has a multi-year collaboration with the Center for Atomic-scale Materials Design (CAMd) in Lyngby, Denmark to increase the scalability of the GPAW code. As a result of the collaboration, the code now scales to 8-racks on VN (virtual node) mode (i.e. all cores within a node are running an MPI process). In addition, an Argonne summer intern from CAMd is working with ALCF applications staff to integrate HDF5 into GPAW. This work will enable the restart of large calculations with wave functions.

VASP

ALCF worked with an INCITE user to identify a bug in the VASP NEB algorithm.

cobalt-subrun

ALCF staff developed the cobalt-subrun feature to support ensemble runs (bundling multiple, smaller runs into one large job). The launch customer for it was Benoit Roux's GateMechProtein, but since then it has been heavily used by some projects, for example:

- Optimization on Nonlinear Dynamics for the APS Upgrade (Discretionary)
- Direct Multi-objective Optimization of Storage Ring Lattices for the APS Upgrade and Beyond (ALCC)
- Computational Protein Structure and Protein Design (INCITE)
- Climate Modeling Uncertainty Quantification (Discretionary)
- Large Scale Hurricane Simulations (Discretionary)

The code has been shared with other sites running Cobalt.

GAMESS

ALCF staff made performance/scalability improvements to GAMESS (a quantum chemistry package). Further improvements are under development. As a result of this work, water cluster benchmarks with GAMESS/FMO on Intrepid show a marked improvement in performance with cumulative measures to reduce I/O overheads. The table below shows the performance of GAMESS with each improvement developed by ALCF staff.

Job Size (No. Waters)	Processor Count	Wall Time (Minutes)	Comments
1,024	2,048	36.9	Original FMO code (v3.3) for cluster MPP
		14.6	FMO v3.4.11 with densities stored in memory
		12.2	Use of RAMDISK feature on Intrepid
2,048	32,768	35.3	FMO v3.4.11
		11.4	RAMDISK

CERFACS

The AVBP code is a combustion CFD code from CERFACS, France. Originally, the code was unable to scale beyond 1,024 nodes, with a strong scaling speedup 4.21 from 128 cores to 1,024 cores (ideal 8). CERFACS visited ALCF for intensive collaboration December 6-10, 2009. The ALCF staff, working with the CERFACS team to analyze performance bottlenecks, determined that the code was using a very inefficient point-to-point communication implementation, inefficient implementation of global norm calculation, inefficient collection of data for I/O, inefficient implementation of collection convergence, stability, and visualization information. After major modifications of the code with functions kpl_update, par_norm, kpl_write, postproc re-written, the code now scales to 8,192 cores. The current state of the code has a speedup of 7.58 on 1,024 cores, and 46.91 on 8,192 cores. ALCF staff performed additional analysis and determined that significant additional improvements can be gained by utilizing appropriate compiler options. In addition, major changes are required for gather/scatter routines (STREAM type operations, 30% of the run time). The code is not threaded and has no use of double FPU/optimized libraries explicitly.

GFDL

GFDL Held-Suarez atmospheric benchmark is the major benchmark for the climate GFDL code. In past years, the ALCF staff have made numerous improvements, including compiling the code using a higher level of optimization, threading incorrectness, and core tuning. This year, they have concentrated on analyzing the scaling of the benchmark to investigate potential additional improvements. A single-thread baseline code, with 6 integration steps, developed by ALCF staff, shows the strong scaling speedup 1.95 from 1,020 nodes (4-way parallel) to 8,160 nodes. An ALCF study showed that the code is performing numerous unnecessary I/O-related functions. These can be removed without degradation of the code's functionality. Additionally, the study found that the performance of posix-type operation could be improved in netcdf code, the I/O code, which provided to GFDL benchmark an extra functionality. After additional tuning, a run on 8,160 nodes went from 1,075.6 seconds to 274.46 seconds, with a speedup of 7.09 from 1,020 to 8,160 nodes. GFDL won a 2010 ALCC award.

Nektar

Nektar code from Brown University is a core of the Nektar-G software—a code used for simulating the arterial blood flow. Nektar can simulate a single flow channel, while Nektar-G is a combination of several channels into a single flow. ALCF performance work was focused on tuning Nektar performance on Blue Gene/P by introducing double FPU intrinsics to the internal BLAS-type routines, improved instruction scheduling, cache reuse, performance counters and flop rate measurements, increasing messaging rate and DMA hardware programming, topology and task mapping issues. Overall, even starting from a well-tuned code, we were able to improve the performance of the full runtime by approximately 30%, as well as identifying many directions for possible future improvements. By aggregating the point-to-point communication updates, we were able to increase the performance by an additional 5%. The Brown University team visited ALCF December 15-18, 2009 and has submitted a proposal for a 2011 INCITE award.

NWChem

ALCF staff improved the efficient and scalability of NWChem significantly over the past year by tuning I/O, communication, adding multithreading and making numerous other changes the NWChem source code. Specifically, the job start time at scale (>16 racks) was reduced from 20 minutes to 45 seconds by eliminating redundant file access during initialization. The communication system underneath NWChem (ARMCI) was redesigned for the Blue Gene/P network using a thread-driven progress engine, rather than OS interrupts. To improve efficiency in the presence of a communication thread, bottleneck kernels in the coupled-cluster module were multithreaded, which also alleviates local memory bottlenecks. Other improvements include removing unnecessary branching in the integral code and tuning the build system to make better use of IBM compiler optimizations. All of these developments have been contributed back to the main branch of NWChem and will be available in the upcoming release.

ARMCI and Global Arrays

ALCF staff identified significant performance problems in ARMCI in the fall of 2009 due to a change in the system software that improved MPI performance, but reduced effective bandwidth in ARMCI by more than 100-fold. Working with IBM, ALCF fixed the issue and enabled ARMCI to achieve peak bandwidth. Subsequently, the progress engine inside of ARMCI was redesigned to better suit the network hardware and lightweight OS kernel. The overall improvement of application codes (e.g. NWChem) was a 50-100% reduction in wall time. A peripheral benefit of the work with DCMF was it allowed ALCF staff to assist LBNL in maintaining UPC/GASNet for the latest BG/P driver release.

The knowledge gained from incremental improvement to ARMCI led ALCF staff to rewrite ARMCI from scratch and integrate it into the MPICH distribution. This will allow ALCF to work closely with IBM to deliver a Blue Gene-optimized version of Global Arrays and will provide a huge performance benefit for current and future ALCF systems.

As part of an ongoing collaboration with Paratools Inc. and PNNL, ALCF staff developed profiling capability within TAU for ARMCI and performed a detailed performance analysis of NWChem on Blue Gene/P. This capability is being integrated into TAU and Global Arrays for distribution in upcoming releases. Further collaboration to add TAU profiling capability for other one-sided messaging layers (e.g. GASNet and DCMF) is planned.

MPQC

Additional quantum chemistry capability has been provided to users in the form of MPQC, which provides core simulation methods (e.g. DFT and MP2). ALCF performance optimizations enabled MPQC to scale to 40 racks.

Combustion Physics

The combustion physics application code was seeing long checkpoint times on Intrepid. The code is strong scaling, and the percentage of time spent in I/O went up as the number of nodes increased. ACLF staff were able to improve the I/O time significantly by moving from per-rank POSIX I/O to collective MPI I/O, reducing the time spent in I/O by almost an order of magnitude. This work was performed at the Leap to Petascale workshop in conjunction with I/O experts from MCS. The project has subsequently been accepted into the Early Science program, and has submitted a proposal to INCITE 2011.

Nek5000

The ability to simulate fluid flow at unprecedented scale and resolution was enabled through work done by ALCF staff; this work enabled the largest-ever simulations with the Nek5000 solver. The Blue Gene/P systems at the Argonne Leadership Computing Facility and Juelich Supercomputer Center were used to perform simulations using over 294 thousand processors, the largest processor count available worldwide on a single system, and employing a significantly larger grid than ever before—over 7 billion grid points. In addition to significant increases in processor count and grid size, the simulations demonstrated the high level of performance the solver can attain at these scales, achieving over 70% parallel efficiency across a near order of magnitude increase in processor count, over 20% of system peak performance, and a sustained flop rate of over 172 TeraFlops. This level of scaling and performance was the result of collaborative work between the ALCF, MCS and IBM. Working with the MCS MPICH team and IBM, an MPI implementation optimized for memory usage was developed to meet the requirements of the Nek5000 multi-grid solver and allowed it to fit the required number of MPI communicators into the available memory. In addition, the larger grid size required the Nek5000 central communication kernel data structures to be converted to 64 bit integers in order to allow global addressing at this scale. Finally, optimized assembly kernels were utilized to enhance the floating-point performance of the solver.

Lattice QCD

ALCF staff have been collaborating with people from Boston University and elsewhere on multigrid solvers for the lattice QCD Dirac equation. The collaboration has a working implementation for one particular discretization (Wilson-clover) that gets up to a 20x improvement in solution time. This discretization is used by the group at JLab (who are running at ORNL) and the "Pion Mass" discretionary project at ALCF, though it hasn't been integrated into their production codes yet. ALCF staff are also working on the domain wall discretization (used on Intrepid by Columbia/BNL) and staggered (used by MILC), this is a work in progress.

Baseline Area 2: Business Results

Charge Question 2: *Is the facility maximizing resources consistent with its mission?*

FY '10 Operational Assessment Guidance: *Business results measure the performance of the facility against operational parameters. For each of the following metrics, report results and provide projections using a methodology developed with the concurrence of the Federal Program Manager: Resource Availability, Resource Utilization, and Capability Usage.*

ALCF Operations Overview

The Argonne Leadership Computing Facility (ALCF) operates a production 557-TF Blue Gene/P supercomputer called Intrepid and an associated production visualization cluster called Eureka. Both systems have direct access to a single, facility-wide file system. HPSS provides archival services to the users. The Table below provides details on both production resources.

System Name	System Type	CPU Type	Clock Speed	Nodes	Cores per node	Total cores	Agg. RAM	RAM per core	Interconnects
Intrepid	IBM Blue Gene/P	PPC 450	850 MHz	40,960	4	163,840	80 TB	512 MB	Proprietary 3-D Torus
									Tree/Collective
									Barrier
									RAS (1G Ethernet)
Eureka	Graphstream COTS cluster	Intel XEON E5405	2.0 GHz	100	8	800	3.2 TB	4 GB	Myricom 10G
	(GPUs)	NVidia Quadro FX5600	1.35 Ghz	200	128	25,600	300 GB	12 MB	PCIe x16 V2.0

Metric 2.1 – Resource Availability

ALCF resource availability will be measured on major systems and reported quarterly. In FY10, the system will be assessed at 85% scheduled availability and 80% overall availability based on our standard method. Our standard method of determining availability: beginning the start of the quarter following general availability, new or upgraded systems will meet 75% scheduled availability and 70% overall availability; systems in operation more than one year since general availability or the last upgrade will meet 85% scheduled availability and 80% overall availability.

Metric 2.1 – Results

Intrepid has been in full production since February 2009, so must achieve 85% scheduled availability and 80% overall availability. For the period of August 1, 2009 through July 31, 2010, Intrepid had 97.9% scheduled availability and 94.6% overall availability, with an average of 2.1% unscheduled downtime.

Metric 2.1 – Projections

Intrepid's availability for this year was within 1% of last year's availability. There is a multi-day electrical outage that may be scheduled before the end of the year; this should be at most a < 2% impact over 6 months. Therefore, the projected overall availability for the entire year is 94% if the power outage does not occur and 91% if the power outage occurs. Scheduled availability projected to remain in the 96% - 98% range. Full details about outages, as well has how calculations are performed are provided in the *Metric 2.1 – Additional Details* section below.

Metric 2.1 – Additional Details

The figure below shows availability data by day from 8/1/2009 to 7/31/2010.

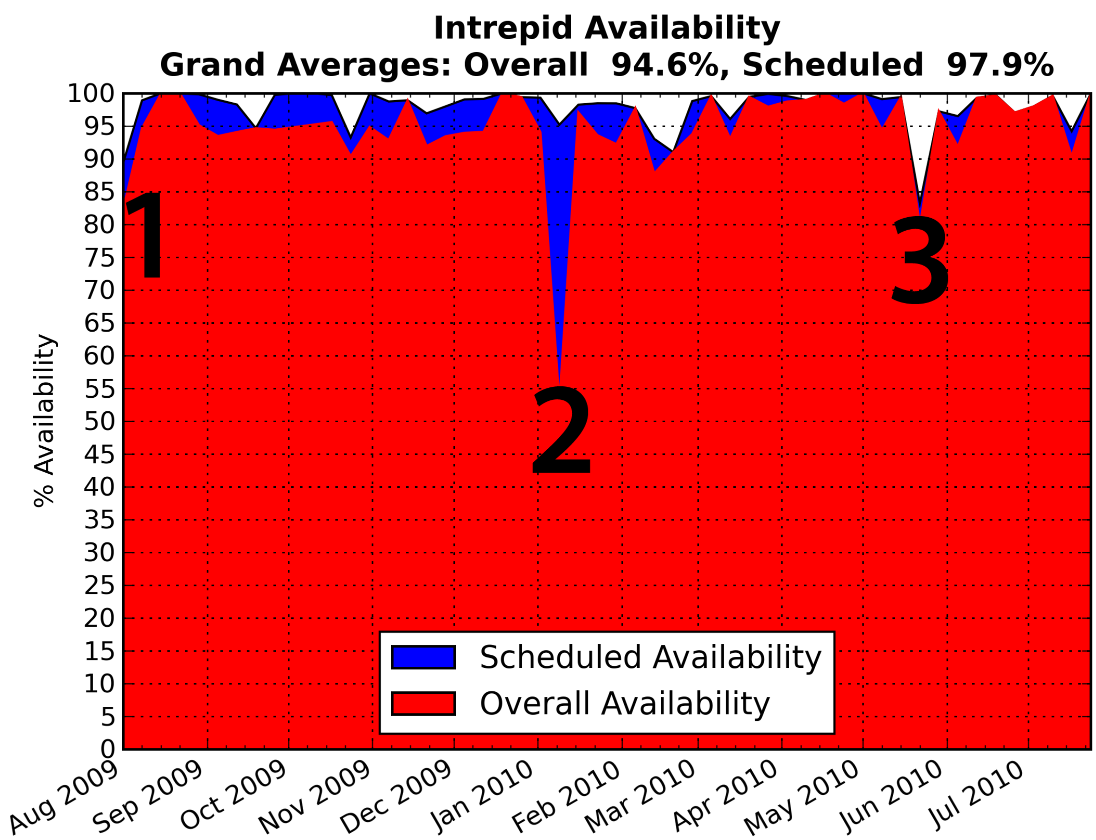

Explanation of availability losses > 10%

A brief description is provided below for the causes of major losses of availability as shown in the annotated figure above

Items 1&2: As noted in the ALCF 2009 OA Report, the Myricom equipment that provides the data and access network for all ALCF resources was experiencing frequent component failures. One of their suppliers, Zarlink, had a batch of contaminated optical transceiver wafers. The ports prematurely aged and began to send spurious data out into the network, causing CRC storms. Several attempts were made to resolve the problem, the first during a maintenance outage in August 2009 where a portion of the transceivers were replaced. After the August outage, it became clear a partial replacement wasn't sufficient. In January, the ALCF took an extended outage to replace every Zarlink transceiver. Since the full replacement, no more failures of this type have occurred, and we believe the problem has been resolved.

Item #3: Included in this item are a number of unrelated outages, identified by their event IDs. The event IDs are unique numbers generated by a new system for automatically tracking and reporting outages that was developed by ALCF staff. The outages are listed, with a brief description, below.

- On May 13th the Blue Gene reported a bridge communication error in the control system (event ID #1620). The system began trying to run jobs on top of other jobs and drained the job queue completely. A control system reset resolved the problem, and it has not occurred since this outage.
- On May 24th, the first, and longest, of 3 cooling outages (event ID #1730). The entire facility was down for almost 25 hours. Root cause analysis has determined that the outage was the result of a combination of high outside temperatures (it was one of the hottest days of the year) and cooling tower maintenance that had the facility supported by only one of the three (when in N+1 configuration) cooling towers. The ALCF was not notified of the maintenance, or we could have operated in a reduced power mode to reduce the risk of a failure.
- On June 25th, a brief chiller outage, caused by a pump tripping a circuit, occurred. The only impact to the ALCF facility was the shutdown of Eureka, which was proactively shutdown to reduce load in the facility.
- On July 1st, another chiller outage occurred. This outage was the result of electrical work that included a scheduled switch over of power from one transformer to another at the sub-station. The switch over, which was expected to have no impact, triggered an unexpected power outage for the entire chiller plan for approximately five (5) seconds. This outage caused a 1-Amp fuse to blow. Locating and the fixing the fuse took nearly three hours. Service was returned to normal after the fuse was repaired.

The ALCF recognized that the initial cooling outage was partially the result of a significant communication problem between FMS and the ALCF. We have since jointly developed and implemented a communication plan to address the problem. As a result, communication between the two divisions has clearly improved. In addition, the ALCF and FMS have developed a plan for improving the chilled water service to the ISSF (Interim Supercomputing Support Facility) where the bulk of the ALCF resources are located. This plan has both short-term (install a temporary air cooled chiller that will be connected to an different power source from the core chillers) and long-term improvements, and is in the process of being implemented.

Availability Calculation Details

Availability is calculated as follows:

$$Max\ Core\ Hours\ Possible = 40\ racks \times \frac{4096\ cores}{rack} \times 24\ hours = 3{,}932{,}160\ Core\ Hours$$

$$Max\ Core\ Hours\ Scheduled = Max\ Core\ Hours\ Possible - \sum scheduled\ losses$$

$$Max\ Core\ Hours\ Available = Max\ Core\ Hours\ Possible - \sum (scheduled\ losses + unscheduled\ losses)$$

$$Overall\ Availability(\%) = \frac{Max\ Core\ Hours\ Available}{Max\ Core\ Hours\ Possible} \times 100$$

$$Scheduled\ Availability\ (\%) = \frac{Max\ Core\ Hours\ Available}{Max\ Core\ Hours\ Scheduled} \times 100$$

These values are calculated on a daily basis, and then arithmetically averaged over 7-day intervals. The data in the charts is the 7-day average. So, the first data point is the average of days 1-7, the second data point is the average of days 8-14, etc. If the number of days is not an even multiple of 7, the last data point is handled as follows: if there are more than half (4 or more) of the data points, a final data point is calculated from those values. If not (3 or fewer), those values are included in the previous data point, which becomes an average of between 8 and 11 data points. This is to avoid significant deviations of the last point due to a small average.

Scheduled availability is defined as the percentage of time that a resource is available to the users, excluding downtime announced to the users with more than 24 hours notice. Overall availability is the percentage of time a resource is available to the users with no exclusions. For example, if there were 14 hours of scheduled maintenance, which took 16 hours, there were 8 hours actually available (24 -16) resulting in 33% overall availability (8/24). There were 10 hours scheduled availability (24-14), and there were actually 8, resulting in 80% scheduled availability. The overall averages for the entire period are a straight average of the daily results.

Recent improvements in availability and failure tracking

In the past year, the ALCF has implemented some significant improvements in how downtime and failures are tracked. In the new process, each week the previous week's failure data is fully processed (instead of processing each month). The weekly data is initially automatically built and preprocessed from various logs and the RAS database using scripts that analyze every job that reported a non-zero exit code, all reservations specially coded to indicate a maintenance reservation, Blue Gene service actions, etc. Once the data has been preprocessed, the team meets to review the data, correlating all data with out of band data such as email trails and IM transcriptions, and do root cause determination where possible. The weekly team processing was a best practice adopted from Los Alamos

National Laboratory. A possible improvement to the process, increasing the review frequency to daily, is currently under evaluation.

Metric 2.2 – Utilization

Utilization will be tracked and reported for Intrepid. This is a reported number, there is no target result associated with the metric.

Metric 2.2 – Results

Average overall utilization on Intrepid was 81.1% for the period January 1, 2010 to July 31, 2010, with an INCITE utilization of 63.7%. The figure below shows the utilization by allocation category, averaged by week. Details on how utilization is calculated are provided in the *Additional Details* section.

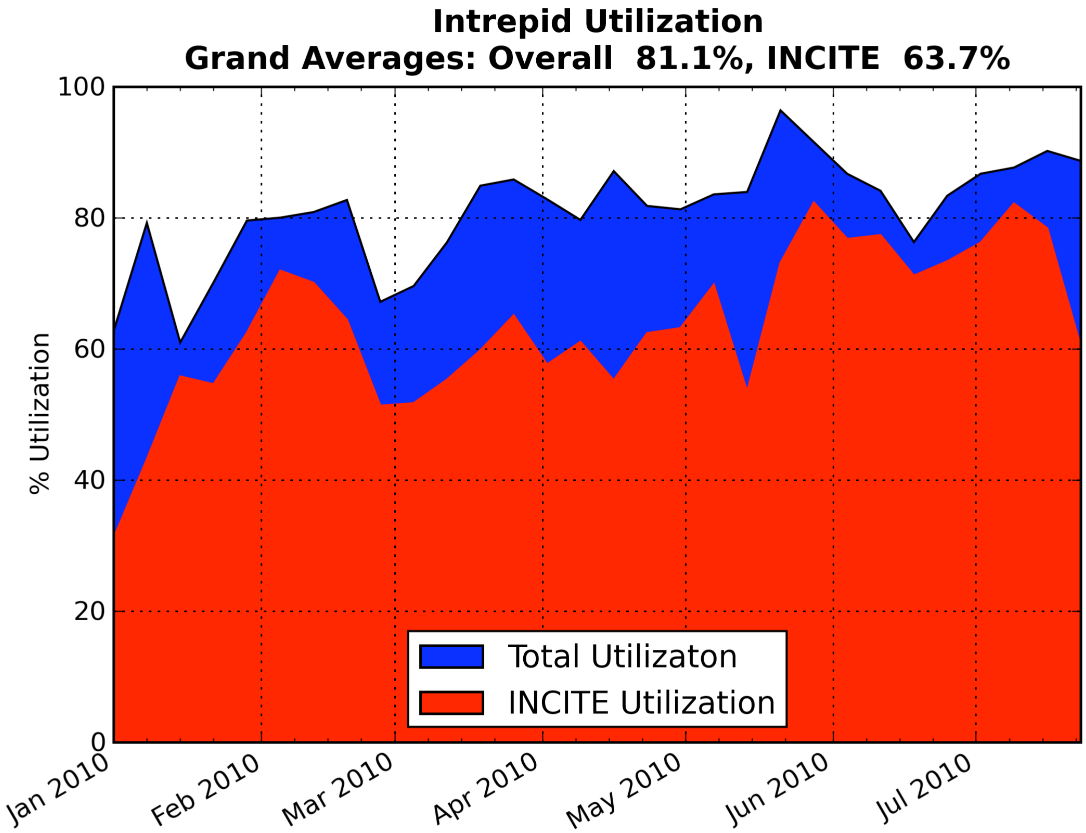

A concern about low utilization was expressed in the 2009 OAR response. Some improvements had already been made at that time, although too late in the year to overcome the initial numbers. These improvements, fixed back fill bug and scheduling policy modifications, were outlined in the *Responses to 2009 OA Recommendations* section. An additional cause for low utilization was projects that were running many short (<30mins), very large (40%-80% of the resource) jobs. Each job would require a drain of the system to free up enough nodes to run the job, which would immediately be followed by other smaller jobs. This drain cost would be paid for each job, resulting in a significant amount of lost system time. The ALCF catalysts have worked with these projects to tie runs together with a single Cobalt job submission using the script mechanism. Finally, a number of back fill improvements were made. A clear trend upwards for utilization is visible in the figure below, which shows utilization for the period August 1, 2009 to July 31, 2010. We believe the most recent six-month average of 81.1% will hold. As a leadership class system, our mission is to support large jobs which will always have a drain cost. Because of this drain cost, we believe that 80% is a good utilization objective.

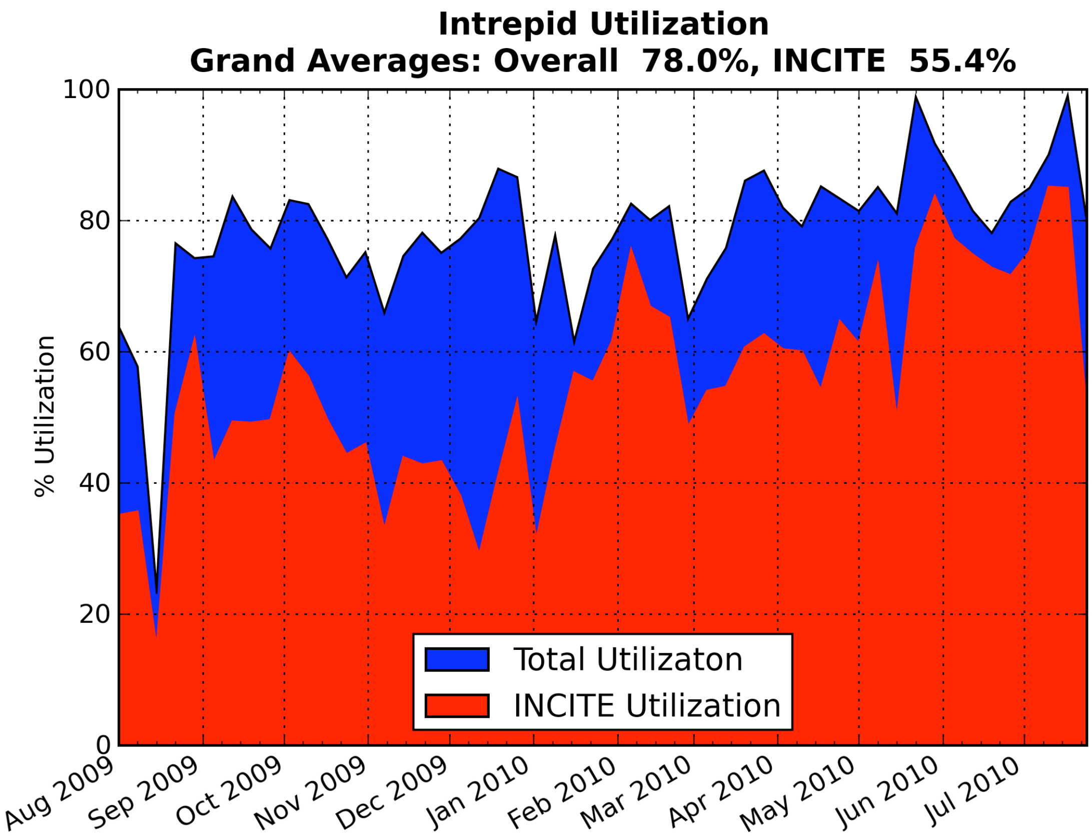

The following pie chart shows the utilization by science domain. The Oak Ridge Leadership Computing Facility (OLCF) and the ALCF jointly selected these science domains for the INCITE program. The ALCF is using this categorization for its entire workload.

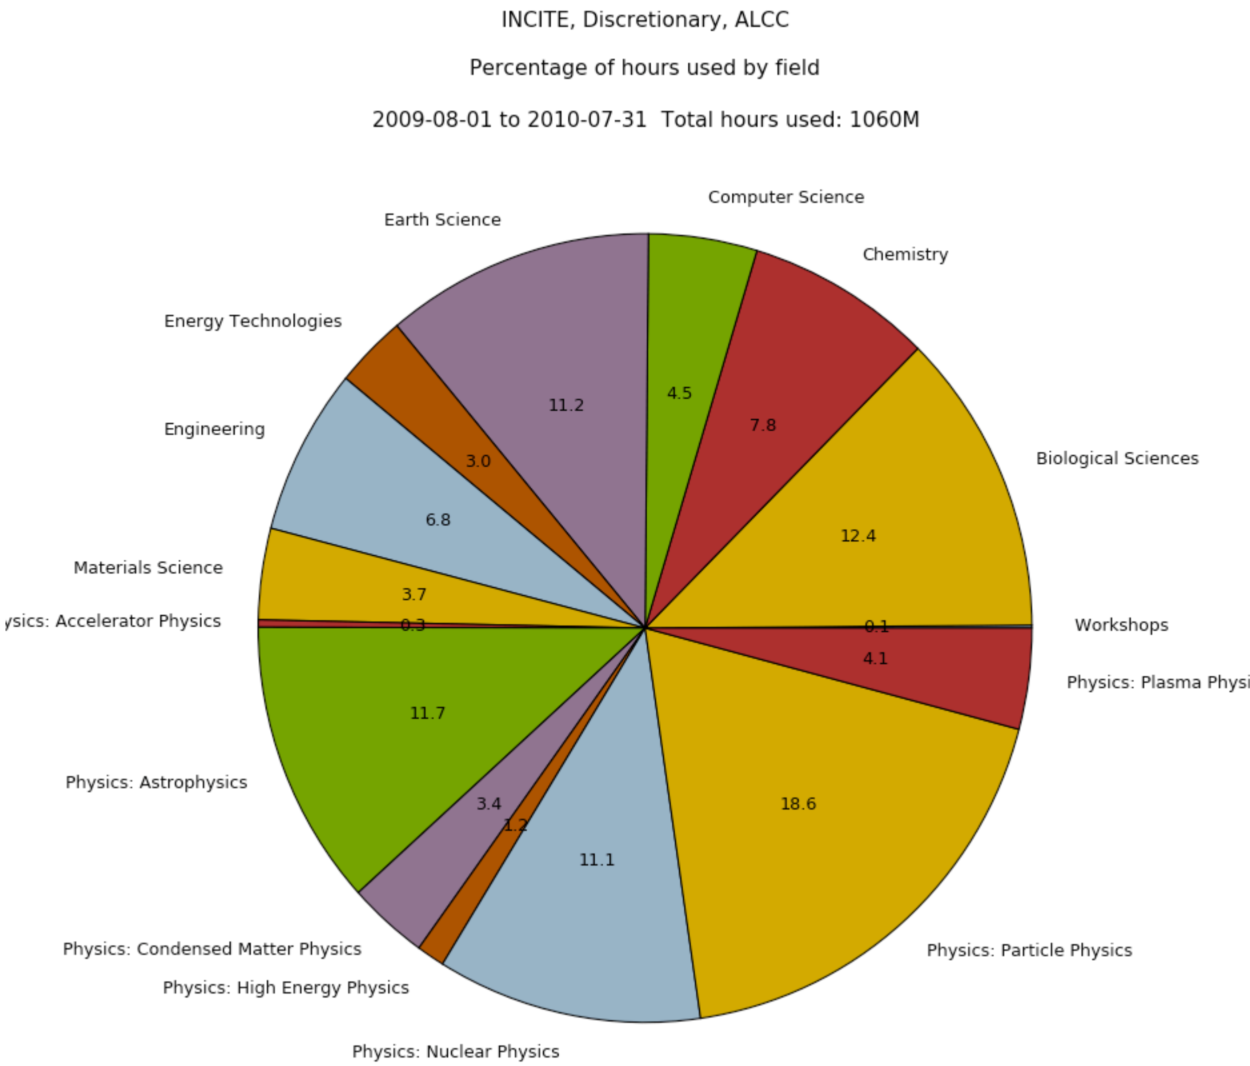

Metric 2.2 – Projections

There are currently no significant scheduling policy changes planned for the next year, policies may be adjusted to accommodate project changes and/or job mix changes. The new backfill algorithm (described in the Baseline Area 2 section of the Responses to 2009 OA Recommendations) is scheduled to be deployed later this year; in addition, adjustments may be made to accommodate the "end of year" rush. Even with these potential adjustments, we expect a similar rate of utilization for the remainder of the calendar year and project overall utilization will be in the 78% to 82% range.

Metric 2.2 – Additional Detail

INCITE core-hour commitment

The ALCF committed to provide 646 million core-hours to INCITE. On July 31, 471.9 million core-hours had been delivered to INCITE projects. If usage continues at the same rate, we should deliver more than 800 million core-hours to INCITE by the end of the 2010 INCITE year.

Utilization Calculation Detail

Job records in the standard PBS format. Each night at midnight, a script runs and processes the day's records and imports the data into the *Clusterbank* accounting package. Clusterbank records the time, date, duration, user, project and various other system parameters for every job run in the facility into a SQL database. Utilization is calculated by pulling the daily total core-hours used, as well as the core-hours used by the job project's category (INCITE, Discretionary, and in the future ALCC). Jobs that cross day boundaries have the hours appropriately apportioned to the days involved in the job. The data is then combined with the availability data (see *Metric 2.1)* and an average computed per day using the following equation:

$$Utilization\ (\%) = \frac{Core\ Hours\ Delivered}{Max\ Core\ Hours\ Available} \times 100^{1}$$

For the graph, the daily values are averaged over a seven-day period, as described in the *Resource Availability* section.

Metric 2.2 –Mean Time To Interrupt (MTTI)

Mean Time To Interrupt (MTTI) will be tracked and reported for Intrepid. This is a reported number, there is no target result associated with the metric.

Metric 2.2 – Results

The overall MTTI for Intrepid is 27.5 hours (see the figure below). For Blue Gene only hardware, the MTTI is 4.6 days. These numbers are significantly lower than we would like and we are working to resolve the issues. The problems are primarily in three categories:

1. *Blue Gene XFPs failures.* It appears that a bad batch of (non-Myricom) XFPs was installed on Intrepid. IBM has run diagnostics, and 80% of them now fail the same diagnostic. IBM is in the process of completing their evaluation. We are in discussions with IBM to replace all failed XFPs free of charge, ideally with ones from a different vendor.
2. *Myricom network failures.* While the optical transceiver failures have been resolved, there remain a significant number of failures within the Myricom network. We are working with Myricom and IBM to determine root cause and to resolve them.
3. *GPFS failures.* We have not yet determined the root cause of the failures of GPFS (IBM's Global Parallel File System). Data on the failures has been provided to IBM and an analysis by the GPFS team is underway.

MTTI data is calculated by taking all interrupts recorded in a week, calculating the intervals between each interrupt, and averaging the interrupt interval for the week. The figure below shows the weekly averages from 8/1/2009 to 7/31/2010. From September 2008 thru December 2009, only automatically generated data was available for determining interrupts. Unfortunately, this data is not very accurate and reports a lot of false positives. The weekly reviews begun in March 2010 (described in *Metric 2.1 – Additional Details)* have greatly improved the accuracy of the numbers. In addition, out of band data sources were mined for data back to January 2010. Data from January 2010 on should be accurate.

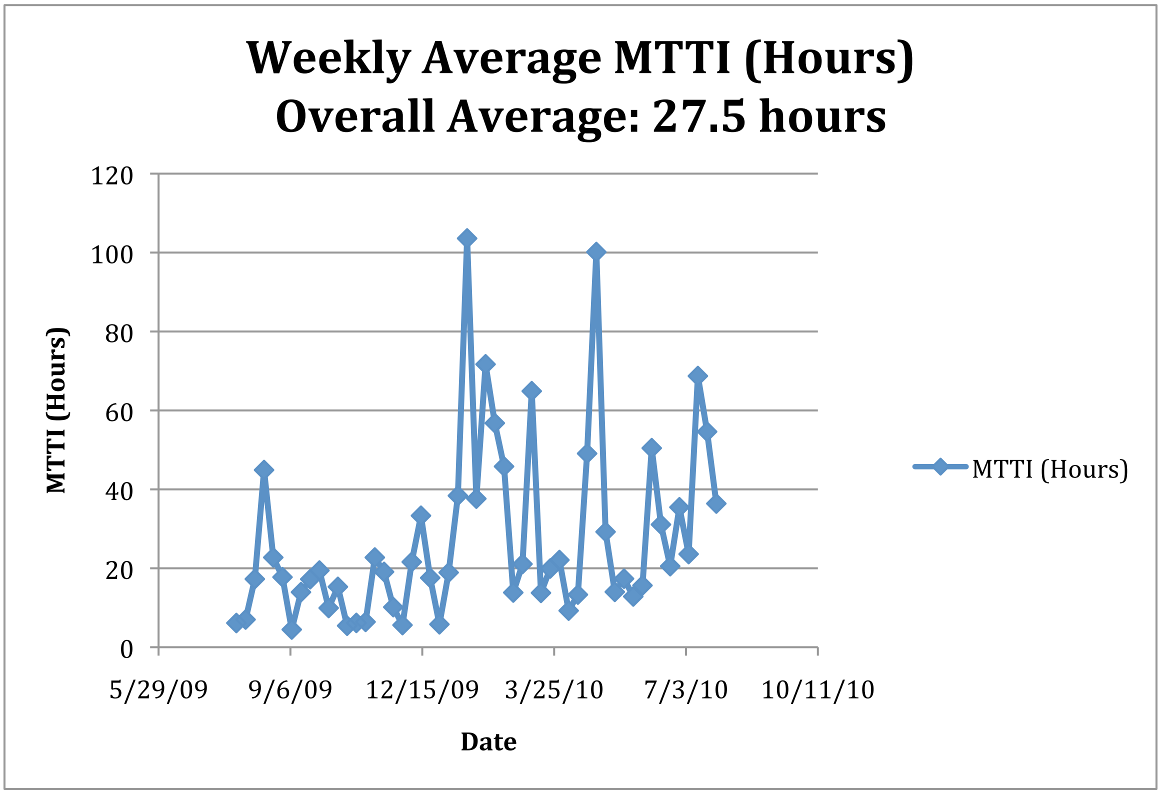

Metric 2.3 – Projections

The resolution to the XFP failures is expected to be completed in late fall, resolution to the GPFS issues may take longer (root cause analysis and resolution of the underlying problem is expected to be time and labor intensive). Therefore, the MTTI is projected to remain at similar levels between now and the end of the year. Once these two issues have been resolved, MTTI is expected to improve significantly.

Metric 2.2 – Additional Details

Disk storage failures are not reported separately as the full resource is considered down when the file systems are down.

Metric 2.3 - Capability Usage

ALCF capability usage of the major systems will be measured quarterly and reported annually. Beginning FY10, 250 million CPU-hours will be delivered to jobs 8,192 nodes or larger.

Metric 2.3 - Results

During the metric target period, August 1, 2009 to July 31, 2010 (inclusive), there was a total of 368.3 million core-hours used by capability jobs. A total of 1,059.6 million core hours were delivered during this time period, with capability jobs being 30.9% of the total (see the figure below). The two cumulative charts show cumulative daily hours used.

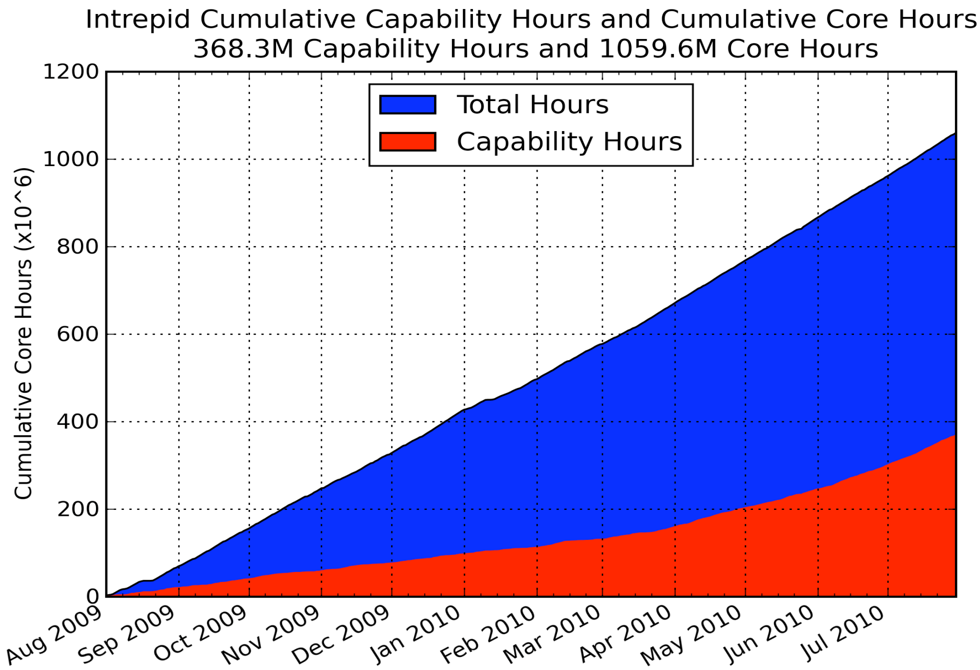

For January 1, 2010 through July 31, 2010, a total of 271.9 million core hours were used by capability jobs, 42.8% of the total 635.0 million core hours delivered (see figure below).

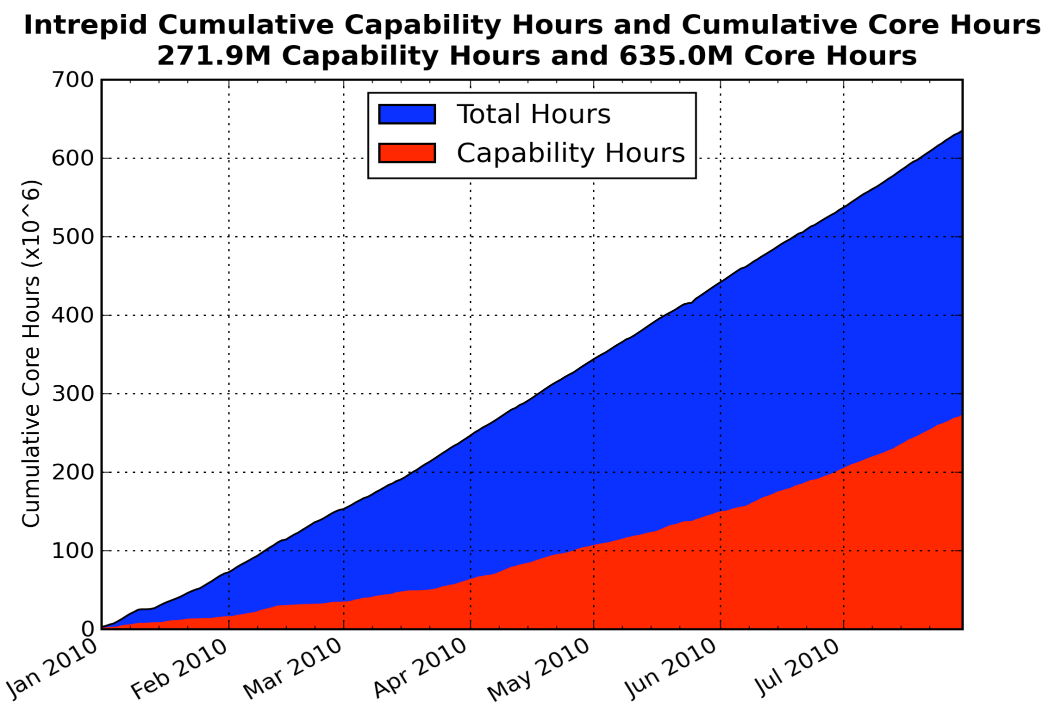

The figure below shows data for all jobs, by month, binned by job size. Smallest jobs are at the top; largest jobs are at the bottom. The 'K' represents 1024 nodes, so 1K is 1024 nodes, 8K is 8,192 nodes (8,192 nodes == 8 racks == 20% of the system == 32,768 cores), etc. Capability (8K and larger) jobs are colored blue, green, yellow, orange and red (i.e. blue and below).

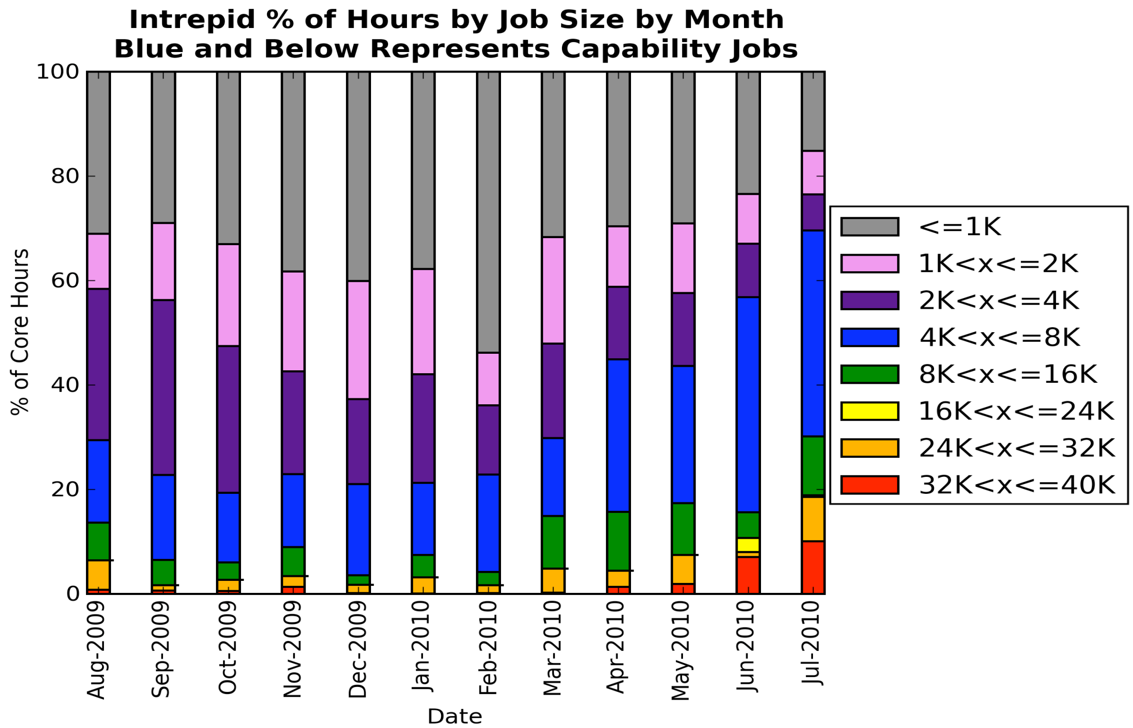

Metric 2.3 – Projections

Using simple extrapolation across two years of data, from 8/1/09 to 7/31/10, results in a projected 352 million core-hours for the full 2010 calendar year. Because of recent improvements on the system, we believe the number will be higher and project a total of between 375 and 400 million core-hours used by capability jobs during the 2010 calendar year.

Metric 2.4 - Job Throughput

Job throughput will be tracked and reported for Intrepid. This is a reported number, there is no target result associated with the metric.

Metric 2.4 – Results

The expansion factor[1] by job size is shown in the figure below; the expansion factors range from nominal for very small jobs, and full system jobs, to 140 for jobs using 32,768 cores. A leadership class facility would expect all expansion factors to be under 80. The expansion factor is calculated as run time - queued time (i.e. holds are not excluded). As discussed in the *Baseline Area 1 sub*section of the *Responses to 2009 OA Recommendations* section, the

[1] The expansion factor is defined to be $(Queue\ Wait\ Time + Job\ Run\ Time)/Job\ Run\ Time$

reported numbers are artificially high due to a lack of state change data. We expect to resolve this problem shortly and should have more realistic data by next year's OAR.

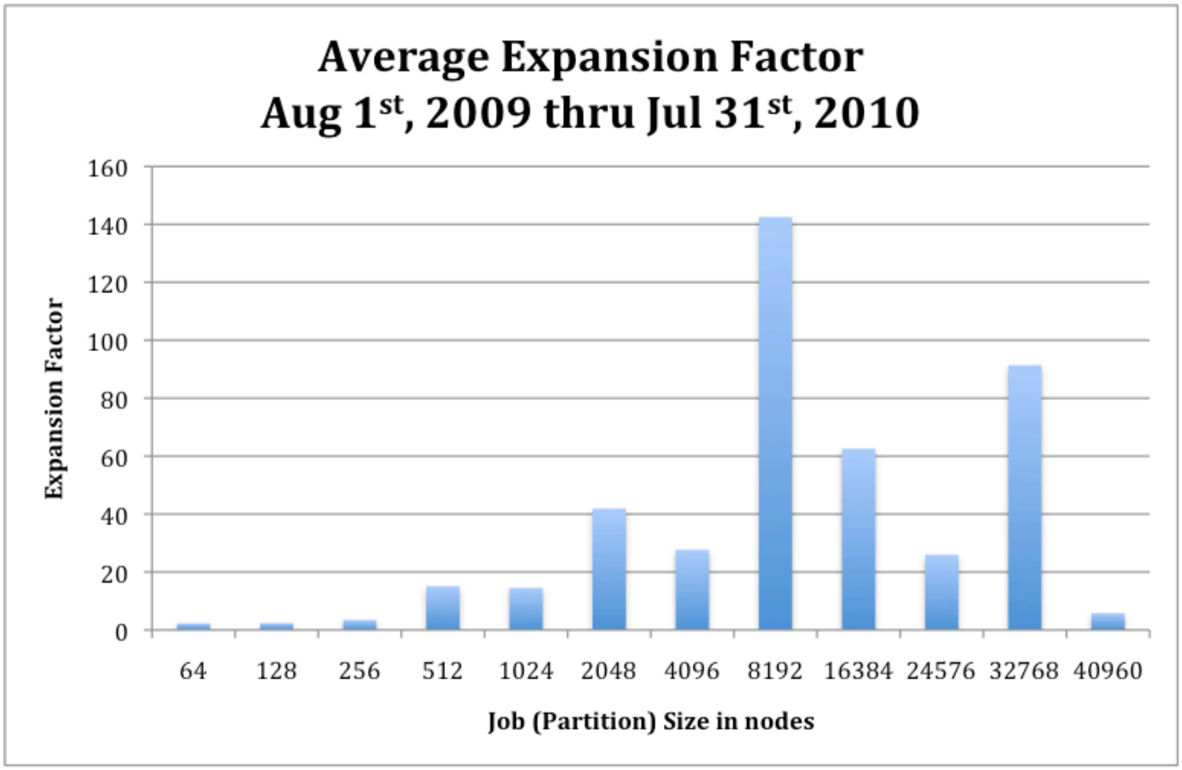

Metric 2.4 - Projections

No projection at this time, due to the expected significant impact of improvements in our ability to calculate accurate expansion factor data.

Metric 2.4 – Additional Details

Monthly reviews of resource manager new feature and bug fix requests also include a review of the wait times by queue, category (INCITE, Discretionary, ALCC), job size, etc. This review provides an opportunity to address potential issues with the expansion factor. User feedback is gathered from the users, and used to inform scheduler policy and cost functions.

Metric 2.5 - Job Exit Data

Job exit data will be tracked and reported for Intrepid. This is a reported number, there is no target result associated with the metric.

Metric 2.5 – Results

From August 1, 2009 through July 31, 2010, a total of 134,679 jobs were submitted on Intrepid. Of these, 82,286 jobs (61.1%) ran to successful completion and 52,393 jobs (39.9%) did not complete successfully. Of the jobs that failed to complete, 48,305 (92.2%) of them were due to user error, and 43,781 (7.2%) due to system failures.

Failures due to user error

Roughly 60% of the user error failures were one of the following:

1. Jobs that ran beyond the requested time limit.
2. Jobs deleted from the queue by the users before the job was started.
3. Jobs killed by users while the job was running.

Category 1 (jobs running past their time limit) includes hung jobs. Jobs are categorized as system errors, until evidence can be found to the contrary.

Failures due to system error

Seventy percent (70%) of the system failures fall into two categories:

1. Blue Gene hardware failures (30% of the system failures). The bulk of these are the XFP failures discussed in *Metric 2.3 – MTTI.*
2. Storage failures (40% of the system failures). Included in this category are file system failures (both GPFS and NFS) and failures of the Data Direct Network raid arrays (DDN). We believe the NFS failures should be resolved with the installation of some equipment currently on order. The DDN failures appear to have been resolved with a firmware upgrade performed in January 2010.

The figure below shows a Pareto chart of all failure causes, system and user failures are labeled.

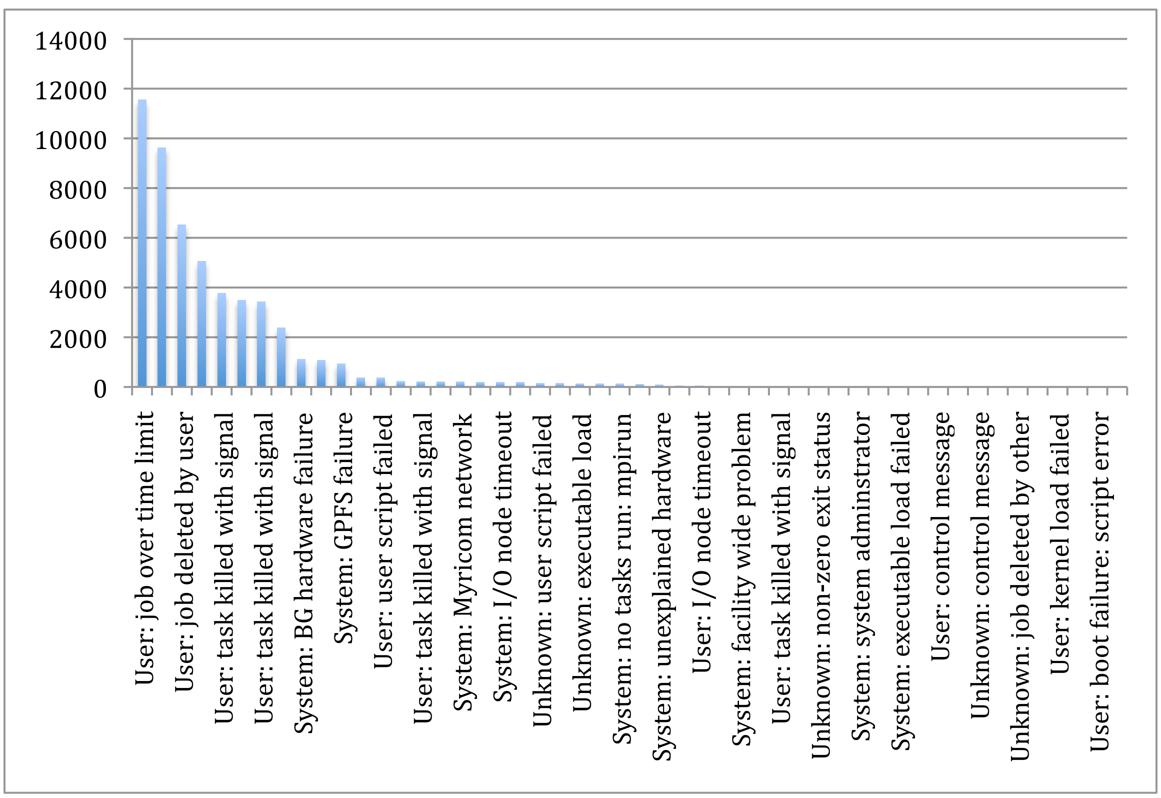

Baseline Area 3: Strategic Results

Charge Question 3: *Is the facility meeting the Department of Energy strategic goals 3.1 and/or 3.2?*

FY '10 Operational Assessment Guidance: *Strategic Results measure the effect the investment has on the performing organization, and should provide a measure of how well the investment contributes to achieving the organization's strategic goals. For each of the following metrics, report results and provide projections for 2011 using a methodology developed with the concurrence of the Federal Program Manager: Science Output, Scientific Accomplishments, and the Allocation of Facility Director's Reserve computer time.*

Overview of ALCF Science Impact

The past year saw the ALCF serve its two billionth processor-hour on its production resource, Intrepid. Interest in using the ALCF resources has continued to increase, and users new and old grow more sophisticated. Considerable progress was made in capability usage on Intrepid (Intrepid capability use is defined as running simulations that span at least 20% of its processors). Recently, over 200 million core hours were awarded under the ALCC program, and we look forward to reporting their accomplishments during the next year. As in the past, Director's Discretionary time allocations continued to play an important role in fostering INCITE (and ALCC) proposals. Results from INCITE, as well as some Discretionary, projects are described in this section.

Research highlights came from diverse fields of interest to DOE, and the ALCF staff has played important roles in achieving these science results. The results are showcased here and in a companion slide presentation, which also calls out ALCF's participation.

During the past year, the ALCF conducted four workshops focused on science applications: "INCITE Getting Started/Introduction to Blue Gene/P," "Leap to Petascale," "INCITE Proposal Writing Workshop," and the "Tau Workshop". The Leap to Petascale workshop played a role in five 2011 INCITE proposals and one Early Science Program proposal.

At SC09, ALCF staff taught a tutorial on Python for scientific and high-performance computing, ran a Birds-of-a-Feather (BOF) session to discuss the challenges and successes using Blue Gene/P, presented a poster on scalability of quantum chemistry codes on Blue Gene/P, led sessions in a BOF on Python for high-performance and scientific computing, presented a poster on the development of ALCF status web site, participated in a Data Center for the Future Challenge, and gave a plenary talk at the Workshop on High Performance Computational Finance.

Additionally, ALCF Catalysts and Performance Engineers gave presentations at the following meetings:

- SciDAC's CScADS summer workshops (Snowbird)
- SciDAC 2010 (Chattanooga); DOE-CSGF HPC Workshop (Washington, DC)
- SIAM Conference on Parallel Processing for Scientific Computing (Seattle)

- MPI Forum (Portland)
- Swiss National Supercomputing Centre Users' Day (Manno, Switzerland)
- American Chemical Society National Meeting (Washington, DC)

Staff also presented posters at many workshops around the world:

- SAAHPC 2010 (Knoxville)
- Kavli Institute Workshop on Lattice Quantum Chromodynamics (Bejing)
- XXVII International Symposium on Lattice Field Theory (Peking U., Bejing)
- NSF-NAIS Workshop (Edinburgh)
- MILC Collaboration Winter Meeting (BNL)
- XXVII International Symposium on Lattice Field Theory (Sardinia, Italy)
- MILC Collaboration Summer Meeting (U. Utah)
- SciComp 16 (San Francisco)

In addition, ALCF Catalysts and Performance Engineers presented seminars at Argonne at the Materials Science Division (MSD), Laboratory for Advanced Numerical Simulations (LANS), Math and Computer Science Division (MCS), at onsite ALCF workshops, and in the MCS lecture series for summer students. They also presented seminars at U.C. Berkeley, Fritz Haber Institute of the Max Planck Society (Berlin), Stanford University, and the Computation Institute (University of Chicago).

The INCITE project teams made significant progress on code development. William Tang's tokamak microturbulence project incorporated the use of ADIOS for their I/O, which gave a 10x performance improvement; Tang's team also got the GTC code scaling efficiently up to 32K cores on Intrepid. GE Global Research's fluid dynamics code was scaled up from 512 to 8K cores on Intrepid with ALCF assistance. Jeff Greeley's team, including Nick Romero of the ALCF, introduced a new layer of parallelization and expanded production of the GPAW code up from 512 to 16K cores on Intrepid. ALCF efforts assisted Theresa Windus' project scale up the GAMESS code from 4K cores to 64K cores on Intrepid, and wrote a new parallel implementation of an important molecular force. John Tonge's Fast Ignition project implemented a new hybrid method on Intrepid that gave a performance boost for some of their calculations of up to 30,000 times faster. The ADLB load-balancing library from Argonne now allows Steven Pieper to run on 32K cores on Intrepid for David Dean's computational nuclear structure INCITE project. The Nek code, used by Paul Fischer's INCITE project, has been scaled efficiently up to the full 160K cores on intrepid, and 294912 cores on Juelich's Jugene Blue Gene/P machine. Don Lamb's project modernized the Flash code's parallel I/O approach. ALCF contributed expertise in scaling up Thierry Poinsot's flow/combustion code from 256 to 16K cores on intrepid. The CCSM project (Warren Washington, PI) introduced threading with OpenMP.

The ALCF staff played a role in preparing the multi-institutional proposal for DOE's new Climate Science for a Sustainable Energy Future (CSSEF) program, including attending a writers' workshop at LBNL, writing a white paper on computational and numerical methods, and collaborating with the other proposal authors at Argonne and elsewhere.

ALCF staff also assisted with various multi-institutional proposals for DOE's Exascale Co-Design Center program, including:

- Flash High Energy Density Physics Exascale Co-Design Center, contributing authorship (two ALCF staff on the proposal).
- Chemistry Exascale Co-design Center (CECC), contributing authorship (two ALCF staff on the proposal)
- A Novel I5 Approach to Exascale Computing and Co-Design for Energy-Related Science, contributing authorship (ALCF staff member is the Argonne PI)
- Exascale Center for Earth System Simulation (ExCESS), contributing authorship (two ALCF staff on the proposal)

Metric 3.1 – Tracking of science output and accomplishments

The LCF will track the science output and accomplishments for each project, including milestone reports, presentations, publications, journal covers, and awards, and will provide quarterly reports on the project results. The LCF will also track technology accomplishments, such as development of reusable code that results in a new tool for its discipline and new algorithm design ideas or programming methodologies.

Metric 3.1 – Results

Science Output

Publications derived from research on our facility are tracked and entered into a SQL database. An automatically generated bibliography of publications with abstracts and online access to some papers is published on the ALCF web site (http://www.alcf.anl.gov/publications/paper_list.php). Details for tracked publications are provided in *Strategic Results* quarterly reports to DOE.

During this past year, 81 new publications from ALCF projects were added to the database, including papers in *Physical Review Letters*, *Science*, and *Nature Physics*.

There were 74 major presentations reported by the project teams for research done with ALCF resources.

Science Accomplishments

This section highlights some of the ALCF scientific accomplishments in the time period of August 1, 2010 through July 31, 2010. Also see the companion slide deck that provides slides for these accomplishments, as well as others.

Accomplishment: DDES, IDDES Simulation of Landing Gear

Discretionary Project: "Improved Boeing IDDES: Delayed Detached Eddy Simulation (IDDES) of Turbulent flow over a Generic Aircraft Landing Gear"

PI: Philipe Spalart (Boeing)

BANC is the Conference on Benchmark Problems for Airframe Noise Computation. One problem category for the conference is unsteady wake interference between a pair of inline tandem cylinders. The paper represents the collaboration with Dr. Philipe Spalart (Senior Technical Fellow, Boeing) and his team of researchers at St. Petersburg, Russia (Prof. Michael Strelets, Dr. Michael Shur, Dr. Andrey Garbaruk, and Dr. Andrey Travin). They used 9M hours (of the 10M discretionary allocation) to run an overlapping grid code on 60M grid points, on 8 racks of Intrepid (8,192 nodes) in MPI+OpenMP mode (8,192 nodes with 4 threads per node) to do the Delayed Detached Eddy Simulation (DDES) and Improved Delayed Detached Eddy Simulation (IDDES) of flow past a generic landing gear geometry (modeled as tandem cylinders). They have presented a paper at the BANC-I workshop that was held June 10-11, 2010, immediately after the 16th Aeroacoustics Conference in Stockholm, Sweden. This paper will appear in the proceedings of the workshop. They are working on a different problem, with the same tuned code; a proposal for this work was submitted for an INCITE 2011 allocation. (Ramesh Balakrishnan of the ALCF worked with Philippe on the INCITE proposal.)

IMPACT: These calculations will spawn a realistic aeroacoustic noise calculation.

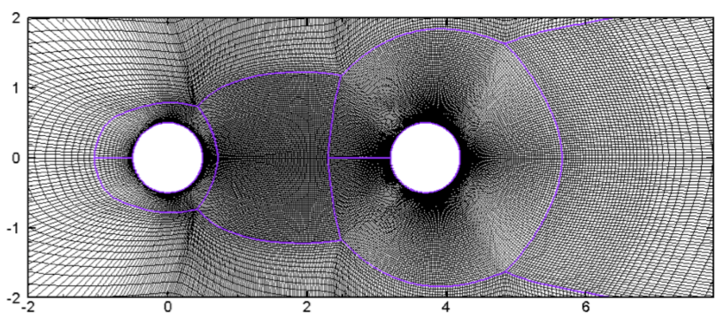

Figure 3. Zoomed grid subset.

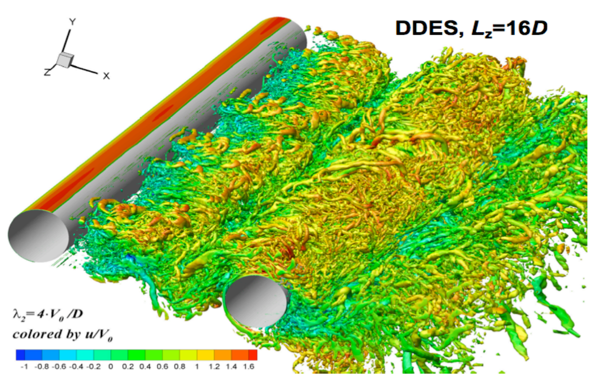

Figure 4. Isosurface of swirl lambda (2nd eigenvalue of velocity gradient tensor).

Accomplishment: Large-Eddy Simulation of Wind-turbine Airfoil

INCITE Project: "Overcoming the turbulent-mixing characterization barrier to green energy and propulsion systems"

PI: Anurag Gupta (GE)

This aerodynamics/aeroacoustics research is studying noise generation by wind turbines with the goal of discovering how to reduce it. The primary investigators are at GE Global Research, working with ALCF staff member Ramesh Balakrishnan. The work was done by Anurag Gupta (PI, GEGR), Hao Shen (co-pi, GEGR) on Intrepid before their INCITE 2010 award of computing time on Intrepid. This work was done in the June - September 2009 time frame. They are getting ready to continue the same work on a much larger number of grid points, utilizing up to 8 racks (32,768 cores) of Intrepid. Their code has been worked on and it now scales to 8 racks with a parallel efficiency of 60% when compared to a similar run on 512 nodes (2,048 cores) of Intrepid. They have done a high-resolution simulation of laminar-to-turbulence transition along the turbine blade that was achieved on 200M grid points on 8,192 cores (pure MPI) without tripping the flow, and at a Reynolds number of $2.5X10^6$. The project expects to get much improvement on this (i.e. they are working to get the correct skin friction at the trailing edge) as they start burning INCITE time (19M hours) on an even more refined grid (esp. refined at the trailing edge).

IMPACT: Noise reduction is critical for next generation propulsion systems, and energy systems such as wind turbines.

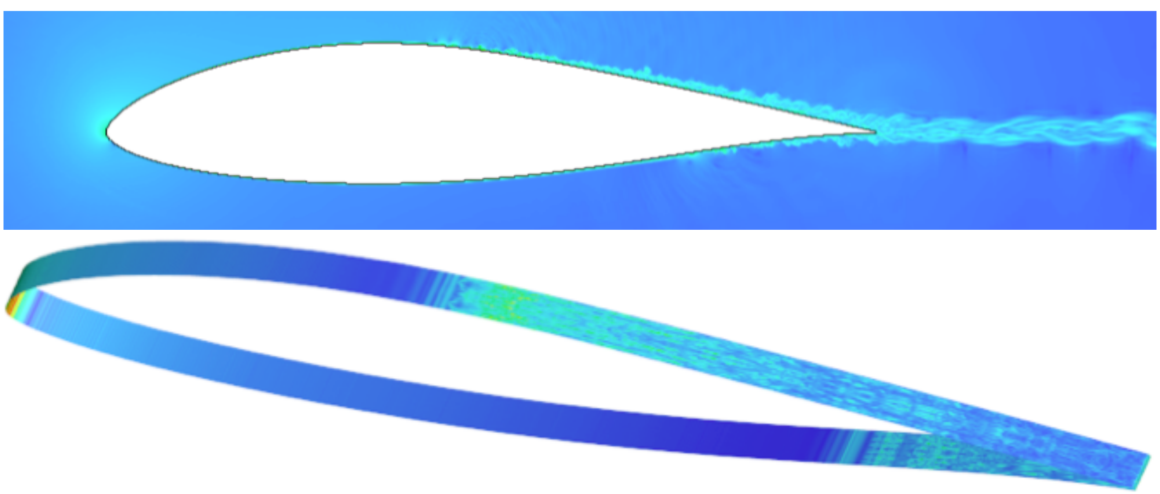

Figure 5. Laminar-to-turbulent transition.

Accomplishment: Reactive MD Simulation of Nickel Fracture

Discretionary Project: "Hierarchical Petascale Simulation Framework for Stress Corrosion Cracking"

PI: Priya Vashisht (USC)

Researchers from USC (Hsiu-Pin Chen, et al.) published a paper in Physical Review Letters (16 April 2010) on simulations of how sulphur particles in nickel grain boundaries alter its fracture properties, making it more brittle. This is the largest-ever (48 million atoms) chemically reactive molecular dynamics simulation on 65 K BG/P processing cores. A related press release is available at http://www.anl.gov/Media_Center/News/2010/BreakSpecial190609.html.

IMPACT: This embrittlement is an important issue in next-generation nuclear reactors.

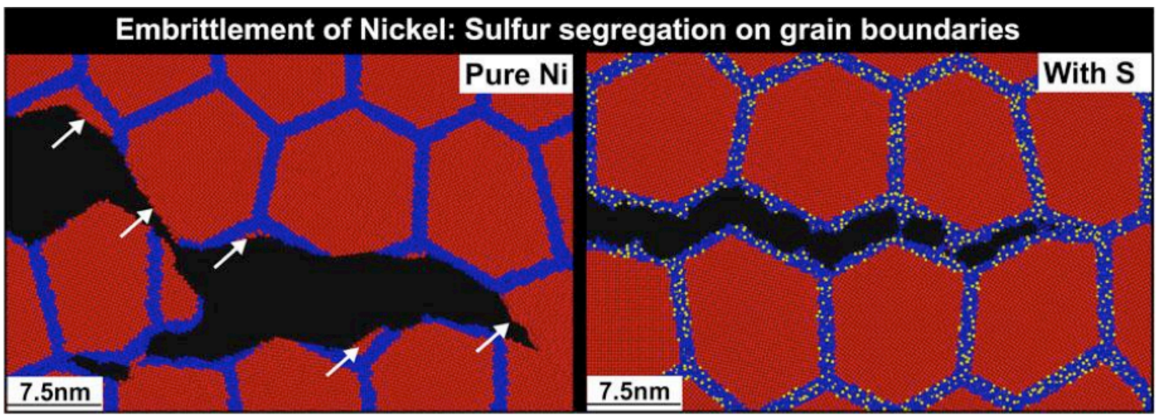

Figure 6. Close-ups of fracture simulations for nanocrystalline nickel without and with amorphous sulfide grain-boundary phases, where red, blue, and yellow colors represent nickel atoms inside grains (>0.5 nanometers from grain boundaries); nickel atoms within 0.5 nm from grain boundaries; and sulfur atoms, respectively. The figure shows a transition from ductile, transgranular tearing to brittle, intergranular cleavage. White arrows point to transgranular fracture surfaces.

Accomplishment: Electronic Structure of Gold Nanoclusters

INCITE Project: "Probing the Non-scalable Nano Regime in Catalytic Nanoparticles with Electronic Structure Calculations"

PI: Jeffrey Greeley (Argonne National Laboratory)

Dr. Greeley's team has investigated the electronic structure of gold nanoclusters varying in size from 13 to 561 atoms—a year-one milestone for this INCITE project. They have determined the stability and adsorption properties with respect to O and CO of the particles. They have also analyzed the global contraction of the particles' surfaces, from which the surface tension is extracted. This contraction influences the adsorption properties of the surfaces, which still have not converged to the corresponding properties of idealized single-crystal surfaces, even for the 561-atom particle. These results have motivated the researchers to study even larger gold nanoclusters (923 to 1,415 atoms). This increase will push the limits both of the GPAW code and the Blue Gene/P. For example, even in ideal weak/strong scaling, the proposed calculations would require running at 32 racks (80% of the full Blue Gene system) for an entire week. An article about these accomplishments is in preparation and will be submitted to a high-profile journal.

IMPACT: Catalytic nanoparticles may revolutionalize electrocatalysis for fuel cells, and metal catalysis for methane conversion.

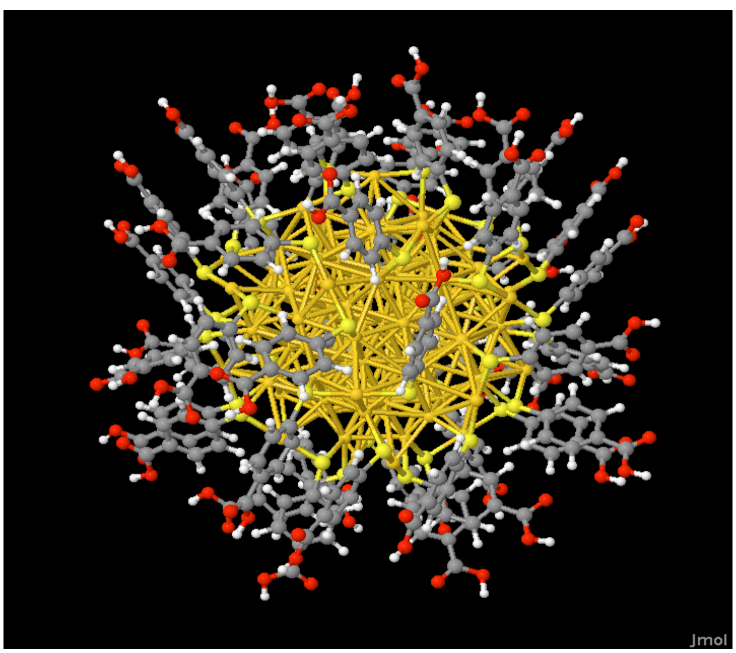

Figure 7. Researchers have investigated the electronic structure of gold nanoclusters varying in size from 13 to 561 atoms. Results have motivated them to study even larger gold nanoclusters (923 to 1,415 atoms).

Accomplishment: Recovering the Dispersion Interaction in GAMESS Molecular Simulations

INCITE Project: "Prediction of Bulk Properties Using High-accuracy Ab Initio Methods Interfaced with Dynamical Calculations"

PI: Theresa Windus (Ames Lab)

The code (GAMESS) uses a Fragment Molecular Orbital (FMO) approach to *ab initio* electronic structure methods. Below, benchmark timings show scaling of GAMESS using the "MP2" forces, which recover the so-called "dispersion interaction"—a quantum effect important for hydrogen bonding. Hydrogen bonding is relevant in that it explains such things as how the DNA double helix stays together. Protein conformation also is relevant in hydrogen bonding. These timings demonstrate the capability to properly account for quantum effects in biological-scale systems. Previously, dispersion was used only for modest-size systems of approximately 64 water molecules. The significance of these latest results is the scaling up of this technique to really large systems that could be regarded as being on the biological scale: replacing water molecules for amino acids results in a sizeable protein.

IMPACT: GAMESS is now capable of modeling biology-relevant molecules.

		Racks:	1	2	4	8	16
		Cores:	4096	8192	16,384	32,768	65,536
Waters	Atoms	Basis Functions	Wall-time-to-solution (minutes)				
128	384	5504	9.7	5.5	3.1		
256	768	11,008		11.6	6.6	4.1	
512	1536	22,016			16.8	10.1	7.4
1024	3072	44,032				24.1	16.8

Figure 8. *Ab Initio* Water Cluster Calculation Timings Using GAMESS on the Blue Gene/P. Method = Fragment Molecular Orbital-2- with calculation of MP2 forces. Atomic Basis Set = aug-cc-pVDZ.

Accomplishment: Computational Protein Structure Prediction

INCITE Project: "Computational Protein Structure Prediction and Protein Design"

PI: David Baker (University of Washington)

Starting with inexpensive backbone-only NMR data, the Rosetta protein-folding application is able to predict the detailed structure, showing good agreement with expensive full NMR experimental results. Results published: Raman, S., et al., "NMR Structure Determination for Larger Proteins Using Backbone-Only Data," *Science*, **327**, 1014 (February 19, 2010).

IMPACT: For systems with more than 150 amino acids, full experimental NMR is difficult and data is therefore scarce. Backbone data is much easier to obtain. This work demonstrates that backbone input is sufficient to get simulated energetics with acceptable accuracy.

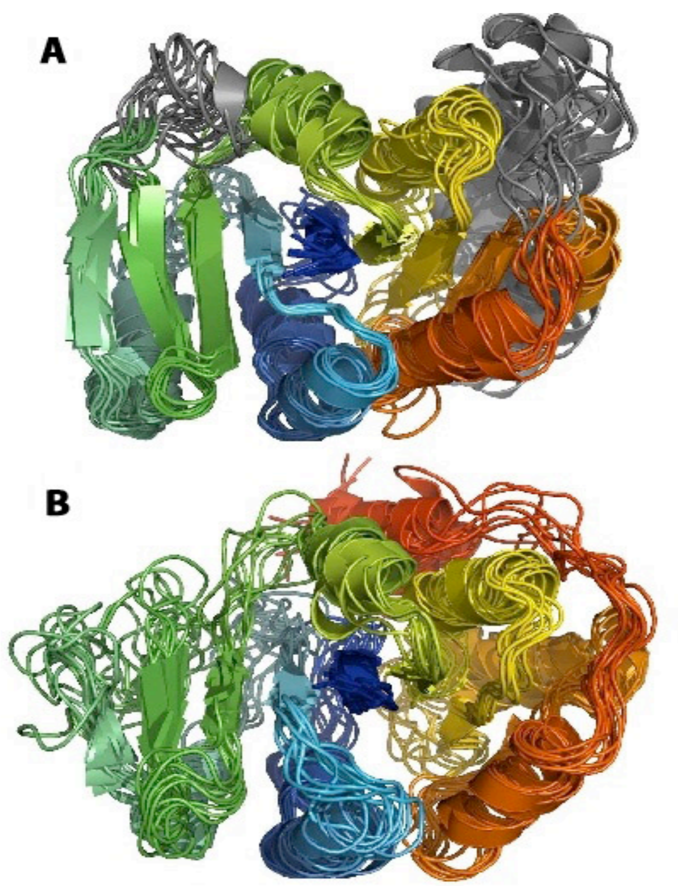

Figure 9. Large protein structures now can be computationally determined by incorporating backbone-only NMR data into Rosetta. Shown here is the structural comparison of ALG13 (201 amino acids) determined (A) computationally using RDCs and backbone NH-NH NOEs, and (B) experimentally by conventional NMR methods. These figures show ensembles of 10.

Accomplishment: UPOs in Navier-Stokes Equations, Whole-Brain Blood Flow Simulations

INCITE Project: "Large Scale Condensed Matter and Fluid Dynamics Simulations "

PI: Peter Coveney, (University College London)

Using a novel space-time variational approach, project researchers have identified the two first Unstable Periodic Orbits (UPOs) in the Navier-Stokes equation. This work was done on Intrepid using their HYPO4D lattice-Boltzmann code and a gradient descent algorithm for the numerical relaxation.

Patient-specific simulation of brain blood flow involves a number of steps (comprising the work-flow) which involve acquiring angiography data, its transfer to local resources, pre-processing locally, staging to remote resources for simulation, and finally reporting (using interactive steering and visualization). The sparse-geometry optimized lattice-Boltzmann code, HemeLB, can be used in two modes: (1) standalone mode where brain blood flow snapshots are dumped to disk and post-processed (no live user interactivity); and (2) real-time steering and visualization mode, where a user can interact with the simulation as it is running on Intrepid. Progress has been made in using HemeLB in standalone mode on Intrepid for the study of three patient-specific internal carotid artery aneurysms.

IMPACT: UPOs are a unifying concept in dynamical systems; improved understanding of fluid UPOs may lead to new statistical descriptions of turbulence. Patient-specific brain blood flow simulations will improve diagnosis and treatment of aneurysms.

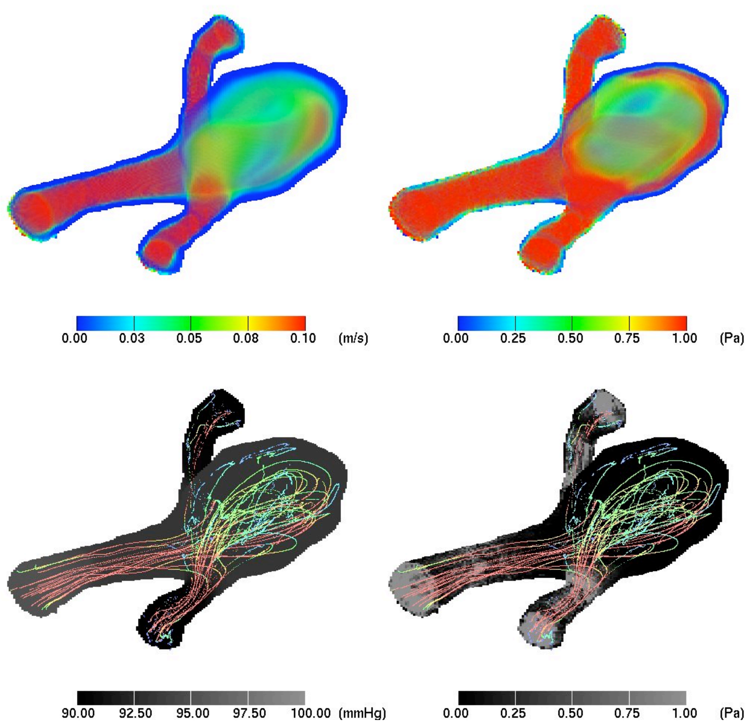

Figure 10. A visualization of one aneurysm. Top Left: volume rendered velocity. Top Right: volume-rendered von Mises stress. Bottom Left: external pressure and streaklines. Bottom Right: external von Mises stress and streaklines.

¬

Accomplishment: Laser Plasma Interaction in NIF

INCITE project: "Simulations of Laser-Plasma Interactions in Targets for the National Ignition Facility and Beyond "

INCITE: PI Denise Hinkel (Lawrence Livermore National Laboratory)

Part of this project is a large-scale particle-in-cell (PIC) simulation of laser backscatter off plasma in a NIF target. Recent experimental results showed unexpected importance of multi-beam interactions, so the planned simulation domain was increased to include multi-beam interactions. The investigators have completed the first three of O(5) 12-hour runs on all 40 racks of Intrepid (160 K cores). This whole set is a single problem.

IMPACT: Final validation of code to be used to optimize targets for future NIF shots.

Figure 11: A NIF ignition target. Ninety-six NIF beams enter the target through both the top and bottom laser entrance holes. The hohlraum is composed of a high-Z cylinder filled with helium gas and contains a capsule of DT in the center.

Accomplishment: Global Simulation of Plasma Microturbulence

INCITE project: "High Resolution Global Simulation of Plasma Microturbulence"

PI: William Tang (Princeton Plasma Physics Laboratory)

The I/O in the GTC particle-in-cell tokamak plasma simulation was recently re-implemented using the ADIOS package, created by S. Klasky of Oak Ridge National Laboratory and collaborators at Georgia Tech and Rutgers. This resulted in I/O on Intrepid that is ten times faster overall. A study of the performance of this I/O on Intrepid is the subject of a paper by Yuan Tian *et al.*, now in review. The GTC code now scales to 128K cores on Intrepid, and the project has plans for upcoming production runs on 64K and 128K cores. This scaling has set the stage to allow researchers to run some calculations on a scale relevant to ITER international fusion project, which they will complete with their INCITE allocation on Intrepid this year.

IMPACT: Being able to run on 128K cores will allow the project's primary simulations to be run on a physical ITER-size plasma.

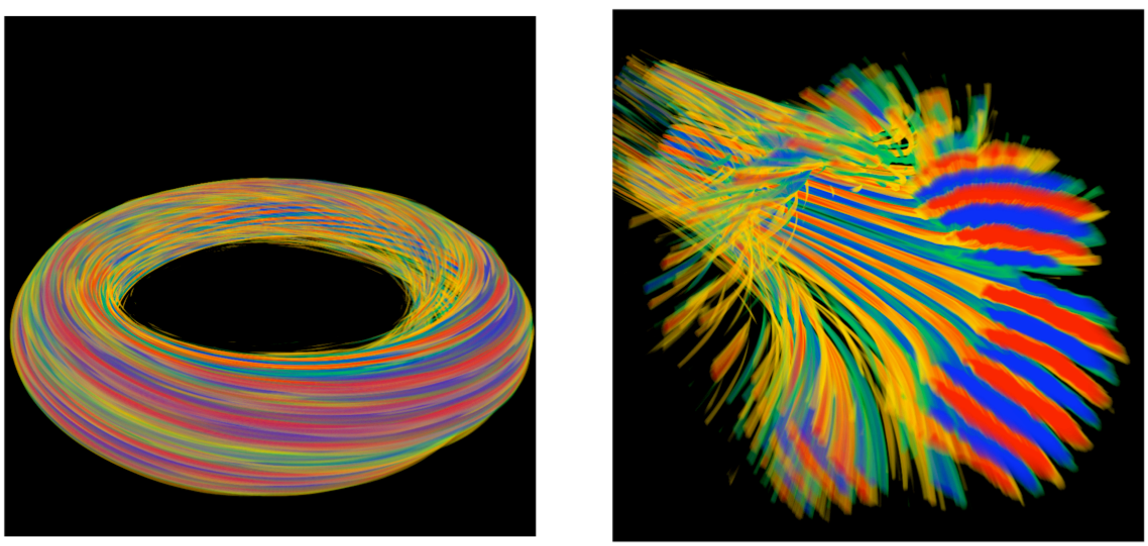

Figure 12. Volume rendering of the electrostatic potential field created by the plasma particles in a GTC simulation. Left: whole volume. Right: cross-section through a poloidal showing elongated eddies of the turbulence.

Accomplishment: Particle-In-Cell/Hybrid Simulations of Fast Ignition

INCITE Project: "Petascale Particle-in-Cell Simulations of Fast Ignition"

PI: John Tonge (UCLA)

The investigators postponed their planned full particle-in-cell (PIC) runs while implementing a newly invented PIC/MHD hybrid method (B. Cohen and L. Divol, to appear in *J. Comp. Phys.*, 2010). This method accurately captures the physics in some regimes of interest to the project and is 300 to 30,000 times faster than full 2-D PIC. This is now implemented and running on Intrepid. The group plans three hybrid runs of isolated targets. Some fully kinetic PIC/transport are still planned; these will take advantage of what's learned in the hybrid runs.

Fast ignition is an alternative to conventional NIF-style inertial confinement fusion. This new scheme separates the compression and heating phases, much like a petrol combustion engine. In the petrol engine, the fuel is compressed by the piston and then ignited via the spark plug. In the case of fast ignition, the driving lasers are the pistons, compressing the fuel to high density around the tip of a gold cone.

IMPACT: While it's believed that fast ignition could have 300x higher energy gain than conventional "hot spot" laser fusion, fast ignition's physics is less fully studied. The much greater density range and high execution speed the hybrid method allows will help map out much more parameter space and better understand the physics.

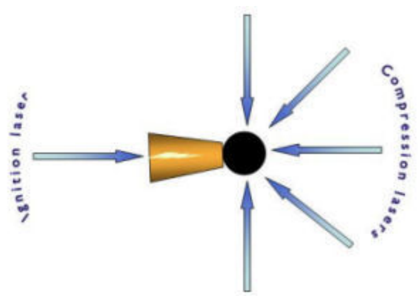

Figure 13. Fast ignition schematic.

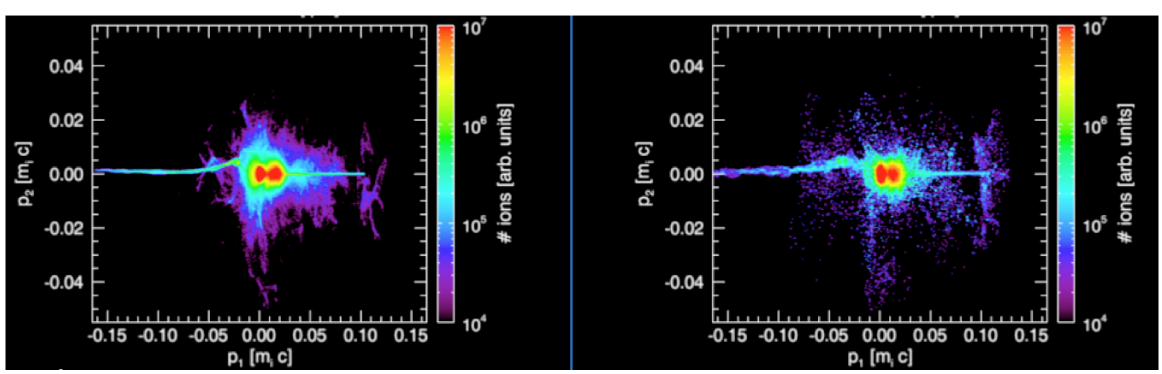

Figure 14. Full 2-D PIC results (left) and hybrid results (right). The hybrid run was 300 times faster.

Accomplishment: Quantum Turbulence in Bose-Einstein Condensates

Discretionary Project: "Quantum Lattice Algorithm for Quantum Turbulence"

PI: George Vahala (College of William and Mary)

The study of coupled Bose-Einstein condensates (BEC) is important for atomic interferometry and holds promise for extremely sensitive detection of gravitational and magnetic field variations. The investigators have developed a novel mesoscopic, reversible algorithm to study quantum turbulence and solitary waves in these condensates. Recently they implemented an algorithm for two coupled BECs. This method, sometimes referred to as a "quantum lattice gas method," is highly scalable, running efficiently on all 40 racks of Intrepid. Researchers have run systems with grids of size 5760^3 and larger. The PI has submitted an Early Science proposal and will submit an INCITE 2011 proposal. The PI and co-authors recently published a SPIE conference paper on this research. The physics of vortex nucleation and the statistics of quantum turbulence are new territory; this work is among the pioneers.

IMPACT: These runs discovered a multi-cascade in the turbulent energy spectrum.

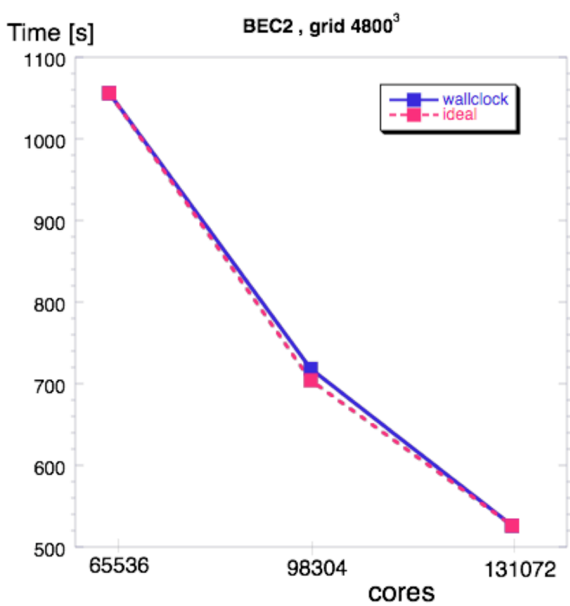

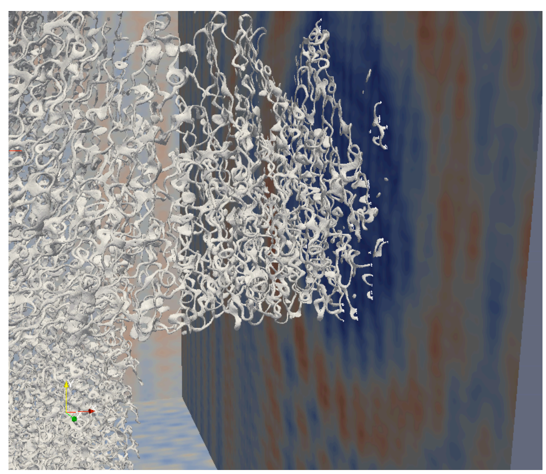

Figure 15. Strong scaling on Intrepid.

Figure 16. Vortex loop/filaments achievable on a 4032^3-grid. This is a blow-up of a sub-region showing degenerate filament vortices after a short time. This image was produced using a new parallel visualization system developed by a collaborator.

Accomplishment: Quantum Chemistry for PhotoCathodes

Discretionary Project: "Electron Aceleration in Laser Wakefields with High NumericalResolution"

PI: Karoly Nemeth (Argonne National Laboratory)

Ultra-thin MgO films on Ag(001) surfaces constitute an example of how ultra-thin surface layers on metals can be used to control the emittance properties of photocathodes. In addition to substantially reducing the work function of the metal surface, the MgO layers also favorably influence the shape of the surface bands, resulting in the generation of high-brightness electron-beams. As the number of MgO surface layers varies from 0 to 3, the emitted electron beam becomes gradually brighter, reducing its transverse emittance to 0.06 mm-mrad. Researchers suggest the use of such photocathodes for the development of free-electron x-ray lasers and energy-recovery linac x-ray sources.

IMPACT: Enables study of phenomena not experimentally accessible today.

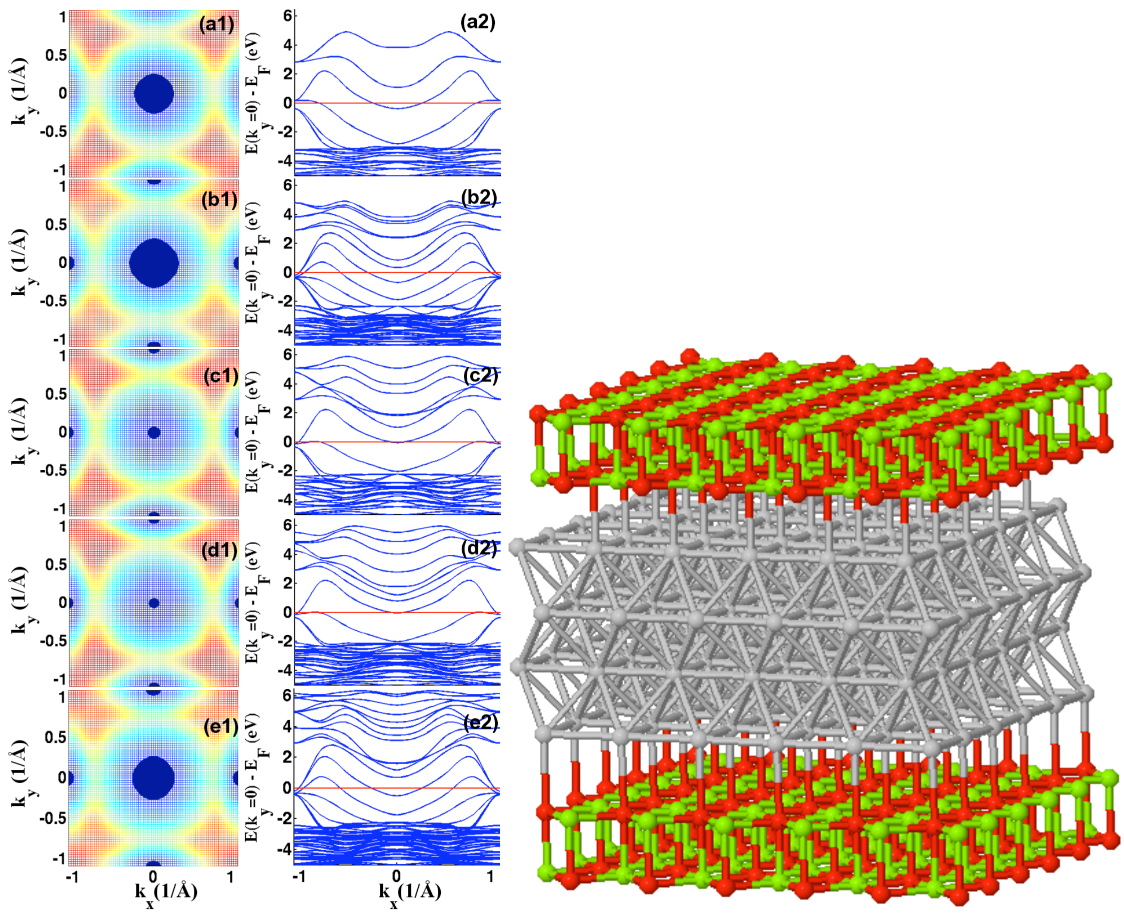

Figure 17. The number of the MgO layers on the Ag(001) surface controls the angular distribution of photo-emitted electrons (the brightness of the electron beam emitted), as indicated by the radius of the occupied region of the surface bands (dark blue spots) in the momentum space. Ultra-low transverse emittance of 0.06 mm-mrad can be achieved with the deposition of 2-3 layers of MgO. Such a low transverse emittance is critical for the development of x-ray free-electron lasers and energy-recovery linac x-ray sources.

Accomplishment: Turbulent Multi-Material Mixing

INCITE Project: "Turbulent Multi-Material Mixing In The Richtmyer-Meshkov Instability"

PI: Sanjiva Lele (Stanford University)

Three-dimensional simulations of the planar Richtmyer-Meshkov mixing have also been performed with the FDL3DI-SU code, for comparison with the experiment by Vetter and Sturtevant (*Shock Waves*,1995). In these experiments a shock of Mach 1.5 was incident on a perturbed planar interface separating air and SF_6 with the shock initially in air. The interface is subjected to a re-shock at a later time giving rise to a chaotic turbulent flow. The total mixing width which is a global measure of the mixing between the two fluids in the flow is calculated and shown in Figure 18.

A qualitative plot of the interface which results from an initial interface having 8 egg cartons along the cross-sectional directions is shown in Figure 19. A cross sectional cut along the center plane is also shown in the Figure. At 4.5 ms (just after reshock) we are still able to see the remains of the 8 egg cartons indicated by the 8 bubble-spike like structures, while at later times the flow is more mixed (turbulent mixing) indicated by the absence of any bubble-spike like structures.

IMPACT: First late-time turbulent multi-material mixing study with accurate compressible turbulence representation. R-M instability occurs in supernovae, inertial-confinement fusion, and hypersonic propulsion systems.

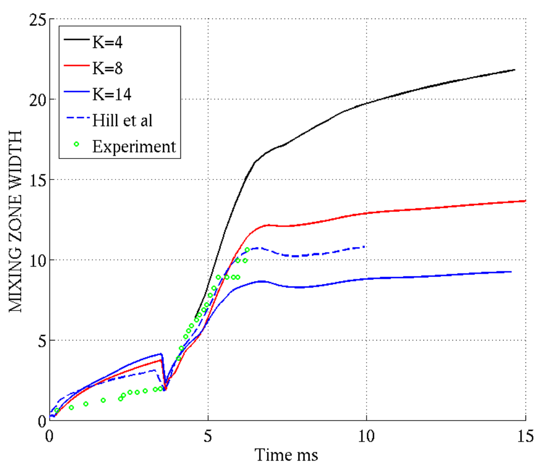

Figure 18. Evolution of the mixing zone between air-SF_6. Reshock seen to occur at ~3.5ms. K is the wavelength of interface perturbation. Results from experiments and previous numerical simulations are shown.

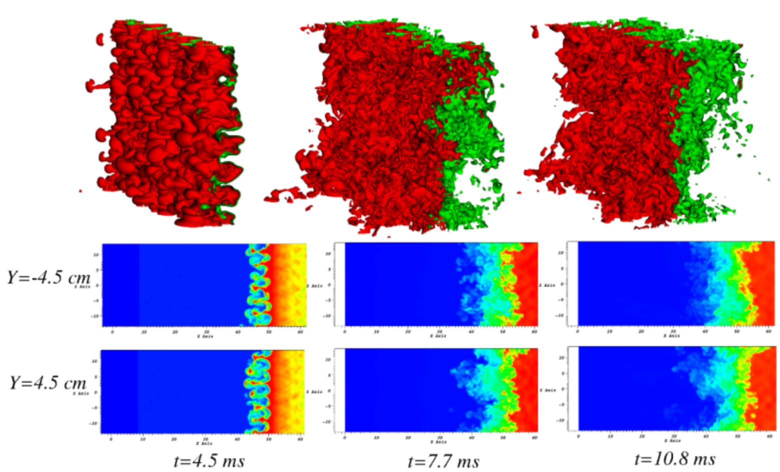

Figure 19. Snapshot of the interface (YZ plane) at three different times.

Accomplishment: ^{12}C Ground and Excited State Calculations; One-Sided MPI in GFMC Computational Nuclear Structure Calculations

INCITE Project: "Computational Nuclear Structure"

PI: David Dean (Oak Ridge National Laboratory)

Having completed ^{12}C ground-state calculations, the team has continued developing and testing various formulations of starting wave functions for the first excited 0^+ state of ^{12}C (The Hoyle or triple-alpha burning state). So far they have not found a starting wave function that remains orthogonal to the ground state wave function during Green's Function Monte Carlo (GFMC) propagation. They are continuing this work.

They have been using ^{12}C GFMC calculations to test a new version of the Automatic Dynamic Load Balancing (ADLB) library that uses the MPI one-sided puts and gets. This has revealed two bugs in the Blue Gene/P implementation of these so far and possibly a third. These bugs have been fixed and did not impact numerical results. They are now getting essentially perfect scaling up to 32,768 cores (code scales to 65,536), which is used for their production runs.

Co-PI Steven Pieper spoke about this INCITE project at the ALCF Leap to Petascale Workshop in May 2010. His talk was entitled *Using Really Big Computers to Study Rather Small Nuclei.*

IMPACT: Completed ground-state ^{12}C calculation was a key milestone and prerequisite to the current excited-state calculations. The ground state represents the best converged *ab initio* calculations of ^{12}C ever.

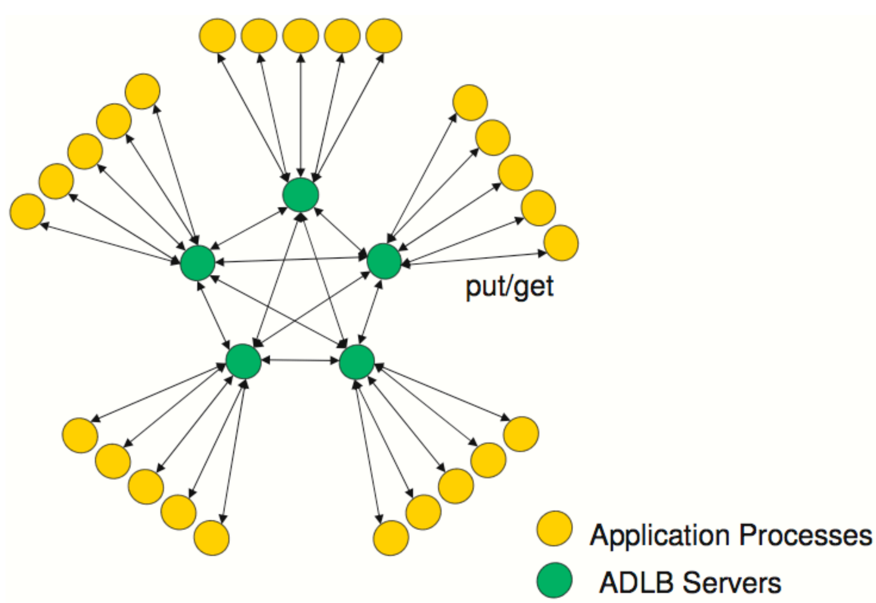

Figure 20. Workflow in Automatic Dynamic Load Balancing. No master required; slaves make calls to ADLB to offload or get work.

Accomplishment: Global Cloud-Resolving Climate Simulation

ALCC Project: "A Proposal from the Geophysical Fluid Dynamics Laboratory to Perform Prototype Ultra High-Resolution Climate-Weather Modeling Studies at Argonne National Laboratory"

PI: Shin-Jian Linn (GFDL)

Investigators completed runs of non-hydrostatic doubly-periodic experiments at 5km resolution with full-physics. Each experiment was run for 60 days at 0, 10, 20, and 30 degrees. Following these runs, they successfully ran the Held-Suarez experiment at 5km resolution with the cubed-sphere non-hydrostatic model for 60-days. With the completion of the Held-Suarez experiment, the projects has begun experiments of the 12km HIRAM non-hydrostatic, cubed-sphere-grid, finite-volume atmospheric model with full physics. The model is expected to run for 5-years and they have already completed several months of the experiment. Primary validation will be based on hurricane hindcast experiments, working toward a global high-resolution simulation. ALCF performance engineers have made several key optimizations to HIRAM on Intrepid.

IMPACT: Held-Suarez test was the key validation of HIRAM's non-hydrostatic dynamical core. Now physics-parameter tuning for hindcast experiments is underway.

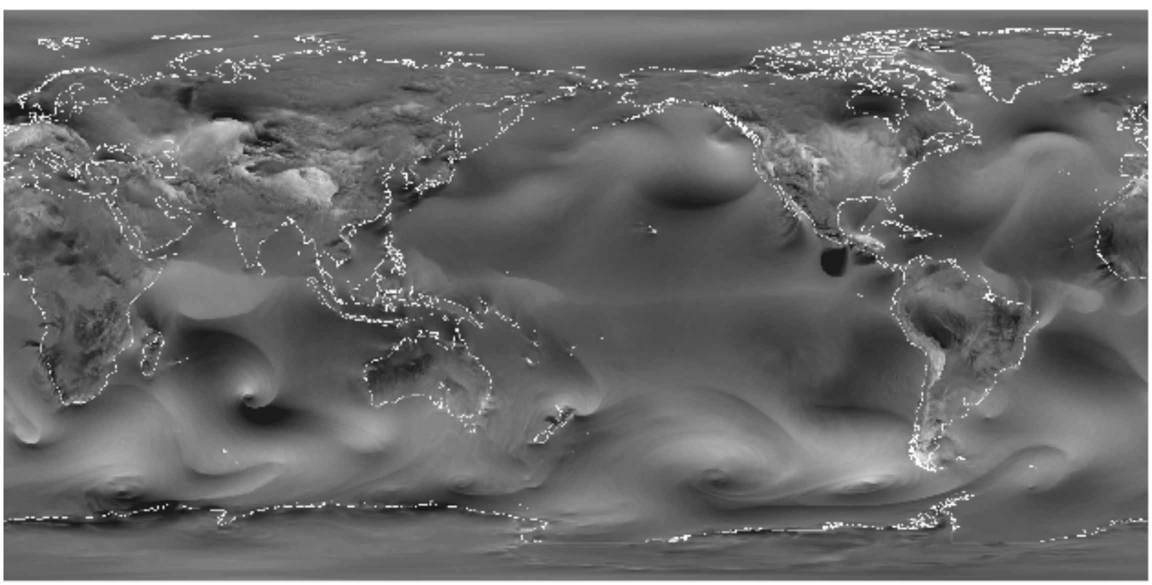

Figure 21. Shown in the figure is output from running a standard Held-Suarez test using a 4.5 km grid and the HIRAM dynamical core. This was an important step in validating the new dynamical core code.

Accomplishment: Hurricane Bill (Aug. 18-23, 2009)

Discretionary Project: "Large Scale Hurricane Simulations"

PI: John Michalakes (NCAR—now NREL)

This project made progress in closing the gap between weather and global models. Investigators ran a large-area simulation using the Weather Research and Forecast (WRF) model, using a moving, vortex-following nested mesh. The configuration was taken from data used in the real-time hurricane forecast system that ran during the 2009 hurricane season. The test case was Hurricane Bill. Using 16K cores on Intrepid, they were able to expand the finest mesh to a 1,000-km square to cover the entire hurricane, not just the eyewall, and expand the domain of the largest coarse-grain mesh to get improved large-scale forcing. Their study analyzed how such high-resolution forecasts could improve hurricane tracking capabilities and help mitigate risks.

IMPACT: Accurate hurricane tracking to mitigate risks.

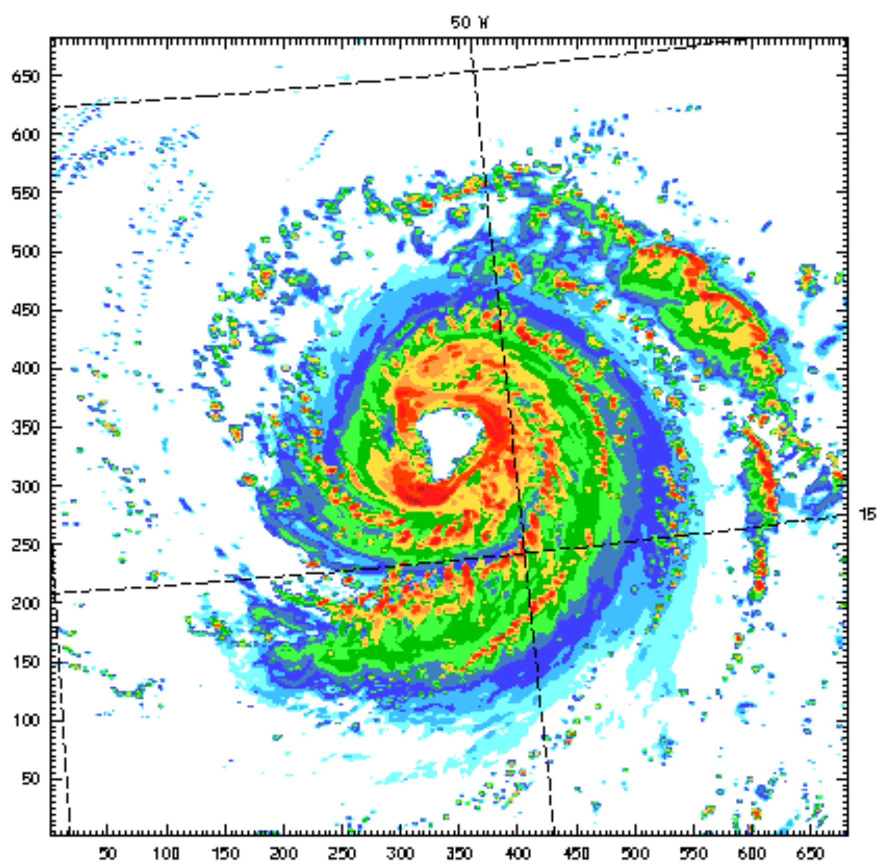

Figure 22. Investigators ran a large-area simulation using the Weather Research and Forecast (WRF) model, using a moving, vortex-following nested mesh.

Accomplishment: Turbulent Thermal Transport in Sodium Cooled Reactor

INCITE Project: "Advanced Reactor Thermal Hydraulic Modeling"

PI: Paul Fischer (ANL)

Paul and the Argonne SHARP project team model turbulent transport of energy away from fuel rods by coolants such as liquid sodium in next-generation nuclear reactor designs. One calculation was for a 217-pin wire-wrapped subassembly, computed on P=65536 cores on Intrepid using the Nek5000 spectral element code to run a large-eddy simulation. The Reynolds number for this flow is Re~10500, based on hydraulic diameter. The simulation is a watershed computation for Nek5000 as it is the first to exceed one million elements (2.95 million) and the first to use in excess of one billion gridpoints.

IMPACT: Liquid-metal-cooled fast nuclear reactors are a key technology for the next generation of safe, clean, nuclear energy.

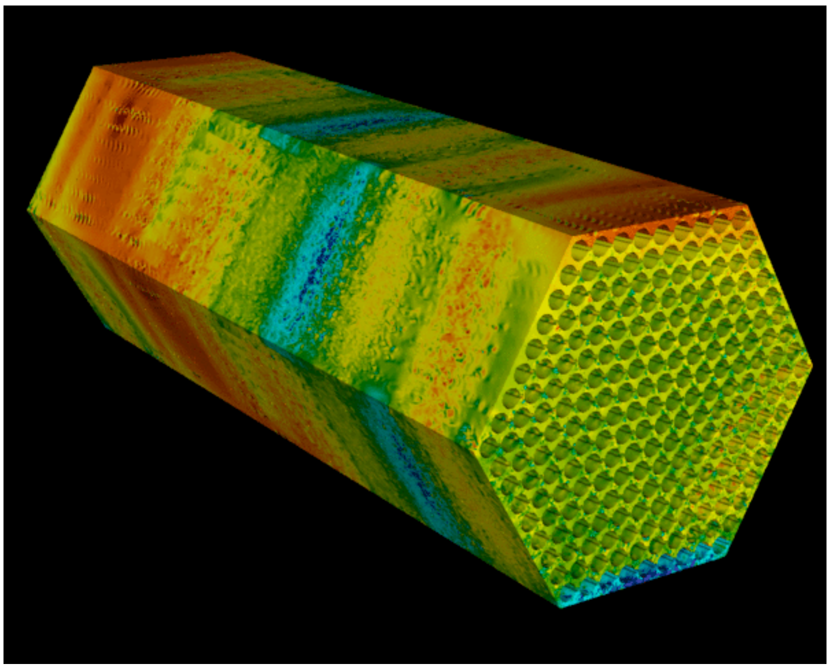

Figure 23. Coolant-flow pressure distribution in the large-eddy simulation.

Accomplishment: Two-Phase Flow & Combustion in Gas Turbines

INCITE Project: "Large Eddy Simulation of Two Phase Flow Combustion in Gas Turbines"

PI: Thierry Poinsot (European Centre for Research and Advanced Training in Scientific Computation (CERFACS))

Since May 2010, in collaboration with Argonne (Dr Balakrishnan) and Stanford (Summer Program 2010), CERFACS has performed the highest refinement simulations performed so far on a real combustor setup. As expected during the INCITE award, scalability production runs have been extended to 16k cores with development tests already reaching 32k cores. Further optimization on I/O and communication will allow further improvements in the short run.

The LES combustion model used in the industrial setups by CERFACS and its mesh independency have been validated on extreme conditions (combustion instabilities) and in a pure two-phase flow case. Validation of the injection modelisation strategy has been performed and a paper is currently finalized for publication (Sanjose et al 2010).

As part of this project, Poinsot and team have conducted the first-ever simulation of a full annular combustion chamber, using the ABVP Large Eddy Simulation code.

IMPACT: Simulations have identified instability mechanisms reducing efficiency in gas turbine combustors. Project goal is increasing efficiency.

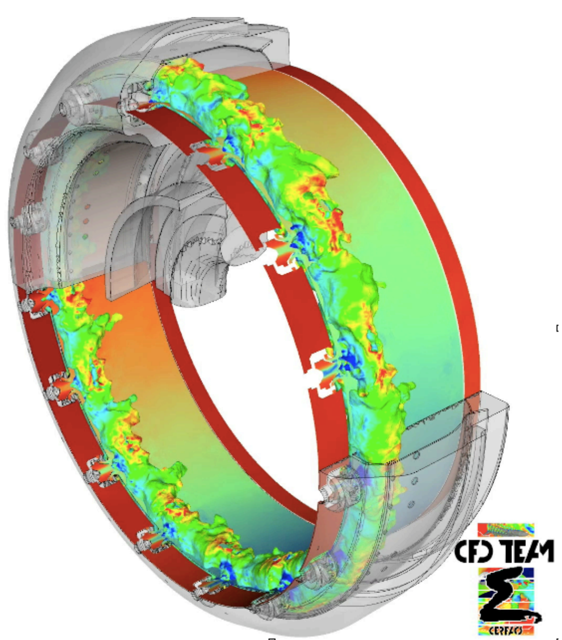

Figure 24. Large-eddy simulation of full annular combustion chamber.

Accomplishment: Protein-Interaction Simulations

INCITE Project: "Protein-Ligand Interaction Simulations And Analysis"

PI: Andrew Binkowski (ANL)

The project team has implemented the newly available "cobalt-subrun" feature on the ALCF BlueGene/P to reorganize how we organize our job submissions, which has improved their throughput and job submission efficiency.

The scientific results produced by runs on Intrepid are currently being validated in biochemical experiments and x-ray crystallographic studies. Preliminary results show that the project's NAMD-based computing pipeline can produce estimations of ligand binding affinity that can be correlated to experimental evidence.

One of the project aims is a proteome-wide analysis of binding domain and small molecule interactions. To maximize the knowledge from this analysis the project team selects targets from the 1,548 proteins (as June 29, 2009) of unknown function from the ProteinStructure Initiative (PSI). To date, the project has analyzed approximately 200 proteins of unknown function.

IMPACT: The project is publishing interaction data for the growing set of unknown-function proteins analyzed in a community repository.

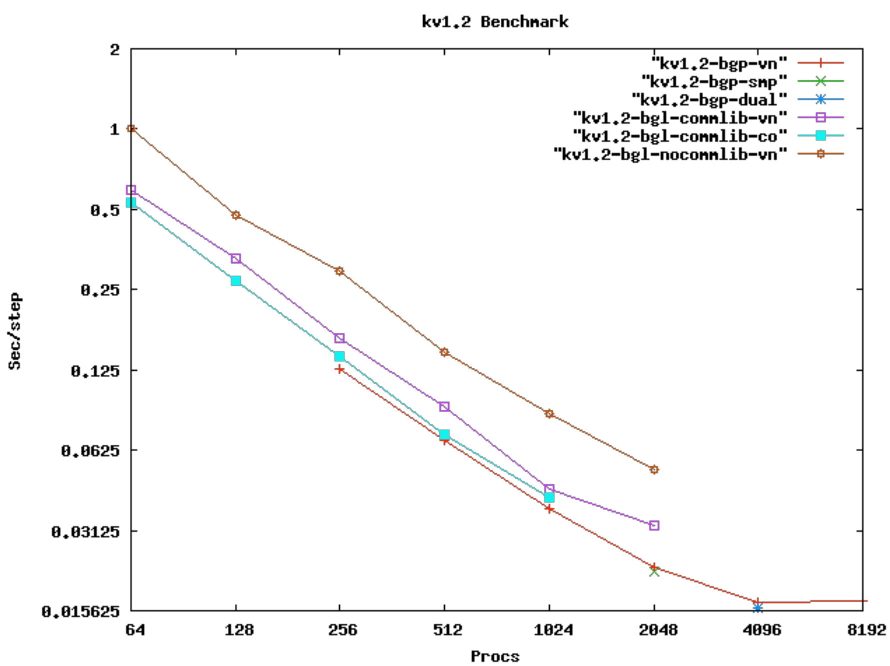

Figure 25. Scaling of NAMD on Kv1.2 Benchmark (352 atoms).

Technology Accomplishments

This section highlights some of the ALCF technology accomplishments during the period from August 1, 2010 to July 31, 2010.

Computational Protein Structure Prediction and Protein Design (INCITE, PI: David Baker)

Ray Loy (ALCF) implemented a new cobalt-subrun system for automating ensemble runs, which was instrumental in the work described in Accomplishment: Computational Protein Structure Prediction. This system is now in use by multiple projects using the ALCF resources.

Protein-Ligand Interaction Simulations and Analysis (INCITE, PI: Andrew Binkowski)

A single ligand binding site FEP calculation generally scales up to 2,048 processors or 1 rack, depending on the system size (number of atoms) and/or number of FEP windows in use. The new colbalt-subrun mode (see above) has been 7implemented that allows a single job to calculate 16, 32, or 64 ligands at the same time. The corresponding processors in use can utilize 8, 16, or 32 racks. This mode significantly simplifies the computational scheme and improves the throughput of ligand scoring.

In addition, with the help of ACLF staff, the project researchers were able to compile the Nucleic Acid Builder (NAB) suite of tools for biomolecular modeling. This allowed them to rewrite a significant portion of our pipeline and streamline the I/O profile by eliminating the need to read/write temporary files to pass data between applications.

Ray Loy (ALCF) set up a framework for running HTC mode jobs under Cobalt (the job scheduler used on Intrepid) and worked with the Cobalt team to improve the system environment that ensures post-job cleanup. As opposed to the normal HPC mode for MPI programs, HTC is suitable for running large sets of serial jobs at the same time. This framework is now used by multiple ALCF projects.

Research Into the Systematics of Type Ia Supernovae (INCITE, PI: Alan Calder)

The project code team successfully implemented a turbulence-flame interaction model in the version of the Flash Code we use. The prescription is that of Colin et al. (Physics of Fluids **12** 1843, 2000).

Large-scale Condensed Matter and Fluid Dynamics Simulations (INCITE, PI: Peter Coveney)

Project researchers have been refactoring the codebase of HemeLB in preparation for adding the efficient, elastic pipe algorithm they have recently published[2].

Computational Nuclear Structure (INCITE, PI: David Dean)

This project has been using ^{12}C GFMC calculations to test a new version of ADLB that uses the MPI one-sided puts and gets. This has revealed two bugs in the BG/P implementation of these so far and possibly a third. Nonetheless they are now getting essentially perfect

[2] G. M. Doctors et al., *Comp. Phys. Comm.*, **181**, 1584–1692 (2010).

scaling up to four racks. ADLB is a general purpose library, free for anyone to use; access it at http://www.cs.mtsu.edu/~rbutler/adlb .

Modeling the Rheological Properties of Concrete (INCITE, PI: William George)

The code team has added to their simulator the ability to model viscoelastic fluids. This is in addition to the ability to model suspensions with shear-thinning and Newtonian fluids.

Probing the Non-Scalable Nano Regime in Catalytic Nanoparticles with Electronic Structure Calculations (INCITE, PI: Jeffrey Greeley)

GPAW code developments, working with Nichols Romero (ALCF):

1. Completed parallelization of dense linear algebra -> can now always run in VN mode with state-parallelization and will never run out of memory.
2. Discovered new efficient BG/P mappings.
3. Further general tuning of the parallel performance. Working more closely with folks at IBM to optimize the code further.
4. Discovered a scheme that corrects for numerical grid effects (egg box effect) which can sum up to several eV's for large systems – and which furthermore influence the relaxation of clusters (grid-snapping to hot-spots of the egg-box effect)
5. Reduced memory footprint
6. Introduced a new layer of parallelism
7. Increased production use from 512 to 32K cores
8. Improved GPAW startup from 41 minutes to 3 minutes, opening the door to scaling up. This was done by a number of innovative methods involving python libraries and Blue Gene ramdisks.

Petascale Adaptive Computational Fluid Dynamics for Applications with High Anisotropy (INCITE, PI: Kenneth Jansen)

The Phasta_In project focused on one science application (aerodynamic flow control) and one development (improved I/O performance). Though not yet complete, the aerodynamic flow control application is the first one ever attempted on a full swept wing configuration. The I/O development now provides two alternatives that are 20-10,000 times faster than our original I/O. The student leading this effort has taken a summer internship at Argonne to transition this capability into Nekton which is an excellent example of the reusability of our developments.

Deterministic Simulations of Large Regional Earthquakes at Frequencies up to 2Hz (INCITE, PI: Thomas Jordan)

The improved scalability of our AWP-Olsen code affects many research groups within SCEC that either use the AWP-Olsen, or use the simulation results obtained with it. The structural inversion work done on ALCF provided an update to a SCEC community velocity model used extensively by a community of over 600 scientists in SCEC.

Lattice QCD (INCITE, PI: Robert Mackenzie)

James Osborn (ALCF) has been working with people from Boston University and elsewhere on multigrid solvers for the lattice QCD Dirac equation. They have a working implementation for one particular discretization (Wilson-clover) that gets up to a 20x improvement in solution time. This discretization is used by the group at Jlab (who are running at ORNL) and the "Pion Mass" discretionary project here, though it hasn't been integrated into their production codes yet

BG/P Plan 9 Measurements on Large-Scale Systems (INCITE, PI: Ronald Minnich)

IBM made an official release of the Plan 9 support code, and this code is now available at the ALCF. This release also included the KittyHawk code.

Simulation of "High" Reynolds Number Turbulent Boundary Layers (INCITE, PI: Robert Moser)

Developed a new plane/pencil decomposition and communication pattern that minimizes the transpose work required for DNS in domains in which the spatial extent in one direction is much longer than the others.

Large Eddy Simulation of Two Phase Flow Combustion in Gas Turbines (INCITE, PI: Thierry Poinsot)

The AVBP code is available for research purposes to a whole variety of research facilities in the world. The new optimizations present in the code allow for the accessibility to the AVBP community to a highly parallel tool working on all accessible architectures so far. All developments during this award are or will be transferred to the official AVBP branch.

High Resolution Global Simulation of Plasma Microturbulence (INCITE, PI: Wiliam Tang)

One of the members of the project team (S. Klasky) leads a parallel I/O software development project called ADIOS. It uses a very flexible runtime adaptive method to optimize the reads and writes on current parallel filesystems at high concurrencies. Tests of the ADIOS software have been carried out during this quarter to gather initial performance on both the GPFS and PVFS filesystems at different concurrencies and message sizes. The results were compared to the well-known parallel I/O benchmark, IOR, which normally gives the highest rates. On PVFS, ADIOS shows a significant performance improvement over IOR without any special optimizations or data alignment. A 1MB boundary alignment of the data on GPFS shows a 15% improvement in ADIOS writing rates on large message sizes. Re-implementing the I/O in the GTC code using ADIOS resulted in a 10x speedup.

Petascale Particle-In-Cell Simulations of Fast Ignition (INCITE, PI: John Tonge)

The project team has implemented a hybrid algorithm that makes possible full-scale 2-D simulation of fast ignition targets over a much greater density scale. Previously, for each four-fold increase in density, the simulation power had to increase by eight-fold due to the necessary increase in resolution. With the hybrid algorithm, high-density regions can be simulated with fluid field equations, allowing much lower resolution. The PIC/MHD hybrid method is from B. Cohen and L. Divol, to appear in *J. Comp. Phys.* (2010). For more details, see the Accomplishment: Particle-In-Cell/Hybrid Simulations of Fast Ignition.

Electronic Structure Calculations of Nano Solar Cells (INCITE, PI: Lin-Wang Wang)

The LS3DF code has been used as a Joule-metric code for 2010. The project team has been able to speed up that code by a factor of ~2 by including band index parallelization, DIIS algorithm for wave function iteration, and a better load balance.

Climate-Science Computational End-Station Development and Grand Challenge Team (INCITE, PI: Warren Washington)

The project has contributed several changes to the CCSM4 and CESM1 build systems so that the main configurations can build and run in the production queue of Intrepid.

Both CCSM4 and CESM1 were released to the climate community in Q2 (CCSM4 on April 1st, 2010 and CESM1 on June 25th, 2010). These are significant code bases containing advanced, state of the art climate and earth system models and a freely available and open source. Both codes contain Intrepid as a supported machine because of the work done under this Incite award. Supporting Intrepid within CCSM4/CESM1 will encourage users to apply for time on Intrepid for climate simulations.

The PIO library is also publicly available on Google Code and its support on Intrepid has been improved this quarter.

Prediction of Bulk Properties Using High-accuracy Ab Initio Methods Interfaced with Dynamical Calculations (INCITE, PI: Theresa Windus)

The code (GAMESS) now has a scalable implementation of the 'MP2' forces, which recover the so-called “dispersion interaction”—a quantum effect important for hydrogen bonding.

ALCF staff has also introduced the use of the RAMDISK feature on Intrepid, demonstrating that it can significantly reduce I/O overheads and improve performance:

Job Size (No. waters)	Processor Count	Wall Time (minutes)	Comments
1024	2048	36.9	Original FMO code (v3.3) for cluster MPP
		14.6	FMO v3.4.11 with densities stored in memory
		12.2	Use of RAMDISK feature on Intrepid
2048	32,768	35.3	FMO v3.4.11
		11.4	RAMDISK

Performance Evaluation and Analysis Consortium End Station (INCITE, PI: Patrick Worley)

- The Berkeley UPC and GASNet project released "Berkeley UPC 2.10.2" and "GASNet 1.14.2" on May 20, 2010.
  These releases included bug fixes for the BG/P that were developed using the INCITE allocation on Intrepid.
- Muster clustering library released, including a portable, tool-ready implementation the CAPEK algorithm described in the paper "Clustering Performance Data Efficiently at Massive Scales". (See listing above.)
  The library is available for download: http://github.com/tgamblin/muster

- Existing Libra load balance analysis tool now available for external download: http://github.com/tgamblin/libra

Awards

This section highlights honors and awards associated with the ALCF and its users this past year.

- Two students working on Thomas Jordan's INCITE project "Deterministic Simulations of Large Regional Earthquakes at Frequencies up to 2Hz" presented posters at SC09 on optimizations made to the AWP-Olsen software. These received Best Student Poster awards.
- Steven Pieper (and Robert Wiringa) of Argonne National Laboratory were presented the 2010 Tom W. Bonner Prize in nuclear physics by the American Physical Society in Washington, D.C. in February 2010. (Re: "Computational Nuclear Structure").
- Kenneth Jansen's INCITE project fluid dynamics code was a Gordon Bell finalist at SC09.
- Andrew Siegel, Michael Smith, Dinesh Kaushik (Argonne), and others were Gordon Bell finalists at SC09 with the UNIC neutron transport code (used in ALCF INCITE and ALCC projects).
- An innovative, energy-saving approach to cooling Argonne's Blue Gene/P supercomputer was recognized with an Environmental Sustainability (EStar) award from the U.S. Department of Energy's (DOE) Office of Science.
- The IBM Blue Gene series of energy-efficient supercomputers, central to breakthrough scientific research around the world, was singled out by President Barack Obama as a Medal of Technology and Innovation award-winner on October 7 in Washington, D.C. Computer scientists at Argonne and Lawrence Livermore national laboratories contributed critical input and software components through a DOE research and development partnership with IBM that strongly impacted Blue Gene's extreme-scale design.
- Two visualizations submitted by Argonne were among the top 10 winners at the SciDAC 2010 Electronic Visualization Night on July 14th. "Verification Study of Buoyancy-Driven Turbulent Nuclear Combustion for Three Different Physical Situations" and "Binary Galaxy Cluster Merger, Simulated Using the Flash Code, Mass Ration 1:1, with an Offset Impact, 4 Different Views" received OASCR awards. The visualizations used ALCF resources.

Metric 3.1 – Projections

Science Output and Science Accomplishments

The ALCF has decided not to make projections of science output and science impact. The decision was based on two considerations. First, the ALCF staff has no control over the computational science output; we work at creating the best possible environment for computational science and provide expert support that contributes to the productivity of the projects, as is documented elsewhere in this report. Second, it is not feasible to derive meaningful projections of science output and science results by projects that use ALCF systems for these reasons:

- The aggressive and innovative nature of the research frequently results in much greater variability of effort – and consequently elapsed time – in developing and running the application codes and obtaining science results, compared with mainstream computational science;
- The groundbreaking nature of the research may introduce more variability in the review time of papers submitted to journals compared to research using well-known methodologies;
- The number of projects is small (by design) for the Leadership Computing Facilities, therefore statistical errors are large. For example, in projecting the same number of publications per project based on previous years' results; year-to-year fluctuations are likely; and
- Although scientists typically are highly motivated to publish their results, the funding sources for some of the project teams do not encourage or value writing journal articles. Therefore, the mix of funding sources for projects – which varies year to year -- can affect the number of journal publications, which is a widely accepted measure of science output.

Allocation of Facility Director's Reserve computer time

We will continue to allocation of Facility Director's Reserve computer time at the current levels.

Baseline Area 4: Financial Performance

Charge Question 4: *How well is the program executing to the cost baseline pre-established during the previous year's Budget Deep Dive? Explain major discrepancies.*

FY '10 Operational Assessment Guidance: *In measuring the financial performance of the operating facility, the operational analysis should compare current performance with a pre-established cost baseline. Report results and provide projections using a methodology developed with the concurrence of the Program Manager, demonstrating operational cost effectiveness. The facility should prepare a financial sheet, which delineates effort, lease, operations, and DME that add up to the facility total budget.*

Metric 4.1 – Monthly Financial Reports

The ALCF will provide monthly reports on the steady state (SS) and DME (Development, Modernization, and Enhancement) costs to compare against plans as described in the OMB 300.

Metric 4.1 – Results

Monthly costs are compared against the baseline budget and are reported each month to DOE in the IT Dashboard report. In addition, a detailed budget/cost analysis is reviewed with ALCF management on a monthly basis.

Metric 4.1 – Additional Details

The following figure shows the planned FY10 budget (blue), cost projections for FY10 actuals (yellow), and the planned FY11 budget (red), broken out by the areas as specified for the July 2010 DOE Budget Deep Dive.

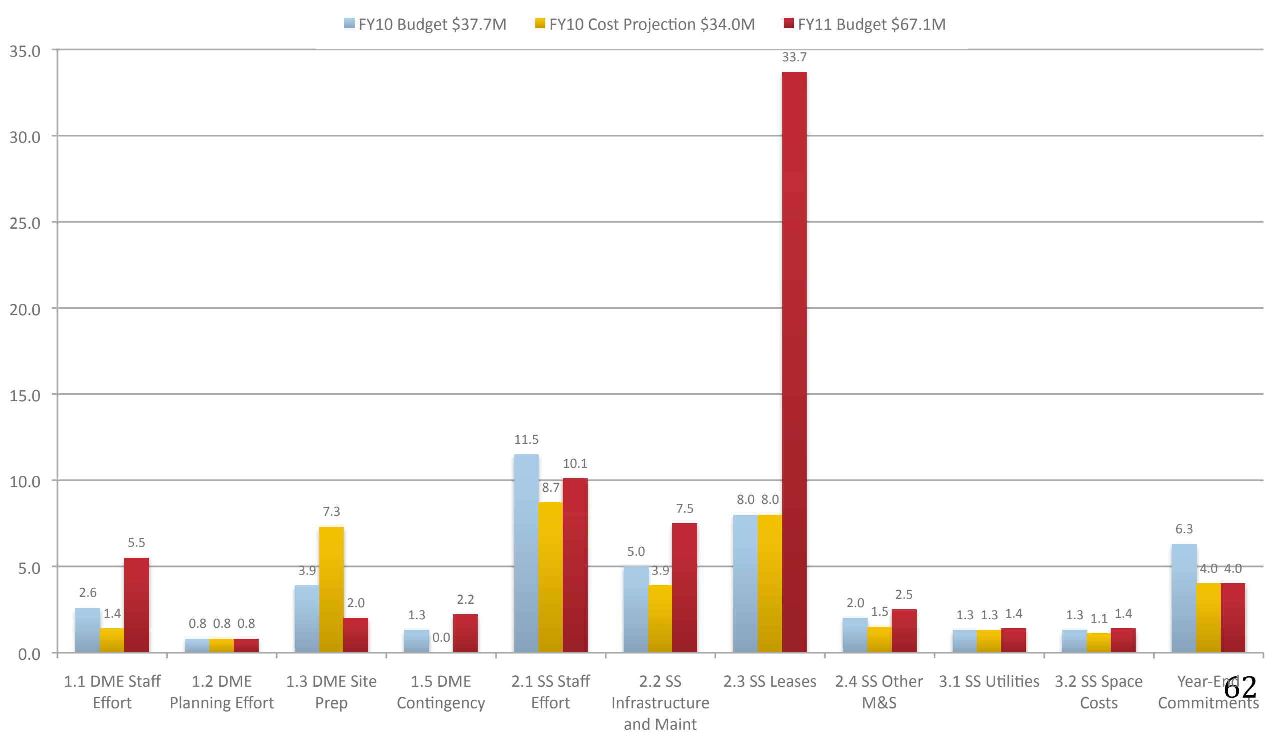

FY10 Budget, Projected Costs, F11 Budget Comparison Chart Notes

Explanations for the variances between the FY10 budget and FY10 cost projections, as well as for anomalies are provided below:

- FY11 large increase in element *2.3 SS Lease*s is the balloon (final) payment for the largest remaining lease.
- Legend totals (top of chart) do not include the year-end commitments shown in the far right columns on the chart.
    - Commitments represent contractually obligated costs in the following year for items such as hardware/software maintenance contracts, contractors, equipment on order, etc.
- The bulk of the FY11 staffing increases are in the *1.1 DME Staff Effort* element. The total FY11 staff on the SS side (*2.1 SS Staff Effort)* will increase from current levels, although the FY11 staff plan is less than the FY10 staff plan. It was expected that SS staffing levels would be reduced slightly during the second year of the ALCF-2 project. However, we are behind our staffing plan, so the net effect will be to increase SS staff in FY11. Total staff effort (1.1 + 1.2 + 2.1) increases in FY11.
- The increase in the *2.2 SS Infrastructure and Maint* element from FY10 to FY11 is due to escalation, projected increase of spares usage, expected increases in maintenance on some software packages, and funds allocated for early access to BG/Q hardware.
- The FY10 Budget vs. FY10 Cost Projection difference in Year-End Commitments, and the difference in *1.3 DME Site Prep* was the result of prepaying for facility enhancements instead of obligating them by year-end. This resulted in planned costs moving from FY11 to FY10.

Financial Details

A financial sheet, which delineates effort, lease, operations and DME, is shown below. The Projected FY10 column is based on actual costs incurred as of June 30, 2010, plus expected costs for the months of July, August, and September. Each category is discussed in detail after the sheet.

(Costs in $ K)	Pre-established at FY09 Deep Dive	Projected FY10	Proposed FY11
Operating Costs:			
2.1 Effort (staff)	11,462.1	8,666.7	10,072.0
2.2 Infrastructure and Maint	5,016.3	3,943.3	7,539.4
2.3 Leases	8,001.1	8,001.1	33,668.9
2.4 Other M&S	1,987.0	1,472.7	2,565.5
3.1 Utilities	1,303.0	1,308.7	1,368.0
3.2 Space Charges	1,315.2	1,112.2	1,381.2
Operating Subtotal	29,084.7	24,504.8	56,595.0
DME Costs:			
1.1 Effort (staff costs)	2,649.5	1,369.7	5,497.6
1.2 Planning (effort)	798.4	774.7	756.8
1.3 Site Prep	3,890.4	7,326.2	1,972.9
1.4 Leases (if any)	0.0	0.0	0.0
1.5 Contingency	1,314.5	0.0	2,252.2
DME Subtotal	8,652.7	9,470.6	10,479.5
Total Program Costs	37,737.4	33,975.4	67,074.5
Carry-Forward to Next Fiscal Year	11,979.2	15,741.2	10,666.7

FY10 Budget and Projected Cost Details

2.1 Effort (staff)

- Includes cost of LCF staff, matrixed staff, consultants, service center personnel.
- Personnel costs are below budget in FY2010, due to the continuing resolution which resulted in a slower ramp-up of staff than planned. Staffing levels have increased toward year-end, in line with the plan.
- Staffing is behind the planned level by an average of 6.7 FTE/month as of 6/30/10. Active recruiting and interviewing are in progress for the remaining hires planned for FY10.

2.2 Infrastructure and Maintenance

- This area includes a number of sub areas, including hardware and software maintenance, capital equipment, spares and hardware supplies. The maintenance includes not just the IBM contracts, but also 3rd party software such as TotalView and Tau.
- Spares and Capital Equipment spending has been limited, due to concerns with FY11 funding levels.

2.3 Leases

- All scheduled payments were made as planned in FY10.

2.4 Other M&S

- Includes cost to support the staff effort (travel, materials and supplies, low-value equipment (such laptops, monitors, disks), outreach, education/training, telephones, etc.)
- M&S estimates are based on existing staff counts, as well as expected number of new hires. Because hiring is behind schedule, staff support costs lag as well.

3.1 Utilities
- Utilities is the electrical power for the data centers and the offices. Projected cost is on budget.

3.2 Space Charges
- Space charges include offices, data center space, conference and meeting rooms, etc. It is partly based on square footage that differ depending on location, as well as number of staff. Because hiring is behind schedule, space charge costs have lagged as well.

FY10 Carry Forward

The FY10 $15.7M carry forward will cover estimated $4.0M year-end commitment amount, which includes $2.0M for DME Chiller equipment and $2.0M for other commitments. (Contractors, Hardware and Software maintenance contracts, etc.). Also included in the carry forward:

- Large carry forward is required so that the scheduled $30.0M lease payment can be made on 2/1/11.
- Carry forward includes $1.3M in DME Contingency that has not yet been used in FY10; Contingency is committed as part of the overall DME baseline budget.
- Carry forward also includes the required funding for Continuity of Ops into FY11.

The table below shows the planned carry forward.

Program Budget for FY10 ($ K)	
FY09 Carry Forward to FY10	7,716.6
DOE Annual Budget	42,000.0
Total Available Funding in FY10	49,716.6
Less Projected FY10 Cost	33,975.4
FY10 Carry Forward to FY11	15,741.2

FY11 Budget Plan

Costs in FY11 are expected to increase in line with the planned funding increase to $62M. The majority (~$25.7M) of the increased funding will be used to cover the scheduled FY11 lease payments of $33.7M, with the remaining increase budgeted to cover staffing increases.

2.1 Effort (staff)

- The ALCF Staffing Plan has 9 additional staff in FY11, bringing the total to 68.7 FTE by year end
  - New hires by group: +5 Catalysts, +3 HPC Operations, +1 Performance Engineering
  - +1 in December, +4 in March, +4 in June
  - Viz & Data Analysis (VDA) matrix staff are increased by 0.6 FTE (in DME)

2.2 Infrastructure and Maintenance

- Maintenance includes increase in FY11 budget to cover early access for BG/Q.
- Capital Equipment includes additional tape archival capacity, additional backup equipment, monitoring system improvements, tape archival performance improvements.
- Spares includes spare Blue Gene parts for Intrepid and Surveyor, increased in FY11 based on projected usage.

2.3 Leases

- "Go/No-Go" decision contractually required in October 2010 for all FY11 lease payments.
- $33.7M lease payments scheduled for FY11; includes $30.0M payment due 2/1/11.
- Final ALCF-1 lease payment scheduled for November FY12.

2.4 Other M&S

- Increase over FY10 budget to account for support of new staff hired during FY10 and FY11, as well as escalation.

3.1 Utilities

- FY11 budget includes provision for increase due to new staff and planned new equipment, as well as escalation.

3.2 Space Charges

- FY11 budget includes provision for new staff and increased machine room space charges, as well as escalation.

Staffing Plan:

The ALCF has an aggressive staffing plan for both FY10 and FY11, supported by the planned funding. ALCF maintains an ability to draw upon contractors and other parts of the Lab to provide support as required. For example, there are close ties to MCS, with a corresponding interchange of intellectual property and use of MCS staff as needed. There is also some reliance on other parts of the lab for operational support, including expertise in cyber security and networking (CIS), facilities (FMS), and technical services (TSD). The table below shows the break out of ALCF staff in FTE by group. In the first column is the current state; the 2nd column shows where we planned to be by the end of FY10 (7 remaining hires); the 3rd column shows where we plan to be by the end of next fiscal year (adding 9.6 new staff).

Group	Current ALCF Staff	Planned Staff @ 9/30/10	Planned Staff @ 9/30/11
Office of Director (OTD)	5.8	5.8	5.8
ALCF-2 Project Planning (PLN)	3.8	3.8	3.8
HPC Operations (OPS)	13.3	16.3	19.3
User Services (USO)	7.3	7.3	7.3
Performance Eng (PNG)	5.0	6.0	7.0
Adv Integration (AIG)	4.5	5.5	5.5
Viz & Data Analysis (VDA)	1.1	1.1	1.7
Catalyst (CAT)	11.3	13.3	18.3
TOTAL	52.1	59.1	68.7

DME staff members are included in the planned FTE numbers above. The DME staff are typically not hired directly into the project, but are either matrixed from other divisions (such as the Argonne project management office), or matrixed from ALCF Steady State. Basically, ALCF back-fills the Steady State positions with new hires, allowing them to come into a robust, well-managed, simpler environment, and shifting the more senior well-seasoned staff over to work on the newer, bleeding edge technology in DME.

During the next year, in addition to normal operations, we are preparing for the new system and supporting the Early Science projects, so the focus is to increase operations and computational science first, then provide additional user support and outreach staff as we prepare to open the new system to users.

As described earlier, the staffing plan from now through FY12 is:

- FY10: 7 additional staff (3 offers are pending, ALCF will continue its aggressive recruiting and interview schedules)
- FY11: Total 9 new staff: +5 Catalysts,+1 Performance Engineer, +3 Operations
- FY12: Total 5 new staff: +5 User Services and Outreach

ALCF Budget Formulation and Monitoring Details

1. The ALCF Staffing Plan and Infrastructure/Maintenance budget are formulated with input from group managers, which is then reviewed, prioritized and fit to the budget plan by senior division management. The ALCF Search Committee, composed of members from all groups and senior division management, meets regularly to discuss staffing plan progress and make adjustments as appropriate.
2. The ALCF Financial Manager and financial assistant carefully monitor all costs and perform a thorough review of the monthly cost reports the first week following the end of each month. The project costs are also monitored, and reviewed each month, by the project controls analyst and the project manager. Project Control Account Managers (CAMS) also review their individual DME WBS costs monthly. Costs are verified for validity and correctness; adjustments are made as needed.

3. Actual cost incurred, for both Steady State and DME, is compared to the budgeted numbers each month, and summary reports are reviewed with senior division management, as well as the project manager. Budgets are adjusted as needed; financial manager and division director approve spending; on the project side, the federal project director and the project manager approves all budgetary modifications thru Project Change Requests (PCR).
4. Effort rates are tracked and reviewed monthly, and updated quarterly as allowed by Lab policy. This allows the ALCF to stay current on rates to avoid surprises at fiscal year-end (variance spreads).
5. Dashboard data (DME and Steady State milestones) is submitted monthly. Reports reviewed with Federal Project Director prior to submission to ASCR.

The figure below shows the planned FY10 Budget versus the actuals and projected costs for FY10. The comparison columns are highlighted in green. The largest variances are in the area of Effort (both in DME and SS), 1.3 DME Site prep, and in 2.2 (caused by capital equipment and spares variances). Details for each are provided in the earlier "*FY10 Budget, Projected Costs, F11 Budget Comparison Chart Notes*" section.

Budget Plan vs. Actual YTD Cost reviewed each month

	FY10 Budget Plan (Budget reviewed with ASCR in Dec FY10)														
	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jun FY10 YTD Plan	FY10 Plan	Jun FY10 YTD Actual	Projection Jul-Sep	Total FY10 Projected	Variance (Plan - Projection)
1.0 DME/Project Costs, if any															
1.1 Effort (Staff costs)	146.0	158.1	137.7	198.6	181.0	243.5	255.8	206.6	260.2	1,787.5	2,649.5	839.3	530.4	1,369.7	1,279.8
1.2 Planning (effort)	76.0	59.1	58.0	63.4	63.4	72.9	69.7	63.4	69.7	595.6	798.4	586.7	188.0	774.7	23.7
FTE Plan	*9.5*	*9.8*	*13.5*	*7.9*	*6.8*	*7.8*	*9.4*	*11.3*	*13.3*			*6.0 June / 7.1 ave*			
1.3 Site Prep	0.0	25.5	0.0	0.0	0.0	0.0	664.3	603.9	664.3	1,958.0	3,890.4	7,228.2	98.0	7,326.2	(3,435.8)
1.4 Leases if any	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1.5 Contingency	39.8	43.5	35.1	46.9	43.8	56.7	177.3	156.5	178.1	777.7	1,314.5	0.0	0.0	0.0	1,314.5
TOTAL DME	***261.8***	***286.2***	***230.8***	***308.9***	***288.2***	***373.1***	***1,167.1***	***1,030.4***	***1,172.4***	***5,118.8***	***8,652.7***	***8,654.2***	***816.4***	***9,470.6***	***(817.9)***
2.0 Operations															
2.1 Effort (staff)	876.8	876.8	876.8	939.5	960.4	960.4	981.3	981.3	1,002.2	8,455.5	11,462.1	5,919.1	2,747.6	8,666.7	2,795.4
FTE Plan	*43.5*	*44.5*	*44.5*	*47.0*	*48.0*	*48.0*	*49.0*	*49.0*	*50.0*			*43.8 June / 40.4 ave*			
2.2 Infrastructure and Maintenance	215.0	215.0	215.0	215.0	215.0	1,433.0	215.0	215.0	215.0	3,153.0	5,016.3	1,417.2	2,526.1	3,943.3	1,073.0
2.2.1 Hardware Maintenance	118.3	118.3	118.3	118.3	118.3	118.3	118.3	118.3	118.3	1,064.7	1,419.8	1,037.6	345.9	1,383.5	36.4
2.2.2 Software Maintenance	26.9	26.9	26.9	26.9	26.9	26.9	26.9	26.9	26.9	242.1	322.5	271.0	90.3	361.3	(38.9)
2.2.3 Capital Equipment	0.0	0.0	0.0	0.0	0.0	641.0	0.0	0.0	0.0	641.0	1,282.0	0.0	720.7	720.7	561.3
2.2.4 Spares	0.0	0.0	0.0	0.0	0.0	519.5	0.0	0.0	0.0	519.5	1,039.0	0.0	756.7	756.7	282.3
2.2.5 Hardware Supplies	0.0	0.0	0.0	0.0	0.0	57.5	0.0	0.0	0.0	57.5	115.0	0.0	108.1	108.1	6.9
2.2.6 Cooling/Chillers	41.0	41.0	41.0	41.0	41.0	41.0	41.0	41.0	41.0	369.0	492.0	41.0	360.3	401.3	90.7
2.2.7 Engineering Studies	28.8	28.8	28.8	28.8	28.8	28.8	28.8	28.8	28.8	259.2	346.0	67.6	144.1	211.7	134.3
2.3 Leases	666.8	666.8	666.8	666.8	666.8	666.8	666.8	666.8	666.7	6,001.1	8,001.1	6,000.8	2,000.3	8,001.1	0.0
2.4 Other (M&S, travel, low value equip., etc.)	115.9	132.5	115.9	165.6	149.0	182.1	165.6	165.6	165.6	1,357.8	1,987.0	463.8	1,008.9	1,472.7	514.3
2.4.1 Professional M&S w/ Travel	17.4	19.9	17.4	24.8	22.4	27.3	24.8	24.8	24.8	203.6	298.0	257.0	144.1	401.1	(103.1)
2.4.2 Low Value Equipment	40.8	46.7	40.8	58.3	52.5	64.2	58.3	58.3	58.3	478.3	700.0	97.1	432.4	529.5	170.5
2.4.3 Staff M&S and Outreach	57.7	65.9	57.7	82.4	74.2	90.7	82.4	82.4	82.4	675.8	989.0	109.7	432.4	542.1	446.9
TOTAL OPERATIONS	***1,874.5***	***1,891.1***	***1,874.5***	***1,986.9***	***1,991.3***	***3,242.4***	***2,028.7***	***2,028.7***	***2,049.5***	***18,967.4***	***26,466.5***	***13,800.9***	***8,282.9***	***22,083.8***	***4,382.7***
3.0 Center charges															
3.1 Utilities (electricity)	108.6	108.6	108.6	108.6	108.6	108.6	108.6	108.6	108.6	977.4	1,303.0	923.8	384.9	1,308.7	(5.7)
3.2 Space (incl. maint., heat, etc.)	109.6	109.6	109.6	109.6	109.6	109.6	109.6	109.6	109.6	986.4	1,315.2	785.1	327.1	1,112.2	203.0
TOTAL CENTER CHARGES	***218.2***	***218.2***	***218.2***	***218.2***	***218.2***	***218.2***	***218.2***	***218.2***	***218.2***	***1,963.8***	***2,618.2***	***1,708.9***	***712.0***	***2,420.9***	***197.3***
TOTAL STEADY STATE (2.0 + 3.0)	**2,092.7**	**2,109.3**	**2,092.7**	**2,205.1**	**2,209.5**	**3,460.6**	**2,246.9**	**2,246.9**	**2,267.7**	**20,931.2**	**29,084.7**	**15,509.8**	**8,995.0**	**24,504.8**	**4,579.9**
ALCF TOTAL (DME and SS)	**2,354.5**	**2,395.5**	**2,323.5**	**2,514.0**	**2,497.6**	**3,833.7**	**3,414.0**	**3,277.3**	**3,440.0**	**26,050.0**	**37,737.4**	**24,164.0**	**9,811.4**	**33,975.4**	**3,762.1**

Baseline Area 5: Innovation

Charge Question 5: *What innovations have been implemented that have improved the facility's operations?*

FY '10 Operational Assessment Guidance: *The Facilities strive to provide high quality, integrated support for scientific research; innovations in support integration should be documented. For each of the following metrics, report results and provide projections using a methodology developed with the concurrence of the Federal Program Manager: Infusing Best Practices and Technology Transfer.*

Metric 5.1 – Tracking and Reporting Innovation

The ALCF will report annually on innovations and best practices developed, shared with the community, and imported from the community, as well as participate in the yearly ASCR-sponsored Facilities Best Practices Workshop.

Metric 5.1 – Results

Innovation: Safety device for passing through openings in machine room raised floors (ALCF Operations Group)

The under floor pedestals, used to hold the tiles, have a very thin "blade-like" plate at the top for the tile to rest on, which is quite sharp. It is in the corner, and is somewhat protected, but was a safety concern for the team. The ALCF operations staff designed an "insert" that drops into the hole and has a lip that rests on the floor and the protruding pedestal, thus preventing a cut. An insert has been created and the team is in the process of testing it at the data center. If it proves to work well, we will provide the design to other facilities.

Innovation: OARtool (ALCF Operations Group)

The ALCF Operations Group developed a new tool called *OARtool*. This tool provides a central repository for availability and interrupt data, along with data manipulation, maintenance and analysis tools. Multiple resources may be tracked, where each resource is a distinct reporting entity such as 'Blue Gene Intrepid' and 'Blue Gene Surveyor'. Analysis mechanics permit additional non-reporting resources that may ordinarily be considered a subset of reportable resources. This facility is useful for internal tracking of disk and other complex subsystems.

Availability tools track overall resource availability, which is then further divided into scheduled and unscheduled periods. Desired capabilities include overlapping entry support, both in system resources and time, fault source tracking, reporting source, and the ability to record 'no impact' records to annotate system events. Overlapping entry support permits multiple availability entries that may wholly or partially represent the same fundamental resource unavailability, which frees administrative staff from determining if another data source (frequently of an automated nature such as system hardware monitors) has already recorded an availability event. Fault source tracking allows deep analysis to spot issues such as systematic problems within a resource's subsystems.

Reporting sources include administratively input data (gathered from operational experience), historical scheduler data and automated software and hardware health monitors such as IBM's Reliability, Availability, and Serviceability (RAS) facility.

Overlapping data presents interesting challenges in ensuring availability data is properly deconflicted prior to analysis. Availability entries include a start and end time along with location data that describes which portion of the associated resources was unavailable. All availability entries, by default, are considered unscheduled. The deconflicting process deconstructs the provided location into the smallest entities that may be scheduled on the resource (for example, psets on a Blue Gene, or a node on clusters). During analysis, OARtool tracks time 'extents' for each schedulable entity. As availability entries are processed, extents are added or extended to reflect a single timeline for each entity. Once processing is complete, another pass is made considering records specially marked as 'scheduled unavailability periods'. A separate set of extents is tracked, and used to mask the first analysis run to determine overall, scheduled and unscheduled availability problems. Treating all data sources as unscheduled, then masking the final analysis with 'scheduled period' records greatly simplifies data record entry.

Interrupt records track data similar to availability records, including resource, fault source, and data origin. Jobs with a common root cause are collected into a single interrupt record, which can be treated either as an instantaneous event, or an event with a start/restore time. In addition, links to other data sources are permitted to associate data such as one or more RAS events with a given interrupt record. Fault sources for interrupts are expanded to include user causes allowing for comprehensive data analysis.

OARTool back end storage is provided by a SQL database. Availability and interrupt events are modeled in a one-to-many fashion where one of either record correlates potentially multiple resource components or jobs together into a single availability or interrupt event. Resources, components and reporting sources are structured such that both the user interface (A PyQT application) and database are sufficiently flexible to permit additions and modifications without changes to the underlying schema. As processing of location information is resource dependent, python fragments are kept in tables, keyed to resource, to allow flexible analysis with little or no change to the application code base.

Future features include an ability to link all data source records in a many-to-many relationship, comparison of automatically reported data and administrative data, and the ability to plug the tool as a whole into other applications. Data linking is particularly intriguing, as it will allow related events to be grouped for speedy root cause analysis. Should a hardware failure trigger a RAS event, with subsequent job interruptions and hardware taken offline, all records involved (automatic reporting, job interrupts and hardware availability) could be linked for ready analysis.

Best Practice: Partnership with Computer Science Developers

For faster application innovation, leading to breakthrough science on the BG/P, ALCF works closely with computer scientists and vendors to provide necessary tools for application developers.

MPI Collaboration Results (Team: MCS and ALCF)

Improved performance for multithreaded MPI: In collaboration with IBM and LLNL, ALCF and MCS worked on improving multithreaded MPI performance by avoiding reference counting of MPI objects within MPICH2 when the application makes some semantic assertions. Preliminary experiments and performance results indicate significant improvements in messaging rates.

Improved performance of one-sided communication operations: Work was done in collaboration with IBM. The previous design introduced high latency making it nearly impossible to overlap communication with computations.

Best P2P communication practice:. The team developed a benchmark that demonstrates the highest messaging rate for random communication pattern on both latency bound and bandwidth bound problems. At least two applications, ABVP from CERFACS and NEKTAR from Brown University, have benefited from these findings for large node counts.

Continued development of fine-grain multithreading capability in MPICH2: Work was done in collaboration with IBM Watson. Included collaborations with IBM on eliminating or reducing the need for reference counting of MPI objects in multithreaded environments. A paper on the work is being submitted to Cluster 2010.

Parallel Virtual File System (PVFS) Collaboration Results (Team: MCS and ALCF):

A combined team formed with members from MCS Systems Software team and ALCF's Advanced Integration Group designed and developed *Darshan* - a lightweight tool for characterizing I/O behavior. This has been enabled on ALCF's BG/P systems to profile I/O characteristics of our applications. A paper titled "I/O Performance Challenges at Leadership Scale," was presented at SC09. The team added members from ALCF's Application Performance Engineering group to analyze INCITE and Discretionary project Darshan data.

The INCITE project EarthquakeSim was identified as not having a good I/O model from early analysis. It was found that 50% of its run time was spent in I/O, and 90% of the I/O time was related to metadata. This I/O signature indicates a potential for significant performance improvement if the issue can be addressed. The team met with the EarthquakeSim application team during the recent Scaling workshop in May to address this issue.

The GFDL code HIRAM was also analyzed with Darshan and the traces revealed a number of file system seek operations. As GFDL code does few if any seeks itself, the culprit must lie with associated libraries. The team found locations in the serial netcdf cache management code where the library called lseek to assert that the current file pointer position matched up with internal bookkeeping. If netcdf is recompiled with assertions disabled, we can save about 30 million system calls on a 2,000 core run.

Best Practice adopted from Los Alamos National Laboratory: Weekly reviews of operational data (ALCF Operations Group)

In discussions with Randall Rheinheimer at Los Alamos National Laboratory, it was learned that they held a regular weekly meeting where they evaluated, confirmed, refined, where possible determined root cause, and recorded the failure and availability events of the last week. As of March 2010, we have adopted this practice and improved upon it by automating much of the data collection and developing database applications that allow us to query all availability and failure events.

Best Practice: Foreign National Cyber Expiration Warning System (ALCF User Services and Outreach Group)

Originally ALCF managed its foreign national cyber access paperwork processing through an administrative assistant responsible for the specific user. While this worked for users internal to Argonne, we were still using significant resources to keep track of when this paperwork expired for each user and we had one administrative assistant whose sole job was tracking these dates and user's paperwork.

Last year we changed the system by which we process foreign national account cyber access paperwork. Instead of the updating of paperwork being initiated by us checking on dates through the lab-wide FAVOR system, we instituted a process where the user and their sponsor received advanced warnings about soon-to-expire paperwork. When the user or the user's sponsor responded to this warning, it triggers the renewal of the necessary paperwork. This warning system was automated programmatically and our corresponding processes were changed accordingly. This new system now allows for user base scaling without requiring an increase in dedicated resources.

Best Practice: Improving utilization and availability through diagnostics scheduling (Team: MCS and ALCF Operations Group)

Blue Gene diagnostics on Intrepid accounted for nearly half of the regularly scheduled maintenance time. MCS and ALCF staff extended the Cobalt scheduler to handle dynamically created partitions and a special diagnostics queue, with a priority of zero, was implemented. Diagnostics jobs are submitted to the queue, and will only run when nothing else can run from any other queue, including backfill. Removing basic diagnostics from scheduled maintenance significantly shortened the regularly scheduled maintenance periods.

Participation in the DOE HPC Best Practices Workshop (ALCF Staff)

Three ALCF staff members participated in the *3rd DOE Workshop on HPC Best Practices: HPC Center Software Lifecycles* September 28-29, 2008 in San Francisico, CA. Susan Coghlan as the *Day 2: Software Stages* Chair, Charles Bacon as a co-chair of the *Sustainment (long term issues)* breakout, and Bill Allcock as a co-chair of the *System Management* breakout.

Baseline Area 5: Risk Management

Charge Question 6: *Is the Facility effectively managing risk?*

FY '10 Operational Assessment Guidance: *Each facility should utilize a risk management plan to document operational risks. The facility should provide information about the development, evaluation, and management of the top five to seven operating and technical risks encountered during the previous year. Develop methodology, with the concurrence of the Program Manager, to evaluate and report on resolution of the major operating risks in CY 2010 and provide projections for 2011.*

Metric 6.1 – Risk Management Process

The ALCF will use a documented risk management process.

Metric 6.1 - Results

The ALCF Risk Management Plan has been in place since November 2006. This plan was developed based on a review of past risk strategies at Argonne and consultations with experts in the field of risk and risk management software and other DOE facilities. In September 2008, ALCF was actively involved in the DOE Best Practices workshop on Risk Management Techniques and Practice. Jeff Sims was a member of the steering committee. Susan Coghlan and Ira Goldberg were both workshop breakout co-chairs. The workshop assessed the current and emerging techniques, practices, and lessons learned for effectively identifying, understanding, managing, and mitigating the risks associated with acquiring leading-edge computing systems at HPC centers. During the past year, ALCF has continued to expand on this proven approach to managing risk by updating and incorporating risk management best practices in to our strategy. The ALCF Risk Management Plan (RMP), Version 3.3, updated June 23, 2010, provides guidance for implementation of the ALCF's Steady State and DME risk management strategy.

Metric 6.1 – Additional Details

Risk Management Strategy

The goal of the risk management strategy for the ALCF is aimed at maximizing the likelihood of success within the overall cost, technical, and schedule envelope. The RMP outlines the risk management process and describes how risks are identified, assessed, tracked, updated, and retired. The risk management tool is used to track risks, quantify and qualify risk impacts, and provide a risk register. In the past year, the ALCF has moved from the *Risk Radar* risk management tool, to the risk management tool, *PertMaster®.* This tool is more flexible, has more extensions, and allows DME project risks to be tied to actual tasks in the WBS with a direct connection to the WBS management tool *Primavera*. The ALCF team continually evaluates operational risks through monthly risk status meetings, and shares major risks with the Program Office. These meetings enhance the effectiveness of risk management by evaluating impacts, contingencies, and mitigations as they develop. Some of the impacts and mitigations become supporting evidence in developing relevant lessons learned. When analyzing steady state risks, we apply the criteria shown in the following figures:

Impact / Risk Area	Low	Moderate	High
Cost	≤ $100K	$100K ≤ $500K	>$500K
Schedule	Impacts current FY OMB E300 specified INCITE production requirements by <5%	Impacts current FY OMB E300 specified INCITE production requirements between 5% and 10%	Impacts current FY OMB E300 specified INCITE production requirements by >10%
Technical	< 10% (Negligible, if any, degradation to computation, data security, storage, or customer support)	10% <= 20% (Significant technical degradation to computation, data security, storage, or customer support.)	> 20% (Severe technical degradation of computation, data security, storage, or customer support.)

Impact / Probability	Low	Moderate	High
High (p >75%)	Low	Moderate	High
Moderate (25% < p < 75%)	Low	Moderate	High
Low (p <25%)	Low	Low	Moderate

Current Steady State Risk Status

The following figure shows our current risk status. To date, we have had a total of 85 steady state risks, 43 of which have been retired; the distribution of the 42 active risks is listed in the above table. The three High (red) risks are:

- Funding availability in FY11 (this is a yearly encountered risk for ALCF)
- Lease payments for BG/Q (far term, will not become active until FY12/FY13)
- Memory costs for the BG/Q (far term, will not become active until FY12/FY13)

Impact / Probability	Low	Moderate	High
High (p >75%)	(1)	(4)	
Moderate (25% < p < 75%)	(2)	(6)	(3)
Low (p <25%)		(23)	(3)

Review of Top Risks Encountered in the Last 12 Months

The top risks encountered in the last 12 months are listed below, along with the risk owner, its projected probability and impact, a description of the actual problem that occurred, the mitigations that were used, and an evaluation of the issues and the management that was used.

1. **Major component failure**
   Risk Owner: Bill Allcock, Director of Operations

   Probability: High

   Impact: Medium

   Description: There was a substantial increase in the failure rate of the XFPs (optical transceivers) that are part of the network between the Blue Gene node cards and the data network. When the failure occurs, there are CRC error storms on the data network that causes a GPFS outage, which causes a system wide outage.

   Mitigation:

   - Leveraging good relationship with system vendor, and high visibility of the project, to reach a no-cost to Argonne solution.
   - Scripts that monitor for early symptoms, to allow pro-active replacement before GPFS and system outage.

- Working closely with the system vendor to manage the problem, as well to assist with the root cause analysis and testing of solutions.

Evaluation and Management: Doing a root cause analysis of a sudden increase in the network and job failures, the team was able to determine the failures were related to the optical transceivers on the data network (XFPs). Further analysis led the team to the realization that the majority were from one specific supplier. The team began working with the system vendor to resolve the problem. The system vendor has performed deep diagnostics on a sample of the XFPs provided by the ALCF team. All samples failed in a consistent manner, leading the team and system vendor to believe the failure is a bad batch of XFPs. The team and system vendor have developed a plan of action to address the problem. The plan is managed through a weekly call. Currently, negotiations are underway between the system vendor and the supplier, we are expecting to resolve the problem in Q4 CY2010.

2. **INCITE Users are not Provided Adequate Support by ALCF**
Risk Owner: Katherine Riley, Catalyst Manager

Probability: Medium

Impact: Medium

Description: The number of INCITE projects may be greater than anticipated, resulting in less support from ALCF to the users. An insufficient number of staff, or a largely overworked staff, due to growing demands and/or loss of personnel will decrease the overall scientific productivity.

Mitigation:

- Use management reserves to hire additional staff (permanent ALCF, temporary contractor, etc.)
- Matrix MCS personnel temporarily to ALCF
- Develop, fund and implement an aggressive Steady State/DME Staffing Plan

Evaluation and Management: A labor analysis was performed by ALCF that resulted in a determination that additional resources were needed to meet the science user community needs. The ALCF staff assists and provides a fully integrated service to the users; oversubscription of the staff has a large impact on the users' ability to produce new science. In response to the staffing needs identified, an aggressive Steady State/DME staffing plan was developed and implemented. Management reserves were used to hire several new ALCF employees to support the users in multiple areas (e.g. Catalysts, Operations, Project Support.) The implementation of the mitigation approach resulted in the science users receiving the necessary support from ALCF and science was not hindered.

3. **Tape archives Lack Performance**
Risk Owner: Bill Allcock, Director of Operations

Probability: High

Impact: Medium

Description: The performance of the tape archive system was lower than anticipated. This resulted in longer user times and ultimately decreased the users' desire to use tape archives. The issue ultimately affected HSI and HTAR clients.

Mitigation:

- Develop scripts to improve performance of tape archive
- Implement tar specification to resolve file size and path length issues
- Assist Science users with tape archive and adjust approaches as necessary to maximize performance

Evaluation and Management: The performance of the High Performance Storage System (HPSS) software has resulted in a "bottleneck" in the HSI and HTAR clients. A single host limits both, and the HTAR could not handle the large file sizes and path names that are sometimes used by the users. Bugs in GridFTP caused instabilities with HTAR when used with HPSS, this has been corrected, resulting in improved performance. New scripts have been written to further increase the performance while running the HSI agent. The file size and path length issues have been improved based on work with Mike Gliesher and the current tar specification in addition to efforts from the HPSS team.

4. **Stability Issues on the DDN9900s**
   Risk Owner: Bill Allcock, Director of Operations

Probability: Low

Impact: Medium

Description: Stability issues with the Data Direct Networks (DDN) caused by bugs in the firmware when write-back cache is enabled on replacement controllers.

Mitigation:

- Institute HSM to HPSF
- Enable mirroring in GPFS
- Implement latest firmware updates from DDN and utilize services contract

Evaluation and Management: By closely monitoring the issue and coordinating with DDN the impact of the failures was reduced. We developed site-specific procedures to avoid known problems when using the DDNs. For example, the write-back cache, known to trigger a firmware bug, is enabled by default on replacement controllers. The team worked with DDN to develop a procedure to ensure the controller was properly configured prior to connecting to the system.

This issue was eventually corrected with a firmware update and the stability has improved. Users have been encouraged to archive their data to tape regularly.

5. **Network Stability Problems**
Risk Owner: Loren Wilson, HPC Network Administrator Sr.

Probability: Low

Impact: Medium

Description: Network stability problems are created through component/hardware issues or firmware bugs in new releases.

Mitigation:

- Advanced scripts were developed for hardware components
- Implementation of new firmware process was developed
- Contract with vendor for parts and/or services

Evaluation and Management: The issues associated the network stability have been resolved primarily through the full Myricom transceiver replacement described in the *Baseline Area 2* section. These transceivers caused service interruptions that led to instabilities and failures throughout the network. We now monitor the network to identify potential issues early and implement advanced scripts to prevent hardware instabilities. Firmware bugs in new releases have the potential to create new stability problems. As such, ALCF has taken the approach of delaying the implementation of software releases unless the features or fixes are explicitly needed. Additionally, substantial maintenance time is allocated to allow for testing and if required, rollback to a previous version of firmware prior to the implementation.

6. **An extended, or full year, Continuing Resolution occurs in FY10, and results in delayed or reduced funding.**
Risk Owner: Pete Beckman, ALCF Division Director

Probability: Medium

Impact: High

Description: Funding is determined politically through Congressional action and is implemented through the DOE. Funding changes and continuing resolutions have occurred historically.

Mitigation:

- Renegotiate the leases to reduce/delay scheduled payments
- Freeze the current hiring plan and release non-essential contracted labor
- Reduce services to the user community.

- Use all management reserves and carry over from FY10.

Evaluation and Management: A continuing resolution did occur in FY10. As a result the ALCF delayed hiring of new personnel and non-contractual payments where possible. Management reserves were used to cover the standing costs of the project. The continuing resolution did not extend into a time frame that required renegotiating leases.

7. **Interruptions to the facility cooling and power**
   Risk Owner: Pete Beckman, ALCF Division Director

   Probability: Medium

   Impact: Medium

   Description: Electrical power and water-cooling for the air handlers used in the ISSF equipment room is not available or below capacity for ALCF operational needs.

   Mitigation:

   - Develop notification procedures for FMS to provide warning of scheduled maintenance or service outages.
   - Communicate schedule with users and assist with planning and scheduling around potential outages.
   - Argonne FMS provides rental chillers for 6-12 months as back up capacity
   - Interconnect site chilled water piping to enable chiller plant redundancy
   - Replace antiquated chiller and control systems

   Evaluation and Management: Two incidences were experienced this year related to facilities not being available to support ALCF operations. One involved cooling tower maintenance that created an issue with the chiller plant not being able to create enough chilled water to cool the ISSF. The ALCF monitoring system automatically sent out temperature warnings to ALCF operation staff and the BG/P was shut down to prevent damage. Some minor components had to be replaced as a result of this unscheduled outage. A second incident was a result of switching electrical power over between substations. During this switch, power was lost to the chilled water plant for 5 minutes. It was partially brought back on, but did not have full power for several hours. The result was not being able to provided chilled water to the ISSF. The ALCF monitoring system automatically sent out alerts to the ALCF staff; the staff were able to put Intrepid into a standby mode, thus preventing any component damage. When full power was restored at the chilled water plant, normal operations at ISSF commenced.

   As a result of these incidents, a formal notification and planning approach was developed between FMS and ALCF upper management to ensure ALCF would be aware, and able to plan for any activities that had potential to adversely affect operations. Additionally, FMS has agreed to provide temporary rental chillers to the ISSF during the next 6-12 months of modernization/construction in the area to provide back up capacity. An IGPP project is being developed to interconnect the laboratories chilled water piping system so that the two chilled water plants that exist in the 300 Area of Argonne and the new 200 Area Central Chilled Water Plant can provide redundancy in the event one of the facilities goes down in the future. In the next one to two years, the antiquated chiller will be replaced and the control systems updated as part of a separate Major Repair funded project.

Projected Major Operating Risks for FY11

The top operating risks projected for FY11 are listed below, along with the risk owner, its estimated probability and impact, a description of the projected risk and potential mitigations that could be used.

1. **An extended, or full year, Continuing Resolution occurs in FY11, and results in delayed or reduced funding.**
   Risk Owner: Pete Beckman, ALCF Division Director

   Probability: Medium

   Impact: High

   Description: Funding is determined politically through Congressional action and is implemented through the DOE. Funding changes and continuing resolutions have occurred historically. Due to the high impact and medium probability this risk has been closely monitored. The current BG/P lease payments are tied to the FY10-11 funding and a substantial $33M balloon payment is scheduled in February 2011.

   Mitigation:

   - Renegotiate the leases to reduce/delay scheduled payments
   - Freeze the current hiring plan and release non-essential contracted labor
   - Reduce services to the user community.
   - Use all management reserves and carry over from FY10.

2. **Electric costs could increase beyond the budget.**
   Risk Owner: Ira Goldberg, Senior Financial Structures Analyst

   Probability: Low

   Impact: High

   Description: The electric rates may increase beyond the anticipated financial plans or the usage increases beyond that projected resulting in higher electrical costs to the project. Currently, the electrical rates were locked at a favorable rate due to economic factors. That combined with the energy efficiency initiatives that were part of the BG/P costs have not exceeded budgeted values.

Mitigation:

- Relock electrical rates at the current favorable rate.
- Use management reserves for increased electrical costs
- Incorporate additional energy saving technologies/approaches inside the computer room

3. **Oversubscription of staff.**
Risk Owner: Pete Beckman, ALCF Division Director

Probability: Medium

Impact: Medium

Description: The "bottoms up" labor analysis performed by ALCF may not yield appropriate resources to meet the science user community needs. An insufficient number of staff, or a largely overworked staff due to growing demands and/or loss of personnel, will decrease the scientific productivity. ALCF staff assists and provides a fully integrated service to the users; oversubscription of staff has a large impact on the ALCF users' ability to produce new science.

Mitigation:

- Develop, fund and implement aggressive Steady-State/DME Staffing Plan
- Matrix underutilized staff from other areas within CELS
- Reduce effort on non-priority tasks
- Use management reserves to hire additional staff (permanent ALCF, temporary contractor, etc.)

4. **INCITE users are not provided adequate support by ALCF.**

Risk Owner: Katherine Riley, Catalyst Manager

Probability: Medium

Impact: Medium

Description: This risk is related to the previous oversubscription of staff risk. The number of INCITE and ALCC projects may be greater than anticipated, resulting in less extended support from ALCF to users.

Mitigation:

- Use management reserves to hire additional staff (permanent ALCF, temporary contractor, etc.)
- Matrix MCS personnel temporary to ALCF

5. **Interruptions to the facility that provided cooling and power.**
   Risk Owner: Pete Beckman, ALCF Division Director

   Probability: Low

   Impact: High

   Description: Electrical power and water-cooling for the air handlers for the computer room floor is not available or below capacity for ALCF operational needs.

   Mitigation: The probability for this event has been decreased due to the expectation of a major cooling tower enhancement project reaching completion in early FY11. This will provide n+1 redundancy to the chilled water plant's cooling tower profile. Additionally, the communications plan and POCs for ALCF will have been in effect for several months, which should decrease the potential for an unknown maintenance event from occurring without ALCF's knowledge and ability to plan.

6. **Major component failure**
   Risk Owner: Bill Allcock, Director of Operations

   Probability: High

   Impact: Medium

   Description: A substantial faction of a non-redundant, system critical component may experience a greater than anticipated failure rate.

   Mitigation:

   - Leveraging good relationship with system vendor, and high visibility of the project, to reach a no-cost to Argonne solution.
   - Working closely with the system vendor to manage the problem, as well to assist with the root cause analysis and testing of solutions.
   - Utilize full service contracts to replace components as they fail.
   - Tracking of all parts failures over time to provide statistical proof of systemic failures may be used to lead vendor to pro-actively replace all of the failing components in the system.

Baseline Area 7: Cyber Security

Charge Question 7: *Does the facility have a valid cyber security plan and authority to operate?*

FY '10 Operational Assessment Guidance: *Each facility should provide information on its approved Cyber Security Program Plan and approved Cyber Security Certification and Accreditation, in accordance with DOE Orders and Federal Regulations.*

Metric 7.1 – Approved CSPP and C&A

The ALCF will maintain an approved Cyber Security Program Plan and an approved Cyber Security Certification and Accreditation.

Metric 7.1 – Results

The ALCF operates its cyber security program in conjunction with Argonne's Cyber Security Program Office. The ALCF leverages the central cyber security services provided by the Laboratory and is currently accredited at a FIPS 199 Moderate level under the Laboratory's Certification and Accreditation envelope. Accreditation was recently extended through January 31, 2012. Copies of the 2008 *Authority to Operate* and the recent *Extension of Authority to Operate* memos are attached.

Metric 7.1 – Additional Details

There were no reportable incidents at the ALCF Facility for the period of August 1, 2009 thru July 31, 2010.

Department of Energy
Argonne Site Office
9800 South Cass Avenue
Argonne, Illinois 60439

2 8 MAR 2008

Dr. Robert Rosner
Director, Argonne National Laboratory
President, UChicago Argonne, LLC
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Rosner:

SUBJECT: ACCREDITATION OF ARGONNE NATIONAL LABORATORY (ANL) CYBER ENCLAVES – AUTHORITY TO OPERATE (ATO)

References:
1. Expansion of Argonne National Laboratory (ANL) Cyber Enclaves – Authority to Operate (ATO) to Include the Advanced Leadership Computing Facility (ALCF), C. Catlett to R. Lutha, February 25, 2008
2. Accreditation of Argonne National Laboratory (ANL) Cyber Enclaves – Authority to Operate (ATO), January 23, 2008

The submitted certification documentation (Ref. 1) has been reviewed for the General Computing Enclave and its major applications:

- Accelerator Control Systems (APS, IPNS)
- Business Systems
- Sensitive Information
- Advanced Leadership Computing Facility

This enclave and its associated major applications include all IT investments at ANL. The attached list designates the systems included in each enclave that have been reported under the requirements of the Federal Information Systems Management Act (FISMA).

Based on
- the testing performed by the July 2006 DOE Office of Science Site Assistance Visit (SAV),
- the FY2007 independent Security Test and Evaluation (ST&E) performed by Grant Thornton,
- the close-out of the DOE-IG's Certification and Accreditation finding, and
- the review of the Advanced Leadership Computing Facility performed by the DOE Chicago Integrated Support Center, Safeguard and Security Division.

A component of the Office of Science

Dr. Robert Rosner - 2 - 28 MAR 2008

I am approving Argonne National Laboratory's General Computing Enclave and its major applications to operate at a FIPS-199 level of moderate through January 23, 2011 in accordance with NIST Special Publication 800-37, "Guidance for the Security Certification and Accreditation of Federal Information Systems". This Approval to Operate (ATO) supersedes the Approval to Operate issued on January 23, 2008. Further this ATO is contingent on the Laboratory satisfactorily addressing the weakness in contingency planning identified by the DOE Chicago Integrated Support Center review of the Argonne Leadership Computing Facility (ALCF).

This approval remains in effect as long as:

1. The security status of each enclave (to include all changes and events that effect security) is reported on a quarterly basis and there are no changes that result in significant increased risk to enclave, laboratory, or DOE assets; and
2. Vulnerabilities confirmed through laboratory monitoring processes are appropriately addressed and do not result in increased risk to enclave, Laboratory, or DOE assets.

The Laboratory should retain a copy of this approval letter with all supporting security certification documentation.

If I can be of any assistance, please contact me or have your staff contact John Kasprowicz at (630) 252-2621 or Francis Healy of DOE Integrated Service Center at (630) 252-2827.

Sincerely,

Ronald J. Lutha
Site Manager

Enclosure:
As Stated

cc: C. Catlett, ANL-CIS/222, w/encl.
C. Woods, SC-31.3/GTN, w/encl.

OFFICE OF THE DIRECTOR
ARGONNE NATIONAL LABORATORY

2010 JUN 2 AM 9:18

Department of Energy
Argonne Site Office
9800 South Cass Avenue
Argonne, Illinois 60439

Action to	CEC
Copy to	BA, JLC, MS
Due Date	—
Return Copy	—
OTD File #	

28 MAY 2010

Dr. Eric D. Isaacs
Director, Argonne National Laboratory
President, UChicago Argonne, LLC
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Isaacs:

SUBJECT: EXTENSION OF ARGONNE NATIONAL LABORATORY (ANL) UNCLASSIFIED CYBER AUTHORITY TO OPERATE (ATO)

Reference: Letter, Skwarek to Livengood, dated May 20, 2010, Subject: Request for Extension of Argonne National Laboratory Unclassified Cyber Authority to Operate (ATO)

0610-002

I am approving ANL's request that its Unclassified Cyber Security ATO be extended from January 11, 2011, to January 31, 2012. This extension applies to ANL's General Computing Enclave and the following major applications; Accelerator Control Systems (APS, IPNS), Business Systems, Sensitive Information, and the Argonne Leadership Computing Facility.

This approval remains in effect as long as:

1. The security status of each enclave (to include all changes and events that effect security) is reported on a quarterly basis and there are no changes that result in significant increased risk to enclave, Laboratory, or Department of Energy assets; and
2. Vulnerabilities confirmed through Laboratory monitoring processes are appropriately addressed and do not result in increased risk to enclave, Laboratory, or Department of Energy assets.

The Laboratory should retain a copy of this approval letter with all supporting security certification documentation. If you have questions, please contact Nancy Oetter, of my staff, at 630-252-2325 or Francis Healy of the Department of Energy-Office of Science-Chicago Office-Safeguards and Security Services at 630-252-2827.

Sincerely,

Joanna M. Livengood

Dr. Joanna M. Livengood
Acting Manager

cc: C. Catlett, ANL
M. Skwarek, ANL

A component of the Office of Science

Argonne Leadership Computing Facility
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

www.anl.gov

U.S. DEPARTMENT OF ENERGY

CONTACT Argonne Leadership Computing Facility | www.alcf.anl.gov | (877) 737-8615