U.S. Department of Energy
FIELD WORK PROPOSAL

<table>
<tr><td>1. WORK PACKAGE NUMBER
10053</td><td>2. REVISION NO.
0</td><td>3. DATE PREPARED
03/31/2010</td></tr>
<tr><td>4. WORK PACKAGE TITLE
SLAC National Accelerator Laboratory Collaborative DOE Enterprise Network Monitoring Deployment</td><td colspan="2">5. BUDGET AND REPORTING CODE
KJ0101040</td></tr>
<tr><td>6. WORK PROPOSAL TERM
10/01/2010-09/30/2013</td><td colspan="2">7. IS THIS WORK PACKAGE INCLUDED IN THE INSTITUTIONAL PLAN?
Yes</td></tr>
<tr><td>8. DOE PROGRAM MANAGER
Dr. Thomas N'dousse-Fetter
Phone: 301-903-9960
E-Mail: tndousse@ascr.doe.gov
Network Research
SciDAC,ASCR
Office of Science
US Department of Energy
100 Independence Avenue
Washington, DC 20585</td><td>11. HEADQUARTERS ORGANIZATION
Office of Advanced Scientific Computing Research (ASCR), U.S. Department of Energy (DOE)</td><td>14. DOE ORGANIZATION CODE</td></tr>
<tr><td>9. OPERATIONS OFFICE WORK PROPOSAL REVIEWER
Paul Golan</td><td>12. OPERATIONS OFFICE
SLAC Site Office</td><td>15 DOE ORGANIZATION CODE</td></tr>
<tr><td>10. CONTRACTOR WORK PROPOSAL MANAGER
Dr. Persis Drell
SLAC National Accelerator Laboratory
650-926-8704</td><td>13. CONTRACTOR NAME
Board of Trustees
Stanford University
SLAC National Accelerator Laboratory</td><td>16. CONTRACTOR CODE
55</td></tr>
<tr><td colspan="3">17. WORK PROPOSAL DESCRIPTION (Approach, anticipated benefit in 200 words or less)

Providing a coherent view of the status of the DOE enterprise network is a critical requirement in identifying performance bottlenecks. This proposal seeks to deploy a distributed monitoring infrastructure that will facilitate performance tuning and diagnosis of network problems by integrating information from the multi-domain, heterogeneous network infrastructure interconnecting DOE laboratories into a simple enterprise-level view, easily accessible by a variety of users. The proposed work leverages existing successful network monitoring projects such as the ESnet network weathermap. Furthermore, the proposed work will complement the E-Center project by deploying and adapting the necessary data collection tools and E-Center software, establishing and enforcing data access policies, and educating networking staff, system administrators, and users on the new capabilities. The infrastructure will be deployed at DOE laboratories, ANL, BNL, FNAL, NERSC, and ORNL and SLAC, which represent an optimal blend of network-intensive, advanced computing facilities within the DOE SC community. The proposed system will act as a proxy to many existing network monitoring projects, and distill and bring the most critical network monitoring information to application users, thus playing a transformative role in bridging DOE advanced network technologies with science applications in a transparent way.</td></tr>
<tr><td colspan="2">18. CONTRACTOR WORK PROPOSAL MANAGER:
[signature] for PP 4/16/10
(Signature)</td><td>19. OPERATIONS OFFICE REVIEW OFFICIAL

(Signature)</td></tr>
<tr><td colspan="3">DETAIL ATTACHMENTS:
Refer to related Field Task Proposals</td></tr>
</table>

SLAC NATIONAL ACCELERATOR LABORATORY

SLAC National Accelerator Laboratory

FWP #: 10053 March 31, 2010

Proposal to the U.S. Department of Energy, Office of Science
Advanced Scientific Computing Research
Program Manager: Dr. Thomas N'dousse-Fetter
Program Announcement: N/A

Collaborative DOE Enterprise Network Monitoring Deployment Proposal

Principal Investigator:	Dr. Les Cottrell
Position title of PI:	*Networking & Telecommunications* Group Leader
Mailing Address:	SLAC, 2575 Sand Hill Rd, M/S 97, Menlo Park, CA 94025
Telephone:	650 926 2523
Facsimile:	650 926 3329
E-mail:	cottrell@slac.stanford.edu

Senior Investigator: Dr. Yee-Ting Li, SLAC

Other Investigators: Graduate Student

Use of Human Subject: None
Use of Vertebrate Animals: None

Laboratory Official 7/16/10
Dr. Persis S. Drell
Director, SLAC

Les Cottrell
Principal Investigator

Table of Contents

1.1. Budget Explanation

A. Senior personnel:
Les Cottrell (2% in Years 1, 2 and 3), the PI, will oversee the overall planning and progress of the project, including assigning the proper resources when required and recruiting the necessary staff. He will ensure that project activities and expenditures are in compliance with Stanford University and Department of Energy policy. He will lead the effort related to dissemination, including submitting papers for consideration in conferences and journals, and assisting science projects interested in reusing the software developed through this proposal.

Yee-Ting Li (10% in Year 1, and 8% in Years 2 and 3) will lead the design and implementation of the project and liaise with the various stakeholders in order to capture requirements and testing of solutions. He will also supervise and mentor the Graduate Student.

B. Other personnel:
Graduate Student (new hire, 50% Years 1, 2 and 3) will do conduct implementation and coding of deliverables and perform testing and analysis of data derived from the tools.

Note about salaries and inflation:
We assumed a 2.5% salary increase in FY'11, 3.5% in FY'12 and 3.5% in FY'13.

C. Fringe benefits:
Current SLAC fringe benefit rates for regular staff (30.6% of salaries) where applied.

D. Permanent equipment:
None

E. Travel:
The travel budget will fund:

- 1 domestic trip in the 1st and 2nd years for the PI, Lead and Graduate student.
- 1 domestic trip in the 3rd year

For each of these trips we assumed one traveler. SLAC's per-diem ($132 lodging, $56 M&IE) was assumed.

F. Trainee/participant costs:
None

G. Other direct costs:

None

H. Total direct costs: indirect costs:
Current SLAC indirect costs rates for labor and travel (52%) and procurements (9.42%) were applied.

The total cost of the project: $299,610 is requested under this proposal.

2. Abstract

Collaborative DOE Enterprise
Network Monitoring Deployment Proposal

Site	***PI/Co-PI***	***CIO***
Fermilab	*Wenji Wu (Lead PI)/Phil DeMar (Co-PI)*	*Vicky White*
Brookhaven National Laboratory	*Dimitri Katramatos (PI)/Dantong Yu (Co-PI)*	*Thomas Shlagel*
Argonne National Laboratory	*Scott Pinkerton (PI)*	*Charlie Catlett*
Oak Ridge National Laboratory	*Susan Hicks (PI)*	*Chris Kemper*
NERSC	*Brent Draney (PI)*	*Rosio Alvarez*
SLAC	*Les Cottrell (PI)*	*Donald Lemma*

January 7, 2010

Abstract

Providing a coherent view of the status of the DOE enterprise network is a critical requirement in identifying performance bottlenecks. This proposal seeks to deploy a distributed monitoring infrastructure that will facilitate performance tuning and diagnosis of network problems by integrating information from the multi-domain, heterogeneous network infrastructure interconnecting DOE laboratories into a simple enterprise-level view, easily accessible by a variety of users. The proposed work leverages existing successful network monitoring projects such as the ESnet network weather map. Furthermore, the proposed work will complement the E-Center project by deploying and adapting the necessary data collection tools and E-Center software, establishing and enforcing data access policies, and educating networking staff, system administrators, and users on the new capabilities. The infrastructure will be deployed at six DOE laboratories, ANL, BNL, FNAL, NERSC, ORNL, and SLAC which represent an optimal blend of network-intensive, advanced computing facilities within the DOE SC community. The proposed system will act as a proxy to many existing network monitoring projects, and distill and bring the most critical network monitoring information to application users. Such information can be integrated into Grid monitoring tools (e.g., Open Science Grid's Google map based MyOSG monitoring system, and Earth System Grid's ESG Visualizer) to "connect the individual dots", i.e., complement data from distributed computing and storage sites with data from monitored network links. The integrated framework will thus "fill the gaps" within the existing Grid monitoring systems and provide users with a complete picture. The outcome of this proposal has the potential of playing a transformative role in bridging DOE advanced network technologies with science applications in a transparent way.

3. Narrative

1. Overview, & Significance

The goal of this proposal is to deploy an end-to-end DOE enterprise network monitoring infrastructure to assist the wide spectrum of DOE users in tracking current network conditions, identifying performance bottlenecks, and providing a coherent view of distributed network information to facilitate performance diagnosis and tuning. While this proposal does not directly deal with performance diagnosis issues, it does seek to equip DOE users with a simple, integrated view of the heterogeneous, multi-domain network infrastructure that supports their applications.

1.1. Background

Over the past several years, perfSONAR has emerged as a common software framework to support ubiquitous gathering and sharing of network information. The perfSONAR framework has three distinct layers. At the bottom is data collection, both passive monitoring and active measurement, of network infrastructure. On top of that is a service layer for administrative controls over the collected data. These services include authentication & authorization of access to the data, as well as lookup & topology services to facilitate location of the data. Finally, there is a user layer for presentation of the data to perfSONAR customers. High-level goals for perfSONAR development include simplifying management & troubleshooting of network infrastructure, enhancing understanding of network conditions & performance characteristics, and supporting next-generation applications that seek to develop network-awareness capabilities. In particular, perfSONAR offers exciting opportunities to collect & analyze monitoring and performance measurement data along network paths that cross multiple, administratively-distinct network domains. Difficulties in monitoring and troubleshooting cross-domain network paths have historically been recognized as limiting factors in optimizing performance of distributed computing applications. PerfSONAR is intended to overcome many of these difficulties.

In its current state of development, perfSONAR supports the following general categories of network service:

- Network interface counters
- Latency measurements (OWAMP; pingER)
- Achievable bandwidth results (Iperf)
- Network diagnostic measurements & tests (NDT/NPAD)
- Publication of available perfSONAR services (lookup service)
- Publication of network topology data

Looking beyond those basic services, perfSONAR has been designed to accommodate easy extensibility for new network metrics and to facilitate the automatic processing of these metrics as much as possible.

In terms of deployment, perfSONAR services have gained a strong foothold within the major US Research & Education (R&E) network backbone providers, in particular ESnet and Internet2. ESnet has a full deployment of perfSONAR servers at its backbone hub locations. The ESnet servers offer ad hoc network performance measurements, as well as regular active network performance measurements between backbone locations. Internet2 has a similar deployment of perfSONAR services across its backbone. Several large science collaborations have begun to evaluate and deploy perfSONAR servers to monitor the network paths used in their distributed computing applications. In particular, the US-LHC experiments have begun site-to-site performance monitoring, using perfSONAR. A number of the National Laboratories have also deployed perfSONAR servers, either in conjunction with the US-LHC monitoring efforts, or other inter-Lab projects.

In summary, perfSONAR is advancing as the de facto network data collection and measurement platform within the US R&E network community to support large scale distributed computing applications.

1.2. Problem space

Despite the progress that's been made in the development and deployment of perfSONAR services within the US R&E community, the current status of its capabilities would have to categorized as limited and fragmented. Problems exist in several dimensions. First of all, end-to-end network path data, even when it's available, remains disjoint. Monitoring and measurement data within one domain must be collected independent of other domains, with each domain potentially presenting different authentication and access issues. There is also a lack of coherence for users in terms where domain-specific monitoring data and tools are located. In some instances, path performance measurements may already exist, by virtue of ongoing active monitoring, but the measurements probably don't match the specific path of interest to the user, and thus may be of little use, if not misleading, to him.

A second dimension to the problem space is partial availability of perfSONAR-provided data along the path of interest to a user or application. While transit network backbones are reasonably well along in their perfSONAR deployments, end site deployments tend to be very spotty. However, data movement is an end-to-end issue; having network monitoring & measurement information on only a part of the path, may not be particularly useful. Note that end-to-end network path means between source system(s) and destination system(s), not just between end site boundaries. The more extensive the perfSONAR monitoring & measurement capabilities within the end sites involved, the more useful the data will be in gaining an understanding of what's happening from and end-to-end perspective.

Finally, there is a knowledge base issue, particularly at the end site. PerfSONAR is a new and emerging network technology. Developing understanding and expertise within site network staff on perfSONAR capabilities and functions will be a requirement for effective end-to-end perfSONAR services. Note that this expertise needs to extend well beyond simply deploying perfSONAR servers and enabling monitoring & active measurement. Among issues that need to be addressed are site authentication & access policies and implementations, development and support for local topology service configurations, and wide-area performance measurement coordination and configuration. Lastly, the site network staff will typically be the conduit for educating users on perfSONAR's capabilities. The deeper and more extensive the staff's knowledge on perfSONAR, particularly in terms of interpreting what perfSONAR data indicates or means, the likelier some benefit from that data will reach the end user.

There is already ongoing effort under way to alleviate the concerns about perfSONAR data being disjoint from a user perspective. The Network Weather and Performance Service E-Center (E-Center) is a facility under design & development that would serve as a central location for users to find monitoring and measurement data for their wide-area network paths. Network path data of interest would be collected from available perfSONAR services by the E-Center, and presented in a manner that users would readily comprehend. The scope of the E-Center Project would be collaborations and experiments supported by DOE Office of Science funding. As such, the focus of the E-Center's initial effort and instrumentation would be on network paths crossing the ESnet infrastructure, and make extensive use of the deployed perfSONAR services already in place across the ESnet backbones. While the E-Center Project will seek to make opportunistic use of end site perfSONAR services as well, it is not a funded part of the project.

1.3. Related Work

ESnet Network Weathermap (http://weathermap.es.net) is a WAN network monitoring tool that allows the visualization of network topologies and their statistics geographically using Google Maps. It reports bandwidth utilization, in terms of inbound/outbound traffic, of network segments along the ESNet backbone, hubs, and routers. It provides the network operator and administrators with the necessary statistics for performance diagnosis and troubleshooting. In contrast, the proposed system focuses on end sites and applications, and provides a holistic view of end-to-end network paths. Core differences are the following:

- User communities and stakeholders: ESnet network weathermap targets WAN network engineers and network operations center administrators who need to know which network domains are involved and what the performance of each involved network component is, while our system provides network monitoring service to application scientists and data transfer users who are more interested in overall performance measurements between their computing and storage servers rather than the performance of the individual network segments involved along an end-to-end network path.
- Targeted network stacks: ESnet weather map monitors layer 1, layer 2 and layer 3. Our system is particularly focused on layer 4.
- Scope: ESnet weather map provides traffic involving only one domain, i.e. traffic entering or leaving, or within the ESnet domain. The proposed system involves multiple network domains between two end sites.

Nevertheless, the two systems have good synergy in adopted technologies and roles within DOE's network monitoring efforts. They can be complementary to each other: on one hand, the E-center system can identify a potential end-to-end performance problem, which can trigger further investigation and assist network engineering in narrowing down the problem area; on the other hand, the ESnet network weathermap can provide WAN network information to our E-center deployment project. Furthermore, the weathermap's Google-based user interface provides a clean, simple and effective view that can be leveraged in our project.

1.4. Proposed solution

We propose a coordinated effort to develop end site network data collection and presentation service capabilities at five ESnet backbone sites, Argonne NL, Brookhaven NL, Fermilab, the NERSC Center at Lawrence Berkeley Laboratory, Oak Ridge NL, and SLAC CA. The proposed effort would extend beyond simple deployment of perfSONAR servers at the participating sites. Each site would be expected to conform to a common perfSONAR deployment path, one that would be closely integrated with ESnet and the other participating sites. One of our overall goals would be the establishment of an end-to-end perfSONAR service infrastructure. E-Center network data access and presentation capabilities would be adapted at each participating end site for local deployment, functionally providing a site gateway for access to locally generated network data. The local site E-Center adaptation would include providing user access to other local network monitoring and performance data as well, although perfSONAR-collected data would serve as the base-level framework for data presentation and analysis.

We propose to achieve our goals by meeting the following set of objectives at each participating site:

- Deployment of perfSONAR Measurement Points (MPs) & Measurement Archives (MAs)
- Establish active performance monitoring with other sites participating in the effort
- Adapt the E-Center platform for local site deployment and support
- Quantify, resolve, and document site security issues involving site perfSONAR & E-Center support
- Integrate local AA support with E-Center data collection & other WAN perfSONAR efforts
- Train local staff to use and support perfSONAR resources, including user education
- Contribute site perspective into the strategic planning of E-Center capabilities & services

Our proposal seeks to address the three facets of the end-to-end network data problem space previously noted. It directly deals with the concerns about partial network monitoring data services coverage by seeding extensive deployment within end sites, the places where coverage tends to be most lacking. While we recognize this proposal only seeds perfSONAR and adapted E-Center deployment at a handful of end sites, it establishes an appropriate test environment to facilitate development of end-to-end network path services, which would be expected to encourage additional deployment at other sites.

The collective effort to follow a common perfSONAR and adapted E-Center deployment roadmap will help to develop an extensive knowledgebase at participating sites. The common deployment roadmap will necessitate site network staff dealing with a broad range of deployment and operational issues, including performance analysis and troubleshooting. In addition, we propose to draft and publish an End Site PerfSONAR & E-Center Deployment Guide, based on our collective experiences. Dissemination of this document to other end sites will facilitate knowledge transfer beyond the participating sites. The Guide is envisioned to be an iterative document, initially published after year one accomplishments, and updated as our collective experiences broaden.

Finally, our proposal seeks to facilitate availability of site network data and put coherence to it. Providing this information to on-site and off-site users, in a controlled manner consistent with local site security policies, will establish a more complete set of network data on the end-to-end path conditions and performance characteristics.

1.5. Significance

We believe that DOE Office of Science research programs will benefit from this project in the following ways:

- Advance the evolution of perfSONAR-based network services by developing the end site component needed to establish a true end-to-end perfSONAR monitoring & measurement capability
- Development of an effective, secure means, applicable to DOE SC environments, for making network monitoring and performance measurement data of local sites more available to the DOE user community
- Enhance the development of the E-Center facility by providing end-site network data information, which isn't within the scope of that project
- Provide end-to-end perfSONAR services for the ESnet prototype 10GE backbone, and other DOE/SC ASCR research projects involving advanced network technologies & infrastructure.

1.6. Participating Sites

The participating national laboratories represent an optimal blend of network-intensive, advanced computing facilities within the DOE SC community. Argonne (ANL) and NERSC will lead DOE's scientific cloud computing investigations through the Magellan Project. Oak Ridge (ORNL) and ANL host the Department's two leadership computing facilities, NCLF and ACLF respectively. Brookhaven (BNL) and Fermilab (FNAL) host the extremely network-intensive US LHC Tier-1 computing centers for the ATLAS and CMS experiments, respectively. ANL, BNL, NERSC ORNL and SLAC are all scheduled to be connected to the ESnet prototype 100Gb/s network backbone, when that facility is deployed. FNAL is co-leading the effort to develop the E-Center, which would be integrally intertwined with the end site perfSONAR efforts. Combined, the collection of participating sites would provide a diverse and well suited environment for developing and advancing end site perfSONAR services.

2. Description of Work

The primary focus of the project will be local deployment of perfSONAR and other site-specific monitoring tools, as well as adaptation and deployment of a local E-Center system to provide controlled access to data generated by those monitoring tools. The proposed work consists of five primary tasks:

- Collaborate with the E-center project to specify the required E-center gateway components

- Deploy a set of network monitoring tools: PerfSonar, and site-specific tools, such as Iperf, nuttcp, and Traffic Flow Tools
- Work with the site cyber-security to classify the network monitoring information, and assign appropriate ACL to each category
- Customize the E-center software by each site to enforce the ACL and conform to the local site security requirements
- Publish network information to its peering E-center system per users' request, and make that information available to the upper level DOE Enterprise E-center.

2.1. E-Center Architecture

The E-Center Project's objective is to provide a coherent, user-understandable perspective of network conditions across a specific network path of interest. The E-Center natively interfaces to perfSONAR monitoring and measurement infrastructure of each network domain along that path to collect the data needed to provide the user the end-to-end perspective he desires. A design objective for E-Center is the capability to function in federated mode, enabling an E-Center to limit its scope of direct data collection to a defined set of boundaries, and rely on peer E-Centers for data outside of those boundaries. For example, the E-Center project will focus on the ESnet network environment. It is envisioned that Internet2 would support a comparable E-Center function to serve its sphere of interest. End sites are likely to have heightened security and policy concerns about the collection and dissemination of network data pertaining to their local network infrastructure. This project proposes to deploy an end site E-Center gateway to deal with these concerns, and provide adequate local control on the accessibility and dissemination of local network data. The E-Center gateway would be an adaptation of the full function E-Center platform for local network monitoring and performance measurement data collection and distribution. Figure 1 displays the relationship between a central E-Center for the DOE Enterprise, and E-Center gateways for individual sites.

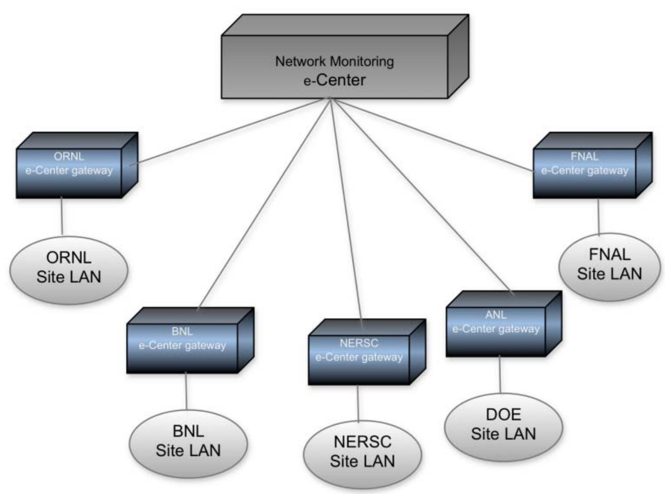

Figure 1: E-Center Architecture

2.2. E-Center Gateway Functionalities

The E-Center gateway will serve two major functions:

- Site Mode: as shown in Figure 2-[A]

- Provide scientists in the local site with LAN monitoring, fault diagnostics, network status, etc
    - Enable in a given site to authenticate to other Office of science E-Center network performance clients, facilitate network monitoring communication with the Office of Science
- Enterprise Mode: as shown in Figure 2-[B]
    - Subscribe to E-Center to publish and site network monitoring information
    - Authenticate users to access LAN performance information
    - Provide authenticated access to other site E-Center gateways
    - Perform scheduled monitoring test as requested by the E-Center

Enterprise mode would be capable of supporting a diverse set of network information access and dissemination configurations. The DOE Office of Science Enterprise would be a common access configuration for all site participants. Virtual organizations (VOs), such as the LHC experiments, would be examples of an access configuration that a specific site might support, if it were part of a specific collaboration. Finally, there would likely be a default access configuration, customized to the security and policy controls for the site, which would provide for general off-site access to the site's network information.

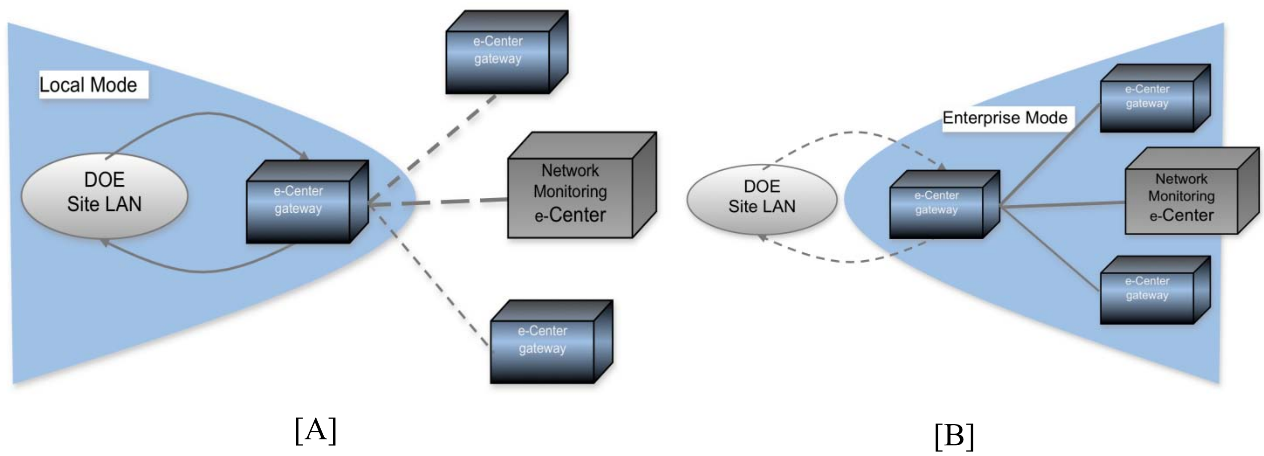

Figure 2: E-Center Local Mode [A] & Enterprise Mode [B]

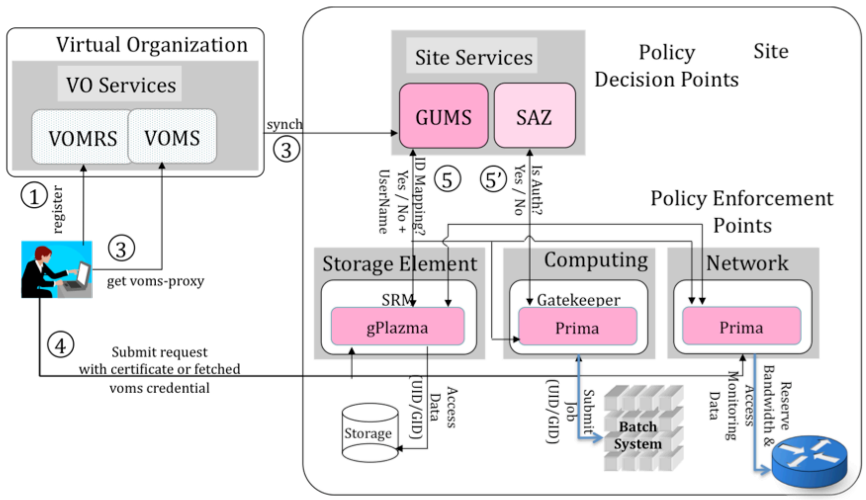

Figure 3: GUMS based User Authentication and Authorization.

2.3. DOE Trust Model

To deal with the multitude of trust issues arising in such an environment where site information needs to be shared among participating sites, we propose to follow a Virtual Organization (VO) – based trust model and the utilization of the Grid User Management System (GUMS). A separate VO specifying all users of this deployment project would include all DOE high performance computing stakeholders: 1) DOE users who need to move their scientific data, 2) site and ESnet network administrators, 3) site information officers responsible for IT services, 4) DOE program officers who need to evaluate the overall network quality of service. For each participating site, the site's GUMS service would perform authentication and authorization by mapping members of this VO to local identities and policies thus enabling controlled access to the site's information depot. GUMS is comprised of web services, web pages for GUMS administration, and command-line tools which interact with the web services. It is a key service in the Open Science Grid (OSG) software stack. Figure 3 shows a GUMS workflow in the OpenScienceGrid environment.

2.4. End Site Objectives

Our proposal is based on coordinated effort at participating sites to achieve a specific set of perfSONAR deployment and service objectives. We envision accomplishment of these objectives will not only establish a comprehensive level of perfSONAR support at the participating sites, but also result in creation of a general template for end site perfSONAR deployment. In addition, the combination of comprehensive end site perfSONAR deployment, along with the existing comprehensive deployment across the ESnet backbone, will combine to create a DOE network environment that facilitates development of true end-to-end network services for both users and applications. Beyond the base-level perfSONAR infrastructure deployment, the end site will be responsible for its own E-Center gateway deployment. The functions of the site E-Center gateway were described in section 2.3. Implementation would involve deploying an E-Center gateway system, adapted to the local site's network monitoring and performance measurement capabilities, and conforming to its security and policy controls on that data. A more detailed description of the individual objectives is listed below:

2.4.1. Deployment of perfSONAR MPs/MAs:

Sites will initially deploy and support one or more perfSONAR Measurement Point(s) (MPs) and Measurement Archives (MAs) along their local site perimeter, near the demark point(s) for their wide area network service providers. As the project moves forward, sites will add MPs/MAs within the local campus network infrastructure, with the objective of extending the monitoring & measurement coverage to reach the site's major distributed computing resources. In addition, sites will develop local network topology configuration support within their perfSONAR services, and contribute to the evolution of a perfSONAR-based topology service.

2.4.2. Active performance monitoring:

Sites will configure their MPs to perform persistent active performance measurements, following project-defined recommendations & configurations. Sites will also provide support for on-demand active measurements from both their perimeter and internal MPs. Finally, sites will collectively investigate and support persistent on-demand active measurements over short and intermediate periods, as resources allow. A model for support of persistent on-demand active measurements will be collectively developed.

2.4.3. Site Deployment of an E-center gateway and archive

Sites will be responsible to deploy the E-Center gateway and archive. The E-Center gateway acts as the site security policy enforce point. It collects the monitoring and performance measurement data via the underlying PerfSONAR infrastructure. Other existing local monitoring or performance tools would also provide network information to the local E-Center gateway. The gateway controls access to the data, based on ACLs and other approved AA mechanisms, and archives the data for future access. The gateway would also cache and archive active performance measurement results obtained per users requests. This caching mechanism can shorten the response time and mitigate the active monitoring load when a similar type of performance measure request is received. To accomplish all these functions, the E-Center gateway requires a state-of-art archival storage and parallel processing capacity. A suggestive set of configurations is:

- CPU: quad core (AMD/INTEL) system
- Memory: 8 GB RAM
- Expandable disk configuration provided by a combination of local disk and direct attached storage system. Initially, 2-4 TB will be required, but with a capability to expand to 50 TB as the project evolves.
- High performance bus: such as hyper transport bus in AMD and Intel QuickPath Inter-Connnect.
- NIC: 2 x 10GigE slots,

2.4.4. Development of local AA support:

Sites will provide authentication & authorization (AA) support for their offered E-Center gateway and perfSONAR services. Initially, this AA service is expected to be based on simple AA mechanisms, such as access control lists. As the AA model for the E-Center, perfSONAR, and related projects evolve toward more robust and distributed technologies, sites will adapt their local services AA support accordingly.

2.4.5. Address site data collection and access security issues:

Sites will work closely with their local computer security support group(s) to ensure that local E-Center gateway and perfSONAR deployments conform to local security policies in a manner that does not hinder development of an integrated research & education network monitoring framework. Each site will develop a risk analysis of its E-Center gateway and perfSONAR services. Sites will collectively draft a guideline on security implications for local campus support of E-Center gateways and perfSONAR MPs/MAs, including risk analysis templates.

2.4.6. End-to-end perfSONAR-based performance analysis:

Sites will actively utilize local and remote perfSONAR services to analyze network performance of local distributed computing applications. Sites will document and share their efforts in this area, and collectively work to advance performance analysis techniques & methodologies. Finally, sites will incorporate availability of local distributed application logs, such as GridFTP, into their perfSONAR services, with an objective of extending the usefulness of perfSONAR in wide area performance analysis.

2.4.7. Develop local perfSONAR knowledge base:

Sites will establish local training programs to educate local network staff, system administrators, and users on the capabilities of E-Center gateway and perfSONAR services. Sites will collectively work to share and jointly develop such training information.

2.4.8. Site contributions to E-Center development:

Sites will actively contribute to the development of the E-Center. They will work closely with E-Center developers to ensure continual access and appropriate responsiveness of their local E-Center gateway services to the central E-Center facility. Sites will also provide feedback in the form of their local site perspective on existing and proposed E-Center services.

2.5. Scope

The DOE Enterprise Network Monitoring Project is intended to be a three year project, roughly coinciding with the three year E-Center project. In the first year of the end site project, effort will be focused on integrating the existing site monitoring capability, particularly the site perimeter perfSONAR monitoring measurement & monitoring with the perfSONAR service capabilities that already exist across the ESnet backbone. The site can select appropriate monitoring information to the web services for publication. Persistent active monitoring will be configured with adjacent ESnet perfSONAR servers, and with other participating sites. An on-demand active measurement capability covering the same site/ESnet and site/site scope will be implemented.

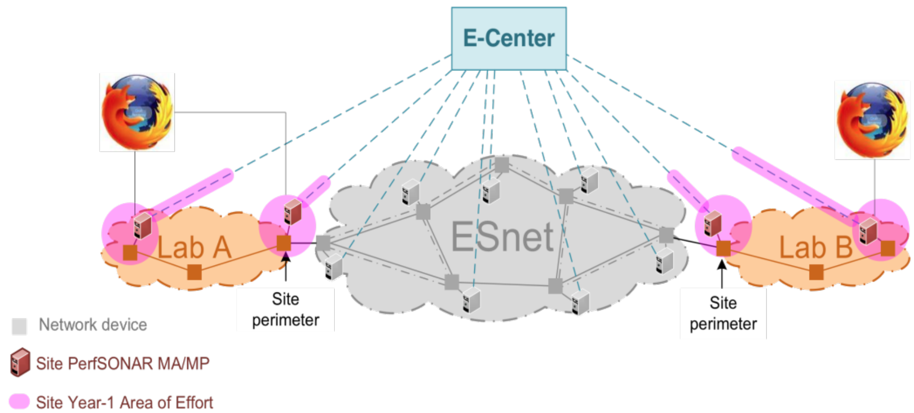

Figure 4: Year 1: Leverage the existing monitoring infrastructure, & adding PerfSONAR as appropriate

In the second year of the project, the focus of the site effort will be extended to include deployment of an E-Center gateway, and adaptation to the site's local perfSONAR monitoring & measurement services. Active performance measurement, both through local perfSONAR infrastructure and other local network measurement tools, will be available for both persistent and on-demand measurements. Local topology information will be incorporated into the perfSONAR topology service.

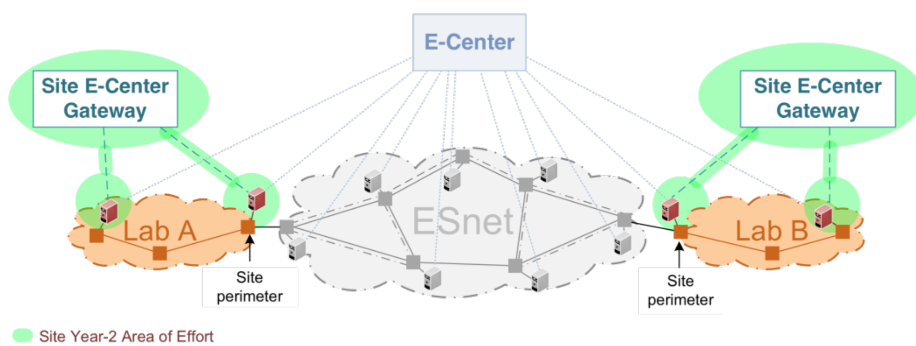

Figure 5: Year 2: Deploy E-Center local mode to collect site monitoring and measurement information

Year three of the project will focus on adapting local E-Center gateway services to supporting the Enterprise mode services described in section 2.3. Direct access to local perfSONAR services will be phased out in favor of controlled access through the site gateway.

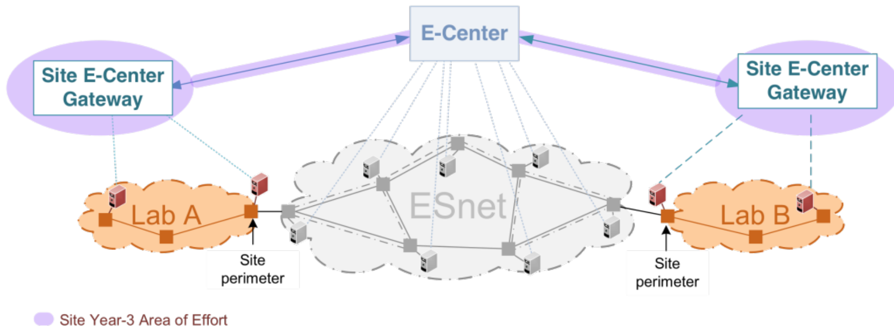

Figure 6: Year 3: Interface for E-Center gateway to establish DOE enterprise E-Center infrastructure

2.6. Future Work of Integrating E-Center Infrastructure into Grid Monitoring Framework

Distributed monitoring must span multiple autonomous networks, computing and storage systems. Existing Grid monitoring systems, such as Open Science Grid's Google map-based MyOSG monitoring system, and Earth System Grid's ESG Visualizer provide monitoring for "individual dots", i.e., individual computing and storage sites, but lack sophisticated network monitoring technologies to address connectivity monitoring and service level measurement issues. An integrated, distributed monitoring system has to provide information about components at all levels (backbone, edges, end point computing and storage sites, layers, dedicated and shared circuits, etc.) to meet federated monitoring needs. Each autonomous system may adopt its own monitoring technology, while a federated monitoring system has to ensure the interoperability of the individual monitoring frameworks. We will participate in working groups in organizations such as the Open Grid Forum to define interoperable protocols and facilitate the integration of network monitoring into Grid/Cloud monitoring systems, and ensure that network is an integral piece of computing resources. In the late stage of the project, we plan to publish network monitoring information into Grid monitoring systems, such as the Open Science Grid's Google map-based MyOSG monitoring system, and Earth System Grid's ESG Visualizer, and use the same Google map technology as user interface to maximize monitoring information flow within the future federated monitoring framework.

3. Breakdown of Work & Project Management:

3.1. Program of work breakdowns and milestones

Our breakdown of project work and deliverables is listed below, based on a three-year project. The breakdown lists yearly technical objectives and associated milestones to achieve the high-level objectives described in Section 2.

<u>Year 1:</u>

PerfSONAR MA/MP	Deploy PerfSONAR MPs/MAs on site perimeter, and within internal site network to support key local computing resources

Deployment & Configuration	
Monitoring & measurement capabilities	Provide SNMP interface counters for site border router(s)
	Configure persistent OWAMP & BWCTL measurements to ESnet from site perimeter
	Support on-demand OWAMP & BWCTL capability from perimeter
Local AA support	Support ACL-based access to network monitoring services and monitoring data
Site Security Issues	Draft risk assessment for monitoring services
Develop local network monitoring knowledge base	Draft network monitoring training document for local staff Deploy Web Interface to expose Local area network monitoring information, with integrated ACL support
Contribution to E-Center	Monitoring Data Mirror/Replicate to the E-center local database Draft the user requirement for E-Center gateway software 1) To support Implement OSG/Grid based cross-site authentication with GUMS 2) To coordinate data collection of site border data by E-Center

Year 2:

Local Network Monitoring Integration	Leverage and adapt other existing network monitoring tools that monitor network status and measure performance at LAN, and site perimeter. (candidate tools are IPERF, and NUTCP)
	Configure PerfSONAR topology service for internal network
Monitoring & measurement capabilities	Provide SNMP interface counters for site internal router(s)
	Implement full mesh persistent OWAMP & BWCTL measurements from local perimeter to other participating sites
	Implement persistent OWAMP & BWCTL measurements from internal MPs/MAs to adjacent ESnet perfSONAR PoP
	Support on-demand internal OWAMP & BWCTL measurement support to authorized off-site users
Local AA support	Adapt AA support as necessary
Site Security Issues	Document site security policies on data & topology information
Development of local monitoring knowledge base	Draft monitoring user support document targeted at local sysadmins
Contribution to E-Center	Coordinate data collection of site internal data by E-Center
	Collaborate with E-Center and joint develop site level E-Center gateway tools 1) Feedback on site perspective of prototype E-Center,

	2) Implement OSG/Grid based cross-site authentication with GUMS 3) Customize the view of network monitoring. 4) Implement the remaining local mode E-center gateway

Year 3:

Monitoring & measurement capabilities	Provide application log transfer measurements to local E-center gateway monitoring data archive
Local AA support	Adapt AA support as necessary
Site Security Issues	Draft security guidelines for campus E-Center gateway and perfSONAR perimeter & internal MP/MA support
PerfSONAR-based performance analysis	Instigate improvements in performance analysis methodologies
Development of local monitoring knowledge base	Complete (collective) Site E-Center gateway and perfSONAR MP/MA Deployment Guide
	Draft monitoring user document targeted at users
Contribution to E-Center	Collaborate with E-Center and joint develop site level E-Center gateway tools 1) Recommendations on capabilities of site instance of E-Center 2) Upload the E-center local monitoring data to E-center enterprise server. 3) Implement the remaining enterprise mode E-center gateway

3.2. Project management structure

We will adhere to standard Project Management Body of Knowledge (PMBOK) methodology, including drafting a project definition document, and based on that document, a detailed project execution plan. The execution plan will specify deliverables, milestones, and reporting mechanisms.

3.3. Resources Requested

Personnel expenditures are estimated at approximately ~ 1/3 FTE/year per site, over three years. This estimated effort is based on site PI effort of 0.1 FTE, and site staff effort of 0.25 FTE. A small amount ($5k/yr per site) of material resources is requested each year for perfSONAR MP/MA and E-Center gateway systems. A very modest travel budget ($1k/yr per site) is requested. The project will assume twice yearly face-to-face coordination meetings, held in conjunction with Joint Techs/ESCC meetings. The travel cost for these meetings is assumed to be a prorated portion of overall cost of the Joint Techs/ESCC trip. The proposed budget is essentially flat over the course of the project. Any salary adjustments for inflation in years 2 and 3 will be offset by slight reductions in effort. Total cost of the project over three years is listed below, as is a breakdown of the per site costs, using Fermilab as a template.

Total:

Budget Item	Year 1	Year 2	Year 3	Total
ANL	$100K	$100K	$100K	**$300K**
BNL	$100K	$100K	$100K	**$300K**
FNAL	$100K	$100K	$100K	**$300K**
NERSC	$100K	$100K	$100K	**$300K**
ORNL	$100K	$100K	$100K	**$300K**
SLAC	$100K	$100K	$100K	**$300k**
Total	**$600K**	**$600K**	**$600K**	**$1800K**

Individual End Site Breakdown (using FNAL as an example)**:**

Budget Item	Year 1	Year 2	Year 3	Total
PI	$8K	$8K	$8K	**$24K**
Developers	$24K	$24K	$24K	**$72K**
Fringe	$19K	$19K	$19K	**$57K**
M & S	$5K	$5K	$5K	**$15K**
Travel	$2K	$2K	$2K	**$6K**
Indirect Costs	$42K	$42K	$42K	**$126K**
Total	**$100K**	**$100K**	**$100K**	**$300K**

4. Principal Investigators (PIs) and Supporting Staff:

Each participating site will designate a PI to coordinate and lead the site's effort and contribution to the project. The PI would nominally be expected to be the ESnet Site Coordinator for the site, or someone designated by the Site Coordinator. The rest of the site effort is assumed to be provided by existing staff, familiar with and currently responsible for supporting networking at the site. Overall PI for the project would be Wenji Wu of Fermilab, who will serve a leadership role in coordinating the overall effort. Dr. Wu is the lead scientist in network research at the Fermi National Accelerator Laboratory. He specializes in network communication protocols and network I/O investigations. He is directly involved in the E-Center Project, and will serve as the technical bridge between the two projects.

5. Biographical Sketches

ROGER LESLIE ANDERTON COTTRELL

SLAC Computing and Computer Services
Stanford Linear Accelerator Center
2575 Sand Hill Road
Menlo Park, CA 94025
cottrell@slac.stanford.edu

EDUCATION

1962-1967 Manchester University, UK
Ph.D: Thesis title – Interactions of Deuterons with Carbon Isotopes
1959-1962 Manchester University, UK
B.Sc.: Physics

PROFESSIONAL EXPERIENCE

2008-Present Stanford Linear Accelerator Center, USA
Head of Network and Telecoms: Management of computer networking services, telecommunications and networking research
1997-2008 Stanford Linear Accelerator Center, USA
Assistant Director SLAC Computing Services: Management of computer networking services, telecommunications and networking research
1995-1997 Stanford Linear Accelerator Center, USA
Acting Director SLAC Computing Services: Management of all SLAC's computing services
1982-1995 Stanford Linear Accelerator Center, USA
Assistant Director, Computing Services: Management of networking and Computing services
1980-1982 Stanford Linear Accelerator Center, USA
Computer Network Manager: Management of SLAC's computer Network activities
1979-1980 IBM U.K. Laboratories, UK
Visiting Scientist: Graphics and intelligent distributed Workstations

SELECTED PUBLICATIONS

Evaluation Of Techniques To Detect Significant Network Performance Problems Using End-To_End Active Measurements, R. L. Cottrell, C. Logg, M. Chhaparia, M. Grigoriev, F. Hara, F. Nazir, M. Sandford. 2006 IEEE/IFIP Network Operations & Management Symposium.

A Hierarchy Of Network Performance Characteristics For Grid Applications And Services, B. Lowekamp, B. Tierney, R. L. Cottrell,, R. Hughes-Jones, T. Kielmann, M. Swany, GGF document GFD-R-P.034, 24 May, 2004, also see SLAC-PUB-10537.

Pathchirp: Efficient Available Bandwidth Estimation For Network Paths, Vinay Ribeiro, Rudolf Reidi, Richard Baraniuk, Jiri Navratil, Les Cottrell, SLAC-PUB-9732, published at PAM 2003, April 2003.

Experiences And Results From A New High Performance Network And Application Monitoring Toolkit, Les Cottrell, Connie Logg, I-Heng Mei, SLAC-PUB-9641, published at PAM 2003, April 2003.

YEE-TING LI

SLAC Computing and Computer Services
Stanford Linear Accelerator Center
2575 Sand Hill Road
Menlo Park, CA 94025
ytl@slac.stanford.edu

EDUCATION

2001-2005 University College London, UK
Ph.D: Thesis title - An Investigation into Transport Protocols and Data Transport Applications Over High Performance Networks
1997-2001 University College London, UK
M.Sci.: Physics

PROFESSIONAL EXPERIENCE

2005-Present Stanford Linear Accelerator Center, USA
Network Specialist: Operations and Research on High Performance Networking technologies and solutions
2005-2005 Hamilton Institute, Ireland
Researcher: Simulation and real-life studies of TCP congestion control algorithms
2004-2004 EGEE, JRA4, UK
Software Engineer: Design and implementation of network monitoring middleware

CURRENT RESEARCH INTERESTS

Distributed systems, network monitoring architectures and schemas, high performance networking, TCP congestion control algorithms, MPLS and Diffserv implementation.

SELECTED PUBLICATIONS

Experimental Evaluation Of TCP Protocols For High-Speed Networks, Y. Li, D. Leith and R. Shorten, IEEE/ACM Transactions on Networking, October 2007

Evaluation of TCP Congestion Control Algorithms on the Windows Vista Platform, Y. Li, SLAC Publications SLAC-TN-06-005, June 2006

Bringing High-Performance Networking To Hep Users, R. Hughes-Jones, S. Dallison, N. Pezzi and Y. Li, Computing in High Energy and Nuclear Physics 04, September 2004

Systematic Analysis Of High Throughput Tcp In Real Network Environments, Y. Li, S. Dallison, R. Hughes-Jones and P. Clarke, Second International Workshop on Protocols for Long Distance Networks, February 2004

6. Facilities and Resources

SLAC was the home site of the BaBar High Energy Physics (HEP) experiment. It is now the home site of the Linear Coherent Light Source (LCLS) experiment and Stanford Synchrotron Radiation Laboratory that includes the SPEAR-3 photon source. SLAC is also a member of the Large Hadron Collider (LHC) experiments at CERN, Switzerland and is a combined US ATLAS Tier 2/3 site housing compute, storage and user facilities. SLAC is also a major contributor to the Large Synoptic Survey Telescope (LSST).

All of these experiments have or will have challenging data network needs both nationally and internationally that we hope to partially address in the current proposal.

We have demonstrated utilization of 35Gbit/s (in both direction) using only two 10Gbit/s connections as part of our record breaking Bandwidth Challenge at the SuperComputing 2005 conference. Contributing with Caltech and Fermilab, we managed to transfer real physics data at a rate of 150Gbit/s peak during a two hour window.

SLAC has hosts dedicated to network measurement that run perfSONAR - collecting and presenting performance data between US ATLAS sites. SLAC has two GPS aerials and connections to provide accurate time synchronization.

The SLAC data center contains many Linux clusters, totally over 6,000 CPUs. For data storage there are 5 PetaBytes of online disk, and automated access tape storage with a capacity of 8 PetaBytes across two Sun SL8500's. In terms of networking, SLAC's datacenter has 17 Cisco 6509's interconnected with multiple 10Gbps connects to provide networking for the compute and storage needs. For outbound traffic, SLAC has two 10Gbit/sec Internet connections to ESnet's Bay Area Metropolitan Area Network (MAN), and also two 10Gbit/sec connections to Stanford University and thus to CalREN/Internet2.

Cluster	Nodes	Hardware	CPU	Memory
hequ	174	Dell R410	dual quad-core 2.93GHz Intel Xeon X5570	24GB
cob	250	Sun Fire V20z	dual dual-core 2.0GHz Opteron 270	4GB
boer	132	Sun Fire X2200M2	dual dual-core 2.6GHz Opteron 2218	8GB
yili	156	Sun Fire X4100	dual dual-core 2.2GHz Opteron 275	4GB
bali	252	Sun Fire X2200M2	dual dual-core 2.6GHz Opteron 2218	8GB
fell	368	Dell Poweredge 1950	dual quad-core 2.66GHz Xeon	16GB

Current and Pending Support

Les Cottrell is currently supported by:

- Department of Energy – 100%

Yee-Ting Li is currently supported by:

- Department of Energy – 100%

The new hire listed on this proposal will be supported 50% through this proposal.

The DOE funds supporting all investigators listed on this proposal are coming under the contract DE-AC-03-76SF00515 with SLAC National Accelerator Laboratory.

There is no pending funding for any of the investigators involved in this proposal. We *expect* to receive funding from the LSST Corporation in FY'11 and FY'12 comparable to the funds provided by LSST Corporation in this fiscal year.

NATIONAL ACCELERATOR LABORATORY

OFFICE OF THE CHIEF INFORMATION OFFICER

Thomas N'dousse-Fetter
PROGRAM MANAGER, Network Research
Computational Science Research and
Partnerships (SciDAC) Division (SC-21.1)
Office of Advanced Sci. Computing Research
U.S. Department of Energy
SC21 Germantown Building
1000 Independence Avenue
SW Washington, DC 20585

Dr. Donald Lemma
Chief Information Officer
SLAC MS 97
donald@slac.stanford.edu
+1-650-926-2944

24 March 2010

RE: UNSOLICITED FIELD WORK PROPOSAL

Dear Dr. N'dousse-Fetter,

Please accept this letter as documentation of our support of the attached Enterprise Network Monitoring proposal. SLAC is enthusiastic about this effort and welcomes the opportunity to join with the other 5 national labs engaged in this proposal.

The proposal leverages and extends the existing network monitoring projects that have emerged into the perfSONAR framework that is being developed for academic and research networks such as ESnet, Internet2 and GEANT in Europe, and does so in an effective manner. I am strongly in favor of establishing an appropriate test environment to enable easy to use, yet powerful remote access to distributed multi-domain network performance and topology information seen from multiple viewpoints. Such access is increasingly critical for identifying and diagnosing performance bottlenecks and to support modern science applications. SLAC currently maintains the largest scientific database in the world and our site would be uniquely qualified to assist with this effort.

Thank you in advance for your kind consideration.

Sincerely,

Donald Lemma, B.Sc., MPA, Ph.D.

Chief Information Officer and Computing Division Director
SLAC National Accelerator Laboratory

[Enc.]