Leveraging Your Local Resources and National Cyberinfrastructure Resources without Tears

Barbara Hallock
Indiana University Pervasive Technology Institute
2709 E 10th Street
Bloomington, IN 47408
bahalloc@iu.edu

Richard Knepper
Indiana University PTI
2709 E 10th Street
Bloomington, IN 47408
rknepper@iu.edu

James Ferguson
National Institute of Computational Sciences
PO Box 2008, Bldg 5100
Oak Ridge, TN 37831-617
jwf@utk.edu

Craig A. Stewart
Indiana University PTI
2709 E 10th Street
Bloomington, IN 47408
stewart@iu.edu

ABSTRACT

Compute resources for conducting research inhabit a wide range, including researchers' personal computers, servers in labs, campus clusters and condos, regional resource-sharing models, and national cyberinfrastructure. Researchers agree that there are not enough resources available on a broad scale, and significant barriers exist for getting analyses moved from smaller- to larger-scale cyberinfrastructure. The XSEDE Campus Bridging program disseminates several tools that assist researchers and campus IT administrators in reducing barriers to the effective use of national cyberinfrastructure for research. Tools for data management, job submission and steering, best practices for building and administering clusters, and common documentation and training activities all support a flexible environment that allows cyberinfrastructure to be as simple to utilize as a plug-and-play peripheral. In this paper and the accompanying poster we provide an overview of campus bridging, including specific challenges and solutions to the problem of making the computerized parts of research easier. We focus particularly on tools that facilitate management of campus computing clusters and integration of such clusters with the national cyberinfrastructure.

Categories and Subject Descriptors

C.1.4 [**Parallel Architectures**], C.2.4 [**Distributed Systems**] K.6.3 [**Software Management**]

General Terms

Clusters, Education, Management, Documentation, Performance, Design, Experimentation, Human Factors, Standardization

Keywords

Campus Bridging, XSEDE, High-performance Computing, Usability, Supercomputing, Academic Computing, Research Computing, Documentation

1. INTRODUCTION

For scientists at the university level in the United States, new avenues of experimentation have meant a constant and ever-growing demand for cyberinfrastructure (CI) resources to complete their experiments. For the purposes of this discussion, these CI resources may range in complexity from something as simple as a plug-and-play external hard drive connected directly with the researcher's personal computer to much more complex supercomputing resources maintained and provided via partnerships between top-tier research institutions in the US, such as the NSF-supported eXtreme Science and Engineering Discovery Environment (XSEDE). XSEDE is "the most advanced, powerful, and robust collection of integrated advanced digital resources and services in the world… a single virtual system that scientists can use to interactively share computing resources, data, and expertise." XSEDE collectively allocates and supports several PetaFLOPS worth of computing resources and has already made great strides in usability and in lowering the barriers of access to major CI resources. Researchers report difficulty obtaining sufficient computing resources to support their research [1], (including allocations of resources through national services such as XSEDE), and using their local resources and national resources, even when they have access to the resources they need. In this paper, we introduce campus bridging as a concept and describe the primary initiatives of the XSEDE Campus Bridging initiative. This paper is intended to inform university and college leadership, CIOs, faculty, campus IT support staff, users, and those who utilize cyberinfrastructure in their scholarship about ways in which they can integrate their local CI resources with national systems such as XSEDE. This can make it easier to use campus-based resources and to leverage national resources such as XSEDE. We focus particularly on tools that make training and information transportable between local campus resources and national resources, and that improve the ability to use local and national resources together to achieve optimal use of both. (In other words, make what you have easier to use and more effective in supporting your research, utilizing federally funded CI resources, all without tears.)

2. THE EXTREME SCIENCE AND ENGINEERING DISCOVERY ENVIRONMENT (XSEDE)

XSEDE is a virtual organization funded by a five-year NSF grant award that includes staff from a number of top-tier supercomputing centers across the United States, including (but not limited to) the Texas Advanced Computing Center (TACC), Indiana University (IU), the Pittsburgh Supercomputing Center (PSC), the San Diego Supercomputer Center (SDSC), the National Institute of Computational Sciences (NICS), the National Center for Supercomputing Applications (NCSA), and others. Major CI resources allocated and supported by XSEDE include Stampede (TACC), Blacklight (PSC), Comet (SDSC), and Wrangler (TACC/IU); major affiliated systems such as the Open

SIGUCCS '14, November 2-7, 2014, Salt Lake City, Utah, USA.
ACM 978-1-4503-2780-0/14/11.
http://dx.doi.org/10.1145/2661172.2661202

Science Grid and Blue Waters; and local resources such as the IU Mason system and other campus-based CI resources.

3. CAMPUS BRIDGING

3.1 Defining Campus Bridging

H. Edward Seidel coined the term "campus bridging" in 2009, when as the Director of NSF's Office of Cyberinfrastructure, he charged six task forces of the National Science Foundation (NSF) Advisory Committee for Cyberinfrastructure (ACCI). That task force developed this description of campus bridging:

"Campus bridging is the seamlessly integrated use of cyberinfrastructure operated by a scientist or engineer with other cyberinfrastructure on the scientist's campus, at other campuses, and at the regional, national, and international levels as if they were proximate to the scientist, and when working within the context of a Virtual Organization (VO) make the 'virtual' aspect of the organization irrelevant (or helpful) to the work of the VO." [1]

More colloquially, the goal is to create "virtual proximity"—everything from a modest local campus cluster to the largest systems should feel as easy to use as a peripheral attached directly to the user's computer. The term "campus bridging" reflects the fact that at the time this taskforce was convened, the national, NSF-funded supercomputing resources seemed to the user to be completely isolated from campus resources. One needed a bridge to get from one to the other. (The "virtual proximity" analogy was coined by Von Welch of the IU Center for Applied Cybersecurity Research.)

A number of challenges exist within the discipline of academic CI. Campus Bridging focuses on a select few. In this section, we present the major challenges campus bridging attempts to resolve and mitigate. In the next section we describe the tools and services being deployed by the XSEDE Campus Bridging team to address these challenges.

3.2 Where Am I and What Commands Work Here?

Since the time Beowulf computing clusters were introduced in 1995 [2], clusters have become extremely popular as a way to accelerate scientific computing applications. Because such clusters can be built locally and easily, they often are created and set up locally with configurations that are highly variable from one cluster to another. This means that software and libraries available on one cluster may not be available on another, or if they are, the versions might not be the same or they might not be accessible in the same location of the system hierarchy.

Learning to interact with a small number of applications may not be an insurmountable challenge for those with little technical expertise. But the difficulty of effectively leveraging the available CI increases significantly in correlation with the number of resources the researcher has to interact with in order to complete his or her experiment. For those who must operate within multiple systems concurrently to achieve results, this disparity can mean frustration and lost productivity due to the cognitive costs of context switching, especially when not much documentation is available at the researcher's level of technical proficiency.

3.3 Documentation

Another significant challenge that campus bridging aims to mitigate is the scarcity of clear, usable documentation available for users who are unfamiliar with the system(s). Where documentation does exist, it is often incomplete or written for a highly technical audience (e.g. system administrators), making it less useful for researchers without a background in the technical discipline.

3.4 Insufficient CI Resources and Allocation "Lag"

A major issue with the research CI ecosystem is the relative scarcity of resources available with which to complete experimental work. XSEDE accepts requests for allocations on a quarterly basis [3]. For researchers on a tight deadline, this allocation timetable can be a challenge. Furthermore, the resources awarded during an allocation process may be as low as half those requested. Information provided by XSEDE [4] makes it clear that preliminary work on XSEDE systems or similar systems in the form of scalability studies provides some performance data that aids the likelihood that a proposal will lead to an allocation of XSEDE resources.

3.5 Moving Data and Other Challenges

Most researchers who use national resources such as XSEDE wish to analyze data that is held someplace other than an XSEDE system – either on local storage resources or on some other national repository. Thus means that those who use XSEDE (either in isolation or in concert with collaborators) must move data from one CI system to XSEDE, and then move the resulting output from XSEDE back to local resources.

For a complete examination of the community-driven use cases that defined the XSEDE Campus Bridging program, see [5].

4. XSEDE CAMPUS BRIDGING SOLUTIONS

In this section we provide an overview of tools and services (including consulting) that are offered by XSEDE through its Campus Bridging program to aid campus IT organizations, researchers, educators, and those who support researchers and educators.

4.1 XSEDE-Compatible Basic Cluster (XCBC) & Documentation Portability

There are tens of thousands of clusters installed on campuses and in research labs in the US. There are a handful of such clusters within XSEDE. In order to increase the compatibility of software between campus resources and XSEDE, the XSEDE Campus Bridging program created the "XSEDE-compatible Basic Cluster" software stack (XCBC). This consists of a definition of a suite of software that represents an open-source subset of the software and tools one can expect to find on any XSEDE cluster, mechanisms for installing that software [6], and tools for installing that software [7, 8, 9,10].

The software definition will evolve over time, and current XSEDE plans call for additions of specific sets of tools (e.g. a set of tools for bioinformatics). Today, one could install the XCBC software stack and have a cluster with an open-source job management system, a standard set of compilers, and basic software and application libraries.

The XCBC can be installed in two ways. If you are building a cluster from scratch, or are so unhappy with your current cluster that you are ready to start over from scratch, you can install XCBC using a cluster installation tool called ROCKS. Instructions are available [9]. Or you might have a cluster that works well and wish to add the ability to run the XCBC software, so that commands that work on an XSEDE cluster also work in the same way on a local cluster. This capability can be implemented by installing sets of packages called RPMs, or Red Hat Package Managers [11], which allow the installation of specific packages or software modules. XSEDE distributes these software modules through a YUM Repository. ("YUM" stems from the Yellowdog Linux distribution, although the current use of YUM repositories goes well beyond the intent of the initial tool that was part of the Yellowdog Linux project [12].) YUM makes it easy to install a particular set of modules from the XSEDE Campus Bridging software repository on a local cluster. You can even sign up for automatic notifications when software in the repository is updated.

This software has a number of benefits. It makes it easier for local IT staff to manage local computing clusters so they can spend more time addressing the specific needs of local researchers and students. It makes it easier for researchers to learn and use computing commands because the same commands work on local clusters and on XSEDE. This approach also means IT staff, researchers, and educators can leverage training resources created by XSEDE and by other educators and staff using the XCBC cluster build. Over time, use of the XCBC can also improve efficiency within a given institution by creating consistency in the setup and management of the disparate clusters that institution operates.

4.2 Globus Transfer

Globus Transfer is a fast, reliable, high-performance service that facilitates secure data movement. Designed specifically for researchers, Globus Transfer provides easy, fire-and-forget features. It automates the activity of managing file transfers between any two resources, whether between two XSEDE resources or to/from XSEDE and another machine, such as another supercomputing facility, cloud resource, campus cluster, lab server, desktop, or laptop [13]. The interface is intuitive, providing a drag-and-drop method of file transfer that is similar to graphical FTP software or the interface for user-friendly cloud storage applications such as Dropbox. Globus software has been available on XSEDE resources for some time. During the first three months of 2014, users of XSEDE transferred a collective 1.32 PB to various XSEDE endpoints and an additional 1.30 PB from various XSEDE endpoints [14].

4.3 XSEDE Execution Management Service (EMS) – Based on UNICORE 6

XSEDE Execution Management Service (EMS) is a tool for seamless and secure job management, based on an XSEDE-customized implementation of UNICORE 6. With the EMS command-line client tool, users can manage jobs, submissions, and associated data movement tools. Access to resources is provided via shell or scripting environment. Users can run jobs, monitor their status, and retrieve generated output, either in single-job mode or in a powerful and flexible batch mode for multiple jobs. XSEDE EMS services can also be incorporated into a graphical interface using the Eclipse integrated development environment [15]. XSEDE EMS supplies web-accessible architecture that can be incorporated into science portals and gateways. All of the XSEDE EMS architecture utilizes a set of common standards for job creation and control, in a workflow made up of a services layer that handles the job definition and management and a system layer that controls the interface to particular grid resources.

The UNICORE software has been deployed to all XSEDE Service Providers (SPs) as of spring 2014, and is useful in creating automated workflows that require the use of multiple disparate SPs with the XSEDE ecosystem. The benefit is that users trade a significant effort in creating, maintaining, and facilitating those workflows for a smaller effort in learning to interface with UNICORE.

4.4 XSEDE Global Federated File System (GFFS)

The XSEDE Global Federated File System (GFFS) is being put into production in 2014. In its initial implementation within XSEDE, GFFS will allow geographically distributed sharing of files independent of physical location. In this regard GFFS has similarities with the well-known Andrew File System (AFS). As XSEDE's implementation of GFFS expands, it will allow file sharing, submitting computing jobs, and sharing computing resources in the form of virtual clusters. The initial deployment of GFFS in 2014 will allow researchers to link, store, and share files on local resources and on XSEDE resources (reducing the need to synchronize multiple copies of an individual file kept in different locations) and to share files with users specified by the file owner [16].

In the XSEDE Global Federated File System, users can use client software to manage files and permissions as they would in a normal file manager. Users can also export local file systems or directories into the GFFS in order to make files at their site available to XSEDE resources, to collaborators at other institutions, or to wherever a client can browse the GFFS. The exported file systems or directories can be used to provide data for analysis, can be shared between collaborators, or can represent job queues via integration with the XSEDE EMS framework. The XSEDE operations team provides the root of the GFFS tree and manages the linking between containers that are exported into GFFS, making them available to be used across XSEDE resources.

4.5 Documentation and HPC University

A wide variety of user guides to advanced computing, parallel computing, and using XSEDE are available online. Using the XCBC software build makes it easier to modify these resources for use in local user support documentation and instructional materials. Major sources of documentation available through and affiliated with XSEDE include the XSEDE web site (xsede.org), XSEDE Knowledge Base [17], and HPC University [18].

4.6 Consulting Services

Campus bridging is a new area. We hope this paper provides some sense of the opportunities available to campus leaders, campus IT organizations, and research labs. However, there is still a fairly large knowledge gap between "I'd like to make better use of my campus resources by integrating them more effectively with national cyberinfrastructure such as XSEDE" and "I understand how to use YUM to download and install RPMs on my local cluster, and where to look within XSEDE web sites to find documents I can modify easily for local use." XSEDE campus bridging staff are available to help with answers to questions, extended discussions of technical capabilities, instructional sessions by videoconference, and even on-site visits to your location. Contact Campus Bridging staff via email at help@xsede.org, or call the XSEDE Help Desk at 1-866-907-2383 to be transferred to a Campus Bridging consultant (8:00am-5:00pm, M-F, except US federal holidays). A discussion with an expert is often the easiest way to begin using XSEDE Campus

Bridging resources as we are still in the early phase of developing communities of practice.

5. CONCLUSIONS

Everyone working in IT in higher education has more than likely experienced a lack of resources related to research computing and research education – not enough computing resources, not enough storage, and not enough time to support researchers and students. This is a matter of US global competitiveness. In this paper, we have explained the concept of campus bridging and described tools that are currently available to the national research and education community. Much as the NSFNET network aligned national efforts in networking during the 1980s, we hope the XSEDE Campus Bridging program will align tools, usage, and documentation of clusters throughout the US. The requirements for alignment differ in their technical particulars, but we believe that aligning disparate campus resources with XSEDE cluster software standards will lead to a more effective national cyberinfrastructure, increase the ease with which researchers can integrate local and federally funded CI resources, and increase the ease with which researchers and educators can find and customize documentation and training materials for local use. Because campus bridging is still in its infancy, XSEDE Campus Bridging personal support and interactions with the national higher education and research communities are particularly important in enabling the community as a whole to work together to achieve the goals of campus bridging and maximize US global research competitiveness.

6. ACKNOWLEDGMENTS

This research was supported by the National Science Foundation through funding for the Extreme Science and Engineering Discovery Environment (XSEDE) via award ACI-1053575. This work was also aided by the Indiana University Pervasive Technology Institute, which is supported by a grant from the Lilly Endowment, Inc.

7. REFERENCES

[1] ACCI Task Force On Campus Bridging. *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging Final Report.* Apr. 2011. doi=http://hdl.handle.net/2022/13210

[2] Becker, D., Sterling, T., Savarese, D., Dorband, J., Ranawak, U., and Packer, C. 1995. BEOWULF: A parallel workstation for scientific computation. In *Proceedings of the International Conference on Parallel Processing* vol. 95. 1995. http://www.phy.duke.edu/~rgb/brahma/Resources/beowulf/papers/ICPP95/icpp95.html

[3] XSEDE Allocations. https://www.xsede.org/allocations

[4] Hackworth, K. Writing and Submitting a Successful XSEDE Proposal. Apr. 2014. https://meeting.austin.utexas.edu/p4p2yy0hcgg?launcher=false&fcsContent=true&pbMode=normal

[5] What are some examples of Campus Bridging use cases related to XSEDE? https://portal.xsede.org/knowledge-base/-/kb/document/bbsv

[6] http://cb-repo.iu.xsede.org/xsederepo/XSEDE-Campus-Bridging-Cluster-software-strategy-2013.pdf

[7] Brueckner, R. New XSEDE tools to aid campus cluster administrators. Dec. 2013. http://insidehpc.com/2013/12/17/new-xsede-tools-aid-campus-cluster-administrators/

[8] XSEDE Campus Bridging. https://www.xsede.org/campus-bridging

[9] What is the XSEDE Rocks Roll, and how do I use it? https://www.xsede.org/web/xup/knowledge-base/-/kb/document/bdpe

[10] Hallock, B., Knepper, R., Ferguson, J., and Stewart, C. 2014. XSEDE14 Campus Bridging Paper. In *Proceedings of XSEDE 2014 Annual Conference* (Atlanta, Georgia, United States, July 13-18, 2014). XSEDE '14. ACM, New York, NY, in press.

[11] RPM Package Manager. http://en.wikipedia.org/wiki/RPM_Package_Manager

[12] The Yum web site. http://yum.baseurl.org

[13] Globus File Transfer User Guide. https://www.xsede.org/globus/

[14] XSEDE 2014 Q1 Quarterly Report. In Press.

[15] The Eclipse Foundation open source community website. http://eclipse.org

[16] Knepper, R. Campus Bridging Final Report. Sept. 2013. doi=http://hdl.handle.net/2022/17378

[17] XSEDE Knowledge Base. https://www.xsede.org/web/xup/knowledge-base/

[18] HPC University. http://hpcuniversity.org/