User Requirements
Best Practices for Campus Champions

13 July 2012

Version 1.0

Table of Contents

List of Figures

A. Document History

Entire Document	1.00	07/13/12	Baseline	Roger Moye, Don Garner, Kimberley Dillman

B. Document Scope

Campus Champions are expected to assist researchers at their institutions with effective selection, access and use of XSEDE cyber infrastructure services and resources. To that end Campus Champions must determine the service, including technical support, and resource requirements of each researcher in order to provide proper assistance. This document is intended to serve as a general guide to determining such service and resource requirements.

C. Assessing User Requirements

In general researchers require various user services before, during and sometimes after using the networking, hardware and software resources of the cyber infrastructure. Besides systems operations services like software installation and file transfers, such services include information about available services and resources, how to get accounts and allocations, how to access resources, as well as, user guides, training, portals and technical support.

In this case technical support refers to support with the technology at various points in the process of selecting, accessing and using the cyber infrastructure. Of course, the amount of technical support required depends on the researcher's technical abilities and willingness to learn and use various technologies. Other services, such as training, user guides and portals may be sufficient for some researchers, while others may require technical support from cyber infrastructure help desk staffs or directly from Campus Champions.

The actual networking, hardware and software resources that researchers require depends on the computational research that they want to perform. However, they are primarily going to know what research they want to perform in terms of the math, physics, chemistry, biological, engineering, etc. concepts involved. Therefore, researchers will require various amounts of fundamental technical support to define the computational research they want to perform in terms of the resources of the cyber infrastructure.

Accordingly, Campus Champions must develop a useful knowledge of what cyber infrastructure services and resources are available. They must work with researchers to define their computational needs in terms of their field of study, then in terms of the available services and resources of the cyber infrastructure, while determining and applying the amount of technical support that they require. Campus Champions must then work with researchers as needed to select and use the appropriate services and resources of the cyber infrastructure to perform their computational research. These efforts are generally sequential, but one does not necessarily have to be completed before the next begins. Also, it may be necessary to break the overall computational project down and apply these efforts to different parts of the project.

D. Know Your Cyber Infrastructure Services and Resources

Participating in Campus Champions meeting and events are a good way to get to know what cyber infrastructure services and resources are available. The XSEDE website is also a good source of information. There is the annual XSEDE conference and the Campus Champions participation in various other research computing events such as the annual Supercomputing Conference. Various online and in-person training is also available. Take time to review the various user guides and documentation that is available for each XSEDE resource. The best way to learn about the services and resources is to try them yourself.

E. Researcher Cyber Infrastructure Background

To understand the basic parameters of their computational requirements, it would be advisable to first establish a context for these requirements by discovering if the researcher has any experience with research cyber infrastructure. The answers to the questions will help frame the remaining questions more accurately and help the researcher produce useful and accurate answers.

1. Do they currently use local resources, and if so, which resources (if applicable)?
2. Are they familiar with XSEDE and its resources? If not, then a brief overview of XSEDE should be presented along with an explanation of Startup and Research allocations.
3. Do they currently use XSEDE resources, and if so, which resources (and do they have an active allocation)?

4. Have they used other national research cyber infrastructure, and if so, which ones?
5. Later, once an XSEDE resource has been selected then the Champion should ask whether or not the researcher will require help in applying for a Startup or Research allocation.

If the researcher has some experience using research cyber infrastructure then you might expect that the remaining questions will be more easily understood by the researcher and the answers more meaningful. Some questions can even be skipped. If they have little experience, then the remaining questions might require additional explanation from the Champion and the answers might need to be refined. If the researcher does not have an allocation, then they will likely need help in applying for one.

F. Researcher Computing Skills

Further understanding the needs of the researcher is to understand their skill level regarding their ability to use cyber infrastructure. Transitioning from a single desktop system running Windows, for example, to a multiprocessor system running Linux and using a batch scheduler can represent a significant obstacle to the effective utilization of cyber infrastructure. Assess the severity of this obstacle by asking these questions of the researcher:

1. What is their current desktop working environment (Mac, Windows, Linux)?
2. What is their current computational working environment (Mac, Windows, Linux)? Single or multicore? Single user or batch scheduler?
3. Is the target system running an operating system for which they are not familiar? If so, offer them a tutorial (either existing XSEDE tutorial, existing institutional tutorial, or one-on-one training).
4. Are they familiar with the type of system where their jobs will be running? Have they ever used that system or one like it? If not, offer them a tutorial (either existing XSEDE tutorial, existing institutional tutorial, or one-on-one training).
5. Have they ever used a batch system? If not, give them a brief hands-on tutorial using your Champions account on a system like Ranger to demonstrate how to submit a job.
6. Are they familiar with remote access to cyber infrastructure (such as SSH)? If not, offer them a tutorial on remote logins using your Champions account. Show them how to login with the XSEDE portal, and alternatively, GSI-SSH preferably from a system where it is already installed and verified to be working.
7. Are they familiar with parallel processing (if applicable) with OpenMP or MPI? If not, explain what parallel processing means and the difference between OpenMP and MPI.

Some researchers will understand, in extreme detail, their computational requirements. However, running their jobs in a batch environment on a remote system can be confusing and intimidating. So taking the time to offer some one-on-one training will ease this transition. Suggested topics, based on the answers to the above questions, are:

1. Brief anatomy of an HPC cluster.
    1.1. What is a cluster?
    1.2. What is the difference between a login node and a compute node?
    1.3. Explanation of the various file systems.
    1.4. What is a batch scheduler?
2. How to login via SSH or GSI-SSH to the target resource.
3. How to transfer data to/from the target resource.

4. Where to store the data on the target resource.
5. How to construct a job script (batch script), including how to invoke their software.
6. How to submit a job and view the job queue.
7. Where to find and retrieve the results.

Each XSEDE resource has its own Getting Started guide which includes much of the above information. Demonstrating these topics hands-on, with references to the Getting Started guides, would be helpful. Further, the XSEDE website has links to tutorials for topics such as Introduction to Linux, Parallel Programming, MPI, OpenMP, and so on. Offer these tutorials to the researcher as needed as a reference for any hands-on training that you provide.

G. Determine Software Requirements

Once you have determined the researcher's level of experience with research cyber infrastructure, and their comfort level, you should proceed to discussing the type of software that they want to use.

1. Which software do they need for this project?
2. Is it commercial, open source, or their own code?
3. Does it require a license, and if so, do they have one?
4. Do they know if it is installed on XSEDE already?
5. Can they install it themselves or will they require assistance?
6. Do they have access to the source code?
7. Does the software have any software dependencies that they know of (specific library versions or dependent on other software packages)?
8. How much disk space does it take to install it?
9. Does it require a specific operating system and version (Red Hat or SuSE Linux, Windows, etc.)?
10. Have they or anyone else used it on XSEDE before? Do any performance benchmarks exist?
11. Is the software capable of running in parallel (MPI or OpenMP)? This will help determine the software's parallel capabilities, and whether or not the researcher knows how to use these capabilities.

Answers to these questions will help the Champion discover how much the researcher understands the software, how it functions, whether or not they will need help installing it, and will guide the Champion in selecting the cyber infrastructure resource(s) where it might run successfully.

H. Determine Hardware Requirements

Determining hardware requirements primarily involves determining:

- how many processors of an available type and number will be required to complete the required number of computational runs in a given amount of time;
- the amount of memory that will be required to hold the computational software and active data;
- the amount of disk storage that will be required for the input and output data.

The amounts can vary widely. In many cases it might be relatively easy to determine how much disk storage will be required for input and output data. However, you must be careful not to overlook storage

requirements for any intermediary data that might be generated. Also, the maximum amount of storage that will be required for the worse case computational path should be covered. The format that the data will be in should also be considered, for example in flat files, databases or encapsulated in models. Since the format may produce extra storage requirements.

You can get an estimate of the amount of memory that will be required from the size and operations of the computational code, along with the size and number of variables and other data structures that the code operates on.

How much processing resources will be required may be the most difficult to determine. Some requirements can be extrapolated from benchmarks, small test runs or from comparable computations Attempting to overestimate the amount required is another common technique. While the practical approach for some situations may be trial-and-error.

See the section titled "Determine Job Characteristics" for more information.

I. Determine Network Requirements

Determining network requirements mainly involves insuring that sufficient and compatible local hardware, software and Internet access are available to connect to the remote cyber infrastructure resources that will be used for the computational research. Believe it or not, some people still use dialup modems to connect to the Internet, Window 95 is still being used, and some systems never get updated or connect to the Internet.

J. Determine Job Characteristics

Once the software requirements have been established, the next step is to determine the parameters of a typical compute job.

- **Compute Job Characteristics**

1. Will the jobs run on a single processor core, or multiple cores? This will determine which systems can run the jobs (high-throughput system or parallel system).
2. Do the jobs use MPI or OpenMP? This will help determine node interconnect requirements.
3. What is the maximum number of processor cores it is capable of using before performance degrades (how well does it scale)? Have benchmarks been run to determine this number?
4. How much memory per process (or per node) is required?
5. How long will the job(s) run (in hours)?
6. Approximately how many jobs do they need to run?
7. Do the jobs need any special hardware (like GPUs, SSDs, shared memory)?

- **Data Characteristics**

1. How much disk space is required for the input and output files per job or per the entire set of jobs?
2. Will the data need to be archived at the end of the project?
3. How long will the data need to reside on the system where the jobs are running (for post-analysis or running of additional jobs)?
4. Does the job require frequent file I/O, or does it read data at the beginning of the job and write the results at the end? This will help determine file system performance requirements.

5. Is the job capable of checkpointing? Can it be stopped and restarted? Some jobs can be stopped and restarted but reading in the checkpoint and other data takes too long and consumes too much wall clock time to be useful.

Answers to these questions will help the Campus Champion select the resource(s) most suitable for running these compute jobs, or will inform the Champion that some benchmark jobs need to be run to understand how the job will perform. Resource selection is beyond the scope of this document. However, the Campus Champions have resources available to aide in this process:

1. A member of the Campus Champions Leadership Team currently maintains a resource selection presentation and data to support this activity.
2. There is an automated resource selection tool currently being developed.
3. Domain experts employed by XSEDE can assist in resource selection. File a Help Desk ticket for assistance.