

Recommendations for XSEDE Allocation Policy Changes

30 January 2017

[Version 1.0]

Allocations Policy Review Team

Jay Alameda (NCSA), Emre Brookes (UTHSA), Lonnie Crosby (NICS), Ken Hackworth (PSC),
David Hart (NCAR, Chair), Chris Hempel (TACC), Ralph Roskies (PSC), Sergiu Sanielevici (PSC),
Amy Schuele (NCSA), Nancy Wilkins-Diehr (SDSC)

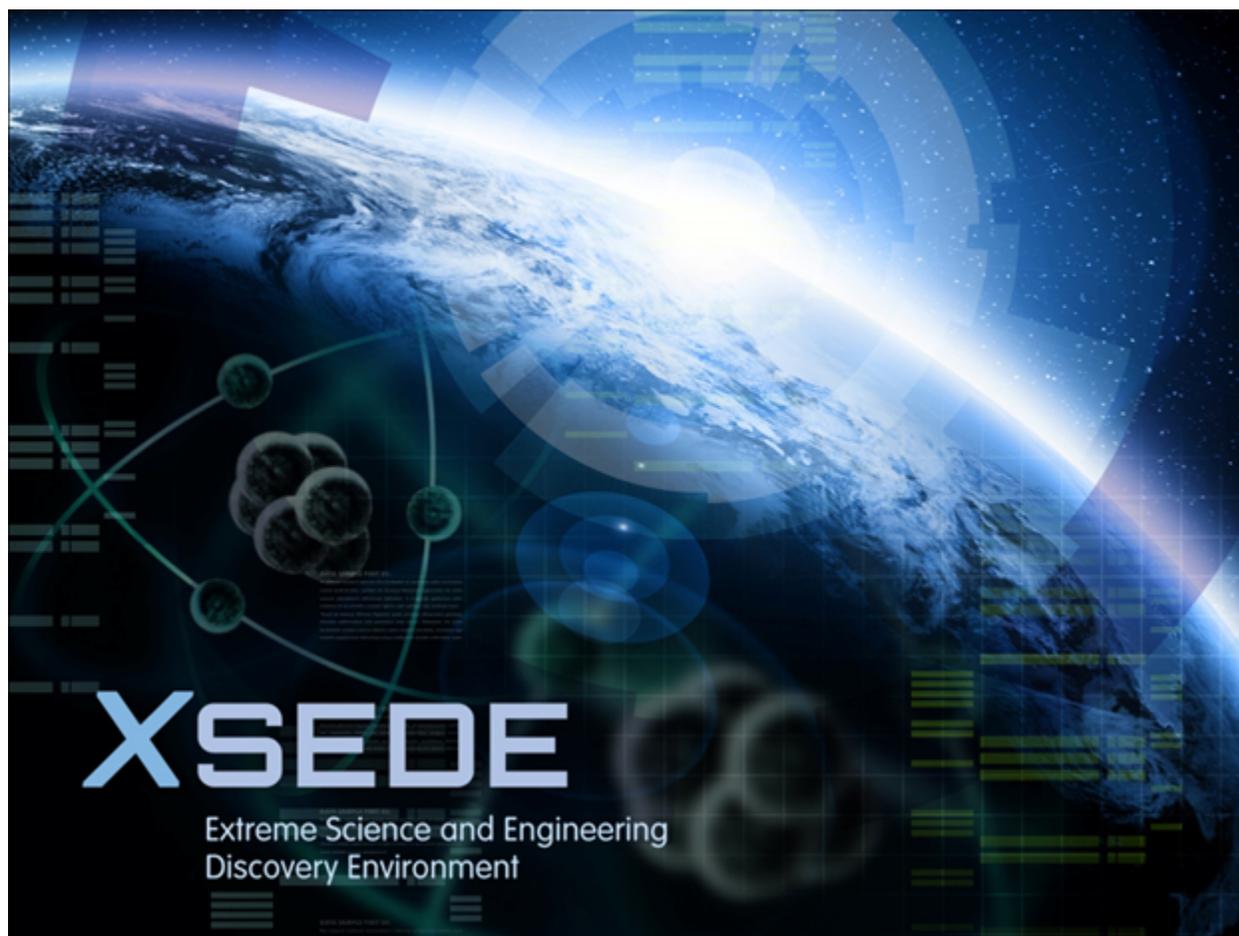


Table of Contents

A.	Document History	iii
B.	Document Scope	iv
C.	Executive Summary	1
D.	Recommendations	2
D.1.	Policy.....	2
D.2.	Procedure.....	8
D.3.	Documentation and Tools.....	11
E.	Actions and Implementation.....	13
F.	Appendix.....	14

A. Document History

<i>Relevant Sections</i>	<i>Version</i>	<i>Date</i>	<i>Changes</i>	<i>Author</i>
<i>Entire Document</i>	<i>1.00</i>	<i>01 Feb 2017</i>	<i>Original submitted for NSF review & approval</i>	<i>Allocations Policy Review Team</i>

B. Document Scope

For this review of the XSEDE Allocation Policies, XSEDE designated a ten-member Allocations Policies Review Team. The group's charter was to ensure that the allocation policies were consistent with the current and future XSEDE ecosystem of resources and aligned with the spirit of the National Science Foundation's principles of merit review, with supporting processes and procedures that had sufficient rigor and necessary structure, while minimizing the barriers to entry and participation by users. The team held biweekly meetings from September 2015 through October 2016. Team members represented XSEDE's Resource Allocations Services (RAS), Extended Collaborative Support Service (ECSS), and Community Engagement and Enrichment (CEE), among others. The team's discussions were informed by regular interactions with a wide range of XSEDE's stakeholder groups, including the User Advisory Committee (UAC), whose chair was part of the policy review team, the XSEDE Advisory Board (XAB), the XSEDE Service Provider Forum (SP Forum), and the XSEDE Resource Allocations Committee (XRAC). Updates on the discussions were presented at XSEDE quarterly meetings and other venues.

Comments on a draft of this recommendations report were solicited from the XSEDE Senior Management Team, the SP Forum, the XSEDE User Advisory Committee, and NSF. The Appendix summarizes the feedback received from each group and how it was addressed. The final recommendation report was submitted to NSF for formal approval. Upon formal approval, XSEDE, via the Resource Allocation Service (RAS), is tasked with defining an implementation plan for the full set of recommendations and then executing that plan to put the updated policies and procedures in place.

C. Executive Summary

In September 2015, XSEDE assembled a team to review the policies used to allocate Service Provider (SP) resources as well as XSEDE's Extended Collaborative Support Services. The team comprised representatives spanning the XSEDE project and the User Advisory Committee. The group's charter was to ensure that the allocation policies were consistent with the current and future XSEDE ecosystem of resources and aligned with the spirit of the National Science Foundation's (NSF's) principles of merit review, with supporting processes and procedures that had sufficient rigor and necessary structure, while minimizing the barriers to entry and participation by users.

To that end, the team solicited input from across XSEDE and from a range of stakeholder groups, including the User Advisory Committee (UAC), whose chair was part of the policy review team, the XSEDE Advisory Board (XAB), the XSEDE Service Provider Forum (SP Forum), and the XSEDE Resource Allocations Committee (XRAC). Updates on the discussions were presented at XSEDE quarterly meetings and other venues.

The team's recommendations fell into two primary categories—policy and procedure—with a general recommendation to improve and update the allocations documentation. On the whole, the team found that the policies were already well-stated to encompass a diverse portfolio of resources. Some policy refinements were recommended with respect to major research facility requests for SP resources, the approach for how XRAC recommendations are reduced to fit within the available resources, reducing the burden on users to continue Startup allocations, in some circumstances, for periods longer than one year, and re-introducing mechanisms to allow multi-year Research allocations.

On the procedure and process front, the team made recommendations to improve processes to ensure that current policies are more effectively implemented. These recommendations include better training for users and reviewers on how different resource types should be requested and reviewed; ensuring more consistent reviews across a large set of reviewers spanning different communities; providing templates to help users improve their submissions; and ensuring that requests without supporting science awards receive appropriate scientific review.

Comments on a draft of this recommendations report were solicited from the XSEDE Senior Management Team, the SP Forum, the XSEDE UAC, and NSF. The Appendix summarizes the feedback received from each group and how it was addressed. The final recommendation report was submitted to NSF for formal approval. Upon formal approval, XSEDE, via the Resource Allocation Service (RAS), has been tasked with defining an implementation plan for the full set of recommendations and then executing that plan to put the updated policies and procedures in place.

D. Recommendations

The recommendations have been organized into three primary categories. Policy recommendations apply to the high-level rules, definitions, and foundations that outline the goals and objectives for XSEDE allocations, including topics related to how current policies should be interpreted or clarified. Procedure recommendations apply to the approaches and practices used to conduct the various allocations processes of submission, review, and administration. The final category, XRAS and Documentation, encompasses specific recommendations related to those aspects of the allocations processes.

D.1. Policy

- 1. After thorough review, the XSEDE allocation policies overall were found to be appropriate and sufficient for requesting and reviewing allocations for a diverse portfolio of resources.**

A primary objective of the group's charter was to ensure that the allocation policies were not inappropriately focused on traditional HPC resources. The group determined that, overall, the XSEDE allocations policies appropriate and sufficient for allocating a diverse portfolio of HPC and non-HPC resources. However, users requesting allocations and allocation reviewers would benefit from improved training on the resource offerings and how to prepare and review requests for non-HPC resources. The group also noted that the policy documentation required minor edits to remove vestigial references or explanations that were overly specific to HPC resources. (See related issues in the Procedure and XRAS and Documentation sections.)

- 2. XSEDE should not implement a cap or limit on the size of the largest requests. Similarly, XSEDE should not try to redirect very large requests to other, non-XSEDE opportunities (like NSF PRAC-allocated Blue Waters or DOE INCITE Resources).**

Supporting very large requests remains part of XSEDE's core mission to deepen access to and use of cyberinfrastructure resources. Similarly, very large requests are often required to pursue scientific objectives that have the potential to produce transformative impacts. Lack of sufficient resources at present to support such demands does not provide sufficient reason to modify the overarching policies. Redirecting very large requests is merely an alternative to enforcing a size limit. Such an approach would be impractical, since it would require XSEDE or the XRAC to serve as an arbiter or interpreter of policies for eligibility, timing, or appropriateness of non-XSEDE allocation opportunities. Coordinating such an approach across just alternative NSF-supported opportunities (Blue Waters, NCAR) would require a strong mandate and agreement across the appropriate NSF directorates.

- 3. Very large-scale requests from major research facilities (e.g., LIGO) should be handled outside of the XSEDE allocation policies and procedures through discussions and agreements between NSF, the major research facility, and the affected Service Providers.**

The allocation policies allow for requests to support large-scale, funded projects that work together as a consortium ("large-scale consortium") or that provide services to a large community of users who are typically not directly collaborating with the project PI (e.g., gateways). While what

constitutes a “large-scale consortium” project may vary, such projects typically operate as described in the policies document: “Often in these cases, a mechanism already exists for allocating community or project resources ... and that mechanism will also be used to make allocations from the time granted to ... individual investigators. Requests for this type of project typically describe the internal processes for managing access of individual investigators” In these situations, the XRAC reviewers cannot assess the individual research activities against the standard review criteria.

While the requests and review criteria have been adapted to these situations, the current policy and processes may not be sufficient when the expectation is that large-scale use of XSEDE-allocated resources will support critical operational needs for the project. For example, LIGO may have the expectation that it will use substantial SP resources rather than, for example, purchasing dedicated hardware. When the size of such an allocation request approaches or exceeds the scale of the largest requests from individual investigators, we recommend a modification to the allocation process.

In such a situation, the parties involved—the Large Facility, relevant SP(s), XSEDE, and the cognizant NSF program officers—need to discuss the expectations of each party and reach agreement on the appropriate path forward. While preparing a standard request may still be the proper course of action, these groups may alternatively agree to reduce the level of resources available to the XSEDE allocations process and devote SP support to the facility directly. Following such an agreement, if NSF and the SPs still recommend that a formal allocation request be submitted by the major research facility, such requests may be accompanied by special instructions for the XRAC.

Individual investigators or small teams who are users of, or associated with, a Major Research Facility may continue to apply through the normal allocations process for resource needs that are separate from the Facility’s critical, operational cyberinfrastructure and to which the normal XRAC review process can apply the standard review criteria without impacting the overall Facility operations.

4. The approach to reducing XRAC recommendations to fit with availability limits following each meeting should not adopt an “all-or-nothing” approach, but should consider incorporating XRAC ratings into the reconciliation formula.

The UAC voiced very strong opposition to adopting an all-or-nothing approach in the post-meeting reconciliation process, in which requests with a nonzero recommended allocation, but whose XRAC ratings fell below some threshold dependent on the available resources would be rejected entirely to allow higher-rated requests to be supported at a greater level. The UAC expressed concerns about putting a halt to ongoing computational work and particularly the impact on New submissions, which already see higher rejection rates.

Analysis of XRAC rating data and this group support further investigation of using the average reviewer rating of requests to modify the post-meeting reconciliation formulas. In such an approach, higher-rated requests would retain more of their recommended amounts, while lower-

rated requests would be reduced by a larger amount. As part of adopting such an approach, this group further recommends that RAS take steps to improve the training and consistency of reviewer ratings. A significant step in this direction is described in the Procedural recommendation regarding the adoption and promotion of a review “rubric” for XRAC members.

5. XSEDE and the Service Providers should adopt policies and procedures that allow investigators to request and receive allocations longer than one year.

Input received by this team indicates exceptionally strong demand from users for allocations longer than one year. NSF (and other agency) awards are typically made for three years with well-defined reporting from awardees during the award period. Such multi-year allocations would significantly reduce the burden on PIs as well as on the XRAC reviewers, who are now regularly faced with more than 200 requests at each quarterly meeting.

While the general consensus of this group favors a policy that allows such multi-year allocations, discussions raised considerable concerns and questions regarding how such a policy might be implemented such that the Service Providers are able to appropriately manage heavily subscribed and over-requested resources. To balance the challenges, allocation PIs should be prepared to accept some limitations and constraints on usage of such allocations, and initial user feedback suggests this is reasonable.

This group identified the following issues and constraints that must be addressed by any successful implementation of such a policy. This list is not exhaustive but does represent some of the major considerations that recurred during discussions.

1. Implementation must avoid the pitfalls of the previously attempted and subsequently discontinued “Multi-Year Request.” Pitfalls included a submission and review process that eliminated most of the desired benefits while increasing confusion on the part of both submitters and reviewers over timing, content, and requirements for interim submissions and reviews. (In practice, the process evolved to be very similar to the normal annual Renewal process, but with more confusion.)
2. The procedure must address issues about resource transitions that may occur during such multi-year allocations. Most likely the burden for handling such transitions (choosing the subsequent resource, asking for a transfer of allocation) should be placed on the allocation PI.
3. The procedure must allow Service Providers to “budget” available SUs across the multi-year period and the intervening Research allocation opportunities. For example, Blue Waters implemented multi-year allocations by setting up a series of annual allocations, but encountered problems when users did not use allocations evenly and asked to move resource units forward or backward in time. One possibility would be to augment allocations that span several years with supplements for future annual increments.
4. Some annual reporting and review may be appropriate, but must not approach the level of effort required for New or Renewal Research requests. Questions to consider are whether

interim reports are required at all or are reviewed as Adaptive requests or by the full XRAC, as well as the timing and requirements for such interim reports.

5. Interim increments (if applicable) should be subject to reduction levels applied at the nearest XRAC meeting. That is, future allocation amounts should benefit or be constrained by future resource availability.
6. The multi-year allocations approach needs to define the eligibility and appropriateness of making supplement requests to increase the originally approved allocations. The need for additional allocation amounts may be due to significant reductions to requested amounts, either by panel recommendation or post-meeting reconciliation, or a change from the originally proposed computational plan.
7. The XRAC should be empowered to disallow the multi-year aspect of a request and recommend only a single-year allocation, for example, if the request does not provide a justification of sufficient quality. Similarly, the XRAC should be allowed to recommend that a PI with a track record of well-reviewed requests and consistent usage of those allocations should consider submitting a multi-year request.
8. Multi-year allocations should track the period of performance of some supporting grant(s). Requests without supporting grants should not be permitted to request multi-year allocations.

Because of these and other issues, the group also suggested that RAS consider a pilot program as part of evaluating possible implementation approaches. Such a pilot may, for example, be limited to NSF-supported projects only, and only permissible for requests up to a modest number of SUs (e.g., a threshold of 1M or 1.5M SUs total).

6. XSEDE should consider capping the size of requests that are permissible without science-reviewed supporting grants.

Because of the challenges of guaranteeing scientific review at the same level of rigor as funding agencies provide for monetary awards, the group reached consensus on supporting a limit on the size of requests that XSEDE should permit. However, further consideration is required to designate a specific limit or method of determining such a limit in a way that is appropriate across the full range of XSEDE-allocated resources and that provides a simple and transparent understanding for submitters, reviewers and RAS staff.

Formulas considered included capping such requests at approximately 1% of a resource's capacity for a given meeting and using some multiple of a resource's Startup limit. However, the values needed for such a calculation are not always easily visible; it would be better to present the number of allocable units to users to avoid confusion. The formula should also apply to requests that ask for multiple resources. The procedural implementation must also consider how and when to verify when the cap applies, ideally prior to being assigned to reviewers. Would requests that exceed the cap be set to the cap? A further consideration is whether and when certain types of supporting grants satisfy the requirement of having had scientific review. For example, awards such as those supporting Research Experiences for Undergraduates (REU) or Integrative Graduate Education and

Research Traineeship (IGERT) activities may not have had scientific review for the work associated with a specific allocation request.

Furthermore, in defining such a limit, it would be advantageous for XSEDE to consider using the same to-be-defined limit or limits to other procedures that apply to requests that are “not too large.” Two examples of such thresholds are (a) the level below which recommended allocations are not cut further by the post-meeting reconciliation formula, and possibly (b) the threshold below which multi-year allocations are allowed.

7. XSEDE should permit, in well-defined situations and for specific purposes, the renewal of Startups allocations beyond one year.

Current XSEDE policies state that Startups are typically not renewed and that, after one year, the projects should submit a Research request to an XRAC opportunity, regardless of the size of that request. The policies also note, in fairly broad terms, that exceptions can be made; the actual granting of these exceptions in practice can be inconsistent and often depends on which person or which Service Provider is involved in the review of a given request. XSEDE needs a clearer policy statement to allow more consistent decision making and transparency.

While the group appreciates the simplicity of the current policy (i.e., that Startups are not renewed), a number of use cases exist for small-scale projects that last longer than a year, but need not rise to the level of requiring Research review by the XRAC process. A common scenario is that of a software or infrastructure developer, for whom access to resources is more important than substantial allocation. Similarly, a small-scale gateway or other project may not anticipate exceeding a resource’s Startup limits in any given year for the foreseeable future.

This group recommends that XSEDE clearly communicate the requirements and permissible exceptions for allowing renewal of Startup projects. Thus, this group recommends that a Startup be allowed to be renewed if they do not fall under the anticipated track for a Research project (evaluation/development, testing, preparation, benchmarking followed by production runs) and are not anticipated to exceed Startup allocation limit(s).

To provide consistent evaluation of renewal requests for Startups, XSEDE should require extra documentation after the first 12- to 18-month period:

- For first time renewal, in addition to the standard project metadata and declaration of grant support (or absence thereof), the PI shall provide a statement of proposed work, including codes to be used and other requirements, a brief justification to explain the small size of allocations requested, and why non-XSEDE resources cannot meet them (maximum of two pages).
- For subsequent renewals, in addition to project metadata updates, the PI shall provide a progress report for the previous year and enter publications into their XSEDE User Portal profile and associate these publications with the project in question.

Because XSEDE may be interested in tracking the number of Startups that fall under certain use cases within these policies, XSEDE and RAS may consider enhancing the metadata collected for

Startup requests to allow projects associated with key use cases—such as Developer, Gateway, or small-scale access to innovative resource—to be tracked and reported on easily.

8. XSEDE should clarify the policies and exceptions when and if to restrict PIs and research groups to only one Research request and allow the XRAC to recommend exceptions on a case-by-case basis.

Current XSEDE policy states that, to minimize the effort required by the submitter and the reviewers, a single investigator should usually only have one active Research project, regardless of the number of supporting grants or the extent of the research activities encompassed by the request. Exceptions are permitted in certain circumstances, including separate projects to distinguish traditional research activities from support for a science gateway, or when the projects are in different fields of science and cannot be reviewed by the same set of experts. Current policy further states that entire research groups “should consider” consolidating requests to minimize the submission and review effort required; this particular point has become a matter of confusion for XRAC members.

XRAS data analyzed for two recent XRAC opportunities indicate that roughly one-third of Research requests have more than one supporting grant. Thus, permitting one request per supporting grant could cause, in the worst case, a one-third increase in the number of submissions—potentially even more since some requests list more than two supporting grants. While, in practice, the actual number would likely be less, with XRAC submissions regularly exceeding 200 per quarter, a potential 33% increase argues for caution.

However, allowing PIs to submit multiple requests had some support from the group, in particular because it would allow better mapping between allocations and supporting grants, and because XRAC reviewers themselves occasionally do complain that too much work is compressed into the document page limits for adequate review.

On this topic, this group recommends leaving the policy largely intact, in support of the original goal to minimize effort for both submitters and reviewers, but the language should be clarified that this policy may be less strictly interpreted for appropriate cause. Two supporting actions are further recommended:

1. The XRAC reviewers should be empowered to identify exceptions to this rule for specific individuals or research groups. For example, the XRAC reviewers may determine that a particular body of work represented in a single request does not permit adequate review commensurate with the work proposed and, while still making a positive recommendation for the current request, may ask the submitter and the XSEDE allocation manager to split future requests into multiple submissions. Similarly, the XRAC reviewers may recommend that multiple requests from the same research group, in the future, be combined into a single request.
2. The XRAC must be given guidance on this policy, which is flexible by design. A discussion at an XRAC meeting would be most appropriate.

In deciding whether to recommend the splitting or merging of requests, the XRAC members should consider the original intent of this policy—to simplify the effort required of both submitters and reviewers. While the tendency may be for maintaining the general policy of one request for one investigator, the policy leaves room for judgment and needs to be interpreted as such.

In recommending that a request be split in the future, the XRAC should provide guidance on the number of expected submissions. Factors that may argue for splitting a request may include a large number of supporting grants; and the inability of the work to be described and reviewed at a level of detail consistent with the scrutiny applied to other comparable requests; clearly separated lines of scientific inquiry, including work that may be more appropriately reviewed by sets of reviewers with different expertise; the ability of junior researchers to initiate their own research programs; or distinct methodologies or contributing sets of team members. In this vein, research groups may choose to submit separate requests with different PIs, for example, to allow junior investigators or postdoctoral researchers to initiate their own lines of inquiry, especially with separate supporting grants.

On the other hand, multiple requests from the same research group or from collaborating researchers may be recommended for merging in the future if the lines of inquiry, methodologies, or contributing staff include significant overlap that can be adequately described in one Research request, again commensurate with the request size or at a level of detail expected of other comparable requests.

D.2. Procedure

The following recommendations and discussion topics apply to procedural matters including implementation of processes to support the allocation policies.

9. XSEDE should formally collect “Final Reports” from projects after their allocations have expired and are not being renewed.

While the allocation policies have long “required” projects to provide a final report, formal mechanisms have not been provided for submission of such reports, format has not been defined, and the policy has not been enforced. This group recommends that XSEDE define the format of such reports and implement a process for accepting such reports in XRAS. Noting concerns from various groups, XSEDE should consider how such reports will be used and keep the burden on users to a minimum. One possibility would be to send Final Reports to the cognizant agency program officers, along with a listing of those projects with overdue Final Reports.

Because of many years with only loosely stated policy, strict enforcement of the final report policy—for example, *requiring* final reports before new allocation requests can be accepted—should be phased in cautiously and gradually.

10. XSEDE should improve its training and handling of the submission and review of requests for a diverse resource portfolio.

While the allocation policies are generally appropriate for managing requests for the evolving resource and service ecosystem, XSEDE should provide better training and guidance for both users and reviewers with respect to communicating the availability of resources, understanding the recommended uses for each resource, helping users get started by selecting the most appropriate resources, preparing successful allocation requests, and reviewing requests in accordance with the unique capabilities and recommended use for each resource. XSEDE is uniquely positioned to provide such multi-resource overviews in training offerings for users and the XRAC.

11. XSEDE and RAS should provide optional document templates to help new users prepare successful Research requests.

Feedback from stakeholder group indicated that new users would appreciate templates for the various documents required for Research requests. XSEDE should work in concert with XRAC to define templates that support the needs of the reviewers, guide users through the necessary preparation, and perhaps most importantly, reduce the chances that requests are rejected for technicalities.

12. XSEDE should define and prepare a “rubric” for XRAC to use when reviewing and discussing requests.

Because the XRAC membership is continually evolving as terms expire and new members join, the XRAC supported the notion of having a rubric or guide to improve the consistency of reviews and to ensure all relevant criteria are addressed. XSEDE should furthermore ensure that the use of such a rubric should be well integrated into the process—that is, made available to submitters, mapped to the review form, and applied in practice by reviewers. Consistency of reviews will support both the rigor of the process, including the science review of requests without supporting grants, and the consistency of request ratings, which are recommended above to be factored into the post-meeting reconciliation algorithm.

13. XSEDE should ensure that all types of allocations requests that do not have merit-reviewed supporting grants receive appropriate review on their scientific or scholarly merits.

The current allocation policies appropriately and sufficiently address the issue of allocation requests without supporting, merit-reviewed awards, namely that all such requests should receive appropriate review of their scientific merits, in addition to the three primary review criteria. However, it is unclear whether such review happens consistently in practice. Data analyzed by this group indicate that up to a third of the requests received at each Research opportunity may not have (or declare) merit-reviewed supporting grants; however, analysis shows the eventual awards to such requests represent 10% or less of the awarded service units. This group emphasizes that XSEDE must remain conscious of the policy and should monitor the process to ensure that it continues to be enforced. When scientific review is warranted, the reviewers should apply the standard intellectual merit and broader impact review criteria already established and documented by NSF.

This group recommends and supports five approaches to improve processing of science review for such requests:

1. Such requests should continue to be reduced by the reconciliation approach to a greater degree than requests that identify merit-reviewed requests.
2. The submission forms in XRAS should encourage more complete data entry. (This has already been completed, in the form of a checkbox that forces submitters to either enter supporting grants or affirm that they have no supporting grants.)
3. XSEDE should consider, in concert with NSF, whether and how to cap the size of allocation requests that can be made without supporting grants. (See Recommendation 6 above.)
4. Allocation reviewers should use a common rubric to ensure review of science, when necessary, is included in reviews, outcomes. (See Recommendation 12 above.)
5. In making review assignments, the allocations team should ensure domain experts on the panel are assigned to requests without supporting grants. In addition, the XRAC members should be given specific notice of requests without supporting grants and reminded to address the science or scholarship.

14. XSEDE should continue to evaluate and adopt practices to manage the XRAC meeting workload in such a way as to ensure all Research requests receive necessary and sufficient review, commensurate with their impact on the overall ecosystem.

This group recognizes that the growing volume of requests, now regularly exceeding 200 submissions to most quarterly Research opportunities, adds to the risk that requests may not receive appropriate levels of review or consistent review. XSEDE must continue to monitor the impact on reviewer workloads and time commitments, and the impact that increased volume places on the time available to discuss requests.

The RAS staff should continue to monitor how meetings are conducted and suggest ways to focus the limited discussion time most productively and efficiently. XSEDE and NSF should re-emphasize the importance of this process and encourage reviewers to stay for a full day of discussion. XSEDE should thus also consider and if necessary budget for the costs of supporting more reviewers for an extra night. At the furthest extreme, if the volume of requests continues to escalate, XSEDE should consider the value or necessity of adding a second meeting day.

The use of plenary and parallel sessions, as well as the use of XSEDE ECSS staff for “adaptive” reviews, should continue, but neither approach is a panacea. Leveraging the recommended rubric for XRAC reviews may offer an opportunity, for example, to further winnow the number of requests that are discussed in face-to-face sessions, by eliminating requests from the agenda if the request is uniformly recommended for rejection.

15. XSEDE should assess and propose methods to smooth out the per-meeting variability in the post-meeting reconciliation process.

The current approach of algorithmically reducing panel recommendations to fit within the available resource limits remains appropriate. However, the scale of such reductions is affected greatly by

the number of requests submitted in a given quarter. As a result, users are inconsistently penalized or rewarded based solely on the number and size of submitted requests at that meeting, not on the overall supply of and demand for resources.

XSEDE and RAS should work with the SPs to modify the post-meeting reconciliation process to use the full range of data available to remove or reduce the impact of the activity level at a particular Research opportunity on the reductions required, e.g., by shifting some of a resource's available Service Units into or from a future quarter. Using available historical data, this group recommends that proposed methods be tested on prior allocation cycles, then applied to upcoming meetings in parallel to the current method to refine and build trust in a proposed approach.

16. XSEDE and NSF, in concert with the SP Forum, should adopt agreements that address the approval and integration of SP-initiated policies and practices into the approved allocation policies and procedures.

The allocation policies implemented by XSEDE are formally “owned” and approved by NSF. Changes and modifications to these policies are conducted through extensive processes involving a wide range of stakeholders such as this one. In turn, XSEDE funds RAS to support ongoing allocation activity and approved plans of action to improve and enhance the infrastructure and process for users and to support a continually evolving ecosystem.

At the same time, separate acquisition solicitations are used to expand the ecosystem with additional resources, SPs, and SP Forum members. At times, these distinct processes can generate SP-proposed changes to allocation policies and supporting infrastructure, which have not been approved or planned for through the allocation policy approval and XSEDE planning processes.

NSF, XSEDE, and SPs need a process to rectify and rationalize the outcomes of these separate processes early in the SP award period, in order to ensure the successful execution of all affected programs and awards. This additional integration phase should include discussion and negotiation on general feasibility of proposed ideas, implementation alternatives or modifications, formal policy change approvals, scheduling timelines, and when necessary, changes to NSF agreements with XSEDE and SPs.

To minimize the impacts of potential disconnects, XSEDE is encouraged to take steps to better document and communicate its allocations policies and infrastructure from an SP perspective. Such documentation could be provided by NSF in support of relevant requests for proposals, for example.

D.3. Documentation and Tools

Along with the preceding recommendations, this group notes that XSEDE and RAS must continue to improve the XRAS interfaces and “radically simplify” the allocations documentation to reduce barriers to entry, including reducing the number of rejected proposals due to non-compliance with requirements, while improving support for and enforcement of existing policies.

This group recognizes that the final implementation plans for the recommendations in this report will eventually evolve into a number of changes to XRAS and documentation. To allow for the most flexibility in the final implementation, we have elected not to detail specific changes to the tools and documentation in this report, though the team does want to strongly communicate that XSEDE should work to simplify its allocations documentation and make it easier for users to find key information.

To exemplify the documentation improvements needed, we are including some comments from this group and the stakeholders who provided feedback. In particular, while the policies themselves are impartial to resources and resource types, the requests and reviews of different resource types need to apply appropriate interpretations of those policies. The general tone of the documentation and video guidance should reflect and encourage the range of activities supported by the variety of resources allocated through XSEDE.

It would be helpful to state the differential interpretation of policy as a logical approach to achieving the high-level goals for the allocation process. Namely, the allocations policies and procedures exist to ensure both that the valuable and oversubscribed traditional supercomputer resources supported for decades by the NSF are put to use in support of the highest priority and most meritorious research activities, and that novel resources and their emerging user communities are supported in the rapid adoption of these resources and in the development of new best practices for requesting and using these resources.

For example, the “Code Performance and Scaling” document has a number of shortcomings. While required by the review process, the document description and name are too focused on HPC-style resources and HPC-style usage modes. The guidance for this document should be revised to address at least two issues: First, “scaling” is relevant to a particular type of resource and usage mode and not the full range of resources available through the XSEDE allocations process. The documentation should clarify the performance or cost information deemed acceptable for different resources and different types of requests (e.g., gateway versus single-investigator). Second, a request for a given resource, per current policies and unchanged in these recommendations, is not absolutely required to show prior experience and benchmark timings on that resource (via a Startup allocation or Campus Champion access, for example). However, requests for highly competitive resources face more intensive scrutiny of cost information, and prior direct experience with a requested resource will help demonstrate the ability to make effective use of that resource. For novel and undersubscribed resources, on the other hand, hands-on access is preferred, but cost estimates based on use of comparable systems or other sources may be offered. In addition, the larger the size of such a request, the greater the expectation for more accurate cost information.

E. Actions and Implementation

Following approval of the recommendations in this report by NSF, the RAS team will define an implementation plan and schedule for updating the policy documentation, allocations processes, and supporting infrastructure. As appropriate, the products and deliverables of the implementation plan will be reviewed and approved by the XSEDE SMT and NSF program officer.

RAS will produce an initial high-level implementation plan and schedule spanning all the actionable recommendations within one quarter of this report being finalized. Implementation effort for specific recommendations will, in some cases, need further discussions within RAS and with external groups to finalize implementation details and schedule. This group will be leveraged as a primary sounding board for advice and guidance, along with external stakeholders.

Some activities in response to these recommendations have already begun. A draft of the reviewer rubric was prepared by a sub-committee of this group and presented to the XRAC at the August 2016 meeting for feedback. Early in the group's discussions, XRAS was updated with a checkbox for users to affirm that they had no supporting grants to list. An initial update to the current allocations documentation, consistent with RAS plans for Program Year 6, was begun in the first quarter of the program year.

F. Appendix

A draft of this document was shared with several stakeholder groups for their comments and feedback. This document represents the synthesis of that feedback, and this appendix summarizes the feedback received from each group.

XSEDE Resource Allocations Committee (XRAC)—While the XRAC was engaged along during the process and that group’s feedback incorporated where appropriate, we received no feedback specific to the draft report.

User Advisory Committee (UAC)—One UAC member suggested a policy change related to handling of gateway requests in the reconciliation phase of the Research allocations process. This group deemed the suggestion a strategic question that was beyond the scope of this group’s charter. The suggestion was raised separately to the XSEDE Senior Management Team for consideration and subsequently not included in the final recommendations.

Service Provider (SP) Forum—We received input from two SP Forum members, one of which identified mainly typographic errors. Of the more substantive suggestions, some were incorporated into the recommendations for documentation improvements. Two elements focused on clarifications for specific fields in the XRAS submissions forms were handled directly, outside of the policy review process.

XSEDE Senior Management Team (SMT)—A limited number of mostly minor edits and comments were provided by the SMT.

National Science Foundation (NSF)—The NSF program officer provided a number of suggestions and recommended edits. Several clarification edits were made, while a number of the suggestions more directly addressed implementation details outside the scope of this report. The implementation concerns will be included in the next phase of this activity.