

XSEDE Capability Delivery Plan

UCCAN-1 Run a Remote Job

Last revised 2016-11-10

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 CAN-1 (“Canonical 1”) describes how computational tasks can be submitted to (and subsequently managed on) XSEDE resources using an XSEDE user identity. Although it is more common for XSEDE users to use the interactive login interface (see previous section) with local job submission interfaces to submit and manage tasks, remote submission is vital for supporting the scientific workflow, science gateway, campus bridging, and federation & interoperability use cases described in the “user needs” section above, as well as several of the high-performance computing, high-throughput computing, big data analytics, and big data visualization use cases.

Use case document: <http://hdl.handle.net/2142/45685>

CDP Summary

The functionality in this use case is **mostly** supported by the operational components listed below.

Gap that we currently plan to address:

- None

Gaps that will not be addressed at this time:

- Access to required data
- Accessibility issue
- Verification of quality attributes

Time and effort summary:

- None

Functionality Gaps

1. Access to required data (suggested priority: medium)

CAN1.b (Access to Remote Data During Execution) describes the need to access job related files on remote systems during job execution. The architectural response recommends use of GFFS to satisfy this, but GFFS is not universally available on XSEDE compute resources, which often do not support external Internet connectivity. This issue could be resolved by making GFFS and internet connectivity universally available on compute resources. Other solutions are possible, though the technical hurdles appear significant.

Plans: There are no plans to address this gap.

2. Accessibility issue (suggested priority: medium)

The current architectural response to CAN-1 (and related HTC-1) requires the scientist end user to produce a detailed description of the tasks to be executed by the system in a language called JSDL. To do this, the user must learn the JSDL language and must understand the details that are important to express in the description. These details may vary from resource to resource. XSEDE provides a graphical, “workbench-like” interface for constructing this description (described in L3D §3.3.2.5), but it requires downloading and installing the Genesis II software to a local computer.

Plans: There are no plans to address this gap.

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

Verifying many quality attributes require 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.

Plans: There are no 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
UNICORE services	Provide the server interface for submitting and managing jobs on XSEDE L1 resources that are accessed via the Genesis II or UNICORE clients.

Genesis II client	Provide the user interface for submitting and managing jobs on XSEDE L1 resources. May be installed on end user systems, campus systems, science gateway systems, or XSEDE login servers.
UNICORE client	Provides both thick client and command line interface to submitting and managing jobs on XSEDE L1 resources. May be installed on end user systems, campus systems, science gateways systems, or XSEDE login servers.
XSEDE Central Database (XCDB)	Provides a database of allocated user projects and project groups that define the users who are authorized to submit jobs to specific resources.
AMIE	A mechanism used by XSEDE and Level 1 and 2 SPs to share project group membership information between XCDB and local authorization systems.