

XSEDE

Extreme Science and Engineering
Discovery Environment

XSEDE Allocations Practices and Procedures

23 October 2019

Version 1.0

Table of Contents

B.	Document History	5
C.	Document Scope	6
D.	Background.....	7
E.	Policies and Related Information	8
F.	XSEDE Allocations Coordinator and Staff.....	8
G.	Service Provider Coordination	8
H.	XSEDE Resource Allocation Committee	9
I.	Allocations in a Heterogeneous Federation.....	10
J.	Research Allocations.....	12
J.1.	Conflicts of Interest	12
J.2.	Basic quarterly cycle	12
J.3.	Quarterly schedule of tasks.....	13
J.3.1.	Before and during submission window.....	13
J.3.2.	Between submission deadline and meeting.....	15
J.3.3.	During the meeting.....	16
J.3.4.	After the meeting.....	18
J.4.	XRAC meeting logistics.....	19
J.4.1.	Site selection	20
J.4.2.	Travel coordination.....	20
J.4.3.	On-site arrangements.....	21
J.4.4.	Meeting communications	21
J.4.5.	Absentee reviewers.....	21
J.4.6.	Meeting budget	22
J.4.7.	Emergency plan	22
K.	Startup Allocations & Educational Allocations	23
K.1.	Validating submissions.....	23
K.2.	Assigning reviewers	24
K.3.	Finalizing awards	24
L.	Other Allocation Types	24
L.1.	Campus Champions.....	25
L.2.	Campus Single-Sign On.....	25
L.3.	Discretionary	25

L.4.	Others	25
M.	Allocation Management Requests	25
M.1.	Extensions.....	26
M.2.	Transfers	26
M.3.	Supplements	26
M.4.	Appeals	26
M.5.	Advances	26
M.6.	Final Reports.....	27
N.	Managing Authorized Persons and Other Users	27
N.1.	Authorized Roles.....	27
N.2.	Verifying Principal Investigators	28
N.3.	Adding and Removing Users	28
N.4.	Re-verifying Users.....	28
N.5.	Changing PIs.....	28
N.6.	Changing Co-PIs and Allocation Managers	29
O.	Appendix: Service Units, Normalized Units and Weighting Factors	29
O.1.	Service Units	30
O.2.	The LINPACK Benchmark	32
O.3.	The “Normalized CPU-Hour” and Weighting Factors.....	32
O.4.	Limitations of the Normalized Unit and Weighting Factor	33
O.4.1.	Application-Specific Weighting Factors	34
O.4.2.	Alternate Weighting Factors	34
O.4.3.	TeraGrid SUs (TG SUs)	34
O.5.	Dollar Cost per Allocated Unit.....	35
P.	Appendix: Reconciliation Formulas and Calculations	35
P.1.	Reconciliation Formula.....	36

List of Figures

Figure 1: A canonical process outline from HPCWorld Consortium (2011), “Handbook of HPC e-Science Infrastructure Allocation Reviewing, Selection and Management.”.....	7
Figure 2: XSEDE has had to adapt its allocation practices and procedures to ensure timely response to requests that in recent years have regularly surpassed 400 Research and Startup requests per quarter.	8
Figure 3. Sample language provided in XSEDE allocation award letters when describing the use and impact of the reconciliation process.....	19
Figure 4. Monthly usage reported to the XSEDE central accounting system across all XSEDE-allocated computing systems over the last five years.....	30

B. Document History

Relevant Sections	Version	Date	Changes	Author
Entire Document	1.0	23 Oct 2019	First version	Ken Hackworth, Stephen Deems, David Hart

C. Document Scope

This document covers the tasks and activities required to administer the XSEDE allocations process. The document focuses on the tasks and efforts required by XSEDE staff to translate, interpret and implement the XSEDE allocation policies and provide researchers with allocations on resources that will help them accomplish their science objectives.

The goal for the XSEDE allocations process can be summarized as ensuring that the cyberinfrastructure portfolio is used as efficiently as possible to produce the best science outcomes. This documents the current procedures XSEDE uses in pursuit of that goal.

The contents of this document are related to and supported by several other documents. Where relevant to this document, these other documents are referenced but not repeated here.

- The [XSEDE Allocations Policies](#) covers the guiding principles and stakeholder requirements for XSEDE allocations. These policies are interpreted, implemented and enforced by XSEDE with respect to allocations on Service Provider resources.
- The [XRAC Reviewer Manual](#)^{*} describes the charter, practices, and tasks required of the members of the XSEDE Resource Allocation Committee (XRAC), from a reviewer's perspective.
- The **XRAS Administrator's Guide** *[document in development]* describes how to use the submission, review, and administrative features of the XRAS system, which is the tool used by XSEDE to support the allocations process.

^{*} <https://www.ideals.illinois.edu/handle/2142/99774>

D. Background

As one of its areas of responsibility, XSEDE manages the processes that support the submission, review, and administration of allocation requests for a portfolio of cyberinfrastructure resources funded by the National Science Foundation (NSF). At a high level, the XSEDE procedures and policies align with the canonical outline of many allocations activities (Figure 1). Thus, as the primary funding agency, NSF lays the groundwork for XSEDE allocations practices by defining the stakeholder goals for the policies and practices (Figure 1, Step 1.1). The XSEDE Allocations Policies are reviewed and approved by NSF, as are these allocations practices and procedures.

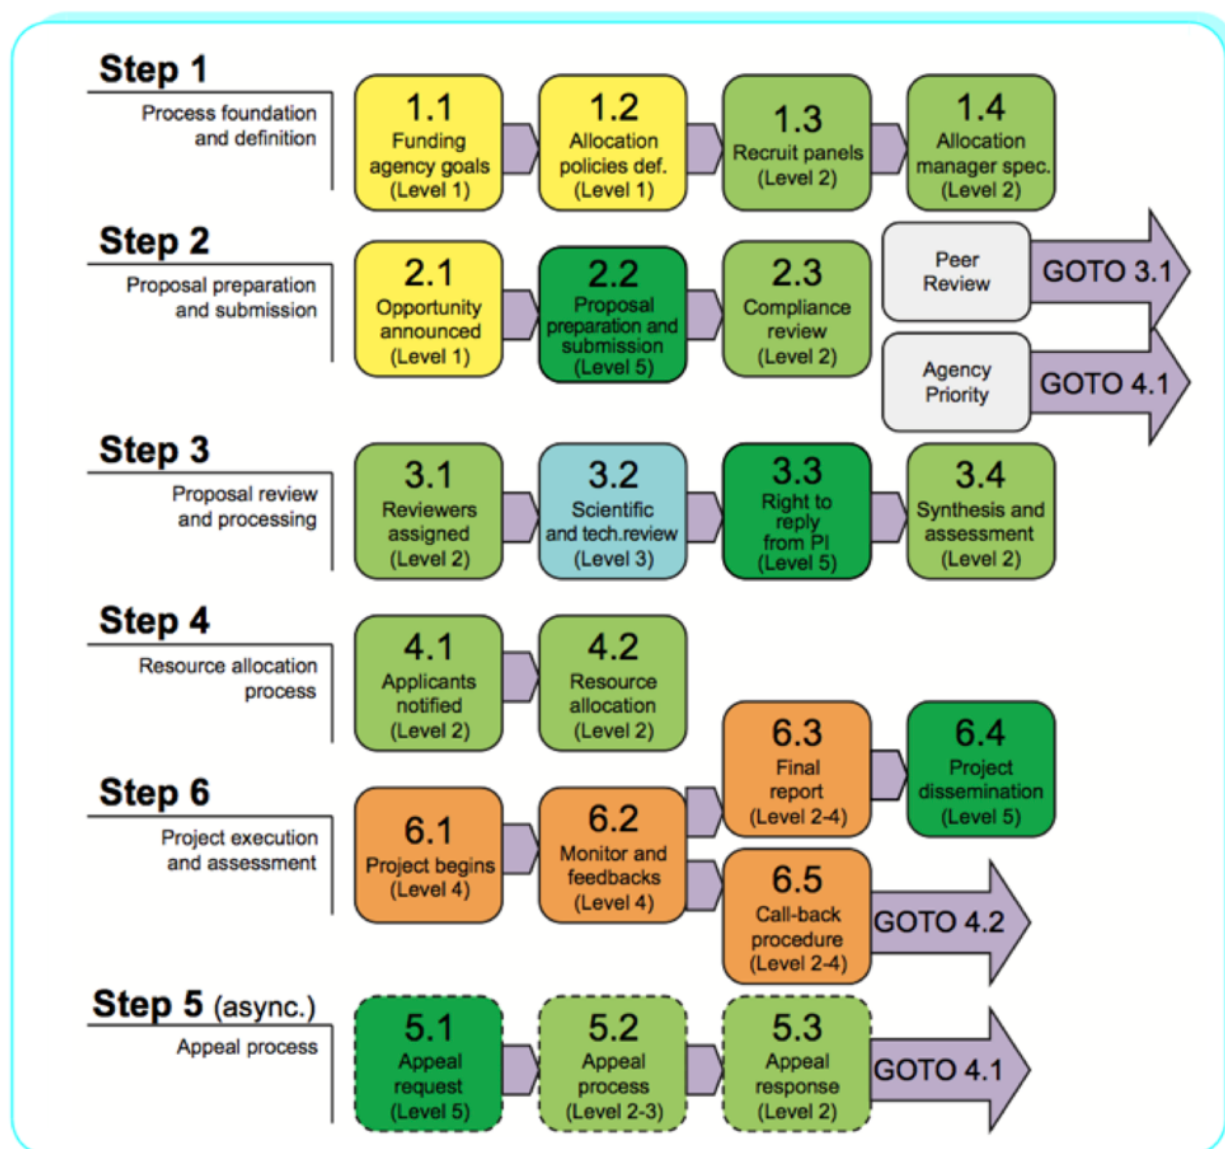


Figure 1: A canonical process outline from HPCWorld Consortium (2011), “Handbook of HPC e-Science Infrastructure Allocation Reviewing, Selection and Management.”

E. Policies and Related Information

The *XSEDE Allocation Policies* are defined in a separate document, which we include here by reference (Figure 1, Step 1.2). In summary, however, XSEDE provides access to resources in high-performance computing, scientific visualization, data storage, and extended collaborative support resources, to enhance scientific discovery. The policies and these procedures are followed to ensure a fair and efficient allocation of these resources. XSEDE encourages and supports projects in all disciplines for a wide range of purposes.

F. XSEDE Allocations Coordinator and Staff

Within the guiding principles established by NSF and the policies approved by NSF for XSEDE to implement, XSEDE staff in the Resource Allocations Service (RAS) area are charged with developing and applying practices to ensure that more than 400 allocation requests each quarter are submitted, reviewed appropriately, and processed in a timely fashion (Figure 2).

The RAS Allocation Policies and Procedures (APP) group, comprises four supported staff (funded at a level of 2.6 FTEs). XSEDE's Allocation Coordinator is the APP manager, supported by three persons in completing day-to-day allocations tasks and allocation meeting support. The Allocation Coordinator has primary responsibility and authority for interpreting the allocation policies. When necessary, policy questions can be escalated through the RAS Director, the XSEDE Principal Investigator, and even up to XSEDE's NSF Program Officer.

G. Service Provider Coordination

Effectively managing the XSEDE allocations process requires effective coordination between XSEDE's Allocation Coordinator and the Service Providers (SPs) whose resources XSEDE allocates.

The XSEDE Allocations Coordinator works directly with SP representatives to review all incoming resource requests. For Startup requests, SP representatives serve as the primary reviewers, to ensure the request is appropriate for the resource. For Research

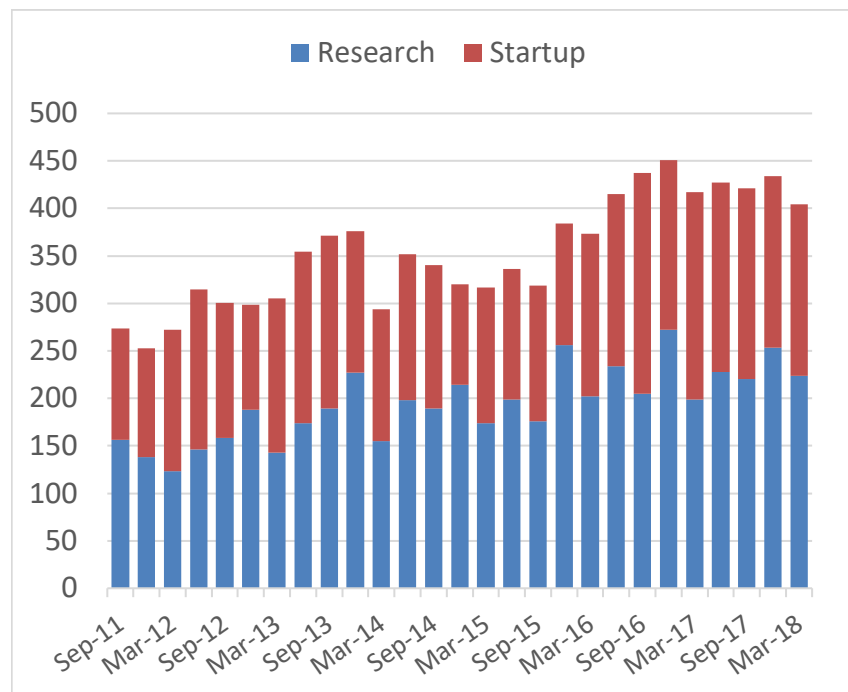


Figure 2: XSEDE has had to adapt its allocation practices and procedures to ensure timely response to requests that in recent years have regularly surpassed 400 Research and Startup requests per quarter.

requests, after each submission window has closed, the Allocations Coordinator engages SP representatives in a correctness and suitability review phase. Due to the changing nomenclatures of resource offerings, PIs will occasionally request resources in the wrong units or ask for resources that are not well suited for their planned work. The allocations team and the SPs will review the submitted documents and accompanying text fields to search for the “true” resource requirements. If additional clarification is needed, the allocations staff or SP representatives will directly contact the submitting research group. Working directly with the SPs also allows the allocations team to raise issues such as large over-request levels and to ensure SPs are engaged in the allocation cycle.

XSEDE’s SP Coordinator also plays a role in supporting SPs’ involvement in the allocations process. Allocations process integration is covered by the SP Coordinator and the SP representatives as part of the annual review of the [Service Provider Checklist](#).

H. XSEDE Resource Allocation Committee

To review larger scale Research requests, XSEDE convenes the XSEDE Resource Allocation Committee (XRAC), which is made up of approximately 60 computational scientists of varying scientific backgrounds from cosmology to computational fluid dynamics (Figure 1, Step 1.3). The panel membership is focused on ensuring that each field of science has a sufficient number of reviewers to handle the requests received during a typical meeting cycle. Thus, areas such as molecular biology, materials research, and chemistry (for example) have more reviewers, while areas with fewer requests (e.g., economics) may be assigned to a group of “general purpose” computational experts. This approach has proven feasible since the XSEDE policies and practices focus on a technical/computational review of the requests. In addition to maintaining a suitable domain balance, the Allocations Coordinator also considers panel diversity more broadly and seeks to recruit members that provide the XRAC with diversity across institutions, institution types, states, academic and non-academic participation, gender and race.

For any given quarterly meeting between 40 to 50 reviewers attend with some reviewing remotely if they are not able to make the trip. In addition to the XRAC proper, XSEDE also enlists the help of staff from the XSEDE Extended Collaborative Support Service (ECSS) to review the smaller requests for XSEDE resources (see Section J.3.2).

XRAC members are asked to serve a three-year term, attending at least three of the four quarterly meetings. Members may be invited to serve more than one term. All XRAC members are volunteers, and XSEDE does not provide honoraria, although travel expenses to attend meetings are paid by XSEDE.

Turnover for panel members varies. New panel members occasionally opt out of serving on the panel due to too many commitments at their home institution or in their personal lives. Occasionally, panel members drop off due to changing institutions or moving from academia into industry. With that being said, some panel members do maintain their activity during such

transitions. XSEDE works to limit XRAC members to no more than two consecutive terms, though exceptions are made.

Recruitment efforts for panel members come in various forms including, but not limited to: recommendations from current/past panel members, recommendations by site representatives and service providers, a call for recommendations by the Allocations Coordinator, direct requests by current PIs, or the Allocations Manager asking a current PI to join the panel (particularly in an area that is currently under-represented in the panel). The XSEDE web site includes a standing invitation to persons interested in joining the XRAC, along with a one-page flyer, which is also distributed at events where XSEDE has a presence. A single presentation slide is also available to all XSEDE staff to include in various talks.

The specific responsibilities and tasks expected of XRAC members are described in the *XRAC Reviewer Manual* and not discussed further here.

I. Allocations in a Heterogeneous Federation

XSEDE allocations span a range of Service Providers, resources types, and resource architectures, providing common policies and practices and permitting financial efficiencies for allocations activities across the NSF-supported resource portfolio. In general, common practices make life easier for users, allowing them to request allocations from a central point, rather than having to apply separately with different Service Providers. At the same time, in managing allocation practices and procedures across a federation, we encounter a number of situations in which we need to assess the size of one allocation request or award to another request or award, to compare the size of an allocation request to some fixed size threshold, or to translate an allocation amount from one resource into an equivalent amount on another resource.

Prior to and through most of the XSEDE program, the inherent challenges in managing allocations across a federation of resources have been greatly simplified since almost all resources were HPC resources, most resources had a homogeneous architecture, and the “processor core” was easy to identify as a common architectural element across all resources. Thus, for several decades, XSEDE and its predecessor programs have dealt with cross-resource allocations by adjusting for performance differences between different types of cores, when needed (see Appendix O). Despite some limitations, this so-called “normalization” approach that has served the community well for decades. It also continues to allow us to make comparisons between today’s resources and historical resources.

In practice, however, most day-to-day activities take advantage of the fact that almost all of the resources being allocated use the notion of “Service Units” (SUs), which over decades of use became synonymous with “core-hours.”² More specifically, SUs tend to be core-hours on a conventional parallel cluster, even as processors have grown from one to dozens of cores.

² Originally, the concept of Service Units was introduced to allow for far more complicated units or formulas for allocation, which may have included memory or disk storage, dating back to the vector machine era. However, with the advent of parallel computers, Service Units became equated with “core-hours,” and that equivalence has proved difficult to walk back.

Thus, most of the time, summing or comparing total SUs has offered a reasonable and simple approach for a person (user or allocations staff) to compare allocations, and when greater precision is necessary—such as when making transfers from one resource to another—normalization factors could be applied. To help users and staff with comparing amounts between two resources, XSEDE provides an [SU Converter](#) tool in the user portal. This tool applies appropriate normalization factors to help users in requesting transfers, for example, and in preparing allocation requests.³ (Exceptions have arisen from time to time, but such systems usually were a small fraction of the overall portfolio and required only minor adjustments to common practices and procedures.)

However, the simple approach relies on one key assumption: namely, that the relative values of allocable units across most of the federated resource portfolio are “similar enough.” This assumption was finally broken with the arrival of the TACC Stampede2 cluster, which, due to its heterogeneous node architecture, has been allocated in “node-hours,” not core-hours. At roughly the same time, the popularity of GPGPU computing within XSEDE skyrocketed, leading to high demand for systems allocated in “GPU-hours.” Most obviously, Stampede2 node-hour requests (as well as requests for GPU-hours) are roughly two orders of magnitude smaller than comparable core-hour requests on other resources.

In fact, in 2019, XSEDE is now allocating resources in more than a half dozen different units:

- Node-hours (Stampede2)
- Core-hours (Bridges, Comet, Open Science Grid)
- GPU-hours (Bridges GPU, Bridges GPU-AI, Comet GPU)
- TB-hours (Bridges Large Memory)
- Virtual CPU-hours (Jetstream)
- Gigabytes (various storage resources)
- “Yes/No” (XSEDE ECSS support, SGCI support)

As of 2019, XSEDE continues to manage its allocations process using *ad hoc* modifications to the longstanding normalization methods for comparing and contrasting allocations. For example, in the reconciliation process (see section J.3.4), GPU-based systems are handled separately from the systems with conventional processors; we assume that requests for GPU systems can be moved to other GPU systems, but not easily moved between conventional core systems.

However, we are taking some early steps to overhaul these practices with a more formal and automatable approach leveraging a more “universal” standard based on the dollar cost of allocated units (see Appendix O.5). We have selected this approach for its practical benefits to the allocations process, including its simplicity in translating complex practical realities for the

³ While the SU Converter can convert to “Normalized Units,” most users are spared that detail. The SU Converter will take basic SUs (allocation units) for Resource A and multiply by the ratio of normalization factors for Resource A and Resource B to produce the equivalent SUs for Resource B.

user community. A future version of this document will be updated to reflect how this dollar cost approach has been implemented and impacted the XSEDE allocation practices and procedures.

J. Research Allocations

The teams associated with XSEDE's Research allocations are typically responsible for 95% of the total system utilization reported by Service Providers. Thus, the bulk of the RAS team's effort is focused on managing the process for these requests. Research requests are reviewed by the XRAC and awarded by XSEDE four times per year.

XSEDE's policies and practice use only a single review phase, with the XRAC members tasked with reviewing the computational appropriateness of each request. The XRAC is specifically instructed to accept the scientific merits of any request that is supported by a merit-reviewed funding award; however, the XRAC does consider the scientific merit of requests without such funding support. For more on the review criteria for the XSEDE allocations process, see the XSEDE Allocation Policies and the *XRAC Reviewer Manual*.

J.1. Conflicts of Interest

Identifying and managing conflicts of interest (COIs) are key aspects of ensuring the integrity of the XSEDE allocations process. XRAC members are deemed to have a COI with a given request if any of the situations that NSF has traditionally defined as a conflict of interest apply between the XRAC member and the personnel associated with a request. The two most common COI types for the XRAC are institutional conflicts and direct participation (as PI, co-PI or major collaborator) in a proposal.

In the case of an XRAC member being the lead or otherwise directly involved in an allocation proposal, that person does not attend the XRAC meeting and does not review requests during that opportunity. Because this policy affects the number of potential reviewers, the Allocations Coordinator does work with XRAC members to balance the number of XRAC member submissions across the four quarterly meetings.

For all other COI types, the XRAC member is not able to view the request nor participate in the discussion. During the discussion of such requests, conflicted XRAC members are asked to step outside of the room. Most such COIs are identified early in the process; for example, the XRAS system identifies institutional COIs automatically as requests are submitted. Less obvious COIs can be reported at any time; the XRAS system provides a way for XRAC members to self-report a COI, and submitting such a report immediately revokes a panel member's ability to view a request and notifies the allocations staff to re-assign the proposal.

J.2. Basic quarterly cycle

XSEDE follows a consistent quarterly schedule to minimize user confusion and staff effort (Table J.1). The schedule is anchored by the four starting dates for Research allocations (April 1, July 1, October 1, January 1). Meetings are held approximately one month prior to the allocation start date; due to the Labor Day holiday, the September meeting has more recently

been held in late August. The submission window lasts one month and opens about 3.5 months before the allocation awards begin.

Upon the opening of a submission window, research groups begin to submit their proposals through XRAS. When the submission window closes, the allocations staff checks all submissions for eligibility and the resource amounts requested for obvious outliers. Reviewers are asked to indicate their proposal preferences and reviews are assigned shortly thereafter. A key element of the guidance to reviewers is that each request should be reviewed on its own merits; the panel is instructed not to attempt to fit their aggregate recommendations into the available resources. Proposals are discussed and award recommendations are made at the XRAC meeting.

After each meeting, the allocations staff and site representatives meet to “reconcile” recommendations with the resources available. This process involves moving allocations from over-recommended to under-subscribed resources, when appropriate, and reducing allocations formulaically, if required. When the allocations staff has returned to their offices, they begin to enter and process awards and rejections. Tickets are fielded from research groups and additional guidance is provided, if necessary.

Table J.1. XSEDE quarterly schedule for Research allocations

Submission Period	Meeting Date	Allocations Begin
Dec 15 thru Jan 15	Early March	April 1
Mar 15 thru Apr 15	Early June	Jul 1
Jun 15 thru Jul 15	Late August	Oct 1
Sep 15 thru Oct 15	Early December	Jan 1

The rest of this section discusses the tasks required to manage this quarterly cycle and the practical aspects of managing the face-to-face XRAC meetings. Note that the quarterly cycles overlap, so the allocations team are typically carrying out different parts of two cycles simultaneously.

J.3. Quarterly schedule of tasks

J.3.1. Because of the fixed quarterly schedule, the responsibilities of the allocation team tend to have a consistent quarterly pattern as well. This section provides an overview of the associated tasks and describes them in more detail. Before and during submission window

Because of the fixed schedule, the Allocations Coordinator can prepare XRAS to handle upcoming Research request opportunities well in advance. The opportunity announcements are typically displayed to users in the XRAS interface for several weeks prior to submissions being accepted.

Also prior to the opening of a Research submission window, the allocations team verifies with SP representatives that they will continue to make allocations on their respective resources for the award period in question. If an SP does not re-allocate their resource, an announcement will be posted within XRAS to indicate that the resource is no longer available for request.

Information about new resources being allocated by XSEDE and decommissioned resources is sent out through email to current users. In addition, announcements are posted on the XSEDE website, sent via XSEDE news, and highlighted directly in the XRAS submission form.

As a submission window opens (Figure 1, Step 2.1–2.2), the Allocations Coordinator hosts a training webinar for those interested in learning more about the proposal writing and submission process. A notification is sent through the XSEDE news system to urge individuals to sign up for the webinar. All registered participants receive a message with a link to the webcast. Participants are encouraged to ask questions during the presentation and follow up via a ticket to help@xsede.org if a more in-depth question and answer is required.

Although the stated deadline for all Research proposal windows is the 15th of the month, in practice the allocations team employs a grace period of one week to accommodate groups that may need a few more days to finalize their proposal due to extenuating circumstances, or make adjustments based on administrative feedback. This grace period is not advertised in any XSEDE documentation. Groups can request this extension through the ticket system.

Table J.2 lays out the calendar of activities for the allocations cycle leading to awards that begin April 1; a similar schedule is followed for the other quarters. As noted, XSEDE has typically been receiving more than 225 requests each quarter.

J.3.2. Before and during submission window

Because of the fixed schedule, the Allocations Coordinator can prepare XRAS to handle upcoming Research request opportunities well in advance. The opportunity announcements are typically displayed to users in the XRAS interface for several weeks prior to submissions being accepted.

Also prior to the opening of a Research submission window, the allocations team verifies with SP representatives that they will continue to make allocations on their respective resources for the award period in question. If an SP does not re-allocate their resource, an announcement will be posted within XRAS to indicate that the resource is no longer available for request.

Information about new resources being allocated by XSEDE and decommissioned resources is sent out through email to current users. In addition, announcements are posted on the XSEDE website, sent via XSEDE news, and highlighted directly in the XRAS submission form.

As a submission window opens (Figure 1, Step 2.1–2.2), the Allocations Coordinator hosts a training webinar for those interested in learning more about the proposal writing and submission process. A notification is sent through the XSEDE news system to urge individuals to sign up for the webinar. All registered participants receive a message with a link to the webcast.

Participants are encouraged to ask questions during the presentation and follow up via a ticket to help@xsede.org if a more in-depth question and answer is required.

Although the stated deadline for all Research proposal windows is the 15th of the month, in practice the allocations team employs a grace period of one week to accommodate groups that may need a few more days to finalize their proposal due to extenuating circumstances, or make adjustments based on administrative feedback. This grace period is not advertised in any XSEDE documentation. Groups can request this extension through the ticket system.

Table J.2. Calendar of activities in support of the Research allocations cycle. Italicized tasks indicate tasks from an overlapping meeting cycle or performed outside of the Allocations team.

Date	Meeting task or milestone
1-Dec	<i>Prior XRAC meeting occurs for awards starting Jan. 1</i>
approx. 15-Dec	Allocation Request Webinar #1
15-Dec	Start accepting Research requests for April 1 start
1-Jan	<i>Prior Research awards start</i>
	Allocation Request Writing Webinar #2
15-Jan	Posted deadline for Research requests
22-Jan	Grace period ends for late submissions (not advertised)
	Administrative review and clean-up of requests, review assignments made
approx. 25-Jan	Reviewers are contacted to optionally enter their proposal review preferences
1-Feb	XRAC reviewers receive assignments
	Requests under review by XRAC
late Feb.	Meeting spreadsheet and other materials prepared by allocations staff
approx. 1-Mar	XRAC meeting
	Complete reconciliation, finalize awards, initiate allocations, prepare admin comments related to declined or substantially reduced awards
7-Mar	Entering/checking final awards for accuracy
15-Mar	Notify submitters of allocation outcomes
15-Mar	<i>Start accepting Research requests for next meeting</i>
1-Apr	Allocation awards begin
7-Apr	Open quarterly allocation user satisfaction survey
30-Apr	Close quarterly user satisfaction survey

J.3.3. Between submission deadline and meeting

Once the submission window (and grace period) has closed, all proposals are vetted for eligibility with regards to position and institution. See the allocation policies for additional details pertaining to eligibility. As they check for eligibility, the allocations staff undertakes a compliance review for all requests (Figure 1, Step 2.3). Any requests that seem to jump out (rather small or extremely large) will be flagged for follow-up. After reading through the submitted proposal it will be determined if the flagged amounts were data entry errors. If a request is unusually small, the PI may be contacted and asked to withdraw the proposal and submit a Startup.

If a request is exceptionally large and poorly justified, the PI may be advised to withdraw their submission due to the high probability that their proposal will be declined. The allocations staff will typically work with such PIs to increase their chances of success in the next allocation cycle.

In light of recent changes to the definition of a Service Unit across the various XSEDE platforms, it has become pertinent that the submitters have requested the resources in the correct type of unit (core-hour, node-hour, GPU-hour, TB-hour, VM-hour, etc.). This responsibility falls mainly on the PI, but the allocations staff and certain SPs may contact PIs for clarification and confirmation prior to assigning reviewers to those requests.

Due to the large number of Research proposals the XRAC receives each quarter (200–250 per quarter), the RAS budget and time limitations prevent the XRAC from reviewing all submitted proposals in one day. XSEDE's experience has been that roughly 150–175 proposals can be efficiently and sufficiently discussed in the course of the XRAC meeting.

To manage the workload and ensure appropriate review for all submissions, some proposals are assigned to XSEDE staff members from the Extended Collaborative Support Service (ECSS) to be reviewed in advance of the XRAC meeting. Typically, the cutoffs for these so-called "Adaptive" proposals are less than 1M core-hours (or equivalent levels of other unit types), and each Adaptive proposal is assigned two ECSS reviewers.

For those requests to be discussed at the XRAC meeting, proposals that request 10M SUs and above are typically assigned 3–5 reviewers and are discussed by the full XRAC during a Plenary session. The Plenary session ensures that the entire panel has a chance to ask questions about the merit of the largest proposals, since XSEDE resources are under high demand and of low supply. The Plenary session also helps establish a consistent approach to discussion and review with the entire XRAC present.

Parallel sessions are utilized to review medium-sized requests and maximize the use of the limited time the XRAC is convened. Three parallel sessions have been implemented, grouping the proposals by the fields of science that typically have the largest number of proposals and shared sets of reviewers. XSEDE currently uses three parallel sessions: PAC

(Physics/Astrophysics), MCD (Biology/Chemistry), and DMR (Materials Research)⁴. Requests from fields of science with fewer proposals are assigned to Parallel sessions to minimize the need for XRAC members to move between sessions. Proposals are categorized as Plenary, Adaptive, PAC, MCD, or DMR by allocations staff as they make their eligibility and compliance review.

When all requests have been vetted, pruned, and/or revised, they are ready to be assigned for review (Figure 1, Step 3.1). Before assignments can be made, the reviewers are given the opportunity to go through the list of requests in XNAS and express preferences for submissions they would like to review as well as those they would not like to review.

Allocations staff try to make review assignments roughly one month before the scheduled XRAC meeting, to ensure reviewers have sufficient time to read and review each of their assigned requests (Figure 1, Step 3.2). Typically, the Allocations Coordinator tries to limit each reviewer to no more than nine proposals to review, though exceptions may need to be made at times. Leading up to the meeting, the Allocations Coordinator monitors the progress of incoming reviews and fields questions from the panel members. The Allocations Coordinator may contact SP representatives and ECSS personnel to urge their peers to complete their reviews in a timely fashion.

J.3.4. During the meeting

All XRAC meetings begin on a Sunday evening and continue the next morning into the late afternoon. Meeting attendees include the XSEDE allocations staff, the XSEDE PI, XRAC members, SP representatives, and XSEDE's NSF program officer (if available). All non-XRAC attendees are considered "observers" and are instructed not to participate in discussions other than to address clarification questions from the XRAC related to their role.

On Sunday evening, the meeting begins with the "caucus" session. During the caucus, the allocations staff welcomes all reviewers, SP representatives and other XSEDE staff present. SP representatives have a chance to address the panel with any changes in the resources they are allocating. The Allocations Coordinator reminds all attendees of the XSEDE Code of Conduct and provides any late-breaking updates. If XSEDE is seeking XRAC feedback on proposed changes to allocation policies or practices, such discussion is usually held during the caucus session.

After the introductions and updates, the caucus begins. Reviewers on the same proposal meet in small groups to clarify their review comments, discuss a consensus recommendation (if possible), and decide who will present the consensus recommendation. During this time, allocations staff and SP representatives remain available to assist in any questions pertaining to specific proposals, allocations policy, and past usage data related to the request.

⁴ XSEDE had only the PAC and MCD sessions for several years, but the volume of materials science requests in recent years led to the creation of the DMR session. Thus, the precise nature and number of the Parallel sessions can evolve over time.

XRAC members can volunteer or nominate others to serve as Session Chairs for the Plenary and Parallel sessions during the caucus session. Meeting chairs may also be recommended by the Allocations Manager when the reviews assignments are being made. Each panel member has the option to volunteer to chair his or her respective session or nominate another reviewer to serve as the chair. Session chairs manage the discussion of requests, keep discussions on topic, and help track the final panel recommendations. The XRAC does not have a formally designated Chair role because the extensive amount of effort leading up to the meeting has been deemed over and above the effort that can be expected of a volunteer panel member.

On Monday morning, the panel reconvenes to begin the Plenary session for discussion of the largest requests submitted for consideration. The process for each proposal is the same. The Plenary session chair announces the name of the proposal, the PI, institution, request, and reviewers. The chair reminds persons with conflicts of interest to leave the room until discussion of the request is completed. After the assigned reviewers have presented a short overview of the proposal (science, computational plan, justification) to the panel, they make a recommendation on the allocations to be awarded based on the justification provided. The XRAC members then have an opportunity to pose questions to the reviewers assigned to these specific proposals. Questions, answers, and comments may circulate the room until the panel has reached a consensus recommendation. This process continues, one proposal at a time, until all Plenary proposals have been presented and discussed. Once all Plenary proposals are discussed, the XRAC members are given the option to go back and revisit any that require additional discussion. For example, a reviewer may feel that a proposal discussed early in the day was handled differently than those discussed late in the session and should be reconsidered in that context. Throughout the session, the allocations staff ensures that XSEDE policy is being followed and applied consistently for all proposals.

If no reviewer is present for a proposal that is discussed at the meeting, the allocations staff will read the reviews entered by the absentee reviewers and may ask other reviewers present to take a look at the proposal (and reviews) to make a second recommendation.

After the Plenary session, the panel breaks up into the Parallel sessions described previously. One member of the XSEDE allocations team is responsible for overseeing and monitoring the review process in each of the sessions. The allocations team member works with the session chair to ensure all members with conflicts of interest are out of the room when required, provide clarity on past allocation history (if requested), and record the recommendations and comments made by reviewers. Once all proposals from a particular session are presented and discussed, the reviewers present are given the option to go back and revisit any that require additional discussion. At the end of the Parallel sessions, XRAC members are thanked for their service and dismissed.

J.3.5. After the meeting

When all the reviewers have been dismissed and the meeting has formally come to an end, the allocations staff and site representatives reconvene to determine if the total Recommended Allocations exceed the total Available Amount on one or more resources. If so, the XSEDE

Allocations Coordinator invokes and manages a “reconciliation” process. The goal of the reconciliation process is to bring the recommended amounts in line with the available (Figure 1, Step 3.4).

The reconciliation process intentionally does not involve the XRAC; reviewers are instructed to focus on the merits of each request. However, during the discussion, reviewers can recommend that a request be exempted from the formulaic reductions. An exemption may be granted when the panel’s recommendation has already reduced the potential award to a level below which the reviewers feel none of the request’s scientific objectives could be completed. Reconciliation exemptions recommended by the XRAC are dealt with separately and on a case-by-case basis.

In the first phase of reconciliation, XSEDE allocations staff and SP representatives are tasked with going back through the requests to identify which proposals could complete their proposed work on XSEDE resources that they may not have requested. If a proposal is deemed “movable,” the recommended award can be transferred in part or in full to another resource, at the Service Providers’ discretion. The transfer may be between similar resources (such as from Stampede2 to Comet), or from more distinct resources (such as from Stampede2 to Jetstream), depending on the nature of the work being conducted. A standard conversion factor, based on the normalization factors of both resources, is applied when moving units from one resource to another (see Section O). While normalization factors can be quite precise, the actual amounts transferred are typically rounded to a few significant digits.

The balancing phase of the reconciliation process continues itself until the excess demand is balanced equally across computing platforms. That is, in this phase, recommended amounts are not being *reduced*; rather, the level of over-recommendation across architecturally similar resources is being equalized. For example, all GPGPU-resources are being balanced to similar over-recommended levels, as are all conventional HPC systems, and so on.

Once the recommendations have been balanced, the reduction phase can be completed. This essential piece of the reconciliation process uses a formulaic solver (within a Microsoft Excel spreadsheet) to reduce the recommended allocations to meet the availability (see Section P). The solver takes into account the level of funding the PI has reported and the funding source or sources. Priority is given to NSF-funded projects. (The solver is also run periodically throughout the balancing process to ensure that we’re reducing allocations by the smallest percentage possible, rather than leaving certain resources under-allocated.)

When the Reconciliation process has been applied after a meeting, the XSEDE Allocations Coordinator includes general information about the Reconciliation reductions in the award notifications (Figure 3).

The total number of SUs requested was more than three times greater than the amounts available on all systems. After the usual merit-review criteria were applied by the assigned reviewers and the panel-recommended allocations were determined, the totaled allocations for the cluster systems were found to be oversubscribed on most systems, but especially so on the GPU resources. It was necessary to adjust the recommended allocations to fit within the available SUs according to the formulation in the XSEDE Allocations Policy document. Every panel-recommended allocation was included in this reduction.

Please note that oversubscription—allocating above the available SUs—causes queue wait times to escalate and significantly slows down research productivity for everyone. Hence, it is an XSEDE policy not to oversubscribe any system. Also, with the allocation of all the available SUs, no surplus remains this quarter for supplements or appeals, and reviewers will closely assess the fairness of additional awards.

Figure 3. Sample language provided in XSEDE allocation award letters when describing the use and impact of the reconciliation process.

Once the allocations staff returns home, they begin the process of entering awards and rejection comments. One member of the allocations staff compiles all of the declined proposals and enters comments with instructions for how the PI should proceed after receiving their rejection notice. Providing specific instructions helps reduce the number of tickets related to declined requests. Another member of the allocations staff enters and saves all of the recommended and approved amounts into the XRAS system, along with any comments noted during the discussion. Once all awards and rejections are entered into the system, a member of the XRAS team generates a report from XRAS to cross-check the values in the system and identify any data entry errors. After the checks have been completed, the allocations staff formally flags each proposal as “approved” or “declined.” Once all proposals are labeled for approval or rejection, the team sends email notifications (via XRAS) to all PIs of their outcomes (Figure 1, Step 4.1) and initiates all approved allocations (Figure 1, Step 4.2), at which time XRAS sends award information to all Service Providers.

Formally, XSEDE allocation policies provide PIs with an opportunity to appeal the outcomes within four weeks of being notified (Figure 1, Step 5.1–5.3). Rejection decisions cannot be appealed, however. In practice, XSEDE does not recommend or encourage appeal submissions due to the already high levels of demand and the need after many meetings to invoke the reconciliation process. For the formal appeals process, see Section M.4.

J.4. XRAC meeting logistics

The following sections describe the process of selecting meeting locations for the XRAC and managing the panel’s travel and on-site logistics. Considerations for sites include, but are not limited to: geography, time of year, price, meeting room availability, proximity to a major airport, ground transportation and AV costs.

J.4.1. Site selection

XRAC meeting locations are selected primarily based on the locations of prior meetings with consideration of any future sites that may already have been booked. Since the panel is comprised of reviewers from across the country, we try to incorporate one meeting in the northern, southern, eastern, and western United States per year. Large, expensive markets, such as Boston, NYC, and San Francisco, are typically avoided. The allocations staff attempts to book sites no later than six months prior to an XRAC meeting. Occasionally sites are booked for multiple years to attain a rooming block discount. Additional consideration is given to ease of travel to the selected cities, so the reviewers aren’t discouraged by difficult flight schedules and total meeting expenses are not inflated by costly flights. Some consideration is also given to the destination’s “appeal” to the panel members. For example, Urbana-Champaign, IL (especially in

December) is not as appealing as Tampa, FL, or Tempe, AZ, both for flight availability and likely weather.

The hotel venue must have at least one meeting room capable of fitting 50-60 persons seated in a U-shape, along with two smaller rooms (up to 20 persons seated in a U-shape) to use for breakout sessions and reconciliation. During the hotel selection process, the allocations team considers the costs for guest room, catering, AV, Internet, and other miscellaneous expenses while narrowing down options. Many selected hotels provide free airport transportation, guest room wi-fi, and breakfast. Free breakfast alone can save XSEDE an additional \$2,000 per meeting. Candidate hotels that are close to the airport help reduce ground transportation costs and maximize the panel members' time in attendance at the meeting.

J.4.2. Travel coordination

Travel for XRAC members is coordinated through the allocations staff, as well as a travel agency contracted by Carnegie Mellon University/Pittsburgh Supercomputing Center (PSC), the site responsible for the XRAC budget. Reviewers who purchase their own airfare may contact XSEDE/PSC for approval of their flights. Reviewers who choose to have the agency book their flights will contact the travel agency directly. The travel agency is responsible for working with the allocations staff to confirm flight purchases. Allocations staff oversee the purchase of each flight to ensure that costs fall in line with the budget.

Allocations staff handle reimbursements for costs that reviewers incur during travel to/from XRAC meetings. In general, no XRAC reviewer should incur any business-related expenses while traveling to serve the panel. Non-XRAC-related costs (extra nights at the hotel, costs of local attractions, etc.) will not be reimbursed. No costs pertaining to alcohol can be reimbursed. Reimbursed expenses include, but are not limited to:

- Airfare
- Lodging
- Meals (itemized receipts only, no per diem)
- Ground transportation (taxis, Uber/Lyft, rental car—if justified)
- Internet access charges
- Mileage
- Parking

SP representatives and XSEDE staff are responsible for arranging their own travel to and from the meeting and for paying their expenses.

J.4.3. On-site arrangements

In general, the allocations staff serves as on-site contacts for the hotel. At a typical meeting, XSEDE provides three meals and two breaks to the XRAC members, SP representatives, and other XSEDE staff present for the meeting. A welcome reception and dinner are provided on Sunday night. Breakfast, lunch, and two breaks are provided on Monday.

The panel sits in a large U-shape with a screen, projector, and microphone at the “top” of the U to be used for presentations and general information. The meeting’s AV needs are relatively modest except for the aforementioned equipment, meeting room Wi-Fi, and power strips for laptop charging. The allocations staff provides an electronic booklet that lists all requests for the meeting and distributes a copy to reviewers, site representatives, and XSEDE staff. Name tents and badges are created for all attendees. XSEDE allocations staff generally collect name tents and name badges after each meeting and reuse them multiple times to further reduce staff time and meeting costs.

J.4.4. Meeting communications

At each XRAC allocation meeting, the Allocations Manager informs the review panel and site representatives in attendance of the upcoming meeting dates and locations. Invitations for specific XRAC meetings are sent out by allocations staff approximately two months before the meeting date. Separate emails are sent to XRAC members and non-XRAC attendees. Both email contacts provide the basic information pertaining to the meeting (dates, location, hotel, address, meeting schedule, and ground transportation info). Both reviewers and non-reviewers are asked to inform the allocations staff of any dietary restrictions to ensure that the meals provided by the XRAC will be substantial enough for their diets.

The email sent to reviewers contains details pertaining to rental car insurance policies, reimbursement policies, travel agent information, and clarification as to whether guests are permitted, or not.

J.4.5. Absentee reviewers

Each meeting typically has a small number of absentee reviewers. XRAC members who indicate that they are unable to attend the meeting in person are asked if they’re still able to contribute reviews ahead of the meeting. If so, they are asked to submit reviews before the start of the caucus on Sunday night. We also ask that, if possible, the reviewer be available by email in case they are needed for caucusing, so that if any questions, comments, or disagreements come up among the reviewers, they may be able to revise their reviews before the meeting begins on Monday morning.

J.4.6. Meeting budget

Table J.3 summarizes the typical budget for an XRAC meeting. Room nights, catering and air transportation comprise most of the budget.

Table J.3. Typical XRAC meeting budget

Budget Category	Typical Cost
Guest Rooms	\$8,600
Catering	\$12,500
Flights	\$10,000
Ground Transportation	<\$1,000
AV/Internet/Room Rentals	\$4,000
Reimbursements	\$1,500

Total	\$37,600
-------	----------

J.4.7. Emergency plan

In the event of an emergency, due to severe weather, natural disaster, or other unexpected disturbances that could prevent all or a large portion of the reviewers and allocations staff from reaching the meeting destination, XSEDE's allocation team has an emergency plan that will be put into place. Due to the number of requests received each quarter and the time required to finalize the requests after the meeting, the emergency plan is guided by the need to complete the review process at or close to the original schedule. The timeline and budget do not allow for the rescheduling of the face-to-face meeting.

The emergency plan primarily consists of transitioning to a set of email or teleconference sessions. Such a transition requires some logistical adjustments to the meeting, and the adjustments may vary slightly depending on which and how many persons are unable to reach the meeting location. The XRAC is regularly briefed on the emergency plans so the process can be effective when needed.

However, once the XSEDE team determines that enough expected attendees will not succeed in reaching the meeting, the following steps are initiated.

First, the XSEDE team sends an email notification to all XRAC members and other meeting attendees that the meeting is going into emergency remote session mode, and they are instructed to monitor their email. In addition, all reviewers will be asked to ensure that their full reviews are entered in the XRAS system, or emailed to the XSEDE staff, if need be. (At a normal meeting, some reviewers may discuss requests based on handwritten notes and enter their review into XRAS only after the proposal is discussed.)

Second, the caucus session will be conducted entirely by email on Sunday evening as planned. Using email contact information available in the XRAS system, reviewers will discuss recommendations and concerns about each of their assigned proposals with their co-reviewers. Critical updates from XSEDE or the SPs will also be shared via email. Simultaneously, the allocations staff will go through all of the submissions, particularly those assigned to XRAC members unable to participate in the remote sessions, to ensure proposals with divergent review opinions or recommendations are flagged for discussion.

The next morning, the Plenary and Parallel sessions will be conducted, assuming a sufficient number of the allocations staff and reviewers are available to participate on the planned meeting day. The allocations staff will assess the attendance to determine if the remote sessions will include face-to-face participation by those on site or if the discussion will be better balanced by having all participants join remotely (e.g., from their hotel rooms).

In addition to the main online sessions, the allocations staff will start and join a separate, chat-only Conflict of Interest "room" and share that connection information with all participants. When attendees must leave the remote discussion session due to a COI, they will join the COI room so they can be told when they can rejoin the discussion session. The COI room will also be

used by the allocations staff to help coordinate when reviewers need to move between parallel sessions.

If the travel disruptions are such that a quorum of reviewers are unavailable on Monday, the remote sessions may be held at separate times throughout the span of the meeting week. A Doodle poll will be sent out to reviewers to see which, if any, times will work best for the online sessions. This option is only a last resort, since the panel had already made arrangements to be free for the meeting.

Once the logistics of each session are finalized, the discussion of requests will proceed as normal. The session chairs will guide the session participants through the proposals and discuss final recommendations.

Following the last of the Parallel sessions, the XSEDE staff and SP representatives will meet via online session for the reconciliation process.

K. Startup Allocations & Educational Allocations

Outside of the XRAC (Research) proposal submission windows, XSEDE always accepts, reviews, and awards allocations for resources through Startup and Educational allocations. The following sections detail the process for handling these types of requests. XSEDE's goal is to have these small-scale awards reviewed and awarded within 10 business days.

K.1. Validating submissions

When a Startup or Educational proposal is submitted, XRAS automatically notifies the allocations staff, and the request is queued for handling. The initial review of such requests begins with the XSEDE allocations staff checking the PI's information to confirm that they are at a U.S.-based institution and they hold an eligible position (i.e., that the PI is not an ineligible graduate student). Next, we check the PI's email to ensure they've submitted with an email that matches their organization; PIs must provide an institutional email address, not a Gmail, Hotmail, or similar generic, free email address. If the PI is a graduate student, we check to ensure that the student has an NSF Graduate Research Fellowship award or honorable mention. Once those checks are completed, we conduct a simple web search to further verify the identity of the PI.

In addition to checking the PI information, we also take note of the resources being requested. If the PI is requesting an amount that doesn't fall within the Startup limits, they will be contacted and asked to update their request.

For Educational requests, a resource justification document and syllabus are required. The resource justification document should detail why the level of resources are required for the class.

Occasionally, requests are submitted to an incorrect opportunity. During an open Research proposal submission window, for example, some PIs may submit their Research requests to the

Startup or Educational opportunity (or vice versa). The allocations staff typically spot these soon after they're submitted as part of the initial validation step.

K.2. Assigning reviewers

New and renewal Startup and Educational requests (along with Supplements) are automatically assigned for review by XRAS. Each SP has a designated default reviewer who is responsible for reviewing Startup and Educational requests for their resources.

To ensure relevant SP staff are immediately aware of incoming Startup and Education requests, XRAS notifications regarding submissions are sent via the “xras-notify” mailing list. All new, renewal, supplement, transfer, and extension requests are sent this list. These notifications alert the allocations staff of actions they may need to review. The xras-notify list includes members of the help@xsede.org address, as well as SP staff responsible for reviewing and handling small-scale requests. The help tickets are queued for XSEDE allocations staff.

XSEDE staff generate a weekly report to highlight requests still under review or pending review of an action (supplement, transfer, extension). If reviews are tardy or approaching tardy a member of the allocations staff will send an email reminder to the reviewer(s).

K.3. Finalizing awards

Startup and Educational requests are typically approved unless some eligibility issue disqualifies the submission. In some cases, requests may be moved or adjusted based on reviewer comments.

Once reviewer comments have been received, XSEDE allocations staff enter the recommended allocation amounts and process the award, notifying the PI and the affected SPs.

L. Other Allocation Types

In addition to Research, Startup, and Educational requests, XSEDE makes available a few other allocation types for more specialized purposes. These other types mainly include Campus Champions, Discretionary, and Staff allocations.

L.1. Campus Champions

In order to be eligible for a Campus Champion allocation, a prospective PI must first be in contact with the Campus Champion team. They provide an overview of the program, what can be expected of a Campus Champion, and any clarification the PI requires before entering the program.

When a new Champion enters the program, the XSEDE Champions team sends an email to welcome the newcomer and copies the allocations team on the email. In addition, the PI is added to the “[Current Champions](#)” web page. This ensures that when a new Champion's

proposal is submitted, the allocations staff can verify that the PI has already been pre-approved for an allocation.

There is no “formal” review process for Champions requests, other than vetting the eligibility of the submitter. When a new or renewal request is received, an allocations staff member will proceed with creating the allocation.

The standard allocation for a Champions award is essentially a Startup-level amount across most XSEDE compute and data resources. An exclusion of the standard allocation is Jetstream Storage, due to the limited disk space for that resource.

L.2. Campus Single-Sign On

This allocation type allows SPs with non-allocated resources to leverage XSEDE’s SSO Hub and single-sign on capabilities on their campus clusters. Such allocations are by invitation-only.

L.3. Discretionary

Discretionary allocations are not often used, but are available to XSEDE SPs who want to leverage XRAS to track awards made to projects as part of the SP-managed discretionary time on their systems. SPs must contact the allocations staff directly (through email or an XSEDE ticket) to let XSEDE staff know who is likely to submit a discretionary request. This allocation type is available as an invitation-only opportunity in XRAS to prevent ineligible submissions.

L.4. Others

Staff allocations are provided to XSEDE management and SP management for limited operational needs or to provide courtesy access for evaluation and related purposes. These generally fall within Startup limits (unless other arrangements are made directly with an affected SP). They maintain their own allocations and add/remove appropriate staff as necessary.

M. Allocation Management Requests

XSEDE continually receives allocation management requests from PIs, co-PIs and Allocation Managers during the course of their allocations (Figure 1, Steps 6.1–6.3). These allocation management requests are submitted via XRAS and processed as described below. When any of these actions are requested, an XSEDE Help ticket is created (via the xras-notify list) and routed to the XSEDE Allocations staff. In a typical quarter, the XSEDE allocations staff will field approximately 500 requests in aggregate for these management actions.

M.1. Extensions

A time extension can be submitted when less than 90 days remain on the allocation and up to 90 days after the expiration of the allocation. Currently two options are available when requesting a time extension: 3 months or 6 months. This constraint ensures that Research awards remain aligned with future Research opportunities. Once an extension is submitted, an

allocations staff member assigns the extension request to the SP representative(s) responsible for reviewing their resources. The SP enters a review with their decision, and once all sites have agreed on a decision, the allocations staff member processes the extension and notifies the PI through XRAS.

M.2. Transfers

A transfer request may be submitted at any time during an active allocation. When a transfer is received, an allocations staff member assigns the transfer request for review to the SP representative(s) that the PI would like to transfer to. If the transfer is approved by the SP(s), an allocations staff member completes the transfer in XRAS, being sure to check the accuracy of the conversions across resources using the XSEDE SU Converter. A notification is sent to the PI through XRAS once a decision has been made on the transfer.

M.3. Supplements

Supplement requests can be submitted up until 30 days remain on the current allocation. If a PI has less than 30 days remaining on their allocation but is in need of a supplement, they must typically first submit and be approved for a time extension. When a Startup or Research supplement is submitted, it is automatically assigned for review to the SP(s) requested. After the SPs have completed their reviews, the allocations staff will process (approve or reject) the supplement request and notify the PI through XRAS.

M.4. Appeals

Appeals must be submitted within four weeks of the date the PI is notified about the review results. In terms of processing, Appeals are treated like Supplements. However, in most cases after an XRAC meeting, there are no additional SUs available for allocation. Therefore, it is difficult to have appeals approved. Occasionally, SP representatives may honor the PI with a small supplement, or ask the PI to submit a supplement later in the allocation period.

M.5. Advances

Once a PI has successfully submitted a Research proposal, the option to request an Advance becomes available. Advances are accepted up until three weeks before the upcoming allocation period; this limitation is in place partly due to XRAS system limitations, but also partly to limit the number of Advance requests from PIs who only want to get a jump on the official allocation start date. A PI can request up to 10% of the SUs requested in the proposal to be used in advance of the start date of the allocation. Once submitted, Advances are assigned for review in the same way as Research/Startup supplements. Once all requested sites have entered reviews, the allocations staff will process the advance in XRAS and notify the PI. Multiple advance requests may be honored, depending on the SP(s). Advances are not applicable to Startup and Educational requests.

M.6. Final Reports

Final reports can be submitted at any time after the expiration of the award. PIs, Co-PIs, and Allocation Managers receive a notification that the option is available for submission. The Final Report helps XSEDE collect additional publications and better understand the reasons that users end their XSEDE-allocated activities.

N. Managing Authorized Persons and Other Users

XSEDE requires a person to have an XSEDE user identity (i.e., an XSEDE User Portal account) before they can submit an allocation request. In turn, the allocations process relies heavily on XSEDE processes related to users and the users associated with authorized roles on projects. In some cases, allocations staff are involved in these processes, so we describe them briefly here.

N.1. Authorized Roles

XSEDE recognizes five roles associated with allocated projects. The first three roles have authority to submit allocation requests and subsequent management actions related to awarded projects, including adding and removing users and being notified of actions requested and awarded to a project.

- **Principal Investigators (PI)**—The PI's identity and eligibility are key to many aspects of the allocations process. A project may have only one PI at a time.
- **Co-Principal Investigator (Co-PI)**—A project may have zero or more co-PIs. No special eligibility rules apply to co-PIs, although this designation typically implies that the individual has some intellectual leadership on the project.
- **Allocation Manager**—A project may have zero or more Allocation Managers. No special eligibility rules apply to Allocation Managers, although this designation typically implies the person is active in the day-to-day oversight of the activity on XSEDE-allocated resources.

The following two roles are also involved in the allocations process. However, these roles are not authorized to submit management actions related to an ongoing project.

- **Submitter**—XSEDE and XRAS allow a New or Renewal request to be submitted by a person other than a PI, co-PI, or Allocation Manager. The Submitter can designate someone besides themselves as the PI on a request. The Submitter's identity is logged by XRAS. However, the Submitter cannot make subsequent management requests for the project (Extensions, Transfers, etc.); thus, in most cases, the Submitter is also associated with one of the Authorized Roles.
- **User**—As a courtesy, XRAS allows a New or Renewal request to include the XSEDE usernames of additional, non-authorized users to be associated with the request. These users are passed from XRAS to the XSEDE accounting service so that accounts can be created on the awarded SP resources.

N.2. Verifying Principal Investigators

Confirming the identity of PIs is key to the allocations process, since PIs can then designate co-PIs and Allocation Managers for their projects without further XSEDE staff involvement. All three roles can then add or remove users without XSEDE involvement. Thus, XSEDE staff diligently follow the PI verification steps described in Section K.1 above.

N.3. Adding and Removing Users

Any PI, co-PI, or allocation manager can add users to their active allocation(s) in the XSEDE User Portal. The “Add User” function is located under the “My XSEDE” tab. The PI, co-PI, or Allocation Manager submitting the add user request is responsible for verifying the XSEDE portal username of the person they are about to add to the allocation. Users can only be added by their XSEDE username.

To remove a user, a PI, co-PI, or Allocation Manager must login to the XSEDE portal and click on “Allocations/Usage” under the “My XSEDE” tab. After clicking on the “Manage Users” button the option to remove users will be available.

N.4. Re-verifying Users

When a PI has a Renewal request pending, they receive a notification asking to re-verify the users they would like to retain from the current allocation. They are directed to the XSEDE Portal to complete this process.

This process was implemented as an additional security measure to ensure that PIs regularly review the users on their projects and that inactive or former members of a research group do not receive continued access to XSEDE resources. If this process is not completed before the start of the Renewal allocations, all users are removed from the project, and the PI, a co-PI, or an Allocation Manager must re-add users.

N.5. Changing PIs

Because XSEDE-allocated Research projects can be renewed indefinitely, the PIs on these longer-running projects may change over time. Typically, XSEDE receives a help request stating that a PI is changing universities or moving to another country, and thus a PI change is needed. Because of the critical role of the PI in verifying all other users on the project, such changes can only be made by XSEDE staff after the PI verification process is completed.

XSEDE staff first make contact with the PI, confirm that this is true, and have the PI designate an eligible person to take over the PI role. Once we have the new PI’s portal account, XSEDE allocations staff make the PI change in XRAS.

N.6. Changing Co-PIs and Allocation Managers

The Co-PIs and Allocation Managers associated with a project can be modified in two ways.

First, during the submission of a New or Renewal request, Co-PIs and Allocation Managers can be designated. When the New or Renewal request is approved, the Co-PIs and Allocation

Managers on a project are updated to reflect the new submission. That is, any prior Co-PIs and Allocation Managers are removed, and the currently listed Co-PIs and Allocation Managers are given their designated roles. If a Renewal request is declined, XRAS makes no changes to the Co-PIs or Allocation Managers.

Second, during the course of an awarded project, the PI can add or remove co-PIs and Allocation Managers through the XSEDE User Portal. These changes take effect immediately.

O. Appendix: Service Units, Normalized Units and Weighting Factors

The NSF XSEDE Program and its predecessor programs (TeraGrid, Partnerships for Advanced Cyberinfrastructure [PACI] and the Supercomputing Centers Program) have allocated and operated since 1985 a variety of high-performance systems in support of national computational science and engineering research. These systems span a wide range of computational performance, and the sites operating them have employed a variety of proposal, usage, administration, and accounting policies over the years. A method was needed, starting with the installation of a second system, to compare the usage between those two systems because one unit of resource usage (e.g., a core-hour) on System A does not represent the same amount of capability as that same unit on System B. This paper describes the requirements for a process to compare performance and usage of dissimilar systems, and the implementation of this process that has been used since the 1980s.

There are basically two processes that drive the need for an inter-system comparison. The first and more fundamental of these is the need to transfer allocated resource units from one machine to another. Such transfers may happen during the allocation process itself, when the allocations staff moves recommended allocations from oversubscribed to undersubscribed resources. While allocations are eventually awarded on specific systems, users may ask for unused allocation amounts to be transferred from one system to another. When a machine is retired, allocated projects may have their allocations on the retired system automatically transferred to a replacement system to minimize interruptions to ongoing projects. Less often, a Service Provider may suggest a transfer to a system more appropriate to the applications being run. In all these scenarios, a conversion factor is used to adjust the transferred allocation to a comparable level of capability on the target system.

The second process requiring comparison of performance is the longer-term reporting and tracking of resource usage. A conversion factor is used, for example, to track total usage of NSF-supported resources over long periods of time, or to track the evolution of usage by a single investigator or discipline. Such tracking can be used to project computing resource requirements into the future, or to aid configuration choices or demonstrate the effect of selected imbalances in system configurations over time. An example of the use of long-range tracking of resource usage is shown in Figure 4, which is taken from a recent XSEDE Interim Project Report.

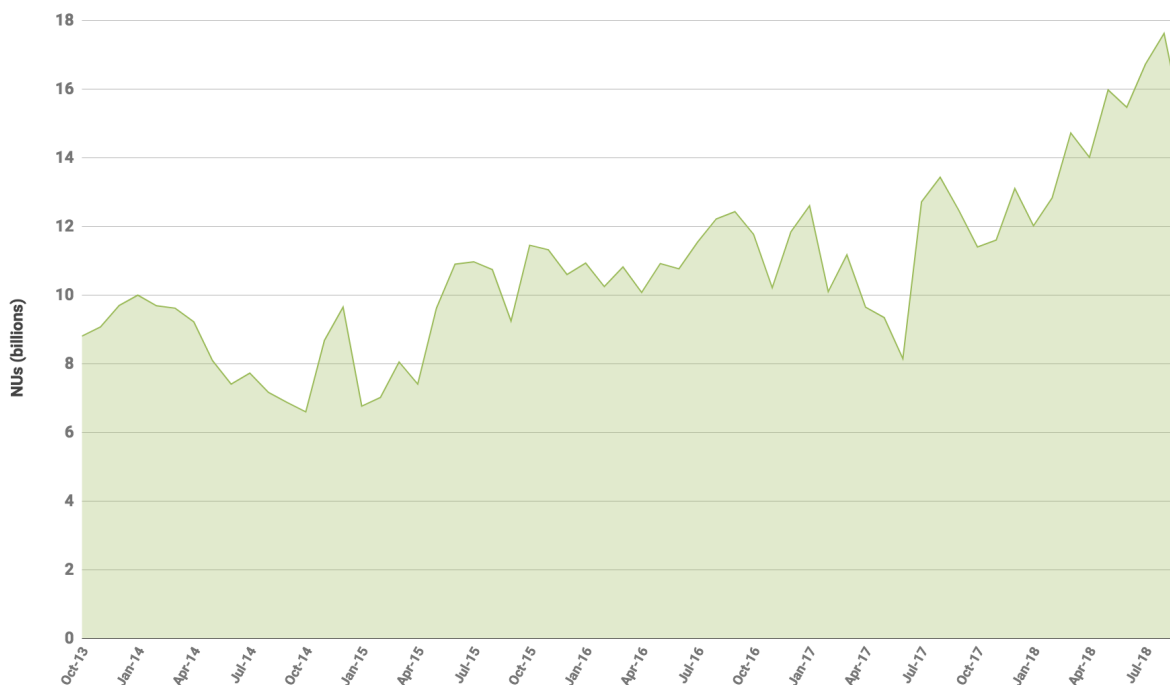


Figure 4. Monthly usage reported to the XSEDE central accounting system across all XSEDE-allocated computing systems over the last five years.

O.1. Service Units

A discussion of the XSEDE allocations process would not be complete without an explanation of Service Units (SUs), even though the rest of this document has studiously avoided using that terminology. As of 2019, the use of “Service Units” is gradually being deprecated in favor of resource-specific unit types; however, the term regularly appears in documentation and common parlance and may still be used by some SPs for their allocated HPC systems.

The term *Service Unit* (SU) has long been used in the allocation process for the amounts allocated and the units of usage reported from the systems. The term was adopted because early vector systems often charged jobs for non-CPU system resources (such as memory or disk), and the more general term Service Unit could encompass all facets of job charges. Over time, however, the term SU became nearly synonymous with “core-hour.” From the late 1980s to the early 2010s, the most common HPC architectures evolved from vector systems to parallel systems boasting single-core processors and nodes, to multi-processor nodes, to multi-core processors and nodes—yet, the base computing element remained the “core.”⁵ The SU concept

⁵ In fact, it took a number of years for the terms “CPU,” “processor,” and “core” to stabilize after being used interchangeably at times. In this document, “core” is used to describe a complete, conventional processing element. A “processor” is the hardware unit that fits into a socket on a motherboard. Modern processors are typically described as “multi-core” or “many-core.” “CPU” is used only when needed to distinguish the conventional processors and cores from nodes with “GPU” hardware.

used in the allocations process thus became tightly linked to the core-hours delivered by the hardware.

However, starting in the early 2000s, the nature of available HPC architectures began to evolve in ways that challenged the commonly held notion of SUs. First, GPUs began to emerge as components of HPC architectures, complicating the understanding of “cores” on a given node. Then, by the mid-2010s, the increasing number of cores on conventional processor chips led to situations where jobs often left cores idle so that the active cores could access more memory; such jobs were still charged for all cores on the node.

Alternate architectures in the NSF-supported portfolio in the latter half of the 2010s further challenged the SU concept. PSC’s Bridges system offers large-memory nodes, specifically designed with memory, not cores, as the component of value. Indiana University’s Jetstream cloud provides access to virtual nodes that may not reflect the number of cores or nodes in the underlying physical hardware. And TACC’s Stampede-2 system offered explicitly heterogeneous types of nodes, with differing numbers of cores.

Under TeraGrid and continuing into the XSEDE era, the allocations process also began allocating resources for which the notion of core-hour and, by extension, the SU had no relevance, including storage systems and user support services.

Thus, up through 2018, the SU concept has remained dominant, even though the allocations process evolved to explicitly use a variety of allocation units that aligned better with the system hardware and SP operational practices. GPU resources are allocated in “GPU-hours,” large-memory systems are allocated in “terabyte-hours,” storage resources are allocated in gigabytes, and XSEDE’s Extended Collaborative Support Service (ECSS) is allocated via “Yes/No” recommendations. The SU concept suffered its heaviest blow when, with the decision by TACC to allocate Stampede-2 in “node-hours,” the largest HPC resource in the XSEDE-allocated portfolio abandoned the base “core-hour” unit.

With respect to our discussion of NUs in this appendix, two points about SUs are important. First, for consistency, allocations have long been made in SUs that equate to actual core-hours, further cementing the conceptual relationship between SUs and core-hours. Second, job accounting practices reflect a challenge in making precise assumptions about the value of a given SU. That’s because, in practice, the notion of SU has always been a “modified core-hour;” HPC system usage is typically reported as the core-hours consumed by a job *modified by the priority factor for the queue in which the job had run*. While jobs in the “regular” or “normal” queue have a queue factor of 1, users often can choose to stretch their allocated SUs by running in lower-priority queues (and being charged less) or accelerating their progress by running in “premium” queues (and being charged more).

O.2. The LINPACK Benchmark

While no single metric can universally represent the performance of a computational system, using individualized applications to handle system-to-system migration conversions would be too cumbersome and costly. More than 2,000 projects are typically active in a given year during

the XSEDE program, making it impossible to use tailored benchmarks for each of them to track long-term usage over 20-year periods. Therefore, a single “representative” benchmark application, or suite of applications, is desired.

This benchmark set must be (a) sufficiently “representative” of a wide variety of applications run on NSF supported systems, (b) open for analysis to be understood by researchers wanting to estimate usage requirements and trends and understand the limitation of “representative” suggested above, and finally (c) easy to run on a wide variety of systems over long periods of time (1985 to the present and beyond).

The benchmark program chosen and used since the Supercomputing Centers era has been the LINPACK benchmark suite and its follow-on, the Parallel LINPACK suite for more recent parallel systems. The suite is a set of well-implemented and well-documented basic linear algebra algorithms that are easy to understand (and apply results to those of your own application) and are relatively straightforward to run on any known system. Although arguably not a perfect choice, LINPACK certainly offers simplicity and ease of understanding, in a “what you see is what you get” sense. Given these factors as well as the long history of LINPACK results on older systems, XSEDE continues to use LINPACK as the basis for these normalizations.

Further supporting this choice has been the observation that while the LINPACK benchmark suite does not necessarily reflect absolute performance expectations for most applications, various real application comparisons have shown that the relative performance of the LINPACK benchmark suite on different platform does reasonably reflect the relative performance of many applications on those same platforms. This is also a critical characteristic of the selected benchmark.

O.3. The “Normalized CPU-Hour” and Weighting Factors

The NSF Centers Program originally defined the base *Normalized Unit (NU)* as one single processor-hour on a Cray X-MP with 9.5 ns clock. The standard LINPACK performance associated with this NU is $R_{max} = 191$ MFLOPS. This value was obtained from the published LINPACK $R_{max} = 427$ MFLOPS for a two-processor Cray X-MP with 8.5 nsec clock, dividing by 2 (processors) and adjusting by a factor of 8.5/9.5 for the difference in clock speed.

Thus, the weighting factor, or conversion factor, for a given resource is calculated as:

$$\textbf{Weighting Factor} = R_{max} / (N \text{ processors} * 191 \text{ MFLOPS})$$

Where N is the number of processors on which the LINPACK benchmark was run, and R_{max} is appropriately converted to MFLOPS.

To compute NUs from the usage reported

$$\textbf{NUs} = \textbf{Reported Usage (in SUs)} * \textbf{Weighting Factor}$$

Weighting factors are also used to calculate the allocation amounts involved in a transfer from System A to System B:

$$SU_B = SU_A * (WF_A / WF_B)$$

Where SU_R is the amount of SUs for resource R and WF_R is the Weighting Factor for resource R . XSEDE provides an SU Converter to help users and allocation administrators with these calculations. In practice, the output values from the SU Converter are rounded to a smaller number of significant digits, recognizing the imperfect nature of the weighting factors.

The Weighting Factors for a set of NSF-supported systems spanning three decades and four program eras are shown in Table O.1.

Table O.1 Weighting Factor calculations and agreed upon values for NSF-supported resources

System	Entered Service	Processors	Calculated Weighting Factor	Agreed Upon Weighting Factor
PSC Cray C90	1993	16	4.483	3.750
NCSA HP Exemplar (SPP-2000)	1997	64	2.255	2.255
SDSC Cray T3E/600	2001	268	2.286	2.286
PSC Lemieux (ES45, AlphaServer SC, 1 GHz)	2001	3024	7.748	7.748
NCSA Tungsten Cluster (Xeon, 3.06 GHz)	2003	2,688	20.081	20.081
NCSA Phase I TeraGrid (Itanium 2, 1.3 GHz)	2003	512	21.576	21.576
TACC Lonestar (Xeon, 3.06/3.2 GHz)	2003	600	21.422	21.422
SDSC DataStar p655	2004	1408	19.711	19.711
TACC Ranger (Sunblade x6420)	2008	62,976	36.015	33.076
NICS Kraken (Cray XT5)	2009	98,928	44.016	44.016
SDSC Comet	2015	46,656	133.016	133.016
TACC Stampede-2	2017	367,024	3,100.881	3,100.881

O.4. Limitations of the Normalized Unit and Weighting Factor

While the NU and the associated Weighting Factors have served and continue to serve their original purposes admirably, the LINPACK-based approach has its limitations. A significant issue is that, because the Cray X-MP is used as the basis for normalization, the weighting factors for modern systems have ballooned, which means that NU values may be three or four orders of magnitude larger than the allocable units presented to users. This section describes some other limitations and the approaches used by XSEDE and earlier programs to work around the limitations.

O.4.1. Application-Specific Weighting Factors

As noted previously, a more accurate measure for converting allocations would be the specific applications being used by a given researcher. In practice, users can perform benchmark runs on the source and destination systems, and the allocations staff will honor that relative performance in calculating the units to be transferred to the destination resource.

O.4.2. Alternate Weighting Factors

The LINPACK benchmark does not always provide a reasonable reflection of relative performance for some machines. In such cases, the Service Providers set the weighting factor for these machines at a more appropriate value, often extrapolating from the LINPACK-based weighting factor of an earlier machine. These “alternate weighting factors” have been applied with greater frequency as novel architectures are deployed and Service Providers adopt allocation units not based on core-hours. When alternate weighting factors are used, the approach is documented in the Resource Description Repository.

O.4.3. TeraGrid SUs (TG SUs)

Originally, NUs closely tracked the size and scale of SUs—that is, weighting factors were close to 1. However, thanks to Moore’s Law, weighting factors increased and the values of NUs and SUs diverged over the years, with NUs gradually becoming orders of magnitude larger than the corresponding SUs, making NUs less suitable as “user-facing” values.

During the TeraGrid program, a semiformal attempt was made to renormalize the NU to the original Phase 1 TeraGrid clusters (Intel Itanium 2, 1.3-GHz cores), which had a weighting factor of 21.576. Allocations on the TeraGrid clusters were made in these Service Units, leading to the term “TeraGrid SUs” or “TGSUs.”

TGSUs were occasionally used in reports or other communication materials because the order of magnitude for TGSU values was better aligned with the numbers of core-hours available on TeraGrid-era HPC systems. With a weighting factor of 21.576, the values of TGSUs were roughly 20x smaller than the comparable NU values.

TGSUs for a given amount of SUs from resource A were thus calculated as:

$$TGSUs = SU_A * (WF_A / WF_{TG})$$

The TGSU notion was embedded in the TeraGrid accounting system, with resource weighting factors all recorded as relative to the TeraGrid cluster, which has a conversion factor of 1 in the database. NUs are calculated from TGSUs by multiplying by the TeraGrid cluster weighting factor of 21.576.

With the passage of time and the end of the TeraGrid program, the TGSU concept has fallen out of favor, because of their tight coupling to a past program and because of the continued performance improvements of modern processor cores. Today, NUs are used by XSEDE in its reporting processes. However, NUs are now roughly three orders of magnitude larger than the SUs on current systems, and XSEDE does not generally use NUs in any user-facing documentation or communications.

O.5. Dollar Cost per Allocated Unit

A final method that XSEDE uses to report on resource allocations and usage across a heterogeneous resource portfolio is the dollar cost. This method is commonly used to provide

project PIs and others with a sense of the scale of an allocation request. One advantage of the dollar cost approach is that it can generally be used regardless of the resource type—compute, storage, or otherwise—as long as a dollar cost can be assigned to the unit of allocation.

The dollar cost is calculated by summing the procurement costs and operational costs over the expected lifetime of the system (often defined by the procurement and operational NSF awards to a Service Provider), and dividing the total dollar value by the number of allocated units expected to be made available over the lifetime of the resource.

For a computing resource allocated in “node-hours,” for example, the allocated units would be estimated by multiplying

- N = the number of nodes in the system
- Y = 8,760, the number of hours in a year
- L = the number of years in the planned system lifetime
- X = adjustment factor to account for downtime, reduced availability, and so forth.

Reasons for “reduced availability” could reflect that the machine is not 100% allocated by XSEDE. For example, NSF-supported resources may be required to allocate only 80% of the system via XSEDE. A campus-supported resource may offer a much smaller fraction of the system to XSEDE users.

With the dollar cost approach, the size and scale of a project that spans multiple resources of different types and using different allocable units can be calculated by the formula:

$$\sum_{i=1}^n R_i * D_i$$

Where R_i is the amount allocated (or used) on resource i , and D_i is the dollar cost per allocated unit for resource i .

P. Appendix: Reconciliation Formulas and Calculations

Before each XRAC meeting, SP representatives determine the amount of available allocation units that can be allotted that quarter to satisfy Research requests. These estimates are called the Available Amounts for each of their resources. In general, approximately one quarter of the yearly available amounts for each system is made available at each quarterly meeting.

As described in Section J.3.5, the first part of reconciliation requires the XSEDE allocations staff and SP representatives move recommended allocations to alleviate oversubscription on individual resources *or to reach a state of equal oversubscription across all comparable resources*. By balancing the oversubscription across comparable resources, any necessary reductions affect all requests approximately equally.

If a condition of general oversubscription exists, even after moving Recommended Allocations to alternate resources, the Reconciliation formula is applied. The Reconciliation formula can be applied across more than one “adjustment group” of comparable resources, if necessary. In

addition, the Reconciliation formula is not applied to requests below a minimum threshold, typically equivalent to a Startup-scale allocation.

P.1. Reconciliation Formula

Once the balancing process is completed, the XSEDE Allocations staff use the following formula, approved by the NSF, to adjust the balanced recommendations to arrive at the Awarded Allocations.

The formula for adjusting an allocation for oversubscription is:

$$\begin{aligned} \text{Award} &= [(1-G) \cdot R_n + (1-G) \cdot F \cdot R_o] \cdot \text{Size}(R,S) && \{\text{for modest oversubscription, } G < \frac{1}{3}\} \\ \text{Award} &= [(1-G) \cdot R_n + (1-G) \cdot F' \cdot R_o] \cdot \text{Size}(R,S) && \{\text{for high oversubscription, } G \geq \frac{1}{3}\} \end{aligned}$$

where

G = Global scaling factor, between 0–1, solved for by non-linear optimization

R = Recommended Allocation

R_n = Component of R that is supported by NSF funding (0–1)

R_o = Component of R that is supported by non-NSF funding ($1-R_n$)

F = Funding priority factor for modest oversubscription, 2

F' = Funding priority factor for high oversubscription, 0.5

S = Size scaling factor, 0.1—small awards reduced less

Three primary factors drive the reconciliation formula:

1. The global scaling factor, G , is used to reduce all Recommended Allocations;
2. The funding factors, F and F' , give NSF-funded requests a priority and in which mixed sources of funds are prorated to their NSF and non-NSF funding ratios; and
3. The size-based factor, S , that linearly reduces large requests more than small requests.

In practice, the funding and size factors are fixed, and the global scaling factor is adjusted by non-linear optimization, using the Microsoft Excel Solver plug-in. The formula works best when applied to sets of resources allocated in common units, of comparable order of magnitude in Recommended Allocation amounts, and oversubscribed to approximately the same degree. In the initial phase of reconciliation, recommended awards are moved among similar systems to achieve comparable oversubscription levels. Thus, the XSEDE Allocations Coordinator may apply the Reconciliation formula to several sets of comparable resources—for example, a set of oversubscribed systems allocated in core-hours, separately to a set of oversubscribed systems allocated in GPU-hours, and yet again to smaller, special-purpose systems (e.g., large-memory systems).

The Allocations Coordinator makes appropriate adjustments for specific systems and uses NUs when appropriate. For example, the solver was originally built with “core-hour” as the definition of Service Unit. When Stampede2 adopted node-hours, XSEDE adapted to use a TACC-

provided formula to convert Stampede2 allocations to “core-hour equivalents” when applying the solver formulas. Similar adjustments are made when dealing with GPU-based resources.