



High Performance Computing Use Cases
(UCHPC 1-3)
Architectural Response v1.1

Table of Contents

[References](#)

[UCHPC-1 User project on a single HPC system](#)

[Summary](#)

[Steps](#)

[Quality Attributes](#)

[UCHPC-2 User project on more than one HPC system](#)

[Steps](#)

[Quality Attributes](#)

[UCHPC-3 Using a scientific workflow system to run on HPC systems](#)

[Steps](#)

[Quality Attributes](#)

References

- [UCHPC] *XSEDE High Performance Computing Use Cases*, June 17, 2013.
- [ARCH] *XSEDE Architecture Overview: A Level 1 and 2 Decomposition for Software Developers and Integrators*, September 1, 2014.
- [L3D] *XSEDE Architecture Level 3 Decomposition*, April 23, 2014.
- [UCCAN-1] *X-WAVE Canonical Use Case 1 - Run a Remote Job, Level 3*, Version 0.94.
- [UCCAN-2] *XUAS Canonical Use Case 2 - Managed File Transfer: Architectural Response*, Version 1.0.
- [UCCAN-3] *X-WAVE Canonical Use Case 3 - Remote File Access, Level 3*, Version 0.91.
- [UCCAN-4] *Canonical Use Case 4 - Interactive Login: Architectural Response*, Version 1.0. May 20, 2014.
- [UCCAN-5] Allocations Management
- [UCCAN-7] Subscribe for Resource Information
- [UCCAN-8] Search for Resource Information
- [UCCAN-9] *Canonical Use Case 9 - XSEDE User Identity and Access Management: Architectural Response*, Version 1.1.

UCHPC-1 User project on a single HPC system

Support for this use case is provided primarily by XSEDE SPs. XSEDE provides two system-wide functions that are used heavily in this use case: identity management (user accounts and accounting) and information resources (documentation and system status). XSEDE and SPs coordinate a third heavily used interface: the interactive login interface. The remainder of the interfaces involved in this use case are provided by the SP that operates the single HPC resource being used. XSEDE offers access-layer interfaces for most of these functions, but because this use case focuses strictly on a single resource and a single user project, users may find it simpler to work with resource-specific interfaces.¹

	Identity Management	Interactive Login	Accessing Remote Files	Submitting and Managing Computations	Transferring Datasets	Discovering and Providing Resource Information
Access	2,3,6,8,9,11,12,13	(2),(3),(6),(8)		(8)	(2),(12)	1,11,13
Service		2,3,6,8,9				
Resource				3,6,8,9	2,12	1,11

Figure 1. This figure shows which elements of the XSEDE architecture [ARCH §D] are engaged during each step of UCHPC-1. Steps numbers in parentheses are alternatives to the primary interface.

Summary

- XSEDE’s access layer provides the primary interfaces for identity management and information services for this use case. (Steps 1,2,3,6,8,9,11,12,13)
- XSEDE and SPs coordinate the service interface for interactive login, which provides significant benefit in this use case. (Steps 2,3,6,8,9)
- The SP provides the primary interfaces for submitting and managing computations, transferring data and application files, and system-specific documentation. (Steps 1,2,3,6,8,9,11,12)
- **If the user anticipates needing to use other XSEDE HPC platforms in the future,** he/she may find XSEDE’s interfaces for interactive login, submitting/managing computations, and transferring data and application files to be helpful. Otherwise, the

¹ As explained in the *XSEDE Architecture Overview*, XSEDE’s architecture is aimed at facilitating user relationships with multiple SP resources and maintaining those relationships over time as each user’s data and computing needs change and as system components change. This particular use case requires neither multiple SP resources nor use over a timescale in which user needs or XSEDE system components change.

user may find the system-specific interfaces offered by the SP to be most relevant/appropriate. (Steps 2,3,6,8,12)

- **XSEDE does not provide significant system features to assist users with steps 4, 5, 7, 10, or 12**, but users should be able to perform these steps without system support.

Steps

- 1. The User studies the Informational Interface to the HPC system to understand the Operational Interface.**

Both XSEDE and the SPs provide websites that contain user documentation, descriptions, and technical specifications about each SP system. The user opens a web browser and navigates to www.xsede.org for basic information about each SP system, and follows links to the SP website for the system he/she is using.

- 2. The User transfers the Application to the HPC System, if not already present.**

We assume that the user's XSEDE identity has already been established [UCCAN-9], and that the user has been granted an allocation to use the SP system [UCCAN-5].

Assuming that the application is not especially large (<1 GB), the user will most likely accomplish this step using the SP's interactive login interface, SSHv2. [L3D §7.3.6] (The SSHv2 interface is an XSEDE service layer interface, specified by XSEDE and operated by each SP. Some XSEDE SPs provide specific SSHv2 services for data transfer that are separate from their login servers. In some cases, these data-oriented SSHv2 services include enhancements to data transfer performance, such as HPN support.) The user opens an SSHv2 client on his/her system, opens a connection to the SP's login server, and authenticates using the SP-specific userid and password. Once SSHv2 access is established, the user can make use of SSHv2's secure copy (SCP) interface to move application files to the SP system.

If the user is already familiar with XSEDE and has used other XSEDE systems, or anticipates using other XSEDE systems in the future, he/she may prefer to use XSEDE's access-layer interfaces instead of the SP-provided SSHv2 interface. XSEDE's interactive login services [UCCAN-4] provide the equivalent of the SSHv2 service interface, and XSEDE's managed file transfer services [UCCAN-2] provide system-wide file transfer services. The primary benefit of these interfaces is that they are consistent across multiple XSEDE systems and are likely to remain the same for future SP systems. However, the SP support staff may not be as familiar with these interfaces as they are with the SP's locally provided SSHv2 interface, described above.

- 3. The User builds, tests, debugs and measures the performance of the Application on the HPC System.**

As in Step 2, the user must login to the SP system, either using the SSHv2 service interface [L3D §7.3.6] or the XSEDE-wide interactive login interfaces [UCCAN-4]. Once logged in, he or she will use the SP system's development interfaces (compilers, profilers, debuggers) to build the application and will use the SP's job submission interface (a resource-layer interface in the XSEDE architecture) to submit computation jobs to the HPC system.

Although XSEDE provides access-layer interfaces for job submission [UCCAN-1], these are most likely not appropriate for application porting, debugging, testing, and performance tuning.

- 4. The User decides which Application tasks (file transfers, jobs submission, output analysis and staging, data management) to perform locally on the HPC system, and which tasks (s)he will execute remotely (e.g. from the XSEDE User Portal, via a service such as Globus Online, or from her/his own workstation).**
- 5. The User decides whether the Application requires any changes for correctness or better job performance ("optimization") on the HPC system.**

The steps above are decisions that the user makes. The user support personnel provided by XSEDE and the SP can assist the user with these decisions, but we assume that the user makes these decisions without direct support by the XSEDE system.

- 6. The User optimizes the Application for the HPC system, until (s)he determines that all the Project goals can be completed within the allocated time and SU budget.**

As in Steps 2 and 3, the user must login to the SP system, either using the SSHv2 service interface [L3D §7.3.6] or the XSEDE-wide interactive login interfaces [UCCAN-4]. Once logged in, he or she will use the SP system's development interfaces (compilers, profilers, debuggers) to optimize the application, using the SP's job submission interface (a resource-layer interface in the XSEDE architecture) to submit computation jobs to the HPC system.

- 7. The User formulates a detailed sequence of Jobs to implement the Project.**

This is an action that the user takes. The user support personnel provided by XSEDE and the SP can assist the user with this work, but we assume that the user takes this action without direct support by the XSEDE system.

- 8. The User begins executing the sequence of Jobs and reports any deviations from the expected behavior to XSEDE user services, who forward them to the HPC SP user support staff if necessary.**

The most obvious choice for users at this point will be to continue logging in to the system and using the local job submission interface as described in Steps 2, 3, and 7. This is simply an extension of the existing use pattern.

However, the user could instead elect to make use of XSEDE's remote job submission interfaces [UCCAN-1]. Assuming that the application is built and optimized on the system and that what remains to be accomplished are a series of application runs, those runs can be submitted from the user's client systems rather than via remote shell. The benefit of this method is that the user's local systems can be involved in submitting, monitoring, and tracking the progress of the application runs and recording any scientifically relevant notes along the way (application run details, configuration notes, etc.). The tradeoff is that this method involves downloading and installing the thick-client software tools for remote job submission, and support staff provided by the SP may not be as familiar with this method as with their local job submission interfaces.

9. XSEDE stakeholders follow up as necessary in diagnosing user problem reports, including conducting tests of their own, forming "tiger teams" with the HPC SP staff, etc.

We assume that XSEDE and SP support staff will use the SSHv2 interface [L3D §7.3.6] for logging in to the SP system and the local job submission interfaces to submit compute jobs when conducting tests in support of the user. The UCCAN-1 interfaces are available, but we assume that support staff will likely prefer to conduct these tests "by hand" rather than involving extra components. The exception to this assumption is if user is using the UCCAN-1 interfaces and the support staff suspects that there may be a problem with the UCCAN-1 interfaces themselves, in which case support staff may choose to replicate the user's complete environment by using exactly the same interfaces as the user.

10. The User cooperates with all HPC SP and XSEDE stakeholder requests to follow up on problem reports, and to investigate specific issues related to the Application and the Jobs which execute it on the HPC system.

This is an action that the user takes. We assume the user takes this action without direct support by the XSEDE system.

11. The User monitors the Informational Interface to the HPC System for any changes in the System that may require changes to the Application or to the Job sequence, including system down times and malfunctions.

As described in Step 1, the websites provided by XSEDE and the SP contain information about system down times, notices about unexpected down time (malfunctions), and other availability conditions.

Each SP system provides local tools for job monitoring and profiling and system utilization and status. (In the XSEDE architecture, these are resource-layer interfaces.) The XSEDE User Portal (XUP) provides a system status interface for monitoring the status and conditions of job queues. (XUP is an access-layer interface in the XSEDE architecture.)

XSEDE also provides a set of information services [UCCAN-7] [UCCAN-8], which may be used for monitoring the status of jobs and the SP system itself. These interfaces are provided primarily to support automated monitoring systems and to facilitate integration with external (non-XSEDE) systems. It is unlikely that scientists on a single XSEDE project will need these interfaces.

12. The User manages the results of the Jobs, and other data relevant to the Project, using the mechanisms provided by the Application.

As stated, this step is largely accomplished using the mechanisms provided by the application. If it becomes necessary to move data from the SP system to the user's local system(s), the same mechanisms described in Step 2 (moving the application to the SP system) can be used to accomplish this data movement.

If data larger than 1 GB must be transferred from the SP system, or if the data is in a form that is difficult to move by hand via SSHv2 clients (e.g., the data is divided into many files), XSEDE's managed data transfer interfaces [UCCAN-2] will be most appropriate for accomplishing this step.

13. The User acknowledges XSEDE and the HPC SP in publications resulting from the Project, as required by the XSEDE User Responsibility Policy. The User cooperates with XSEDE stakeholder requests to generate reports, science stories, etc. based on the results of the Project.

This is primarily an action the user takes without XSEDE system support. However, the XSEDE User Portal (XUP) provides an interface specifically for recording published articles and other works that benefitted from XSEDE services. As described in Steps 1 and 11, the user logs into XUP using his/her XSEDE identity [UCCAN-9] and navigates to the relevant section of the XUP interface to record a publication.

Quality Attributes

The use case description provides several quality attributes. All of them are highly dependent on non-system issues. E.g., selection of well-prepared users and likely-to-succeed projects during the allocation process, availability of the HPC system during operational periods, performance of the user support staff, etc. There do not appear to be any quality attributes in this use case that the system alone could prevent us from achieving.

UCHPC-2 User project on more than one HPC system

This use case depends heavily on SP-provided HPC resources. Due to the involvement of multiple SP-provided systems, however, *all* of XSEDE's system-wide functions are relevant. Interfaces provided by the SPs are used for submitting and managing computations (particularly during the application preparation, debugging, and optimization phases) and discovering and providing resource information (particularly the application preparation and monitoring phases). The remainder of the use case can be accomplished using XSEDE access-layer interfaces, taking advantage of XSEDE-wide functions.

	Identity Management	Interactive Login	Accessing Remote Files	Submitting and Managing Computations	Transferring Datasets	Discovering and Providing Resource Information
Access	2,3,6,8,9,10,13,14	3,6,8,9,13	8,9,13	(8),9	2,9,13	1,9,12,14
Service		(2),(3),(6),(8),10,(13)				
Resource				3,6,8,10	(2)	1,12

Figure 2. This figure shows which elements of the XSEDE architecture [ARCH §D] are engaged during each step of UCHPC-2. Steps numbers in parentheses are alternatives to the primary interface.

Steps

- 1. The User studies the Informational Interface to the HPC systems to understand the Operational Interfaces.**

Both XSEDE and the SPs provide websites that contain user documentation, descriptions, and technical specifications about each SP system. The user opens a web browser and navigates to www.xsede.org for basic information about each SP system, and follows links to the SP website for the system he/she is using.

- 2. The User transfers the Applications to the appropriate HPC Systems.**

We assume that the user's XSEDE identity has already been established [UCCAN-9], and that the user has been granted an allocation to use the SP system [UCCAN-5].

Because this use case involves multiple SP-provided systems, the user will most likely prefer a common interface provided by XSEDE (particularly one that recognizes their XSEDE userid and

password) for accomplishing this step. XSEDE's interactive login services [UCCAN-4] provide the ability to login to each SP-provided system using one's XSEDE userid and password, and XSEDE's managed file transfer services [UCCAN-2] provide system-wide file transfer services, also using one's XSEDE userid and password. The primary benefit of these interfaces is that they are consistent across multiple XSEDE systems and are likely to remain the same for future SP systems. However, the SP support staff may not be as familiar with these interfaces as they are with the SP's locally provided SSHv2 interface, described below.

Assuming that the application is not especially large (<1 GB), the user could also accomplish this step using each SP's interactive login interface, SSHv2. [L3D §7.3.6] (The SSHv2 interface is an XSEDE service layer interface, specified by XSEDE and operated by each SP. Some XSEDE SPs provide specific SSHv2 services for data transfer that are separate from their login servers. In some cases, these data-oriented SSHv2 services include enhancements to data transfer performance, such as HPN support.) The user opens an SSHv2 client on his/her system, opens a connection to each SP's login server, and authenticates using the SP-specific userid and password. (This userid and password will be different for each SP and could even be different for multiple systems provided by the same SP.) Once SSHv2 access is established, the user can make use of SSHv2's secure copy (SCP) interface to move application files to each SP system.

3. The User builds, tests, debugs and measures the performance of the Applications on the HPC Systems.

As in Step 2, the user must login to each SP system, either using the XSEDE-wide interactive login interfaces [UCCAN-4] or the SSHv2 service interface [L3D §7.3.6]. Once logged in, he or she will use each SP system's development interfaces (compilers, profilers, debuggers) to build the application and will use each SP's job submission interface (a resource-layer interface in the XSEDE architecture) to submit computation jobs to the HPC system.

Although XSEDE provides access-layer interfaces for job submission [UCCAN-1], these are most likely not appropriate for application porting, debugging, testing, and performance tuning.

- 4. The User decides which Application tasks (file transfers, jobs submission, output analysis and staging, data management) to perform locally on each HPC system, and which tasks (s)he will execute remotely (e.g. from the XSEDE User Portal or from her/his own workstation).**
- 5. The User decides whether the Application requires any changes for correctness or better job performance ("optimization") on one or more of the HPC systems.**

The steps above are decisions that the user makes. The user support personnel provided by XSEDE and the SP can assist the user with these decisions, but we assume that the user makes these decisions without direct support by the XSEDE system.

6. The User optimizes the Application for the HPC systems, until (s)he determines that all the Project goals can be completed within the allocated time and SU budget.

As in Steps 2 and 3, the user must login to each SP system, either using the XSEDE-wide interactive login interfaces [UCCAN-4] or the SSHv2 service interface [L3D §7.3.6]. Once logged in, he or she will use each SP system's development interfaces (compilers, profilers, debuggers) to optimize the application, using each SP's job submission interface (a resource-layer interface in the XSEDE architecture) to submit computation jobs to the HPC systems.

As in Step 3, it is likely that XSEDE's access-layer interfaces for job submission [UCCAN-1] lack features needed during the application optimization phase.

7. The User formulates a detailed sequence of Jobs to implement the Project.

This is an action that the user takes. The user support personnel provided by XSEDE and each SP can assist the user with this work, but we assume that the user takes this action without direct support by the XSEDE system.

8. The User begins executing the sequence of Jobs and reports any deviations from the expected behavior to XSEDE user services, who forward them to the relevant HPC SP user support staff if necessary.

The most obvious choice for users at this point will be to continue logging in to the system and using the local job submission interface as described in Steps 2, 3, and 7. This is simply an extension of the existing use pattern.

However, the user could instead elect to make use of XSEDE's remote job submission interfaces [UCCAN-1]. Assuming that the application is built and optimized on the system and that what remains to be accomplished are a series of application runs, those runs can be submitted from the user's client systems rather than via remote shell. The benefit of this method is that the user's local systems can be involved in submitting, monitoring, and tracking the progress of the application runs and recording any scientifically relevant notes along the way (application run details, configuration notes, etc.). The tradeoff is that this method involves downloading and installing the thick-client software tools for remote job submission, and support staff provided by the SP may not be as familiar with this method as with their local job submission interfaces.

If the applications running on separate SP systems require access to a common set of files (e.g., for configuration or input data), XSEDE's remote file access interface [UCCAN-3] may be used to provide access to a shared filesystem.

9. **The SPs are as consistent as possible considering compatibility with each individual machine. E.g. it makes no sense to ask all SPs to run the same compiler or MPI library, but it makes sense to ask all SPs to use "modules", to provide consistent path names for libraries and 3rd party applications, to provide consistent end points for GridFTP / Globus Online, etc. ("common user environment").**

The XSEDE architecture provides the following system-wide functions that provide consistency across SP systems.

- XSEDE-wide identities [UCCAN-9], which are used to provide a single userid/password for all of the functions below.
- Remote job submission interface [UCCAN-1], which provides a consistent mechanism for submitting computation/application tasks to SP-systems. (Note that this is an *alternative* to the primary interface recommended for this use case, which *is not* consistent across SP systems.)
- Managed file transfer interface [UCCAN-2], which provides a consistent mechanism for moving files to/from/between XSEDE SP-provided systems.
- Remote file access interface [UCCAN-3], which provides a shared filesystem that can be accessed from XSEDE SP-provided systems.
- XSEDE's SSO (single sign-on) Hub, which provides an interactive login interface [UCCAN-4].
- XSEDE User Portal (XUP), which provides a queue-monitoring interface, the resource allocation system interface [UCCAN-5], user documentation, and an interface for allowing users to register any publications that were supported by XSEDE.
- Information services [UCCAN-7] [UCCAN-8], which provide mechanisms for monitoring and querying the status of SP-provided systems. (Note that these are *alternatives* to the primary interfaces recommended for this use case, which *are not* consistent across SP systems.)

XSEDE SPs also coordinate further details of their systems beyond those listed above, such as the tool used to customize the shell environment on their systems ("modules"), common paths and environment variables, and so on.

10. **XSEDE stakeholders follow up as necessary in diagnosing user problem reports, including conducting tests of their own, forming "tiger teams" with the HPC SP staff, etc.**

We assume that XSEDE and SP support staff will use the SSHv2 interface [L3D §7.3.6] for logging in to the SP system and the local job submission interfaces to submit compute jobs when conducting tests in support of the user. The UCCAN-1 interfaces are available, but we assume that support staff will likely prefer to conduct these tests "by hand" rather than involving extra components. The exception to this assumption is likely to be if the support staff assumes

that there is a problem with the UCCAN-1 interfaces themselves, in which case support staff may choose to replicate the user's complete environment by using exactly the same interfaces used by the user.

11. The User cooperates with all HPC SP and XSEDE stakeholder requests to follow up on problem reports, and to investigate specific issues related to the Application and the Jobs that execute it on the HPC systems.

The user takes this action without direct support by the XSEDE system.

12. The User monitors the Informational Interface to the HPC Systems for any changes in the System that may require changes to the Application or to the Job sequence, including system down times and malfunctions.

As described in Step 1, the websites provided by XSEDE and the SP contain information about system down times, notices about unexpected down time (malfunctions), and other availability conditions.

Each SP system provides local tools for job monitoring and profiling and system utilization and status. (In the XSEDE architecture, these are resource-layer interfaces.) The XSEDE User Portal (XUP) provides a system status interface for monitoring the status and conditions of job queues. (XUP is an access-layer interface in the XSEDE architecture.)

XSEDE also provides a set of information services [UCCAN-7] [UCCAN-8], which may be used for monitoring the status of jobs and the SP system itself. These interfaces are provided primarily to support automated monitoring systems and to facilitate integration with external (non-XSEDE) systems. It is unlikely that scientists on a single XSEDE project will need these interfaces.

13. The User manages the results of the Jobs, and other data relevant to the Project, using the mechanisms provided by the Application.

As stated, this is largely accomplished using the mechanisms provided by the application. If it becomes necessary to move data from the SP system to the user's local system(s), the same mechanisms described in Step 2 (moving the application to the SP system) can be used to accomplish this.

If the applications running on separate SP systems require access to a common set of files (e.g., for configuration or input data or for collecting output data), XSEDE's remote file access interface [UCCAN-3] may be used to provide access to a shared filesystem.

14. The User acknowledges XSEDE and the HPC SPs in publications resulting from the Project, as required by the XSEDE User Responsibility Policy. The User

cooperates with XSEDE stakeholder requests to generate reports, science stories, etc. based on the results of the Project.

This is primarily an action the user takes without XSEDE system support. However, the XSEDE User Portal (XUP) provides an interface specifically for recording published articles and other works that benefitted from XSEDE services. As described in Steps 1 and 12, the user logs into XUP using his/her XSEDE identity [UCCAN-9] and navigates to the relevant section of the XUP interface to record a publication.

Quality Attributes

The use case description provides several quality attributes. All of them are highly dependent on non-system issues. E.g., selection of well-prepared users and likely-to-succeed projects during the allocation process, availability of the HPC system during operational periods, performance of the user support staff, etc. There do not appear to be any quality attributes in this use case that the system alone could prevent us from achieving.

UCHPC-3 Using a scientific workflow system to run on HPC systems

This use case depends heavily on SP-provided HPC resources. Due to the involvement of an automated workflow system, however, *all* of XSEDE's system-wide functions are relevant. Interfaces provided by the SPs are used during the application preparation, debugging, and optimization phases for submitting and managing computations and discovering and providing resource information. The remainder of the use case can be accomplished using XSEDE access-layer interfaces, taking advantage of XSEDE-wide functions.

	Identity Management	Interactive Login	Accessing Remote Files	Submitting and Managing Computations	Transferring Datasets	Discovering and Providing Resource Information
Access	3,5,6,8,10	5,6,8,15	6,8,10,12,15	(6),(8),10,12	3,6,8,10,12,15	1,14,16
Service		(3),(5),(6),(8),12,(15)				
Resource				6,8,12	(3)	1,14

Figure 3. This figure shows which elements of the XSEDE architecture [ARCH §D] are engaged during each step of UCHPC-3. Steps numbers in parentheses are alternatives to the primary interface.

Steps

- 1. The User studies the Informational Interface to the HPC systems to understand the Operational Interfaces.**

Both XSEDE and the SPs provide websites that contain user documentation, descriptions, and technical specifications about each SP system. The user opens a web browser and navigates to www.xsede.org for basic information about each SP system, and follows links to the SP website for the system he/she is using.

- 2. The User selects a Workflow System that will interact programmatically with the Operational and Informational Interfaces.**

This is a decision made by the XSEDE user. However, the XSEDE community and XSEDE user support services have experience with a wide variety of workflow systems, several of which have been used to manage XSEDE user projects. Taking advantage of this professional

experience may be extremely helpful to the user. The XUP's user support interfaces are most likely the best route to obtain advice based on past XSEDE experiences.

3. The User transfers the Application (including the relevant components of the Workflow System) to the HPC Systems.

Because this use case will eventually involve job submission by a workflow system, the user may find it helpful to begin using XSEDE's system-wide interfaces as soon as possible (particularly those that recognize their XSEDE userid and password). XSEDE's interactive login services [UCCAN-4] provide the ability to login to each SP-provided system using one's XSEDE userid and password, and XSEDE's managed file transfer services [UCCAN-2] provide system-wide file transfer services, also using one's XSEDE userid and password. The primary benefit of these interfaces is that they prepare the user for interfaces that will be used in later steps of this use case. However, the SP support staff may not be as familiar with these interfaces as they are with the SP's locally provided SSHv2 interface, described below.

Assuming that the application is not especially large (<1 GB), the user could also accomplish this step using each SP's interactive login interface, SSHv2. [L3D §7.3.6] (The SSHv2 interface is an XSEDE service layer interface, specified by XSEDE and operated by each SP. Some XSEDE SPs provide specific SSHv2 services for data transfer that are separate from their login servers. In some cases, these data-oriented SSHv2 services include enhancements to data transfer performance, such as HPN support.) The user opens an SSHv2 client on his/her system, opens a connection to each SP's login server, and authenticates using the SP-specific userid and password. (This userid and password will be different for each SP and could even be different for multiple systems provided by the same SP.) Once SSHv2 access is established, the user can make use of SSHv2's secure copy (SCP) interface to move application files to each SP system.

4. The User decides how best to use the Workflow System to programmatically launch and manage HPC Jobs on the HPC Systems.

XSEDE's information services [UCCAN-7] [UCCAN-8] provide programmatic interfaces to service availability and status information. These interfaces may be used by workflow systems as inputs to resource selection algorithms.

XSEDE's remote job submission mechanism [UCCAN-1] is the recommended mechanism for remote and/or automated execution of HPC jobs on XSEDE systems. As described in UCCAN-1, this involves installing (or using a previous installation of) the UNICORE or Genesis II client software on the system where the workflow system executes and configuring the workflow system to use the UNICORE or Genesis II client to launch jobs.

Note that the user is free to choose other interfaces (e.g., the SSHv2 service-layer interface [L3D §7.3.6] combined with each SP's job submission interface) to submit computation jobs to

the HPC system. We discourage this practice because the UCCAN-1 mechanisms handle a wide variety of user and system needs that most users will not want to assume responsibility for. Past experience has demonstrated that these user and system requirements are both important and difficult for individual users to manage on their own. It is easy for them to become a distraction from the user's scientific work. We highly recommend leveraging existing interfaces to mitigate these issues rather than taking them on individually.

5. The User configures the Workflow System accordingly.

If the workflow system is executed on the user's systems (personal or shared computers at his/her home institution), the user must accomplish step 5 using locally supplied mechanisms. If the workflow system is executed on an XSEDE system (presumably one provided by an XSEDE SP), the user may use XSEDE's interactive login interface [UCCAN-4] to login remotely to the system and configure the workflow components.

6. The User builds, tests, debugs and measures the performance of the Application on the HPC Systems. This includes the ratio of the number of scientifically significant results produced per unit time with and without the Workflow System.

As in Step 3, the user must login to each SP system, either using the XSEDE-wide interactive login interfaces [UCCAN-4] or the SSHv2 service interface [L3D §7.3.6]. Once logged in, he or she will use each SP system's development interfaces (compilers, profilers, debuggers) to build the application and will use each SP's job submission interface (a resource-layer interface in the XSEDE architecture) to submit computation jobs to the HPC system.

Although XSEDE provides access-layer interfaces for job submission [UCCAN-1], these are most likely not appropriate for application porting, debugging, testing, and performance tuning.

7. The User decides whether the Application (including the Workflow System) requires any changes for correctness or better job performance ("optimization") on one or more of the HPC systems.

The user makes this decision. The user support personnel provided by XSEDE and the SP can assist the user, but we assume that the user makes these decisions without direct support by the XSEDE system.

8. The User optimizes the Application (including the relevant configuration or features of the Workflow System) for the HPC systems, until (s)he determines that all the Project goals can be completed within the allocated time and SU budget.

As in Steps 3 and 6, the user must login to each SP system, either using the XSEDE-wide interactive login interfaces [UCCAN-4] or the SSHv2 service interface [L3D §7.3.6]. Once logged in, he or she will use each SP system's development interfaces (compilers, profilers,

debuggers) to optimize the application, using each SP's job submission interface (a resource-layer interface in the XSEDE architecture) to submit computation jobs to the HPC systems.

As in Step 6, it is likely that XSEDE's access-layer interfaces for job submission [UCCAN-1] lack features needed during the application optimization phase.

9. The User formulates a detailed sequence of Jobs to implement the Project.

The user makes this decision. The user support personnel provided by XSEDE and the SP can assist the user, but we assume that the user makes these decisions without direct support by the XSEDE system.

10. The User begins executing the sequence of Jobs and reports any deviations from the expected behavior to XSEDE user services, who forward them to the relevant HPC SP user support staff if necessary.

As established in step 9, executing the project's jobs is a matter of instructing the workflow system to begin a set of application jobs. The workflow system will use the mechanism set up in steps 4 and 5 to submit (and presumably monitor) each individual job until all of the jobs are complete.

If the jobs running on separate SP systems require access to a common set of files (e.g., for configuration or input data), XSEDE's remote file access interface [UCCAN-3] may be used to provide access to a shared filesystem.

11. If the User has evidence, from the Workflow System's monitoring component, that the Workflow System deviates from its expected behavior, (s)he includes this information in the report to XSEDE user services, as well as to the Support Providers for the Workflow System (if these are external to XSEDE).

The user takes this action without direct support by the XSEDE system.

12. XSEDE stakeholders follow up as necessary in diagnosing user problem reports, including conducting tests of their own, forming "tiger teams" with the Support Providers for the Workflow System, HPC SP staff, etc.

We assume that XSEDE and SP support staff will use the SSHv2 interface [L3D §7.3.6] for logging in to the SP system and the local job submission interfaces to submit individual compute jobs when conducting tests in support of the user.

If the support staff suspects a problem with the UCCAN-1 job submission interface or with the user's workflow system, the support staff may replicate the user's complete environment by

using exactly the same interfaces used by the user (including the user's workflow system), as covered in steps 4, 5, 9, and 10.

13. The User cooperates with all HPC SP, XSEDE stakeholder and Workflow System provider requests to follow up on problem reports, and to investigate specific issues related to the Application and the Jobs that execute it on the HPC systems.

The user takes this action without direct support by the XSEDE system.

14. The User monitors the Informational Interface to the HPC Systems for any changes in the System that may require changes to the Application (including the configuration of the Workflow System) or to the Job sequence, including system down times and malfunctions.

As described in Step 1, the websites provided by XSEDE and the SP contain information about system down times, notices about unexpected down time (malfunctions), and other availability conditions. XSEDE's information services [UCCAN-7] [UCCAN-8] also provide programmatic access to this information.

XSEDE's information services also provide interfaces for monitoring the status of jobs and the SP systems themselves. Some workflow systems can use these interfaces to perform automated monitoring on the user's behalf. This would be done as part of steps 4 and 5. If so, the workflow system may be configured in step 9 to respond to deviations from normal behavior, or it could be configured to simply notify the user of these deviations.

In addition to XSEDE's information service interfaces, each SP system provides local tools for job monitoring and profiling and system utilization and status. (In the XSEDE architecture, these are resource-layer interfaces.) The XSEDE User Portal (XUP) also provides a system status interface for monitoring the status and conditions of job queues. (XUP is an access-layer interface in the XSEDE architecture.)

15. The User manages the results of the Jobs, and other data relevant to the Project, using the mechanisms provided by the Application.

As stated, this is largely accomplished using the mechanisms provided by the application. If it becomes necessary to move data from the SP system to the user's local system(s), the same mechanisms described in Step 3 (moving the application to the SP system) can be used to accomplish this.

If the applications running on separate SP systems require access to a common set of files (e.g., for configuration or input data or for collecting output data), XSEDE's remote file access interface [UCCAN-3] may be used to provide access to a shared filesystem.

16. The User acknowledges XSEDE and the HPC SPs in publications resulting from the Project, as required by the XSEDE User Responsibility Policy. The User cooperates with XSEDE stakeholder requests to generate reports, science stories, etc. based on the results of the Project. In particular, use of an XSEDE supported Workflow System should be acknowledged.

The user takes this action without XSEDE system support. However, the XSEDE User Portal (XUP) provides an interface specifically for recording published articles and other works that benefitted from XSEDE services. As described in Steps 1 and 12, the user logs into XUP using his/her XSEDE identity [UCCAN-9] and navigates to the relevant section of the XUP interface to record a publication.

17. The User will also want to acknowledge any non-XSEDE Workflow System Provider in the resulting publications.

The user takes this action without XSEDE system support.

Quality Attributes

The use case description provides several quality attributes. All of them are highly dependent on non-system issues. E.g., selection of well-prepared users and likely-to-succeed projects during the allocation process, availability of the HPC system during operational periods, performance of the user support staff, etc. There do not appear to be any quality attributes in this use case that the system alone could prevent us from achieving.