Lessons From 20 Continuous Years of Cray/HPC Systems

Liam Forbes, Don Bahls, Gene McGill, Dr. Gregory Newby, Oralee Nudson
Arctic Region Supercomputing Center
University of Alaska Fairbanks
Fairbanks, AK, USA
lforbes@arsc.edu

***Abstract*— The Arctic Region Supercomputing Center (ARSC) was founded in 1992/1993 with a Cray Y-MP M98 (denali) and since then has operated or owned at least one Cray system, including most recently a Cray XK6m-200 (fish). For 20 years, ARSC has shared high performance computing (HPC) experiences, users, and problems with other University HPC centers, DoD HPC centers, and DoE HPC centers. In this paper, we will document and present the user support and system administration lessons we have learned from the perspective of a smaller, regional University HPC center operating and supporting the same architectures as some of the largest systems in the world over that time. We will cover topics such as purchasing and supporting Cray systems at the far end of US transportation, remotely supporting Cray systems from 2,000 miles away, supporting users from eight different time zones, and being part of the process of growing the next generation of HPC users while they develop the next generation of HPC problems. We will also provide some commentary on what happens to a customer as an HPC vendor changes hands three times, changes operating systems three times, changes hardware architecture four or more times, and the HPC market expands and adapts beyond the scope of what a small research University or organization encompasses. Comparisons to experiences with HPC hardware and software products from other vendors will be used to illustrate some of the points.**

Keywords-user support, system administration, center operations, Cray, ARSC

I. INTRODUCTION

Founded in 1992 on a Cray Y-MP M98 (denali) that was initially operating for 9 months in Cray Research's Chippewa Falls facility, the Arctic Region Supercomputing Center (ARSC) at the University of Alaska Fairbanks (UAF) has owned or operated Cray systems for a diverse group of users and funding agencies. ARSC has participated in a variety of high performance computing (HPC) organizations, supported hundreds of users, and cooperated with other University HPC, Department of Defense (DoD) centers, and National Science Foundation (NSF) HPC centers. This year we are proud to celebrate 20 years of operation. In this paper, we hope to present various lessons learned from the perspective of a smaller, regional/academic HPC center owning, operating, and supporting the same architectures and systems as some of the largest centers in the world over that time. We believe it is important for smaller centers to be successful and continue filling a variety of roles in the broader HPC community, or HPC in general will be less successful. We hope to highlight lessons learned from remote deployment and system management, managing complex systems with limited staff, and other insights where UAF/ARSC may represent a minority in the supercomputing community. We also provide a little perspective on working with one HPC vendor for two decades as they have dealt with their own changes.

II. INTRODUCTION TO THE ARCTIC REGION SUPERCOMPUTING CENTER

ARSC is the HPC research unit of UAF. Funding for ARSC operations and acquisitions currently is a combination of direct funding, external grants, and contracts from a variety of sources, including the NSF and Lockheed Martin through the DoD High Performance Computing and Modernization Program (HPCMP). ARSC provides HPC resources for University of Alaska researchers and their partners, including HPC systems, large data storage, and high-speed networks, when individual or departmental resources are not sufficient to efficiently support their research. In addition to compute cycles, ARSC provides networking, storage, visualization, education, and planning services.

ARSC computational systems and resources are designed to meet the academic and research needs of university staff, faculty and students. Since ARSC's first compute cycles were realized in 1993, staff at ARSC have established a reputation for providing outstanding service in all aspects of HPC, massive data storage, and HPC network support. ARSC consultants, systems analysts, and specialists provide direct assistance to scientists who use HPC to conduct or advance their research. Consultants provide training and support to users by phone, email and on-site, as needed, while specialists provide expertise in areas such as code migration and optimization. Research activities and computationally intensive investigations conducted at ARSC include climate and weather modeling, ice sheet modeling, oceanic physical and ecological systems, materials science and engineering. Fig. 1 shows ARSC's allocation percentages in core research areas, demonstrating the diversity of users, projects, and activities supported.

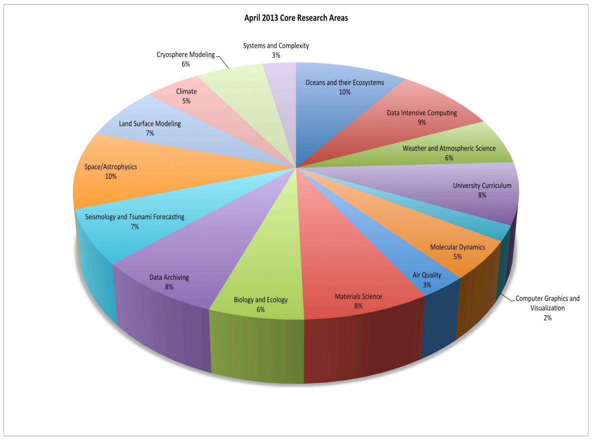

Figure 1. April 2013 ARSC Core Research Areas.

Table 1 identifies the Cray systems owned or operated by ARSC since 1992/1993. The primary host name is provided, followed by the system model and dates of operations. Additional notes about the system, such as its ranking in the Top500 [1] over its lifetime, or its location if not in Fairbanks, are also included.

The number and variety of systems, demonstrates ARSC's success at providing supercomputing resources despite being located in the Alaskan interior. It also demonstrates what we view as a crucial aspect of ARSC's success, which is being able to develop a close and lasting relationship with vendors. In fact, we often joke that we know who really wants to work with us by who is willing to come to Fairbanks in January or February. Truthfully, a good HPC vendor is just as willing to work with smaller university centers as national laboratories. Plus, ARSC stands out among North American centers thanks to geography (and it is easy to come up with fun and unique swag for Supercomputing.)

III. General Lessons for Small High Performance Computing Centers

Vendors are one of three relationship groups that are most important for a small HPC center to track and maintain. Users and funding providers are the other two. By consciously recognizing these groups, and developing different processes to maintain the relationships a center can more effectively achieve its mission and goals.

In the best case, vendors become an extension of the center staff skills and knowledge. When new projects and opportunities arise, it's then possible to draw in one or more vendors as partners rather than just hardware or software providers. Pick one or two vendors that provide products or services closest to the primary center mission and goals, and develop deeper relationships with those vendors. Make center resources and staff available to the vendor. Visit the vendor as frequently as possible and regularly invite them to visit the center. Develop person-to-person communications, not just exchanges of automated support emails. Doing this pays off when a center needs to launch large projects like major resource upgrades, or when it's time to develop the next big proposal.

Having good relationships already in place makes it easier to start new conversations rather than having to build a relationship from scratch while also trying to start something new. A vendor who is viewed as a partner rather than just a storefront can be very helpful to a small center's long-term

Hostname - Model	Dates of Operation	Notes
Denali - Cray Y-MP M98	1992-1997	largest system memory in the world, for one week #251 1993/06, #302 1993/11, #405 1994/06
Yukon - Cray T3D	1994-1996	hosted by Denali #58 1994/06, #55 1994/11, #83 1995/06, #99 1995/12, #127 1996/06, #171 1996/11, #241 1997/06, #344 1997/11
Yukon - Cray T3E	1996-2003	#70 1997/06, #62 1997/11, #67 1998/06, #74 1998/11, #44 1999/06, #56 1999/11, #78 2000/06, #107 2000/11, #131 2001/06, #199 2001/11, #383 2002/06
Chilkoot - Cray J90	1998-1999	
Chilkoot - Cray SV1ex	2000-2003	first customer to receive the SV1ex processors
Rime - Cray SX6	2002-2003	located at ARSC for testing and development by ARSC, Cray, and potential Cray customers
Klondike - Cray X1	2003-2005	#116 2003/06, #71 2003/11, #154 2004/06, #202 2004/11, #353 2005/06
Nelchina - Cray XD1	2005-2008	additional chassis located and operated at George Washington University to support FPGA testing and development
Pingo - Cray XT5	2009-2010	#109 2008/11, #205 2009/06, #290 2009/11, #435 2010/06
Chugach - Cray XE6	2010-present	located at, then transferred to, the U.S. Army ERDC ITL #83 2010/11, #100 2011/06, #142 2011/11, #236 2012/06, #230 2012/11 soon to be combined with two other DoD Cray XE6 systems
Tana - Cray XE6	2010-present	
Fish - Cray XK6	2012-present	

TABLE I. Arctic Cray systems over the years success. It also makes it easier for a vendor to work extra hard on behalf of the center. ARSC considers Cray Inc. such a partner vendor. For example, fig. 2 shows Cray engineer working at negative 22 degrees Fahrenheit outside the UAF data center to crate equipment.

Figure 2. Crating equipment at -22 degrees Fahrenheit.

This is another case where remoteness actually enhances ARSC. The ideal of Alaskan living attracts people, and organizations, with can-do personal attitudes that are reflected in work attitudes and results. This is very useful in a leading edge technology field.

Similarly, good relationships with existing users and funding providers lead to opportunities to expand both. ARSC has a number of past and present user projects that have used the center's resources for a significant number of years. These users are not always the largest users and the projects may be very focused, but by working with the same users over the years ARSC staff have developed additional insights into the users needs and they have been able to point out areas where ARSC can improve. Actively engaging with users is key. If possible, provide new projects some initial focused time and attention. A week of early focus is better than a two-days-every-month support level agreement. Review jobs for those that terminate abnormally and inform the user(s) rather than letting them find out for themselves.

A small center needs to monitor the size of its user base and project collection closely though as too many users or too many funded projects can stretch the center's capabilities, leading to un-manned projects or dissatisfied users. These failures can quickly become the death of a small center, especially in politics-prone academic environments. The best way to maintain good relationships with users is through proactivity. Proactively communicating with users about system issues that impact them (whether they know it or not), or new center developments helps maintain proper expectations and cultivates reciprocal communication as well. There's nothing more frustrating than discovering a user knew about a problem, but rather than informing the help desk or center staff decided to focus on other priorities and just ignore the issue hoping it will eventually be resolved.

The best way to maintain good relationships with funding providers is by meeting and exceeding the goals of the product they are funding. A center that can develop a reputation for delivering on the agreements made to receive funding, for responsibly handling funds, and getting more out of funding than originally expected will be able to maintain and find new funding opportunities. Being responsive to reporting requirements and informal status checks definitely helps this as well. Choose funding opportunities that have a high synergy with the center mission and goals. Often, smaller, targeted funding that helps a center develop distinction in an area is more beneficial than larger blanket funding.

Developing points of distinction allows a small center to be successful when funding or staff size is insufficient. Examples from ARSC's history include developing, writing, and maintaining a newsletter containing good technical data, being the HPCMP storage lifecycle management development and test center, and developing high resolution, geo-local weather forecasts for Alaska. All three efforts raised ARSC's profile and helped us stand out to our local user base, and among peer institutions. They attract new users, new staff, and new funding opportunities. Small centers cannot afford to be all things to all potential shareholders, so strategically identifying and developing distinguishing characteristics, activities, or resources allows for stability and growth over time. A few focused distinctions will get more attention than overall uniformity.

Challenge projects, whether they are "grand challenges" or not, are a great way to really engage staff with users and funding sources. By developing a core body of work that uses significant resources, the center has a base to build from and opportunities to do new things. Research guides how these opportunities develop, but a programmatic approach provides a framework in which staff and users can imagine new potentials.

There is more to HPC technology than the hardware on the floor. Small centers have as much opportunity to dive into code or underlying components as large centers, but they have to plan for it and make it a point of distinction since a small center cannot afford to just throw bodies at technologies. Developing a technology, or at least in-depth skills in an existing technology, can be a center's distinguishing characteristic.

Another way to develop good relationships is to get involved in communities and user groups (such as the Cray User Group) as much as possible. By allowing staff opportunity and time to get involved, they improve their own skills and help build a center's reputation. According to recent research [2], for skilled knowledge workers, these opportunities are more motivating than financial gain. Additionally, technical communities can provide additional support and research beyond a small center or vendor's capabilities. With more technologies coming from or becoming open source, an active, open development community can often provide answers or clues more quickly than a vendor's support or a center's staff. These folks are intimately involved in the technologies, use their knowledge and save your energies.

A center should find ways to engage with key players in the communities important to their resources. Take part in activities like standards development, program reviews, conferences/meetings, and external advisory panels. By putting people forward, a small center can learn more about the larger HPC market and needs. Consider staff exchanges to learn new skills and share local expertise. Do a few good things well, and others will come to you for resources and support. Be willing to get help for the other things, and people will be willing to provide resources and support.

In all these efforts, remote-ness is no longer an issue - once you are far enough away to not bump into somebody on the way to the loo, you are remote. In HPC, remote versus local is an attitude, not a distance. Tools for collaborating remotely and sustaining relationships have vastly improved in the last two decades. There is still a need to make connections in person and travel for outreach / learning / maintain profile, but travel is not required for day-to-day operations and communication.

A bigger challenge for a small center is crossing boundaries such as government vs. university vs. commercial. The most important thing is to develop relationships that last longer than one allocation cycle. A small center can be involved in multiple realms; in fact taking knowledge from one realm to another can be a point of distinction. However, it requires committing people to understanding each realm's operations and communications to be effective in both. The need to do this should grow out of service to the center's primary users or funding providers.

Underpinning the above are centers' needs to develop and look after staff. Low turnover allows for development of success. Keeping and developing the staff that look after the user relationships can mean more to a center's success than anything else.

IV. User Support Lessons

User support is absolutely crucial to user satisfaction, to system productivity, and to an HPC center's reputation and continuation. It should go without saying that all other center activities should start with a user project, issue, or idea. For service oriented organizations, focusing on resources or center meta-activities leads to low utilization and eventually shutting down the organization.

To provide optimal HPC customer service, hire people who want to be in the position of user support instead of using it as a "start here" entry-level position at an HPC center. Find people who are interested in helping others solve problems, learn technology, or perform research. Additionally, find people who are comfortable working on many un-related tasks throughout the day, who can tolerate interruptions, and have no fear of interacting with others, on the phone as well as in person.

User support staff should also look for ways to communicate with other HPC center user support personnel. Sharing ideas of new resources to check out, what works, and what does not is very helpful. Again, technology communities and user groups make ideal opportunities.

If possible, embed users in the user support processes. An advisory council or periodic process reviews that incorporate the most successful users (or at least the most vocal), are outreach opportunities for a small center. These users can become important standard bearers for the center during budgeting or project proposal cycles, especially in academic environments.

In general, large, capacity focused HPC centers have to be unfair to medium and smaller users/projects in order to provide the largest resources to the largest users. When the goal is to make a large resource run a single job or workload, smaller jobs may have to wait for long periods or can be preempted. Small, capability focused centers should concentrate on a few small to medium users/projects, and help them grow to be large users. Sometimes successful users will outgrow a small center and move on. However, successful people usually remember who made them successful and are willing to give back in some way.

This is not to say that a small HPC center cannot find opportunities for users to monopolize an entire resource. Look for and take advantage of events like acceptance testing to conduct additional science for users and try out non-standard run cases. During system installation and initial shake out, one or two "friendly users" can really help to identify configuration issues, software bugs, or hardware problems. Pick users for these events who are willing to contribute back to the center's mission or goals in some way. Then stay in contact with them on a daily basis to help get the most science out of the system before it has to start running the production workload. Most of these opportunities are temporary and eventually have to give way to normal center operations, but some may become successful in their own right and take on a life of their own.

User support staff should also be utilized to provide training. User training events are good opportunities to help users grow and put a personal touch on the center's activities. Such events must be well advertised though, and often need to creatively explain how computing applies to non-traditional areas. Focus the primary user training events around the center's distinguishing characteristics, but leave room to modify or expand the training in the future.

HPC user support is an active and changing arena, but often thankless. In an academic environment, user support career paths can be limited unless one changes jobs or organizations. Good staff often go to work for those they used to help, the vendors whose products they used to train users to use, or larger centers with more focused responsibilities, including leadership opportunities. Smaller centers provide a broader range of experiences for user support staff and should highlight the varied opportunities, as well as celebrate being the foundation for staff to grow from.

V. System Support Lessons

Users are a small center's customers. User support staff provide the center's face and front-line service for those customers. However, when a center is large enough to have separate user and system support staff, user support staff are also customers themselves. System support staff serve as the interface between the user support staff and the center's resources. Through several re-organizations and workload adjustments, ARSC has maintained that system support staff must be responsive to the needs and input of the user support staff; it's even in the system support job description. This ensures the user's experiences and needs guide most of the work performed on a daily basis.

The fundamental systems administration principles for Crays/HPCs are functionally the same as administering any other information system: provide configuration and change management, implement and automate reliable operational procedures, create and follow appropriate policies, and maintain current documentation. Whether system support staff focus on a single resource, like a Cray XE6, or a variety of resources including an HPC cluster, multiple servers and desktops, and network devices, these fundamental principles still apply. Small centers should be looking for, developing, and keeping system support staff who can adapt to a variety of resources and apply the fundamentals in all cases.

Even small HPC centers/customers can improve the largest HPC systems. For example, along with other centers, ARSC has provided Cray:

- opportunities to test new configuration options or explore using a system in a unique way
- critiques and suggestions for external servers ("ES") support
- input on XE power management
- input on methods to operate and administer systems
- porting new applications and models to existing architectures

As with large HPC centers, a small center's ability to improve technology and products is rooted more in staff talents and capabilities than in resource size. Small centers can be more flexible in their focus.

Site preparation is key to successful operations. Spend a lot of time getting ready for new systems to arrive, including talking to other center's technical staff. The more work that can be performed before arrival, the faster a system can be up and running. Develop acceptance test plans that focus on making sure a system delivers and is reliable rather than trying to chip away at the system cost. A good test plan takes careful design and lots of input, including from the vendor. Make sure there's plenty of room for spare parts near where they will be needed. Try to keep spare components in a "warm" state and test them at least once before they are actually needed. This avoids swapping a broken part for another broken part.

As part of site preparation, have a well-developed installation plan that includes proving the system is operating correctly and within expected performance criteria. Also, have verified emergency power down plans with prepared regression tests, based on the installation tests, to confirm a system returns to functioning properly after an uncontrolled shutdown. Efforts spent documenting these site guidelines save a lot of time during and right after operation interrupts.

Remote HPC system administration is eminently possible, with some planning and foresight. A local operations staff (even a single person who can get to the data center in an acceptable time frame) is necessary periodically for hands-on procedures. These break-fix situations can typically be contracted as part of system hosting though. Plan for quarterly data center visits to schedule maintenance or proactive procedures. No users and very few staff are working regularly next to an HPC system in the data center. Users are working remotely at all hours of the day. The difference of being in the next room, the next county, or the next state has no significant.

System support staff should always be considering the next stage of a resource's lifecycle. Often a mid-life upgrade can be more productive than a completely new system. ARSC has had systems for as little as two years, and systems that have lasted a decade or more. Staff should ensure the system configuration and functionality adapts to user's changing requirements, which is much more possible in a small center than a large one.

VI. CONCLUIONS

ARSC has been a proud operator of Cray systems for twenty years. During that time funding has fluctuated as much as any academic research institute. However, we have always maintained a focus on supporting users getting the most out of the resources we support and developing staff (and students) to their highest capabilities. Cray Inc., in its various incarnations, as an ARSC vendor, and informal partner, has been key to our success in those areas.

From a user perspective, "Cray" means great capacity/capability and carries significant cachet. Having Cray systems in Alaska is surprising to students, faculty, researchers, and visitors. Knowing nothing else, it signifies to them that ARSC is a real HPC center. (From the system administrator perspective, a modern Cray is just a strangely modified linux system, but still has very cool cachet.) As a small center, we are able to leverage that prominence in other areas, helping our parent organization improve research and results. However, once people get beyond the name "Cray", we have to be able to deliver the support we promise.

Unbeknownst to the users, good vendor support has a huge impact on good system operations. We rely on our vendor partners for responsiveness, technical expertise, troubleshooting guidance, and solution brainstorming. Cray Inc. has always provided these qualities for ARSC.

To get the best results in these areas, we have to be sure to inform vendor support of any significant system events or changes, in order to maintain that support and prepare them to provide the best possible support afterwards (ex: system moves, network reconfigurations, third party hardware or software installs/upgrades). This two-way relationship provides our users the best resources we can give them.

Our experience indicates that regional or small university centers like ARSC no longer compete for the largest sums of money and the largest capacity systems in the world. However, we can come close, and we can provide the same technologies and architectures to a user base that does not need the largest systems in the world. Further, smaller centers can participate in the development of new technologies that span the entire spectrum of research computing, including HPC. It often requires smaller centers to be able to learn to say no to projects and ideas in order to avoid reaching beyond what's actually possible to do within limited funding and staffing.

Efforts like the NSF's Campus Bridging task force [3], EDUCAUSE's Campus Cyberinfrastructure working group [4], and solicitations from various United States federal government departments (NSF, DoE, DoD, NIH, ...) for focused research resources and support mean there will always be a need for smaller, regional academic HPC centers. Similarly, many commercial and government entities will also always have needs to do their own in-house HPC. Therefore centers like ARSC need to exist in order to support the general HPC environment/market with trained staff and users. Having HPC vendors who are partners rather than just retailers contributes to small, successful centers and ensures a vigorous future of HPC growth and innovation.

ACKNOWLEDGMENT

We thank the following people for their memories and contributions to this paper, as well as their many years of service at ARSC.

- Virginia Bedford, ARSC Associate Director System Services, 1994-2011.
- Barbara Horner-Miller, ARSC Associate Director User Services, 1997-2011.
- John Metzner, Cray Field Support Engineer, 1993-2005, 2008-2011.
- Dr. Greg Newby, ARSC Chief Scientist, 2003-present, ARSC Director 2011-present.
- Guy Robinson, ARSC Research Liaison, 1997-2003.
- Dr. Frank Williams, ARSC Director 1996-2011.

REFERENCES

[1] Top500.org, "University of Alaska - Arctic Region Supercomputing Center." Available from: http://www.top500.org/site/49130, URL last confirmed April 4, 2013.

[2] D. Pink, "Drive, The Surprising Truth About What Motivates Us," New York, NY: Riverhead Books, 2010.

[3] NSF Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging, "Final Report," March 2011. Available from: http://www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf, URL last confirmed April 4, 2013.

[4] EDUCAUSE and CASC, "Developing A Coherent Cyberinfrastructure from Local Campuses to National Facilities: Challenges and Strategies," April 2009. Available from http://net.educause.edu/ir/library/pdf/EPO0906.pdf, URL last confirmed April 4, 2013.