

XSEDE Capability Delivery Plan

HPC-2 User project on more than one HPC system

Last revised 2016-09-21

Background

Use cases describe community needs, requirements, and recommendations for improvements to cyberinfrastructure “CI” resources and services. A Capability Delivery Plan “CDP” is an executive summary of use case support gaps, of plans to fill those gaps with new or enhanced capabilities, and of existing operational components that already support aspects of a use case.

Use Case Summary

Use case HPC-2 describes how a researcher with an XSEDE allocation on more than one HPC resource can learn how to use the systems, get data and code to the systems, adapt and optimize code for the systems, execute jobs on the systems, obtain user support, and report results of their work. A key user need in this use case is for XSEDE SP resources to offer a consistent user environment and a similar set of system tools, software applications, and software libraries. Use case document: <http://hdl.handle.net/2142/45703>.

CDP Summary

With the following exceptions, XCI believes that the features described in the use case are fully supported by the XSEDE system components listed below. XCI does not have technical solutions for any of the exceptions listed below, but we have a suggestion for how XSEDE might address the first one.

Gap(s) that we currently plan to address:

- Inconsistent user environment and software applications/libraries

Gap(s) that will not be addressed at this time:

- Technical expertise required by researchers
- Verification of quality attributes

Time and effort summary:

- None

Functionality Gaps

The features described in the use case are supported by the XSEDE system components listed below. A key element of the use case that is not provided by these components is the consistent

user environment and software applications and libraries on SP resources.

1. Inconsistent user environment and software applications/libraries (suggested priority: medium)

Step 9 expresses the need for consistency of local applications and working environments across the XSEDE HPC resources. User support staff from all of the XSEDE service providers actively work together to identify and acquire these applications and tailor these environments to be as similar as is practical.

Because each SP resource is specified, designed, and delivered in response to a unique award from NSF, it is not possible for XSEDE to control the user experience on each resource. Ultimate responsibility for the software and environments on resources belongs to the SPs, and is driven both by their experience with scientific users on current and past systems and by their responsiveness to specific NSF system solicitations.

Nevertheless, there are several things that XSEDE can--and should--do to ensure a reasonable amount of consistency across resources. Specifically, *XCI recommends that XSEDE's CEE (Community Engagement & Enrichment) team and the XSEDE SPs (Service Providers) work together on a joint activity to accomplish the following tasks.*

1. Identify the most important elements of consistency needed by XSEDE researchers who use more than one SP resource.
2. Develop and carry out joint plans for consistency across independent SP resources.
3. Document the features of SP resources where consistency is actively maintained so that researchers can plan accordingly.

2. Technical expertise required by researchers (suggested priority: high)

All of the HPC use cases, as written, imply strongly that researchers who use XSEDE HPC systems prefer to, or need to, work very closely with the low-level interfaces of the XSEDE resources. There are good reasons for this need or preference, but it nevertheless puts a huge burden on the researcher and his/her associates to acquire expertise about the system and its low-level interfaces and to spend time inspecting, monitoring, and debugging the system and their application. Anything that can be done to reduce the need for scientists to deal with these issues personally would surely improve their productivity as researchers. *There are no current plans to address this issue.*

3. Verification of quality attributes (suggested priority: low)

Verifying quality attributes requires significant one-time and ongoing testing. XSEDE has decided that the costs of this testing would not bring sufficient benefit. Instead XSEDE will monitor user satisfaction, usage, and available performance metrics and address quality issues when raised by users. *There are no current plans to address this verification gap.*

System Components That Support This Use Case

The following XSEDE operational components currently support this use case.

Component	Supported Functionality
XSEDE User Portal (XUP)	The front-researcher interface to the XSEDE system where researchers register with XSEDE, access XSEDE and SP resource documentation, and request allocations to use XSEDE SP resources. XUP also provides a directory for researchers to find their login details on SP resources.
Globus Auth	Provides the authentication service used by researchers to login to XU and obtain an XSEDE OAuth2 token that can be used with other XSEDE services, including Globus Transfer (see below).
MyProxy, MyProxy OAuth2 (OA4MP)	Used by some XSEDE researchers to obtain XSEDE tokens for direct remote login to SP resources (as opposed to SP-specific credentials). Also used by Globus Transfer (see below) to acquire tokens for authenticating to source/destination endpoints.
XSEDE Kerberos	The repository that stores XSEDE usernames and passwords and authenticates XSEDE identities for Globus Auth
XSEDE SSO Hub	Used by XSEDE researchers to login to SP resources using their XSEDE userid and password (as opposed to SP-specific credentials).
XSEDE Information Services	Used by XUP to automatically populate user documentation pages that show the availability of software and services on specific SP resources
UNICORE services	Provides the server interface for remotely submitting and managing jobs on XSEDE L1 resources that are accessed via the Genesis II or UNICORE clients.
Genesis II client	Some researchers use Genesis II clients to remotely submit and manage jobs on XSEDE SP resources. May be installed on researcher systems, campus systems, science gateway systems, or XSEDE login servers.
UNICORE client	Some researchers use UNICORE clients to remotely submit and manage jobs on XSEDE SP resources. May be installed on researcher systems, campus systems, science gateways systems, or XSEDE login servers.
Genesis WS-Trust STS	Translates XSEDE OAuth2 tokens (user identity, group membership) obtained from Globus Auth into the signed SAML chains required by th

	Genesis II client and UNICORE services. Currently in development.
Globus Transfer	Some researchers use Globus Transfer to move code and data to and from XSEDE HPC systems. Globus is a hosted web application accessed by users via a web browser, and it manages quality attribute regarding reliability, performance, scalability, and history accessibility. It also provides an SSH-accessible command-line interface and a RESTful web API for application integration.
GridFTP, Globus Connect	Used by XSEDE SPs, campus IT providers, and researchers to create endpoints from which (and to which) Globus Transfer can transfer files