LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-659251

CY14 Livermore Computing Focus Areas

T. M. Connell, K C. Cupps, T. E. D'Hooge, T. J. Fahey, D. M. Fox, W. S. Futral, M. R. Gary, R. J. Goldstone, P. G. Hamilton, T. M. Heer, J. W. Long, R. J. Mark, C. J. Morrone, J. D. Shoopman, J. A. Slavec, D. W. Smith, R. R. Springmeyer, D. M. Stearman, P. C. Watson

August 25, 2014

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Contents

Purpose ***1***

Executive Summary ***1***

Process ***1***

CY14 Drivers ***1***

LC CY14 Focus Areas ***2***

Improving Efficiency and Service Quality **3**
File System Improvement 3
Virtualization Environment Improvements 4
Lorenz User Portal Improvements 4
Hopper Data Management Improvements 5
Additional Significant Efficiency and Service Quality Improvements 5

Monitoring, Measurement and Automation **7**
Center-wide Monitoring Tool 7
Log Processing using Splunk 8
Lustre Monitoring 8
Archive quotas 8

Next Generation Data and Computing Environments **9**
Next Generation HPC Platform Preparation 9
Exascale Platform Preparation 10
Data Science Initiatives 10

Hiring, Training and Mentoring New Staff **11**
Hiring and New Employee Orientation 11
Training and Mentoring New Staff 11

Purpose

This document identifies and describes four cross-cutting areas demanding the particular attention and focus of Livermore Computing (LC) during Calendar Year 2014 (CY14). Identifying these focus areas allows LC management and staff to make sound decisions about the application of effort and expenditures throughout the year. ***This document is expressly not intended to be a comprehensive description of LC efforts***. Rather, it shines a light on high priority issues involving multiple LC disciplines.

Executive Summary

The LC team undertook a survey of primary Center drivers for CY14. Identified key drivers included enhancing user experience and productivity, extreme-scale platform preparation, best practice process improvement, and business expansion. The team organized critical supporting efforts into four cross-cutting focus areas:

- Improving Service Quality and Efficiency
- Monitoring, Measurement and Automation
- Next Generation Data and Computing Environments
- Hiring and Training of New Staff

In each area the team detailed high level challenges and identified discrete actions to address these issues during the year. Identifying the Center's primary drivers, issues, and plans is intended to serve as a lens focusing LC personnel, resources, and priorities throughout the year.

Process

This document is the result of a two month series of weekly meetings involving representatives from throughout LC. These disciplines included: networking, visualization, security, customer services, development environment, archive storage, systems software, operations, the Advanced Technology Office (ATO), and file systems. The meetings explored each discipline area in turn, focusing on issues and efforts as they related to primary LC drivers. Each area presented to the assembled team allowing for an extensive interchange of information amongst disciplines. The team then identified critical areas requiring focused attention during CY14.

CY14 Drivers

The LC team identified the following primary drivers for Calendar Year 14 (CY14):

User Experience and Productivity: Enhancing the productivity of LC users and providing a positive and productive experience remains a cornerstone of LC's mission. This driver includes providing high availability and reliability in support of LC customers.

Extreme-scale Machine Preparation: The next five years will witness the procurement and deployment of two new machine architectures within the LC, namely the Commodity Technology System (CTS) 1 machines and the Sierra Advanced Technology System (ATS). The Sierra system will strain or break many of today's HPC software products and technologies and will require significant changes to the solutions that LC fields today. Contract selection of both of these machine architectures will take place in CY14. Making these machines productive for LC customers requires significant upfront investment and preparation during CY14.

Process Improvement and Automation: The number and variety of environments and systems supported by LC has grown markedly in the last four years (Collaboration Zone, Restricted Zone, Infrastructure Zone, iHPC, SNSI, Global Security, Green Collaboration Environment, etc.). During this same time both staffing and funding has dropped significantly. Every LC discipline will need to identify ways to improve, streamline and automate internal processes in order to accommodate existing environments and to allow growth into new mission areas.

Data-centric Computing Paradigms: The location and use of Non-Volatile RAM (NVRAM) in upcoming compute architectures and the prospect of mission expansion in the area of Big Data in support of science shine a light on the need for the LC to grow its knowledge-base in this area. The arrival of the Catalyst machine provides a platform for LC staff, CASC researchers and LC customers to delve into data management and operational models possible on these machines. LC is no stranger to massive amounts of data and this expertise can be leveraged in support of related, but dissimilar, data science areas.

Center-wide Resource Management: Today's HPC jobs are scheduled on discrete machines without attention to off-machine network bandwidth and data requirements. A job's impact on these shared assets is ignored as is the possible impact of contention for shared resources on the job being scheduled. This leads to poor job performance and inefficient and contentious use of shared resources in the center. The lack of Quality of Service (QOS) guarantees for all job-required resources regularly leads to poor job performance. The LC needs to move to a model where the entirety of a center is scheduled holistically to maximize user productivity and resource utilization.

Best Practice Measurement and Synthesis Effort: The LC suffers today from not being able to track improvements or degradations after deployments or changes in our environments. This impacts decision-making and may lead to poor investments in hardware, software, and manpower.

Business Expansion (Work For Others - HPC-IC, CRADA, SNSI): Growing Center business is key to the LC's future. We must find ways of leveraging the significant investments that have been made in the SNSI environment, the HPC-IC, and the significant preparation for a CRADA in file system, STAT and Spindle software product areas.

LC CY14 Focus Areas

Four cross-cutting (multidisciplinary) focus areas were called out for the primary attention of LC staff during the coming year. These areas were:

- Improving Service Quality and Efficiency

- Monitoring, Measurement and Automation
- Next Generation Data and Computing Environments
- Hiring and Training of New Staff

Below, each focus area is briefly introduced, followed by a description of the area's key contributing issues and finally, a description of LC's planned approach to address the issues during the year.

Improving Efficiency and Service Quality

Supporting high levels of user productivity is the core of LC's mission. Providing consistent, quality services across a growing number of environments and mission spaces requires development of improved efficiencies and the elimination of wasted effort. A number of areas for improvement have been identified that will enhance our customers' and our own experience and productivity while providing a high return on investment.

File System Improvement

Issue:
The largest challenge, and probably the most well-known pain point for LC users, is the availability, reliability and performance of the LC's Lustre parallel file systems. For some years now the LC has struggled to provide reliable globally mounted parallel file systems in each of its environments. Despite significant improvements and attempts to keep up with the scaling of LC compute resources, LC customers continue to suffer through poor file system availability and insufficient metadata performance. These challenges have directly impacted the ability of LC customers to do work.

Approach:
Work will first focus on improving reliability and availability of individual file systems. While globally mounting file systems directly benefits users by allowing access to data wherever they work, and benefits the Center by minimizing the number of copies of data a user requires, it comes at the price of reliability. Experts in the LC and in the Lustre community agree that complexity of global mounts is at the core of many of the availability issues we have experienced. The LC therefore intends to perform an experiment with its classified file systems, mounting file systems to subsets of machines, and in some cases single machines, in order to try to enhance file system stability. At the same time, the LC will make the significant hardware investments necessary to migrate all Lustre file systems to a single hardware and software base thus minimizing the number of releases and hardware architectures managed and supported by LC staff. The LC will undertake both of these efforts and, after 6-8 months, evaluate the experiment as part of an analysis of how and where to mount file systems in our centers.

File system stability will free LC Lustre developers to focus on the second pain point, Lustre metadata performance. While there are many areas to be worked, the two that represent the highest return on investment are 1) The implementation of the ZFS Intent Log (ZIL) and 2) Large dnode support. The ZIL effort will implement a persistent cache to log and protect metadata changes and to buffer them before writing them to disk. The ZIL work itself should not provide a large performance gain, but will replace a workaround that is currently in place that can lead to compromised data integrity. ZIL work will be started during the year with the hope of production deployment in CY15. Large dnode support will expand the size of an internal data structure

obviating the need for multiple accesses for commonly accessed data. Large dnode coding will be scoped and planned during the year with development scheduled for CY15.

Virtualization Environment Improvements

Issue:
Over the last year the LC successfully consolidated many of its unclassified infrastructure servers onto a single hardware (Cisco UCS 5108) platform via server virtualization. This work drastically reduced the number of servers and physical infrastructure (power, networking, cabling) required. The success of server virtualization led to many critical infrastructure services residing on a single platform. While failures within a single UCS system are well handled, the failure of a full UCS or of the physical infrastructure supporting the hardware has not been adequately planned for. Our initial deployment also taught us a number of lessons and revealed room for improvement in the creation and deployment of new Virtual Machines (VM) environments.

Approach:
The highest priority virtualization environment effort will involve enhancing VM Disaster Recovery (DR) capabilities. LC administrators will, as part of the move of equipment into the B453 USR, deploy a newly architected UCS configuration on a USR UCS platform. VMs will be migrated to this new hardware as appropriate. A secondary unclassified test and development system will continue to reside in a geographically separate building and will be used in the event of a complete system failure of the primary unclassified UCS environment. As part of this DR capability, VM images will be replicated to multiple locations as well as backed up to our tape backup infrastructure. Recovery to the secondary system will occur if the primary UCS system is so badly damaged that recovery cannot be achieved in a timely manner. Only after this work is completed will a full implementation of the SCF virtualization environment and its accompanying DR configuration be completed.

Using lessons learned in our initial virtualization effort, automated tools will be developed in order to streamline VM creation and deployment. A VM "cluster" approach will continue to be used as a management method making our VM TOSS updates match those of production LC clusters. An investigation will be undertaken into the viability of running Red Hat Enterprise Virtualization (RHEV) atop our production TOSS stack. Improvements in these areas may lead to the LC offering VMs in support of LC customer requirements.

Lorenz User Portal Improvements

Issue:
The Lorenz effort is focused on making HPC easier for users through the use of modern web technologies. The MyLC dashboard is a part of this effort, and is a customizable view into the Center from a user's perspective. Users turn to the MyLC dashboard to get the latest information about the Center. Unfortunately, some useful historical information is not currently being displayed by Lorenz, the tool is limited by the inability for it to switch user identities (for use by LC staff and customers with multiple accounts), and users do not have the ability to register for notification of events.

Approach:
The Lorenz team will coordinate and develop a repository for tracking events that change the configurations of the clusters, and will deploy a portlet that displays this information in a calendar format. This will help users and staff better understand the connection between changes in code

behavior and changes that occurred on the clusters. The team will also develop and deploy an 'su' capability throughout Lorenz which will allow LC staff to "become" a user and observe information as the user can see it. This will also allow an end-user to switch between multiple user accounts within a single browser session.

To improve the timeliness of information made available to users, the Lorenz team will add a subscription and notification mechanism, allowing users to subscribe to events (such as job completions or hitting file system quotas) and receive notifications in a number of different ways. This will tie into Flux and system monitoring tool efforts as appropriate.

Hopper Data Management Improvements

Issue:
Hopper is a customer-driven data management tool that performs file transfers, searches, directory synchronization, disk usage exams, and other tasks. A user is able to perform parallel actions, including data transfers, through a simple drag-and-drop interface. While easy to use, Hopper cannot meet our users' requirement for displaying disk usage information from Lustre quickly; this makes file system cleanup more difficult for users, putting stress on heavily used Lustre file systems and increasing the need to take more drastic actions such as purging. Hopper's ease of use can also lead to inefficient and possibly disruptive actions being performed, such as the default choice of a transfer tool that is inappropriate for the number or size of the files being moved.

Approach:
To increase user and resource efficiency, the Hopper team will develop and deploy a file transfer wizard. This wizard will take into account resource usage efficiency, source and destination location, and the number of involved components in order to make appropriate choices for data transfers.

To reduce the time it takes to get detailed disk usage information, the Hopper team will attempt to leverage Lustre's Robinhood Policy Engine. When users are asked to clean up their LC directories, many turn to Hopper for its graphical disk usage view to see how their data are distributed. For large Lustre directories, though, the scan can take hours or even overnight to complete. While Lustre metadata improvements (mentioned above) will help, the Hopper team would prefer to unburden Lustre by leveraging Robinhood for analysis functions involving Lustre metadata. Unfortunately, efforts to deploy Robinhood have revealed critical weaknesses in its implementation, making it unclear if it can function in the LC. If Robinhood proves viable, the Hopper team will use it to speed up metadata queries while offloading requests from Lustre metadata servers.

Additional Significant Efficiency and Service Quality Improvements

Issue:
A number of other significant issues exist which limit the LC's ability to deliver high quality service to its customers. These include:

- *Long job startups on BGQ machines* - Job startups on IBM BGQ are delayed due to slow library-loading performance for dynamically-linked HPC applications.

- *The LC Hotline is not able to perform adequate triage* - Hotline consultants are not able to leverage many system administration tools for initial problem triage/investigation.

- *Time to first archive byte* – With archive disk cache capacity having been stagnant for several years, archive users often experience delays accessing files because slower tape reads are required. In addition, the aging disk subsystems deployed in the archives are not as reliable as newer generation systems, and require high levels of Operations effort to maintain.

- *Debugging at scale issues* - Bugs that occur at large and extreme scales are becoming more common as systems increase in size and complexity. The STAT tool has proven effective for debugging hang cases at even the largest scales of Sequoia. However, tools do not exist to debug other classes of bugs, such as crashes, numerical instability, or faulty execution that occur at such scales.

- *Poorly understood network bandwidth requirements* - Network monitoring has shown that some Center Ethernet and InfiniBand (IB) links/ports are highly underutilized and could be repurposed.

- *Suboptimal communication on Livermore Computing Division (LCD) cross-cutting issues* - Despite now being a single division, it is often the case that LCD staff members are unaware of efforts and projects being performed elsewhere within their division. This limits the ability of personnel to leverage existing code, techniques, and expertise and can lead to redundant work being performed.

- *Facilities challenges* - B115 and B451 are aging facilities where outages and ongoing maintenance requirements are leading to escalating amounts of downtimes to platforms and Center infrastructure. In addition the USR computer room in B453 has serious problems with its floor making it extremely difficult to efficiently work on power and network issues within the room.

Approach:
The issues will be addressed as follows:

- *Long job startups on BGQ machines* - LLNL's Spindle provides a mechanism for efficient and scalable loading of shared libraries, executables, and python files from shared file systems. Spindle has demonstrated excellent scalability on LC machines. Spindle will be ported to IBM BGQ platforms in order to improve startup times on Sequoia and Vulcan machines.

- *The LC Hotline is not able to perform adequate triage* - Hotline consultants will begin training in select areas leveraging LC administrator and developer personnel. Once trained in administration tool use they will be given permissions to use these tools to enhance initial triage and troubleshooting. File system administration tools will be the initial target area. A Lorenz portlet or utility will serve as the user interface and will increase ease of use.

- *Time to first archive byte* – New generation archive disk cache subsystems will be procured and deployed, effectively doubling the size of the current cache and improving archive user experience by providing faster access to data for a longer period of time. In addition, this

will improve efficiency for the LC in terms of maintenance, as the archive will be using the same, highly reliable disk hardware subsystem used by LC file systems.

- *Debugging at scale issues* - The LC has developed a prototype extension to STAT, known as the DySectAPI, that enables STAT to perform more advanced debugging operations and to capture a wider range of bugs while maintaining its proven scalability. The LC will continue to evaluate, design, and implement the Dysect API and move it towards production.

- *Poorly understood network bandwidth requirements* - LC Ethernet and IB experts will survey bandwidth utilization of network switches. Undersubscribed switches and ports will be repurposed as a networking right-sizing effort within the Center.

- *Suboptimal communication on Livermore Computing Division (LCD) cross-cutting issues* - The LC will create an online inventory identifying significant software development projects and products. The LC will also execute a pilot project mimicking StackOverflow which will allow LC personnel to ask and answer questions surrounding programming and scripting challenges. The LC will continue to feature short presentations following morning Tactical meetings detailing selected internal development efforts or challenges.

- *Facilities challenges* - The LC will target moving all but a few essential services out of B115 during CY2014 and will relinquish responsibility for the building to Global Security. Equipment remaining in B115 will remain under a tenant agreement with GS. In parallel with this, initial design and network planning will be completed for a new modular computer facility (B654) which will eventually take pressure off of B451 facilities infrastructure. All equipment in the B453 USR computer room will be moved out, the floor fixed and remodeled, and the room repopulated in an efficient manner with larger cooling aisles, overhead cable runs, and expanded wide and deep racks where required.

Monitoring, Measurement and Automation

The LC environment is complex and is comprised of over 169 different compute and infrastructure systems. Monitoring the health and performance of these systems and their ability to work together is extremely challenging. Providing proper automated response to issues detected via monitoring is even more difficult. In the face of increasing complexity and decreasing or flat staffing levels, the LC finds itself searching for new solutions capable of rapidly detecting failures and performance regressions and taking appropriate actions in an automated fashion.

Center-wide Monitoring Tool

Issue:

The LC utilizes a diverse set of monitoring tools to track the health and status of its HPC systems and infrastructure. Skummee is the "single pane of glass" which presents the best overall picture of Center health. Skummee uses SNMP and other methods to monitor center resources in a scalable manner. Skummee provides threshold status checking and reporting, and historical trending and analysis. While Skummee works well, it is locally developed and supported and is only partially integrated into the myriad of monitoring, logging, and reporting tools in the Center. The LC requires improved monitoring coverage enabling faster and more complete problem discovery and response.

Approach:

LC will sponsor a cross-organization team with the goal of surveying select monitoring solutions that represent best practice monitoring techniques for HPC centers. This team will participate with Tri-Lab colleagues in a "deep dive" on monitoring requirements and solutions. As part of its charter the team will study the Splunk, Zenoss and Lighweight Distributed Metric Service (LDMS) monitoring tools which are used by Los Alamos National and Sandia National Laboratories. The team will identify and recommend a path forward for LC center-wide monitoring.

Log Processing using Splunk

Issue:
The volume of unstructured log data generated by the LC's HPC systems is enormous. As processor and system counts grow, providing post- and real-time log processing tools becomes even more critical. Over the last year the LC has begun to heavily leverage the Splunk log processing product. Splunk is an extremely powerful machine data analysis tool capable of capturing log and other information, correlating it, and aggregating the data into a repository. While these efforts have been successful, they fall far short of leveraging Splunk to its full potential of streamlining problem discovery and providing real-time monitoring alerts.

Approach:
Building upon work done in 2013, the LC will expand its tool chest of Splunk queries and automated alerts and actions. As new operational issues are encountered where automated flagging or responses are appropriate (e.g.; InfiniBand issues), administrators will enhance and adapt Splunk-driven monitoring and response capabilities. System administrators across the LC will enhance their knowledge of the Splunk tool and will be encouraged to grow the LC's toolset of automated responses and alerts.

Lustre Monitoring

Issue:
As mentioned above, I/O issues are one of the highest pain points for LC customers. The LC has had many issues surrounding Lustre performance, availability, and reliability. It is often the case that Lustre administrators are unaware of issues until users complain. Once administrators are aware that an issue exists, it is often very difficult for to determine if the source of the problem is load related, due to the actions of a single "bad actor" with harmful I/O patterns, an issue with hardware, or an issue with Lustre software.

Approach:
File system problems are often related to or triggered by excessive load, oftentimes a result of a single user's harmful I/O patterns. Identifying abusive I/O patterns is extremely difficult. The LC would like to simplify the identification of the source of harmful I/O so that it can engage involved user or application teams in high Return On Investment (ROI) I/O consulting activities. In the LC the Lustre Monitoring Tool (LMT) is currently only leveraged to display Lustre server and router activity. LMT contains valuable historical information that is actively stored in a back-end database. This information will be mined and combined with information from various sources to aid in harmful I/O detection. Once identified, high ROI consulting will be undertaken leveraging the DARSHAN I/O performance analysis tool (where possible) in order to characterize and improve individual application I/O patterns.

Archive quotas

Issue:

HPSS archives currently use the locally developed Annual quota (Aquota) tool to inform users and resource managers of excessive space utilization in LC archives. A small minority of archive users do not pay attention to advisory quotas or to repeated warnings from resource managers. This leads to the overuse of archive resources.

Approach:
Archive developers and administrators will launch a development project to enhance the Aquota tool. After being automatically warned a configurable number of times, users heavily surpassing their archive quotas will be automatically blocked from storing more data within HPSS until they have freed space to an established limit. The enhanced Aquota tool is planned for deployment in CY15.

Next Generation Data and Computing Environments

Extreme-scale and Exascale machine horizons and the rise of big data and data analytics disciplines necessitate that LC invest in preparation efforts including architectural research ,targeted software development, and production data analytics system administration and support. Supporting future machines and their architectures will require scaling efforts on most HPC infrastructure components. The LC's future also involves new mission support areas including data science technologies. These areas require LC to rapidly grow its expertise in non-traditional HPC spaces.

Next Generation HPC Platform Preparation

Issue:
Two major computer procurements will take place in CY2014. The LC will procure the Sierra Advanced Technology System-2 (ATS-2) machine which will be delivered in late CY2017 and Commodity Technology System-1 (CTS-1) machine(s) that will be delivered in CY2015. The procurement of these machines and the preparation of LC infrastructure to support these world-class resources will be complex and will involve every discipline within the LC.

Approach:
The LC will fully engage individuals across all discipline areas in the Sierra and CTS-1 procurements, freeing time and rearranging priorities such that necessary expertise is brought to bear on these important procurements. The Sierra machine will be competed and selected as part of the Collaboration of Oak Ridge Argonne and Livermore (CORAL). Sierra Non-Recurring Engineering (NRE) and build contracts will be authored, negotiated and signed during the year. User outreach and town halls will be conducted in support of the generation of the CTS-1 Request For Proposal (RFP). A CTS-1 usage model document will be prepared. The CTS-1 RFP will be completed and released and vendor selection will take place during the calendar year.

Infrastructure preparation for these machines will include evolutionary improvements across the Center's software and hardware portfolios. Included in these evolutionary improvements will be the development of a five year network upgrade plan and the ongoing development of HPSS version 7.5 (targeted for GA in CY15) which advances HPSS toward the goal of multiple core servers by introducing partitioned metadata.

A number of revolutionary improvements will also be pursued. Flux, an extensible next-generation job and resource management framework, will target management and scheduling of the entire computer center as a single pool of resources. Flux will be prototyped and evaluated for further development.

Looking beyond CTS-1 the possible use of ARM processors in HPC is compelling. To study the efficacy of such an approach, the LC will invest in ARM64 hardware testing and evaluation at scale during CY2014. The LC will also undertake focused efforts, testing and education surrounding the use of GPU accelerators to inform both platform architecture explorations as well as to enhance our visualization capabilities and futures. Toward this end the LC will use GPU accelerators on edge, max, and rzgpu machines for experimentation and development using OpenACC and CUDA languages.

Exascale Platform Preparation

Issue:
At this point it is unclear what HPC architectures will be dominant or even viable in the exascale timeframe (defined here as CY2022). HPC vendors themselves have yet to settle on their targeted architectures for Exascale. Understanding and narrowing the solution space early and identifying key technological hurdles will be necessary so that application and infrastructure teams can prepare for the arrival and efficient use of these platforms.

Approach:
ICCD's Advanced Technology Office (ATO) will team closely with LC discipline experts and customers in investigations and vendor engagements surrounding next generation platform technologies. Investigations will be performed on ARM, Intel Xeon Phi, IBM PPC64, GPU and Intel x86 microarchitectures. The LC will continue active involvement and participation in all five DOE FastForward1 efforts (completing June 2014) and with FastForward2 and DesignForward projects. These efforts all focus on accelerating the deployment of exascale machines.

Data Science Initiatives

Issue:
The last few years have seen a global explosion in the use of data science techniques to analyze large amounts of typically unstructured data ("Big Data" gathered from social networks, sensor networks, etc.) into valuable and actionable data products. These data discovery techniques can be applied to many diverse datasets in a wide variety of fields ranging from national security to sustainable energy. While LC has a long history dealing with lots of data, namely files generated during HPC simulations, it has only recently begun focusing this expertise and pairing it with data analytics.

Approach:
LC will study and continue to prototype ways of leveraging and growing our HPC and data expertise for use in data science frameworks. In particular LC will investigate the use of a range of file systems and architectures as underpinnings to Hadoop. The LC will gain familiarity building and managing Hadoop clusters while providing early access to small number of project teams to understand Hadoop from the user perspective. The LC will then explore opportunities to improve Hadoop performance with HPC technologies including Lustre, Infiniband, and NVRAM. This effort will include the continued development of the Magpie project, a set of scripts that supports launching and running Hadoop jobs atop unmodified HPC clusters and infrastructure. By the end of CY2014 the LC will determine if supporting Hadoop-like big data jobs on an existing HPC infrastructure is viable and what other data analytics support models should be supported by the LC.

In parallel with these efforts the LC will fully integrate and commission the Catalyst machine. Catalyst is a unique machine featuring large amounts of DRAM and NVRAM in each compute node and NVRAM centrally located for burst buffer and other applications. This machine is part of a partnership with Cray and Intel corporations. The LC will host a competition for projects wanting to run on the machine, will select winning projects, and will support these projects in their investigation of new approaches including application check-pointing, in-situ visualization, out-of-core algorithms and big data analytics. The LC will also actively pursue projects surrounding NVRAM use cases in support of extreme-scale and exascale burst buffer architectures.

Hiring, Training and Mentoring New Staff

LC staffing has decreased dramatically over the last three years. Fortunately the LC's budget is able to support substantial strategic hiring during CY14. It is vital that the LC hires well, trains well, and successfully integrates new employees in a manner that ensures prolonged productivity, retention and personal and career development.

Hiring and New Employee Orientation

Issue:
Hiring has long been a challenge for the LC and LLNL in general. Many eager and qualified candidates are lost due to slow offer turnaround times. Once hired new employees often are hindered by the time necessary to get adequate office space, desktop equipment, and new accounts.

Approach:
The LC will aggressively pursue targeted hiring opportunities. The LC will actively advertise positions externally in forums such as HPCWire's Job Bank. Accelerated Candidate Events (ACE), which target hiring decisions and offers within 24 hours, will be leveraged by LC staff. Hiring managers and their interview teams will endeavor to quickly phone screen candidates and bring them onsite at the earliest possible opportunity. Hiring decisions will preferably be made on the day of the interview and hiring managers will expedite hire packages and remain in continuous contact with candidates. In advance of a new employee arriving, hiring managers will ensure that appropriate office space is secured and desktop resources are prepared in advance. Hotline account specialists will streamline the generation of initial user accounts for new employees. Hiring managers will ensure that Q-clearance paperwork is rapidly submitted.

Training and Mentoring New Staff

Issue:
It is challenging for new employees to physically participate with their coworkers and managers on projects while awaiting their Q-clearances. Rapid integration of new employees into groups and projects is also hindered a lack of spaces and tools for collaboration. Mentoring of new employees is difficult when mentor and mentee are physically apart.

Approach:
In order to foster quick integration of new employees into their groups and projects the LC will leverage available collaboration methodologies. Project teams will use collaborative tools including the Atlassian software development toolset with a focus on providing helpful documentation for both new and existing employees. The LC will expand the number of physical collaboration spaces allowing new employees to meet with their colleagues more easily and with greater frequency.

Hiring managers will select appropriate mentors or mentor-pairs for newly hired employees. Mentors will be prepared and provided with expectations and the importance of their mentoring contributions will be communicated. Good mentorship performance will be highly valued by LC management. New employees will also be paired together whenever possible for companionship, camaraderie, and to best take advantage of educational and training opportunities.

Fostering cohesiveness with a new employee's organization is vital. Hiring managers will encourage new employees to take advantage of tours, lectures, 9AM Tactical meetings, hackathons and other opportunities to learn about the LC, LLNL, mission, and the work being performed. Forums and techniques that address LCD communication on cross-cutting issues (see above) will be made fully available to new employees.