Deploying a High Performance Computing Cluster at a Predominantly Undergraduate Institution

Dr. Evan C. Lemley, University of Central Oklahoma

Professor Lemley teaches thermo-fluid engineering and works with undergraduates to perform fluid dynamics research that is mostly focused on small scale flow problems. He is currently an Assistant Dean of Mathematics and Science and a Professor of Engineering and Physics at the University of Central Oklahoma, his home institution for more than fifteen years. Previously, Professor Lemley worked as a mechanical engineer in the power industry. His bachelor's degree is in physics from Hendrix College and his M.S.M.E. and Ph.D. were earned at the University of Arkansas.

Deploying a High Performance Computing Cluster at a Predominantly Undergraduate Institution

Abstract

This paper is focused on the deployment and use of a high performance computing HPC cluster at the University of Central Oklahoma, which is a predominantly undergraduate institution (PUI) and a Carnegie Master's Level - Larger Programs Classification. In January 2015, UCO received a National Science Foundation (NSF) Major Research Instrumentation (MRI) grant to fund an HPC cluster to enhance and extend computationally oriented research and education. By July of 2015 the *Buddy* Cluster (so named for the school mascot) was deployed and ready for early users. In September 2015 there were only a handful of users compared to January 2017 when a total of 150 users have been registered on the system. The Buddy Cluster has been used for a wide range of engineering and other applications such as bioinformatics, computational fluid dynamics, heat transfer, biomedical applications, image processing, and statistical analysis. Note the bulk of the use of Buddy is by the Department of Engineering and Physics Department at UCO. Although the topics discussed are general, the paper addresses how engineering physics programs can deploy and make use of similar clusters.

This paper will discuss the impacts to computational research and the research infrastructure at UCO. The process to acquire and deploy the *Buddy* Cluster will be detailed including the process to solicit proposals from and choose a cluster vendor. The process to get users working in the cluster environment - including internet browser-based access and use of certain cluster software and the use of more standard command line access and use - are discussed.

Introduction

There is a recognition of the national need for developing high performance computing (HPC) resources, which include human resources that will use this technology. In March 2007 the National Science Foundation (NSF) released a report entitled "Cyberinfrastructure Vision for 21st Century Discovery"[1] that addresses how high performance computing (HPC) is necessary to science and engineering disciplines to answer the most basic research questions and to solve technical problems of national need. In July of 2016 the White House released a report called the National Strategic Computing Initiative[2] which is a call to "maximize the benefits of high performance computing (HPC) research, development, and deployment." These reports and others indicate the need for increased computational power and technological training for science and engineering research and innovation and for national security. The increased use of HPC will require training more scientists and engineers to appropriately use HPC in their work and more generally in deploying and using cyberinfrastructure (CI).

Certainly all academic institution types have a role to play in HPC and CI education. This paper is focused on the deployment of a HPC cluster at a Predominantly Undergraduate Institution (PUI). At a PUI the engagement of students in research projects is mostly from the undergraduate ranks which is markedly different from research intensive university. The idea of undergraduate research (UGR), if unfamiliar, may seem like a stretch, but Kuh[3] and others[4-5] have observed the positive effect of improved persistence in college courses generally and in STEM (Science, Technology, Engineering, and Mathematics) courses in particular, for those students involved in UGR.

This paper is in a series of reports about progress of high performance computing at UCO. A previous ASEE paper discussing the use of Buddy in an engineering physics course on fluid dynamics was documented in 2016[6].

Background

The University of Central Oklahoma (UCO) is a PUI, with some Master's degrees. The College of Math and Science (CMS) at UCO houses the only STEM departments on campus: Biology Dept. (BD), Chemistry Dept. (CD), Computer Science Dept. (CSD), Engineering & Physics Dept. (EPD), and Mathematics & Statistics Dept. (MSD). CMS has seen over 54% growth in majors from 2006 to the present[7] with increases in nearly all CMS majors.

For over five years UCO has made undergraduate research (UGR) a top priority across campus based on the knowledge of the very high impact UGR has on persistence in college courses.[4] Efforts to embed UGR at UCO are led by the Provost, the Office of Research and Grants (ORG), and by college Dean's. The results of these efforts to embed UGR have led to a quickly evolving research-intensive environment, where faculty and students are supported internally and externally to perform undergraduate research. The focus and scale of the UGR effort have made UCO a nationally recognized model for implementation of UGR, by the Council on Undergraduate Research[8].

Prior to the award of the NSF-MRI grant in January of 2015, significant momentum of research projects had been achieved with nearly 20 computational projects that were in-flight across the CMS. These projects range across many disciplines including bioinformatics, computational fluid dynamics, heat transfer, biomedical applications, electromagnetics, image processing, mathematical biology, and statistical analysis. UCO's efforts to implement a UGR program had supported these projects in terms of small equipment, supplies, and student employment. The need for higher end computational modeling and simulation had become necessary in these projects. Although some access to off-campus HPC resources was available to the faculty and the students involved in these projects, that was not enough. These researchers needed more support. Although the NSF-MRI grant funded HPC equipment, just as importantly, it caused the campus environment for support and assistance for performing HPC to be transformed.

Cluster Acquisition and Deployment

The NSF-MRI grant was funded in early 2015. A set of requirements for the HPC cluster were determined by an internal committee with members including co-authors (Lemley and Chen) and other faculty from across CMS with strong interest and backgrounds in computing. A Request for Proposal (RFP) was developed and distributed to potential vendors. Vendors' responses for the RFP were rated and assessed by the internal committee, resulting in the selection of a vendor that best met the RFP: *Advanced Clustering Technologies (ACT)*. The *Buddy* cluster was delivered in mid-June of 2015 and on-site set-up and training occurred in late June 2015.

As potential guidance to other PUIs that plan to acquire similar systems the items that made up the RFP are shown below:

1. *System Specifications* (Details on the following items),
    a. Operating System (Linux as is used on almost all of these clusters),
    b. Overall System Components Summary,
    c. Compute nodes,
    d. High-memory nodes,
    e. Graphics Processing Unit nodes,
    f. Management node(s),
    g. Storage node(s) - parallel?,
    h. High-speed network for message passing,
    i. Management network,
    j. Racks and power distribution,
    k. Software
2. *Extended Warranty*: Details of what you and your institution and/or funding agency requires. Carefully consider components that are more likely to fail and if some components are warrantied by the original equipment manufacturer. What is the process for warranty service and what potential impacts could this have on your uptime.
3. *Additional Warranties or support services*: Is there a need for up-front pricing on longer term warranties or other support services,
4. *Delivery*: Requirements that delivery cost are included, the location for delivery, and any other delivery requirements,

5. *Installation/Deployment*: What are the needs for installation, what specifically will the vendor do and what specifically will your institution have prepared for installation in terms of electrical, networking, cooling. What is the target delivery and installation dates.
6. *Training and Other Services*: Your needs and requirements for training for system use and administration.
7. *References*: Obtain the contact information of 3-5 current or recent customers to verify their satisfaction with the vendor and system. Consider the requirement that some references are similar to your institution and that some customers have similar systems to the system you plan to acquire.

In our case some special stipulations were made regarding software to be installed on the cluster. This was related to the fact that it was anticipated that our cluster would 1) have users that were unfamiliar with typical command line HPC environments, and 2) have a considerable commercial software footprint. These needs drove us to request a web-based interface for using our cluster that we could customize and tune to new applications and user needs.

Co-author (Gutierrez) is currently an undergraduate student in computer science and was hired as a Student Cluster Intern (SCI) in July 2015 to assist and learn cluster system administration. All coauthors installed and tested a number of software packages from July - November of 2015. These software packages, both open source and commercial, were ones requested by UCO faculty.

Initial User Intake and Training

By September 2015 system users were given access to the system. The two modes of using the *Buddy* cluster are through secure shell (ssh) command line access and through a web-based front-end created by *ACT* called e*QUEUE*. Documentation was written and distributed to users on how to gain access and use some software on the *Buddy* cluster. By December 2015 the system had over 40 users and over 10 pieces of software had been installed and were in use by faculty and students.

Training has mostly consisted of training and setup for users in faculty offices and student labs by the coauthors. A weekly help session was arranged in Spring 2016, but it was difficult to find a common time when most users could come. So instead, user training and setup is mostly handled now on a one-on-one basis.

Many cluster systems of this type do not require very much setup of a user's machine. Attributes of *Buddy* make setup of the user machines a little more challenging. The reasons for this are described in the section entitled Results and Current Situation.

Results and Current Situation

The *Buddy* cluster has been deployed and in operation since June 2015 with faculty and student access starting in September 2015. A number of changes have been made and achievements accomplished on the *Buddy* cluster, these are listed in Table One.

Table One Changes and achievements on the *Buddy* cluster from Dec. 2015 to Jan. 2017.

	Change and/or Achievement	***Notes***
1	External Access to Buddy	The state's networking organization (OneOCII) enabled a 10 Gigabit per second connection to *Buddy* in May 2016.
2	Public Key Infrastructure (PKI)	In tandem with item 1, so that off campus access was secure, the PKI arrangement was put in place.
3	150 Users	The number of users has steadily increased to 150 users.
4	>10 Research Projects	There are over ten active research projects being performed on *Buddy* by nine faculty and their students.
5	>5 Classes Have Used *Buddy*	Access and flexibility of *Buddy* is being used in classes where computational facilities are needed. At least five different classes at UCO have been using *Buddy*.
6	5 Tours and Outreach	Tours of the data center with *Buddy* and related outreach activities have been done with five groups, including middle school students, high school students, Fulbright scholars, and political representatives.
7	2 Poster Presentations	Poster presentations have been made at one state and one national conference.
8	2 Conference Papers	One conference paper was presented at the 2016 International Mechanical Engineering Congress and Exposition and one conference paper was presented at the 2016 American Society of Engineering Education Annual Conference and Exposition.

9	1 Journal Article	One journal article was published in Colloids and Surfaces A: Physicochemical and Engineering Aspects in 2016.
10	1 Master's thesis	One master's thesis on fuel cell cooling plates was completed in May 2016.
11	Virtualized sessions for certain software that runs in other operating systems..	Certain software runs exclusively in Microsoft Windows ® and now can take advantage of *Buddy* cluster hardware.

This paper has a companion poster in which some of the research results are highlighted. In particular some of the results from classes and research are available on the poster.

The direct impacts on computational research and computationally oriented classes that have been observed are in Table 2.

Table Two Impacts on computational research and computationally oriented classes with *Buddy*.

	Enhanced Ability	***Impact***
1	Ability to Perform Companion Simulations to Experimental Research	The ability to accurately model physical experiments advances research significantly in terms of the impact the results have in the scientific field. There are several cases in which *Buddy* has enabled this effort.
2	Ability to model problems with extremely fine resolution	This allows a model to undergo an extensive "grid independence study" that is required to publish purely computational results. The trustworthiness of the data from these simulations is much higher.
3	Ability to do many simulations quickly via batch jobs using significant computational resources.	This allows a researcher to run a lot of runs where say a single parameter is changed between runs. The computational power of *Buddy* is that one can do a very large range of a parameter or work with several variable parameters and get results quickly.
4	Access to software externally at any time	In the past students have often had to be on campus to use certain specialized science and engineering software, but remote access to *Buddy* makes it possible for students and

		faculty to be productive off campus and at times that computer labs are not open. This has impacted research, but also the classes that are using *Buddy*.

The items in Tables One and Two reflect considerable effort on the part of faculty, students, and the co-authors of this paper. It should not appear as though these results were not “hard-won.” In the conclusions of this paper we list some of the issues that have arisen in this deployment and operation of the *Buddy* cluster in hopes that others can at least be aware of pitfalls.

Conclusions and Discussion

At UCO an NSF MRI grant was competed for and won for a HPC cluster, *Buddy*, to enable and enhance the research computing environment at UCO, which had not had any such facilities before. The competition for *Buddy* took several tries with NSF, which were strongly supported within UCO and externally. This is especially true of the OneOCII organization, which is the statewide community of current and future cyberinfrastructure and HPC providers. This organization seeks to enhance the CI landscape in the state and a large part of that has been through NSF and other funding.

This paper details our process including the acquisition, deployment, and initial operations of Buddy. A companion poster contains some results of research performed using the *Buddy* cluster. The process of acquiring, deploying, and operating *Buddy* has not been smooth as it may appear, but that is likely due to the significant impacts the grant and *Buddy* have had on the UCO campus. In order to assist others that may go through a similar process some pitfalls to be aware of are listed here - although these are not unique to PUI’s, the authors believe the probability of running into these issues is somewhat higher at a PUI. This list is directed at those who might seek funding for a HPC cluster or another similar CI project:

- Communication is key, you will need technical and other information from potential users. The best policy is to write documents such as support letters for them so their time spent is minimized. Since you often need data on memory use and storage, you may need to have multiple discussions with users to reasonably approximate this info.
- Understanding what your needs will be computationally and demonstrating that you know what these are and have calculated these in a reasonable way, is key to convincing reviewers you can be trusted to acquire, deploy, and operate the hardware.

- In the RFP process, you should consult the rules for how this should be done at your campus. Often you will need to bring the purchasing department (or equivalent) and/or legal counsel in very early to ensure the RFP and subsequent vendor selection and purchase are executed in a timely fashion.
- Software that is needed for applications on clusters such as *Buddy* often cannot be purchased via an equipment grant. If you are reliant on commercial software this can be very expensive. It's important to get institutional commitments to purchase this software before the grant is even submitted. A chair, dean, or provost would not like to be surprised at the additional cost just to use the hardware for what it was intended.
- Faculty at PUIs typically have higher teaching loads than more research intensive institutions. You will need to observe when they are most productive (often summer) and plan to be help them the most at this time. These faculty will be sometimes unable for long periods of time to be productive, but then they are able to work on a computational project. Try to stay up on what they are doing and when they may be ready to do work on your hardware. It's important not to pressure them to use the hardware.
- You may have users that are not as technically capable as others to use your hardware, but help them as much as possible. We considered this up front, and a web-front end for some applications was built into the RFP.

Acknowledgements

The authors would like to acknowledge the UCO CURE-STEM program for support. In addition the statewide cyberinfrastructure group, OneOCII, was critical to receiving funding and growing the HPC support infrastructure for *Buddy*. This material is based upon work supported by the National Science Foundation under Grant No. ACI #1429702.

References

1. National Science Foundation Cyberinfrastructure Council, 2007: "Cyberinfrastructure Vision for 21st Century Discovery." http://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf0728.
2. National Strategic Computing Initiative Executive Council, 2016, July, "National Strategic Computing Initiative Strategic Plan."
3. Kuh, George D. Excerpt from "High-Impact Educational Practices: What They Are, Who Has Access to Them, and Why They Matter". Association of American Colleges and Universities, 2008.
4. Springer, Leonard, Stanne, Mary E., and Samuel S. Donovan, 1999, "Effects of Small-Group Learning on Undergraduates in Science, Mathematics, Engineering," Review of Educational Research, 69(1), pp. 21-51.
5. Gregerman, Sandra R., Lerner, Jennifer S., Hippel, William von., Jonides, John., and Biren A. Nagda, 1998, "Undergraduate Student-Faculty Research Partnerships Affect Student Retention," The Review of Higher Education, 22(1), pp. 55-72.

6. Lemley, E. C., & Gillispie, A. M., & Dorety, A., & Meier, A., & Riahinezhad, S. (2016, June), Integration of High Performance Computing into Engineering Physics Education Paper presented at 2016 ASEE Annual Conference & Exposition, New Orleans, Louisiana. 10.18260/p.25426
7. UCO Office of Institutional Research, http://www.uco.edu/academic-affairs/ir.
8. John F. Barthell, Wei R. Chen, Beverly K. Endicott, Charles A. Hughes, William J.Radke, Charlotte K. Simmons, Fall 2013, "Encouraging and Sustaining a Culture of Student-Centered Research at a Predominantly Undergraduate Institution." CUR Quarterly, Council on Undergraduate Research, vol. 34, no. 1.