

**COMMITTEE OF VISITORS
REPORT**

Advanced Networking Research Program

October, 2011

Date of COV: October 10-11, 2011

Program: Advanced Networking Program

Office: Advanced Scientific Computing Research (ASCR)

Agency: United States Department of Energy

Committee Membership:

Prof. Victor Frost, University of Kansas

Prof. Peter Steenkiste, Carnegie Mellon University

Prof. Gregory E. Monaco, Ph.D., The Great Plains Network

Prof. George N. Rouskas, North Carolina State University

Ms. Victoria White, FermiLab

Prof. Ty Znati, University of Pittsburgh

o. Executive Summary

The Committee of Visitors (COV) met to review the management processes for the Next Generation Networking for Science (NGNS) elements of the Office of Advanced Scientific Computing Research (ASCR) program at the DOE Germantown location on Tuesday, October 11, 2011.

The COV is extremely grateful to the program officers and other ASCR staff who gave graciously of their time and knowledge to help the COV in its deliberation. Their support was instrumental in enabling a smooth and effective review process.

Finding and Recommendation:

Based on the presentations and discussion with the NGNS office, the COV considers the Networking programs under review to be generally effective and reasonably well managed.

- **Findings:** The COV found that the process used to generate solicitation appears to be effective and fairly well-administrated. The COV also found that the review process was conducted in accordance with the DOE normal standards of peer review and that the program officers are clearly dedicated and competent, demonstrating significant initiative, leadership and considerable knowledge in managing the program.
- **Recommendations:** The program should move from a passive to active notification mechanism for workshop participation and funding opportunities, broaden participation in workshops and establish predetermined schedule dates for solicitation announcements and proposal due dates. The program should also provide greater visibility into the award process, budget and scope reductions and the tracking of progress for the next COV. The NGNS is encouraged to automate the archiving of interactions between program managers and PIs, reports from site visits, and progress reports in a single easily-accessible repository.
- **Findings:** The NGNS program contributes to DOE's leadership role in the discovery, development, and deployment of forefront computing and networking capabilities; and it is internationally recognized for the impact of its high-quality research results. The NGNS has engaged top-level network researchers and large-scale, high-performance network infrastructure developers both in first-class research and innovations and in persistent development that lead to world-class networking and middleware capabilities.
- **Recommendations:** The NGNS Program should continue to address networking challenges closely associated with the DoE mission, and continue its leadership in high performance networking and middleware. The Program should further engage the broader CS community in its research initiatives and workshops, continue to coordinate their networking research and development efforts with other funding agencies, and continue to expand its collaboration with the international community. The Program should establish clear strategic plans regarding future funding allocations between long-term fundamental research, near-term research and development, and testbed support.

The report discusses several additional recommendations, which are offered in the spirit of improving the Program's processes and to ensure continued leadership in high-performance networking and middleware.

1. Introduction

The Advanced Scientific Computing Advisory Committee (ASCAC) for the Office of Science, United States Department of Energy (DOE), was charged by W. F. Brinkman, Director of Office of Science, with assembling a Committee of Visitors (COV) to review the management processes for the Next Generation Networking for Science (NGNS) elements of the ASCR Networking Research Program. The program is focused on high performance networking tools and middleware to help the ASCR research communities utilize the capabilities of current and future computing infrastructure.

A COV of five members was formed to review NGNS elements. One member of the COV, Dr. Vicky White from Fermi National Accelerator Laboratory (Fermilab), is the ASCAC member on the committee. The list of all participants in the COV is provided in Attachment 1 and the letter charge to the committee from the chair of the ASCAC is provided in Attachment 2.

This report presents the findings and recommendations of the COV. The review covers the operations of the Networking Research Program during the fiscal years 2008, 2009, and 2010, with a focus on (i) the processes used to solicit and review research proposals, recommend awards, and manage research projects and (ii) the breadth and depth of the research and development portfolio to support distributed, collaborative high-end science.

Prior to the meeting, the COV was provided with a link to a secure website with information for the COV members to review prior to the COV meeting. The available information included the ASCR functional organization chart, the ASCR merit review procedures, spreadsheets listing information about the proposals submitted during the fiscal years of 2008-2010, the Funding Opportunities Announcements (FOA) for the different Networking Research Programs under review, including the Early Career Research Program, for both DOE laboratory projects and university projects, and documents related to the proposals submitted to these announcements, including research project descriptions, proposal reviews and recommendations. Due to cyber security concerns, information related to some proposals, such as reviewer names and content of the reviews, was excluded. The COV, however, was informed that full access to all of this information will be provided during the COV meeting.

In preparation for the main COV meeting, the COV held a planning meeting at the Fairfield Inn, 20025 Century Boulevard, Germantown, Maryland, from 4:00 p.m. to approximately 7:30 p.m. on Monday, October 10. Prior to this meeting, the COV members were requested to read the informational documents about the Networking Research Program available on the website and review a set of proposals randomly selected from the different programs.

During the planning meeting, the COV considered the material that was reviewed by the COV members, discussed the logistics and scope of the review process, finalized the overall structure of the COV report in relation to the COV charge, and identified a number of issues that the COV members needed additional clarification about from the NGNS office. A COV member was then assigned to lead the discussion related to a specific set of issues during the COV meeting.

The COV meeting was held at the DOE Germantown location on Tuesday, October 11, 2011. The meeting opened with a series of presentations. Dr. Christine Chalk provided information about the logistics of the meeting. She also provided a summary of the charge and briefly discussed the DOE conflict of interest policy. Dr. Barbara Helland reviewed the functional organization of the ASCR office. She also discussed the ASCR program mission and priorities, emphasizing the role of the Networking Research Program in the discovery and deployment of the forefront computational and networking capabilities needed to support advanced research in the physical sciences.

Following Dr. Helland's presentation, the two Program Managers, Dr. Richard Carlson and Dr. Thomas Ndousse-Fetter, presented an overview of the Next Generation Networks for Science program, focusing on the specific program's goals, the approach used to evaluate proposals and the major accomplishments of the program during the fiscal years under review. The emphasis for the briefing was on the two main elements of the program, namely High-Performance Networks and High-Performance Middleware. The Program Managers described the capabilities of ESnet, a primary provider of network connectivity for the Office of Science, including its core IP backbone and the recently developed Science Data Network (SDN); the latter provides an independent backup to the IP backbone and allows access to dynamically provisioned bandwidth to carry the national labs massive data flows. The Program Officers also discussed the ESnet projected capabilities, the increasing need for advanced middleware that is critical to the productivity of scientific collaborations, and the widening gap between the High Performance Computing and High-speed Network infrastructure. The last part of the briefing elaborated on how FOAs are developed, how proposals are reviewed, how decisions to approve or decline proposals are made and how post-award management and monitoring are carried out. A description of the current portfolio was also provided, including the portfolio distribution of collaborative research projects – involving more than two PIs – and single PI research projects; the affiliation of the funded researchers across national labs, universities and industry; and the portfolio distribution of long-term projects, short-term projects and testbed activities.

Following the presentations, the COV met with the Program Managers to clarify issues related to the website material, discuss preliminary findings and request further information to aid in the development of the COV report. The requested material included reports of the research requirements workshops, the “High-Performance Networks for Distributed Petascale Science” and the “High-Capacity Optical Networking and Deeply Integrated Middleware Services for Distributed Petascale Science” FOAs, an overview of budget data, statistics and information relevant to review panel membership, Letters of Intent, award progress and final reports. The Program Managers provided links to the missing FOAs and workshop reports, produced documents for budget data and discussed additional information related to proposals and awards.

After the program summary, the COV met in executive session to discuss and develop a synopsis of the COV's findings and recommendations. Before the end of the meeting, the COV met briefly with Dr. Helland to briefly summarize their findings. The final report was prepared using e-mail exchanges between the COV members. The COV members were grateful for the active and helpful engagement of the ACSR program managers throughout the review process.

2. CoV Charge

The specific charge to the COV included the evaluation of the following two major program elements:

1. For both the DOE laboratory projects and the university projects, an assessment of the efficacy and quality of the processes used to:
 - a. Solicit, review, recommend proposals;
 - b. Document proposal actions; and
 - c. Monitor active projects and programs
2. Within the boundaries defined by DOE missions and available funding, an assessment of how the award process has affected:
 - a. The breadth and depth of portfolio elements, and
 - b. The national and international standing of the program with regard to other computer science research programs that are also focused on high performance networking tools and middleware for science.

3. Efficacy and Quality of Processes

The COV considers the NGNS program under review to be generally effective and reasonably well managed. The objectives of the High-Performance Networks and High-Performance Middleware Program elements are well-aligned with DoE's priority to deliver the forefront computational and networking capabilities needed to enable world class research in the physical sciences and to facilitate collaboration among scientists across the world. The program officers are clearly dedicated and competent public servants. They have demonstrated significant initiative and leadership in managing the program and have shown considerable knowledge of their respective portfolios and communities of practice.

The NGS program has achieved significant impact in the deployment of the DOE's network infrastructure, including ESnet, DOE Laboratories Local Enterprise Networks, and Host systems. This was achieved through the transfer and successful deployment of major findings of funded research projects. Highlights of funded programs include advances in the design of high-speed transport protocols, dynamic provisioning of high-capacity network bandwidth between end systems, control and monitoring protocols for network performance, and advancing the state-of-the-art in virtual networks.

Charge 1(a): Assess the efficacy and quality of the processes used to solicit, review, and document applications and proposal actions.

The COV found that the solicitation and review processes appear to be effective and fairly well-administrated.

Solicitation Development

Science trends for Network R & D are gathered from the outcome of a pool of workshops in domain science areas (e.g. math, physics). These general trends are then incorporated in a

charge to a formal network workshop committee. The workshop committee incorporates these into the Networking workshop where working groups identify networking needs related to general scientific trends. Relevant findings of the networking workshop are then incorporated into a solicitation or request for proposals to the community of researchers. NGNS uses this process to assess ESNNet requirements, core network research and scientific collaboration. Attendees at these workshops include scientists, network researchers, and industry. Through these workshops a set of Science Networking Requirements is developed, which are prioritized by the program as input to the solicitations.

The workshop committee members invite attendees to participate in the workshop. Program Managers indicated that they are at liberty to accept requests to attend workshops from individuals who show interest. The importance of an invitation to such a workshop is indicated by the fact that over 25% (18 of 65) of the proposals submitted in response to the subsequent September, 2008, solicitation that followed from this workshop were from workshop attendees serving as PI (other proposals that did not have a workshop attendee as PI did include an attendee as a co-PI).

It is not clear if the broadest participation of the network research community is obtained both in terms of requirements development, e.g., timely notification and encouragement to attend workshops, and response to solicitations, e.g., passive notification of new research opportunities. Inconsistent timing of solicitations may also be a barrier to obtaining broadest participation in the program.

Solicitations are posted on-line. They are open for at least thirty days, with a general call of unsolicited proposals that is open throughout the year. Solicitations are also advertised in the Federal Registry and on Grants.gov. The timing for proposal submission, however, is not consistent; it varies considerably, is not uniform, and is driven by varying conditions. In one case, the solicitation required submission of a LOI, prior to submitting a research proposal. Proposals, which did not comply with the LOI requirement, however, were accepted and reviewed, making the role of these letters ambiguous.

Documents on the website describe four basic review criteria that cover technical merit, appropriateness of method and approach, competence of the team, and the budget. Program managers can add additional criteria specific to each solicitation. The COV had access to four solicitations. Two of the solicitations (395, 09-26) added two criteria to the four basic criteria. Both these solicitations were for the Early Career Research Program, and the additional criteria cover relevance to the mission of the program and potential for leadership. One solicitation (246) did not add specific criteria; it did, however, include a broad statement that factors, such as relevance to the mission, would be considered. It also encourages, in a separate section, collaboration between institutions, including collaboration with DOE National Laboratories. Finally, one solicitation did not mention the merit criteria.

DoE publishes annually a general call seeking unsolicited proposals. This call invites PIs in EPSCoR institutions to submit proposals for special consideration. When an EPSCoR-designated proposal is submitted to the program it goes directly to the EPSCoR office and is

then redirected to the specific program area for consideration and co-funding. It is not clear from the program officers how one takes or avoids the EPSCoR path when submitting an unsolicited proposal.

Recommendations:

- The program should move from a passive to active notification mechanism for workshop participation and funding opportunities, broaden participation in workshops and establish predetermined schedule dates for solicitation announcements and proposal due dates. To this end, the NGSN office is encouraged to develop a mechanism to notify the community, including state and regional research and education networks, of the opportunity to participate in all phases from workshop and solicitation development to solicitation announcement.
- The NSGS Program should maintain some consistency in handling the review criteria across solicitations, clarify the intent and role of the LOIs in the review process, and follow that intent.
- The NSGS Program should consider including in the solicitations statements that clearly and explicitly document specific research and development expectations, such as deployment on ESnet and other DOE high-performance networking infrastructure. Such a practice will be very helpful in setting the right scope for long- and short-term research projects and ensure fair and appropriate reviews of the proposals.

Review Process

The COV found that the review process was conducted in accordance with the DOE normal standards of peer review. Proposals are evaluated using a combination of mail-in reviews and panel reviews. Normally, panel reviews are used (possibly combined with mail in reviews from experts), although if a solicitation has few proposals, evaluation may be purely based on mail-in reviews. In both mail-in and panel review cases, a minimum of three reviews for each proposal was the standard across all FOAs.

Program managers select mail-in reviewers and panelists based on their expertise and background in the areas relevant to the solicitations and the DOE mission. Given that the merit criteria are very broad, the reviewer pool is very diverse. Reviewers typically come from academia, government labs, and industry. The information the COV received about the reviewers' pool used to review proposals in the fiscal years of 2008 to 2010 shows that mail and panel reviewers have appropriate expertise and background.

Each reviewer evaluates the proposal with respect to the merit criteria for the solicitation and the reviewer also assigns an overall score for the proposal in the range 1-10. Panels discuss 20-25 proposals, on average, and each panelist typically reviews about 5 proposals. At the start of each panel, the program manager gives a presentation focused on reviewing the merit criteria and discussing the panel objectives and panel process.

The panel discusses each proposal, but does not rank proposals or assign them to specific funding categories. Panelists are encouraged, but not required, to update their proposal ranking

upon listening to the panel discussion. No panel consensus ranking is required. PIs get a copy of the reviews when they are informed of the decision, but they are not given the numeric scores assigned by the reviewers.

The COV found that the panel size seems to be reasonable. The COV also found that the reviewers were well informed and knowledgeable of the state-of-the art related to the proposals they were assigned to review. The ultimate decisions made by the Program Manager were found to be, in general, reasonable and justified. The NGNS Program Managers use the panel review mechanism and in-mail review extensively to make approval or decline decisions, taking into consideration DOE's mission and priorities and the need to maintain a balanced portfolio. The review of a proposal sample shows that the decision of the program manager was reasonably in line with the comments of the reviewers, although in some case the program manager may have weighted the competency and prior successes of the investigators rather heavily.

The COV found it difficult to understand all of the factors that go into an award decision because so much of the negotiation with the proponents is not recorded and a large number of factors, aside from the merit of the proposal as judged by the reviewers, enter into the decision. However, the portfolio of awards appears to be mission driven. The projects directly address issues of concern to the DOE and seek to develop networking and middleware technology that has potential for near term deployment in DOE high-performance network infrastructure.

The COV review of proposal documents revealed differences in the way approved and declined proposal decisions are handled. The unevenness in the level of the documentation associated with declined and approved proposals made it difficult for the COV to gain thorough understanding of the decline process, particularly for highly-ranked, declined proposals.

While the proposal review panels do not have to reach consensus on proposal rankings or produce a panel summary that documents the discussion about a particular proposal during the panel meeting, it appears that the panel discussion plays a significant role in determining the strengths and weaknesses of the proposals. The Program Manager has considerable flexibility in taking the advice of the panel and mail reviewers and formulating a proposal to senior ASCR management on what to fund. As such, the Program Manager has to exercise a significant level of discretion in reconciling disparities between reviewers' panel scores in order to make final decisions. The COV found that reviewer written comments about a proposal, rather than the scores provided by the reviewers, are weighted most heavily. Priorities of ASCR and the need to balance the program both in subject matter, short or long term deliverables and needs of other program offices are also used in making the recommendation. ASCR senior management has the final approval for awards to be made.

In some solicitations, e.g., 08-27 solicitation, Letters of Intent were required: "A one-page Letter of Intent (LOI) is **REQUIRED** and must be received by October 31, 2008, 5:00 p.m., Eastern Time... A response to the Letters of Intent encouraging or discouraging formal proposals will be communicated to the applicants **by November 14, 2008. Formal proposals will be accepted only from those encouraged to submit. No other formal proposals will be considered**". Proposals, from applicants who have not submitted LOIs as required by the

solicitation, appear to have been accepted and reviewed. This, however, may be due to an archiving problem.

Recommendations:

- There is no central repository that lists potential reviewers. Such a repository would constitute an official information system, compliant to federal privacy and information gathering rules. NGNS is encouraged to harness current systems and systems under development to develop such a repository to aid Program Managers in forming high-quality, diversified reviews panels for each solicitation. The reviewers' pool should include both reviewers who are capable of evaluating proposals that are focused on short-term DOE requirements as well as reviewers who are capable of evaluating longer-term proposals with high payoff.
- It would be useful to ask panels to collectively develop a short summary document reflecting the discussion of each proposal. The goal of the summary is not to report a consensus recommendation, since this is not the objective of the panel, but rather to identify points of agreement and of contention that may be useful for the program manager in making the final decision. It may also be useful to make this discussion summary available to the PI as part of the review material, particularly for young investigators.
- The COV was pleased to learn of NGNS participation in the Early Career PI Program, but was disappointed that no proposals were funded under this program. This is mostly due to the failure of these proposals to address DOE's mission and priorities. The COV believes that NGNS will benefit greatly from the Early Career PI Program and encourages the Program Managers to find ways to reach out and clearly convey the objectives and priorities of the NGNS Program to young investigators. The NGNS program management team is also encouraged to periodically revisit the balance between long term and short term research. Longer term research may also provide an opportunity to engage and attract young investigators.
- The rules, with respect to how proposal declinations are handled, have apparently changed over the years. While writing a detailed review analysis for each declination would place too high of a burden on Program Managers, limiting this requirement to highly-rated, declined proposals would be useful in assessing overall funding decisions and helpful to PIs in future submissions.

Charge 1(b): Assess the efficacy and quality of the processes used to monitor active awards, projects and programs.

An important role of the NGNS office is to monitor progress of the research projects funded by the Program. The COV found that the NGNS Program Managers use effective mechanisms to monitor ongoing awarded projects, including progress and final reports, site visits, and PI teleconference meetings. Such a monitoring is essential to ensure that the research goals and milestones of the funded project are achieved in a timely and satisfactory fashion.

The COV recognizes that the NGNS program competes with other programs of ASCR in the internal budget process to determine the amount of money available for new activities, given

prior year commitments. As such, budget cuts are expected, in budget-limited programs. The negotiation of awards in the NGNS Program, in cases where the full budget requested is not granted, is done informally through discussion and correspondence between the lead of the proposal and the program manager. The final revised budget is recorded in the folder along with a revised budget justification statement. Program managers stated that they ensured that cuts in budget were reflected in reduced work and deliverables. In the folders reviewed, however, the COV could not see evidence of proponents being asked to address reviewers concerns prior to award, but this could be because such dialogue is not recorded, as noted above.

Awarding of grants, after decisions are made, is a lengthy process and ASCR must get final procurement paperwork done by June 30th for an award in the current FY. A compressed timeline and process sometimes has to be used in a year of continuing resolution in order to squeeze solicitation/review/award process into the period between funding availability but before June 30th. Awards are usually made within 6 months of the proposal deadline.

Award management and award progress monitoring is achieved based on PI annual progress reports and final reports and on holding regular PI teleconference meetings. The COV found that annual progress reports are not available online for a large number of funded projects. The COV review shows that Lab and University proposals are handled in the same way and held to the same standards. The follow up on addressing reviewer concerns is handled informally, and no notes can be found in the files to assess this aspect of the process. The COV commends the Program Managers for active participation in the monitoring process and for enabling good interactions to take place, not only between program managers and PIs, but equally important between different members of a research project.

Recommendation

- The COV recommends greater visibility into the award process, budget and scope reductions and the tracking of progress. To this end, NGNS is encouraged to automate the archiving of interactions between program managers and PIs, reports from site visits, and progress reports in a single easily-accessible repository.

4. Effect of the Award Process on Portfolio

Charge II: Within the boundaries defined by DOE missions and available funding, assess how the award process has affected the breadth, depth, and national and international standing of the portfolio elements.

The NGNS has engaged top-level network researchers and large-scale, high-performance network infrastructure developers both in first-class research and innovations and persistent development that lead to world-class networking capability to enable unprecedented science critical to DOE missions and priorities.

The research portfolio includes several projects that are unique in addressing challenges related to the DoE infrastructure, including projects on provisioning circuits alongside the IP infrastructure, creating dynamic circuits, on-the-fly, without the need for human involvement,

and co-scheduling of resources. Such projects are at the forefront of research and development and have been adopted or are being considered for adoption by national or international groups outside DoE.

The research portfolio is of high quality and addresses challenges that are distinctly relevant to the mission of the DOE. Although fundamental research comprises a relatively small fraction of the overall portfolio, the selected projects are essential for enhancing the capabilities of the DOE infrastructure and of strategic importance to the DOE mission. Within the budget constraints, the program has been able to identify and fund projects to develop technology that can be readily deployed on the DoE infrastructure so as to address the needs of the science community.

The NGNS program contributes to DOE's leadership role in the discovery, development, and deployment of forefront computing and networking capabilities and the program is internationally recognized for the impact of its research results and the high quality of its visionary workshop reports. The outcome of several of the research projects funded by the NGNS Program, such as PerfSonar, On Demand Secure Circuits and Reservation System (OSCARS), and other network control protocols and middleware, is prominent in the international community and recognized as world-class. Research work produced by these projects is highly cited in prestigious national and international high-performance networking conferences.

Recommendations:

- The solicitations should continue to address networking challenges closely associated with the DOE mission, and funding must continue to be allocated by taking into account projects' impact on the DOE mission and objectives. Given the relatively small size of the program and future fiscal challenges, the NGNS office should establish clear strategic goals, regarding future funding allocations between long-term fundamental research, near-term research and development, and testbed support.
- The Program should further engage the broader CS community in its research initiatives and workshops, and find effective ways to nurture and engage the next generation of leading network researchers and developers in creative and effective network research and development within the context of DOE's mission goals and priorities.
- The NGNS Program is encouraged to continue to expand and enhance international collaborations in DOE's network research and development priority areas and foster leading-edge partnerships throughout the global networking community.
- The NGNS office has played an active role in coordinating their networking research and development efforts with other funding agencies. The NGNS Program is encouraged to continue to develop synergistic and collaborative activities with other federal funding agencies, to leverage resources across all agencies, and expedite discovery and deployment of innovative high-performance networking infrastructure.

NGNS Committee of Visitors

Prof. Victor Frost, University of Kansas
Lawrence, Kansas 66045
phone: (785) 864-4833
fax: (785) 864-7789
Email: frost@eecs.ku.edu
<http://www.ittc.ku.edu/~frost/>

Prof. Peter Steenkiste, Carnegie Mellon University
Voice: 412-268-3261
Fax: 412-268-5576
E-mail: prs@cs.cmu.edu

Prof. Gregory E. Monaco
Director for Research & Cyberinfrastructure Initiatives
The Great Plains Network
Kansas State University
Phone: 785-554-9657
Teleconference: 877-366-0711 Participant Code: 57348838#
Email: greg@greatplains.net
<http://www.greatplains.net>

Prof. George N. Rouskas
North Carolina State University
Raleigh, NC 27695-8206
Email: rouskas@ncsu.edu
Tel: 919-515-3860
Fax: 919-515-7896

Prof. Vicky White, Fermilab
MS 120, P.O. Box 500, Batavia,
Illinois, 60510
Tel : Office Fax: 630-840-2783
Email: white@fnal.gov
<http://www-do.fnal.gov/~white/>

Prof. Ty Znati, University of Pittsburgh
Senott Square
University of Pittsburgh,
Pittsburgh, PA 15260
Email: znati@cs.pitt.edu
Tel: 412-624-8417