

# Thoughts on the NSF Future Directions Interim Report

Glenn K. Lockwood, Ph.D.  
Oakland, California  
January 2015

## Introductory Remarks

---

The National Academies recently released an interim report entitled Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020 that was commissioned to take a hard look at where the NSF's supercomputing program is going. Since releasing the interim report, the committee has been soliciting feedback and input from the research community to consider as they draft their final report, and I felt compelled to put some of my thoughts into a response.

NSF's HPC programs are something I hold near and dear since I got my start in the industry by supporting two NSF-owned supercomputers. I put a huge amount of myself into Trestles and Gordon at the San Diego Supercomputer Center, and I still maintain that job encompassed the most engaging and rewarding work I've ever done. However, the NSF's lack of a future roadmap for its HPC program made my future feel perpetually uncertain, and this factored heavily in my decision to eventually pursue other opportunities.

Now that I am no longer affiliated with NSF, I wanted to delineate some of the problems I observed during my time as an NSF subcontractor with the hope that the community really thinks about how they can be addressed. The report requested feedback in nine principal areas, so I've done my best to contextualize my thoughts with the committee's findings.

## 1. How to create advanced computing infrastructure that enables integrated discovery involving experiments, observations, analysis, theory, and simulation.

---

Answering this question involves a few key points:

1. Stop treating NSF's cyberinfrastructure as a computer science research project and start treating it like research infrastructure operation. Office of Cyberinfrastructure (OCI) does not belong in Computer & Information Science & Engineering (CISE).
2. Stop funding cyberinfrastructure solely through capital acquisition solicitations and restore reliable core funding to NSF HPC centers. This will restore a community that is conducive to retaining expert staff.
3. Focus OCI/ACI and raise the bar for accountability and transparency. Stop funding projects and centers that have no proven understanding of operational (rather than theoretical) HPC.

4. Either put up or give up. The present trends in funding lie on a road to death by attrition.
5. Don't waste time and funding by presuming that outsourcing responsibility and resources to commercial cloud or other federal agencies will effectively serve the needs of the NSF research community.

I elaborate on these points below.

## 2. Technical challenges to building future, more capable advanced computing systems and how NSF might best respond to them.

---

*"New approaches that co-locate computational and data resources might reduce costs and improve performance. Recent advances in cloud data center design may provide a viable integrated solution for a significant fraction of (but not all) data- and compute-intensive and combined workloads."*

This strong statement is markedly unqualified and unsubstantiated. If it is really recommending that the NSF start investing in the cloud, consider the following:

- Cloud computing resources are designed for burst capabilities and are only economical when workloads are similarly uneven. In stark contrast, most well-managed HPCs see constant, high utilization which is where the cloud becomes economically intractable.
- The suggestion that cloud solutions can "improve performance" is unfounded. At a purely technological level, the cloud will never perform as well as unvirtualized HPC resources, period. Data-intensive workloads and calculations that require modest inter-node communication will suffer substantially.

In fact, if any cost reduction or performance improvement can be gained by moving to the cloud, I can almost guarantee that incrementally more can be gained by simply addressing the non-technological aspects of the current approach of operating federated HPC. Namely, the NSF must

1. Stop propping up failing NSF centers who have been unable to demonstrate the ability to effectively design and operate supercomputers.
2. Stop spending money on purely experimental systems that domain scientists cannot or will not use.

**The NSF needs to re-focus its priorities and stop treating the XD program like a research project and start treating it like a business.** Its principal

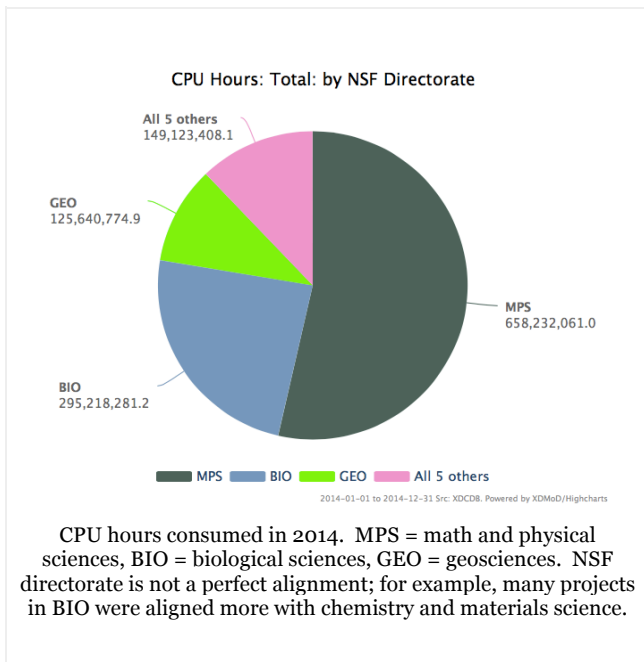
function should be to deliver a product (computing resources) to customers (the research community). Any component that is not helping domain scientists accelerate discovery should be strongly scrutinized. Who are these investments truly satisfying?

*"New knowledge and skills will be needed to effectively use these new advanced computing technologies."*

This is a critical component of XD that is extremely undervalued and underfunded. Nobody is born with the ability to know how to use HPC resources, and **optimization should be performed on users in addition to code**. There is huge untapped potential in collaborative training between U.S. federal agencies (DOE, DOD) and European organizations (PRACE). If there is bureaucratic red tape in the way, it needs to be dealt with at an official level or circumvented at the grassroots level.

### 3. The computing needs of individual research areas.

XDMoD shows this. **The principal workloads across XSEDE are from traditional domains like physics and chemistry, and the NSF needs to recognize that this is not going to change substantially over the lifetime of a program like XD.**



While I wholeheartedly agree that new communities should be engaged by lowering the barriers to entry, these activities cannot be done at a great expense of undercutting the resources required by the majority of XD users.

The cost per CPU cycle should not be deviating wildly between Track 2 awards because the ROI on very expensive cycles will be extremely poor. If the NSF wants

to fund experimental systems, it needs to do that as an activity that is separate from the production resources. Alternatively, only a small fraction of each award should be earmarked for new technologies that represent a high risk; the Stampede award was a fantastic model of how a conservative fraction of the award (10%) can fund an innovative and high-risk technology.

### 4. How to balance resources and demand for the full spectrum of systems, for both compute- and data-intensive applications, and the impacts on the research community if NSF can no longer provide state-of-the-art computing for its research community.

*"But it is unclear, given their likely cost, whether NSF will be able to invest in future highest-tier systems in the same class as those being pursued by the Department of Energy, Department of Defense, and other federal mission agencies and overseas."*

The NSF does not have the budget to support leadership computing. This is clear even from a bird's eye view: DOE ASCR's budget for FY2012 was \$428 million<sup>1</sup> and, by comparison, NSF ACI's budget was only \$211 million<sup>2</sup>. Worse yet, despite having half the funding of its DOE counterpart, the NSF owned HPC resources at seven universities in FY2012 compared to ASCR's three centers.

Even if given the proper funding, the NSF's practice of spreading Track 2 awards across many universities to operate its HPC assets is not conducive to operating leadership computing. The unpredictable nature of Track 2 awards has resulted in very uneven funding for NSF centers which, quite frankly, is a terrible way to attract and retain the highly knowledgeable world-class staff that is necessary to operate world-class supercomputers.

### 5. The role of private industry and other federal agencies in providing advanced computing infrastructure.

The report makes some very troubling statements in reference to this question.

*"Options for providing highest-tier capabilities that merit further exploration include purchasing computing services from federal agencies..."*

This sounds dirty. Aren't there are regulations in place that restrict the way in which money can flow between the NSF and DOE? I'm also a little put off by the fact that this option is being put forth in a report that is crafted by a number of US DOE folks whose DOE affiliations are

<sup>1</sup> *Advanced Scientific Computing Research Funding Profile by Subprogram and Activity.*  
<http://science.energy.gov/~media/budget/pdf/sc-budget-request-to>

<sup>2</sup> *Directorate for Computer and Information Science and Engineering (CISE).*  
[http://www.nsf.gov/about/budget/fy2014/pdf/18\\_fy2014.pdf](http://www.nsf.gov/about/budget/fy2014/pdf/18_fy2014.pdf)

masked by university affiliations in the introductory material.

*"...or by making arrangements with commercial services (rather than more expensive purchases by individual researchers)."*

Providing advanced cyberinfrastructure for the open science community is not a profitable venture. **There is no money in HPC operations.** I do not see any "leadership" commercial cloud providers offering the NSF a deal on spare cycles, and the going rate for commercial cloud time is known to be far more expensive than deploying HPC resources in-house at the national scale<sup>3</sup>.

## **6. The challenges facing researchers in obtaining allocations of advanced computing resources and suggestions for improving the allocation and review processes.**

*"Given the "double jeopardy" that arises when researchers must clear two hurdles—first, to obtain funding for their research proposal and, second, to be allocated the necessary computing resources—the chances that a researcher with a good idea can carry out the proposed work under such conditions is diminished."*

XD needs to be more tightly integrated with other award processes to mitigate the double jeopardy issue. I have a difficult time envisioning the form which this integration would take, but the NSF GRF's approach of prominently featuring NSF HPC resources as a part of the award might be a good start. As an adaptive proposal reviewer within XSEDE and a front-line interface with first-time users, I found that having the NSF GRF bundle XSEDE time greatly reduced the entry barrier for new users and made it easier for us reviewers to stratify the proposals. Another idea may be to invite NSF center staff to NSF contractors' meetings (if such things exist; I know they do for DOE BES<sup>4</sup>) to show a greater amount of integration across NSF divisions.

In addition, the current XSEDE allocation proposal process is extremely onerous. The document that describes the process<sup>5</sup> is ridiculously long and contains of obscure requirements that serve absolutely no purpose. For example, all XSEDE proposals require a separate document detailing the scaling performance of their scientific software. Demonstrating an awareness of the true costs of performing certain calculations has its merits, but a detailed analysis of scaling is not even relevant for the majority of users who run modest-scale jobs or use off-the-shelf black-box software like Gaussian. The only thing these obscure requirements do is prevent new users, who are generally less familiar with

all of the scaling requirements nonsense, from getting any time. If massive scalability is truly required by an application, the PI needs to be moved over to the Track 1 system (Blue Waters) or referred to INCITE<sup>6</sup>.

As a personal anecdote, many of us center staff found ourselves simply short-circuiting the aforementioned allocations guide and providing potential new users with a guide to the guide. It was often sufficient to provide a checklist of minutia whose absence would result in an immediate proposal rejection and allow the PIs to do what they do best—write scientific proposals for their work. Quite frankly, the fact that we had to provide a guide to understanding the guide to the allocations process suggests that the allocations process itself is grossly over-engineered.

## **7. Whether wider and more frequent collection of requirements for advanced computing could be used to inform strategic planning and resource allocation; how these requirements might be used; and how they might best be collected and analyzed.**

The XD program has already established a solid foundation for reporting the popularity and usability of NSF HPC resources in XDMoD<sup>7</sup>. The requirements of the majority are evolving more slowly than computer scientists would have everyone believe.

Having been personally invested in two Track 2 proposals, I have gotten the impression that the review panels who select the destiny of the NSF's future HPC portfolio are more impressed by cutting edge, albeit untested and under-demanded, proposals. Consequentially, taking a "functional rather than a technology-focused or structural approach" to future planning will result in further loss of focus. Instead of delivering conservatively designed architectures that will enjoy guaranteed high utilization, **functional approaches will give way to computer scientists on review panels dictating what resources domain scientists should be using** to solve their problems. The cart will be before the horse.

Instead, it would be far more valuable to include more operational staff in strategic planning. The people on the ground know how users interact with systems and what will and won't work. As with the case of leadership computing, **the NSF does not have the financial commitment to be leading the design of novel computing architectures at large scales.** Exotic and high-risk technologies should be simply left out of the NSF's Track 2 program, incorporated peripherally but funded through other means (e.g., MRIs), or incorporated in the form of a small fraction of a larger, lower-risk resource investment.

A perspective of the greater context of this has been eloquently written by Dr. Steven Gottlieb<sup>8</sup>. Given his

<sup>3</sup> *The Magellan Report on Cloud Computing for Science.* <http://www.alcf.anl.gov/magellan>

<sup>4</sup> *U.S. Department of Energy Office of Science – Materials Science and Engineering (MSE) Division Principal Investigators' Meetings.*

<http://science.energy.gov/bes/mse/principal-investigators-meetings/>

<sup>5</sup> *XSEDE Allocation Policies.* <https://portal.xsede.org/allocation-policies>

<sup>6</sup> *U.S. Department of Energy Leadership Computing - INCITE.* <http://www.doeleadershipcomputing.org/>

<sup>7</sup> *XSEDE Metrics on Demand.* <https://xdmod.ccr.buffalo.edu/>

<sup>8</sup> *Whither the future of NSF Advanced Computing Infrastructure?*

description of the OCI conversion to ACI, it seems like taking away the Office of Cyberinfrastructure's (OCI's) autonomy and placing it under Computer & Information Science & Engineering (CISE) exemplifies an ongoing and significant loss of focus within NSF. This changed reflected the misconception that architecting and operating HPC resources for domain sciences is a computer science discipline.

This is wrong.

Computer scientists have a nasty habit of creating tools that are intellectually interesting but impractical for domain scientists. These tools get "thrown over the wall," never to be picked up, and represent an overall waste of effort in the context of operating HPC services for non-computer scientists. Rather, operating HPC resources for the research community requires experienced technical engineers with a pragmatic approach to HPC. Such people are most often not computer scientists, but former domain scientists who know what does and doesn't work for their respective communities.

### **8. The tension between the benefits of competition and the need for continuity as well as alternative models that might more clearly delineate the distinction between performance review and accountability and organizational continuity and service capabilities.**

*"Although NSF's use of frequent open competitions has stimulated intellectual competition and increased NSF's financial leverage, it has also impeded collaboration among frequent competitors, made it more difficult to recruit and retain talented staff, and inhibited longer-term planning."*

Speaking from firsthand experience, I can say that **working for an NSF center is a life of a perpetually uncertain future and dicing up FTEs into frustratingly tiny pieces**. While some people are driven by competition and fundraising (I am one of them), an entire organization built up to support multi-million dollar cyberinfrastructure cannot be sustained this way.

At the time I left my job at an NSF center, my salary was covered by six different funding sources at levels ranging from 0.05 to 0.30 FTEs. Although this officially meant that I was only 30% committed to directly supporting the operation of one of our NSF supercomputers, the reality was that I (and many of my colleagues) simply had to put in more than 100% of my time into the job. This is a very high-risk way to operate because committed individuals get noticed and almost invariably receive offers of stable salaries elsewhere. Retaining talent is extremely difficult when you have the least to offer, and the current NSF funding structure makes it very difficult for centers to do much more than continually hire entry-level people to replace the rising stars who find greener pastures.

<http://www.computer.org/cms/Computer.org/ComputingNow/docs/CISE-17-02-EIC.pdf>

Restoring reliable, core funding to the NSF centers would allow them to re-establish a strong foundation that can be an anchor point for other sites wishing to participate in XD. This will effectively cut off some of the current sites operating Track 2 machines, but frankly, **the NSF has spread its HPC resources over too many sites at present and is diluting its investments** in people and infrastructure. The basis for issuing this core funding could follow a pattern similar to that of XD where long-term (10-year) funding is provisioned with a critical 5-year review.

If the NSF cannot find a way to re-establish reliable funding, it needs to **accept defeat and stop trying to provide advanced cyberinfrastructure**. The current method of only funding centers indirectly through HPC acquisitions and associated operations costs is unsustainable for two reasons:

- The length of these Track 2 awards (typically 3 years of operations) makes future planning impossible. Thus, this current approach forces centers to follow high-risk and inadequately planned roadmaps.
- All of the costs associated with maintaining world-class expertise and facilities have to come from someone else's coffers. Competitive proposals for HPC acquisitions simply cannot afford to request budgets that include strong education, training, and outreach programs, so these efforts wind up suffering.

### **9. How NSF might best set overall strategy for advanced computing-related activities and investments as well as the relative merits of both formal, top-down coordination and enhanced, bottom-up process.**

Regarding the top-down coordination, the NSF should drop the Track 2 program's current solicitation model where proposers must have a vendor partner to get in the door. This is unnecessarily restrictive and fosters an unhealthy ecosystem where vendors and NSF centers are both scrambling to pair up, resulting in high-risk proposals. Consider the implications:

1. Vendors are forced to make promises that they may not be able to fulfill (e.g., Track 2C and Blue Waters). Given these two (of nine) solicitations resulted in substantial wastes of time and money (over 20% vendor failure rate!), I find it shocking that the NSF continues to operate this way.
2. NSF centers are only capable of choosing the subset of vendors who are willing to play ball with them, resulting in a high risk of sub-optimal pricing and configurations for the end users of the system.

I would recommend a model, similar to many European nations', where a solicitation is issued for a vendor-neutral proposal to deploy and support a program that is built around a resource. A winning proposal is selected based

on not only the system features, its architecture, and the science it will support, but the plan for training, education, collaboration, and outreach as well. Following this award, the bidding process for a specific hardware solution begins.

This addresses the two high-risk processes mentioned above and simultaneously eliminates the current qualification in Track 2 solicitations that no external funding can be included in the proposal. By leaving the capital expenses out of the selection process, the NSF stands to get the best deal from all vendors and other external entities independent of the winning institution.

Bottom-up coordination is much more labor-intensive because it requires highly motivated people at the grassroots to participate. Given the NSF's current inability to provide stable funding for highly qualified technical staff, I cannot envision how this would actually come together.