Abstract

Network Weather and Performance Service E-Center

Historically, user expectations for data movement performance across wide-area networks are rarely met. The user is left wondering what the problem is, or even how to investigate it. The user's general instinct is to suspect that his sub-optimal performance is tied to problems or limitations across the wide-area network path he is using. Unfortunately, his network path is likely to cross multiple, administratively-independent network domains, with network equipment and technology varying by domain. The emerging perfSONAR network monitoring framework offers opportunities to gather network monitoring and measurement data within each network domain along the multi-domain network path. However, each domain's data is disjoint from the data being collected in other domains. The user is likely to have difficulty finding where the data is located, and the face access issues to get at it. What the user really desires is a central location where he can retrieve wide-area network monitoring and performance data of interest to him. The central location needs to present the data in a coherent manner that's understandable him, and providing a level of detail adjustable to his interests.

We propose to design, develop, and deploy a Network Weather and Performance Service E-Center facility. The E-Center will serve as a central location for users to find monitoring and measurement data for their wide-area network paths. Our objective is to make the data easily accessible, and presented in a manner that users will readily comprehend. The scope of our targeted user community will be projects, experiments, and collaborations that are supported by DOE Office of Science funding. As such, the focus of our initial effort and instrumentation will be on network paths crossing the ESnet infrastructure.

Our primary deliverable will be a facility that supports a basic set of network monitoring and measurement services, including current network conditions, active performance measurements along a network path of interest, and a capability to display link counters along that path. This data collection will be based on use of currently available perfSONAR services from ESnet, as well as end sites and other network providers, if available. We will also provide web-based portal services, in the form of blogs, wiki pages, and interactive mashups to enable the user to help and advice on topics or issues of interest. A secondary deliverable will be a software package that can facilitate deployment of E-Centers elsewhere. We envision a federated model of cooperating analysis centers, each serving a local customer base in a highly efficient manner.

Our project team offers an optimal blend of expertise and perspective. Fermilab has extensive experience with collecting and analyzing network data in near real time. LBL (ESnet) similarly has a long history of collecting and analyzing network monitoring data, as well as designing frameworks for exchanging measurement information. Both LBL and Fermilab staff are actively involved in the development effort of perfSONAR. Bucknell University brings statistical analysis and data mining expertise into the project. The ESnet backbones are fully instrumented with perfSONAR monitoring services and will provide an excellent, production-quality environment for development of E-Center services.

Network Weather and Performance Service E-Center

Phil DeMar (PI), Matt Crawford (Co-PI) {Fermilab};

Brian Tierney (Co-PI) {LBL}; Michael Frey (Co-PI) {Bucknell}

May 7, 2009

I. Overview, & Significance

1.1 Background

Since the emergence of the Internet, wide-area networks have traditionally been represented by clouds. From an end user perspective, his data transfers disappear into a cloud and almost magically reappear at the remote end. The network-as-a-cloud symbol hides the underlying complexity of the network technology and infrastructure that moves the data from source to destination. As long as the user's performance expectations are being met, he normally doesn't care about the internal complexity of the wide-area network, or its operational status. However, data movement performance across wide-area networks has historically failed to meet user expectations. Users normally anticipate sustaining data transfer rates that approximate the bandwidth capacity of their underlying network path. In practice, such wide area transfer rates are rarely achieved, particularly across the capacious backbones of the Research & Education (R&E) networks. The user is typically left wondering whether his performance expectations are simply too high, or whether there are underlying problems or resource constraints limiting his network performance. When he attempts to resolve these issues, his perspective of the network as a cloud becomes a hindrance, rather than a convenient simplification. The user finds that he needs tools and services to assist in determining whether his wide-area network performance lies within reasonable expectations, and if not, what is limiting that performance. He needs to be able to find out where those tools and services are located, as well as the scope of their coverage. Finally, the typical user will have difficulty with interpreting the output from those tools and services. We propose to establish a network weather and performance service center to provide the user with the ability to peer inside the network cloud that is carrying his data, and potentially determine whether there are any problems or resource limitations impacting his data movement.

1.2 Problem space

A user that perceives sub-optimal network performance with his wide-area data movement faces a daunting task in identifying where his problem might lie [1]. His problem space is comprised of a highly complex set of layers and domains. On his host system alone, there may be problems within the application software itself, the middleware that supports the application, the underlying operating system, and even the system hardware, such as the NIC card or disk I/O capabilities. Beyond that, his typical R&E network path is a concatenation of local and wide-area network links that cross multiple, administratively-distinct network domains. The general instinct for the user is to suspect that sub-optimal wide-area data movement performance is tied to problems or limitations across the wide-area network path in use. More often than not, this turns out to be an erroneous assumption. However, the strategy to examine the wide-area network path first can still be a prudent course to follow. It becomes the first element in a divide-and-conquer strategy. Upon finding the wide-area network path performance to be within expectations, the user can then move on to investigating the host system, or even the application itself, as a limiting factor in his network performance.

Once a user turns his interest to understanding the network component of his wide-area data movement, he still faces a high degree of complexity. His network path is likely to cross multiple, administratively-independent network domains, with the underlying network equipment and technology varying by domain. Historically, users have at least been able to rely on long-standing layer-3 tools, such as ***ping*** and ***traceroute***, to help understand rudimentary characteristics of their network path. But as circuit-based network technologies emerge, even these basic tools may disclose little or nothing about the wide-area network path in use.

Tools and monitoring data normally exist within each administrative domain that could assist the user with performance measurement and diagnosis. The emergence of perfSONAR [2] as a common network monitoring framework offers opportunities to access network performance measurements provided by the individual network domains within a multi-domain network path. However, each domain's data remains disjoint from monitoring data and measurements being collected in other domains. There is also a lack of coherence for the user in terms where domain-specific monitoring data and tools are located. There will likely be access issues with getting at the data. In some instances, path performance measurements are actively being conducted, but they are unlikely to match the specific path of interest to the user, and thus may be of little use to him.

Even if the user were able to find, collect, and collate performance analysis data from all the different domains along his network path, he would run into the more difficult hurdle of interpreting and understanding the information he's assembled. First, there's the knowledge gap between what a normal user understands and what a network expert is able to decipher from data that is collected for and used by the latter on a daily basis. Second, the amount of collected data is likely to be large and diffuse, presenting difficulties for the user in distilling out only what is pertinent and understandable to him.

What the user really desires is a central location where he can retrieve wide-area network monitoring and performance data of interest to him. The central location needs to present the data in a coherent manner that's understandable to the user, and providing a level of detail adjustable to a specific user's interests. Finally, the central location should provide venues or channels for the user to seek advice or help in understanding the data that's been presented to him.

1.3 <u>Proposed solution</u>

We propose to design, develop, and deploy a Network Weather and Performance Service E-Center facility. The E-Center will serve as a central location for users to find monitoring and measurement data for their wide-area network paths. Our objective is to make the data easily accessible, and presented in a manner that users will readily comprehend. The scope of our targeted user community will be projects, experiments, and collaborations that are supported by DOE Office of Science funding. As such, the focus of our initial effort and instrumentation will be on network paths crossing the ESnet infrastructure. However, we recognize the user perspective as being an end-to-end one, and will incorporate end sites that are not on the ESnet backbone where practicality allows.

We intend to provide a set of services grouped broadly into three service classes, network monitoring & measurement data services, web-based user services, and a portal service to point users to support and other helpful information.

The network monitoring and measurement services will be based on data collected from perfSONAR services infrastructure. In this regard, we will make extensive use of the deployed perfSONAR services already in place across the ESnet backbones. Initially, we will provide three levels of increasingly more detailed performance data on the network path between two sites:

- A high level network weather map that displays current network conditions
- End-to-end performance measurements for latency, throughput, and packet drops
- Hop-by-hop link data across a network path

We will also provide a forecasting service on the predictive behavior of the user's network path. In addition to perfSONAR-based services, the E-Center will provide direct support for and the access to legacy performance monitoring services, IEPM/BW [3] and PingER [4].

Web-based services will provide intuitive, user-friendly access to performance data. It will be a principal design goal to provide the user with a customizable view of the data he is interested in. In addition, the E-Center is intended to serve as an electronic forum for users to raise questions or discuss network performance topics and issues. Blog, Wiki, and Chat capabilities will be supported.

The third service class will be a portal service to expertise and to information related to network performance issues. This would include pertinent network performance documentation stored locally, and pointers to remote documentation. As part of our operational support of the E-Center, we will provide local network wizards expertise, to the extent practicable.

The scope of our project, in the proposed initial phase, will be limited to DOE Open Science community and the network facilities that directly support them. However, the E-Center will be designed with a larger scope in mind. We envision a federated E-Center capability, expanding to include the rest of the US R&E community in the second phase, and eventually extending to cover international R&E communities as well. We also expect to broaden the scope of services provided by the E-Center. The user perspective on wide-area data movement is the observable end-to-end result. To accommodate that perspective, it will be necessary to eventually extend monitoring and measurement coverage into the local network environment, and ultimately into host systems themselves. While these goals are not within the scope of our proposal, we will design our facility to accommodate them later on.

Our primary deliverable will be the facility that supports the services outlined above. Equally important will be the development of a software package to establish E-Centers elsewhere. In that regard, our objective is two-fold. First, encourage the deployment of additional E-Centers, and establish federated services with them. Our long term vision is cooperating analysis centers, each serving a local customer base in a highly efficient manner. Second, make the E-Center portable, so that it can be replicated or relocated elsewhere within the DOE/SC community with minimal effort.

1.4 Significance

We believe our proposed project would provide the following general benefits to the entire R&E community:

– Provide a building block for development of federated network performance services that would enhance the distributed computing capabilities of large-scale, global collaborations

– Advance the use of emerging perfSONAR infrastructure and active monitoring capabilities. We believe our effort will encourage wider-scale deployment of perfSONAR services, as well as encouraging development of enhanced perfSONAR services

In addition, we believe that DOE Office of Science research programs will benefit directly and more deeply from our project in the following ways:

– Provide DOE/SC users with a set of services that would enable them to better understand wide-area network performance aspect of their distributed computing efforts

– Establish DOE/SC in a leadership in the development and advancement of perfSONAR-based network monitoring services.

1.5 Project Team

Our project team offers an optimal blend of expertise and perspective. Fermilab has extensive experience with collecting and analyzing network data in near real time. The Laboratory currently maintains one network traffic analysis system that automatically detects and blocks network scanning, and another that generates a collaboration-scale network weather map. Fermilab also provides an end-to-end network performance analysis service to assist remote collaboration sites in debugging wide-area data movement performance problems. Expertise gained in developing and supporting these services should align very well with the skill sets required for development of the E-Center. LBL (ESnet) similarly has a long history of collecting and analyzing network monitoring data, of designing frameworks for exchanging measurement information, and of deploying advanced wide-area networking technologies. The ESnet backbones are fully instrumented with perfSONAR monitoring services and will provide an excellent, production-quality environment for development of E-Center services. Both LBL and Fermilab staff are actively involved in the development effort of perfSONAR. Bucknell University brings statistical analysis and data mining expertise into the project. They are well suited to developing the predictive services that will be based on current and archived network performance data. In addition, Bucknell will bring a useful outside perspective in the development of user interfaces for E-Center services.

II. Research Design and Methods

2.1 Design Architecture

The architecture of the E-Center is defined by four distinct services, perfSONAR-based monitoring & measurement services, legacy measurement services, portal services to assistance and relevant information, and web-based user services. Figure 1 displays the architecture, and the relationships between the types of services. At the core, a central control module facilitates interaction of the services. The control module provides authentication and authorization services. It also provides an archive function for collected monitoring and measurement data. User interaction with the E-Center is via a Web-services interface. Web-based user services will provide a customization function for monitoring and measurement data desired by the user. In addition, electronic forums for questions and dialog provide the user with venues for interaction with other users. The user interface also provides a channel for federation of E-Center services with other E-Center instances that serve related user communities. Network monitoring and measurement data is collected and provided through perfSONAR-based services, as well as legacy services for older monitoring tools that are not perfSONAR-based. The legacy services will be integral to the E-Center. PerfSONAR-based network services are directly collected from the network infrastructure (principally ESnet) that serves the E-Center community. PerfSONAR lookup services will be supported within the data collection module. The data fusion module manipulates collected perfSONAR data for specific perfSONAR monitoring/measurement services. Finally, the portal services provide access to related useful information and assistance (performance wizard) services.

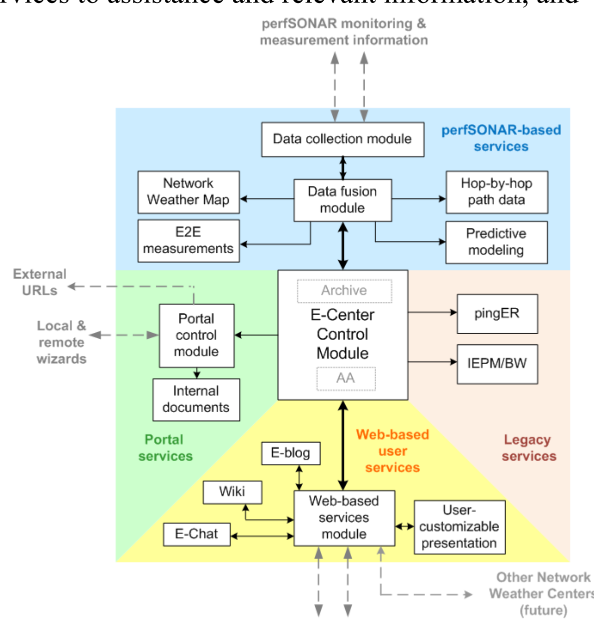

Figure 1 E-Center Architecture

2.2 Design Goals

2.2.1 Visualization

For the Network Weather and Performance Service E-Center, visualization will be a critical design element. It helps users interpret monitoring & measurement data, and thus better understand the current network status. It can also help to discover correlated abnormal network behaviors that are critical to quickly debugging the network performance.

Essentially, our E-Center will be a web application for network status visualization, able to correlate data from different network monitoring infrastructures (e.g. perfSONAR, pingER) and present the result in a readable, easy to understand form. We intend to build an interoperability layer on top of the existing monitoring tools to create a unified presentation system for the different types of monitoring and performance measurement data. The E-Center will provide snapshot views (network status maps) of network performance and traffic flows in near real-time. These network status maps will also enable users to access more detailed statistics by clicking on a link or node.

For the network weather service e-center, the global view will be a dynamic map of the currently monitored physical locations based on geographical representation. The user can then focus his view on specific regions, sites or network paths, and save these views for future visits. By selecting a specific site or connection, the user will be able to view status information and discussions related to network improvements and problem remediation. This should help reduce the loss of information that occurs with unstructured exchange of email. By linking historical performance data to these ongoing discussions, any network modifications and associated performance improvements can be easily correlated.

Google Maps [5] and Google Earth [6] are web mapping service applications and technologies that power many map-based services. They are ideal tools for sharing geographical information in a simple way. It is anticipated that our Network Weather Service will have better visualization effect for end users by referencing real geographical backgrounds. We are planning to develop our E-Center to be based on Google Map and Google Earth technologies.

2.2.2 User Interface Design

All features of the proposed system will be accessible through a web-based interface, using a Web 2.0 interactive [7] approach. This web-based interface will provide maximized user activities, easier communication and collaboration, better usability, and faster performance.

The E-Center's user interface design should be flexible, and able to provide customizable views for end users. This customizable interface design will enhance the user experience and should help users in better understanding the data. Further, the user interfaces should facilitate interactive access to other E-Center services. For example, end users should be able to access more detailed statistics by clicking on a link or node on the network status map, and provide questions, feedback, or comments on what he's observed.

The E-Center will depend on a set of tools and measurements for network performance diagnostics and analysis. Existing tools and measurements provide a variety of useful performance analysis features. Some legacy tools are available via command line only, while others provide graphical interfaces only on the host system running the tool. Some newer tools provide the desired web-based interface, but require a level of expertise that exceeds many users. An analysis tool or service is of little or no value to the average user if it can't be configured and run through a web browser. The web browser is the only interface that is universally available and understood. A fundamental design goal of this project is that full accessibility via web browser is an absolute requirement. The E-Center will support consistent web-based graphical interfaces for all sets of monitoring & measurement tools, and make them understandable to the casual user.

2.2.3 Design Non-goals:

Within the scope of our proposal, our goal is to develop E-Center web services for network status visualization, able to correlate data from different network monitoring infrastructures and present the result in a readable, easy to understand form. We do not envision the E-Center becoming a network trouble-shooting center, in terms of diagnosing or investigating. We limit the scope of our involvement in network trouble-shooting to making monitoring & measurement information available that may be of use in investigating those types of issues or problems. Similarly, the E-Center per se is not intended to provide detailed analysis of poor network performance, although we expect to assist the user in where that type of help might be available. Our longer term strategic vision (Section 3) does include a true end-to-end performance analysis component, but the current proposal is focused on providing monitoring & measurement data on the network path between source and destination sites.

2.3 E-Center Services:

2.3.1 perfSONAR-based Services:

The E-Center will provide a set of services based on monitoring & measurement data collected from perfSONAR infrastructure. We will gather the data from the currently deployed mesh of the perfSONAR network monitoring services within ESnet, as well as DOE Labs and other DOE/SC collaboration sites that have deployed perfSONAR measurement points (MPs). We anticipate these MPs will be based on the perfSONAR NPToolKit [8] package, jointly developed by Internet2, ESnet and Fermilab. The list of currently supported NPToolKit network monitoring services include SNMP-based monitoring, static and dynamic end-to-end circuit monitoring, ping-based (pingER) probes, BWCTL and OWAMP [9] on-demand active measurements, and NDT [10] server data. In gathering this data, we plan on routinely collecting performance measurement data that's already available, and initiating requests for specific performance measurement tests. The data collection will be built on top of the perfSONAR client API. Through perfSONAR's global lookup service, desired network measurements from the identified perfSONAR network monitoring facilities would be collected and passed up for data collating and fusion.

2.3.1.1 perfSONAR-based Services – Current Network Conditions:

Collected perfSONAR data will be used to provide the following services that describe current and historical perspectives:

1. A network weather map depicting current network conditions: SNMP interface counters will be regularly gathered to track utilization and packet discard rates. Initially, this weather map would include the ESnet backbone(s), as well as site connections into ESnet. This service would run in the background continuously, and be instantly available to user requests.
2. End-to-end performance measurements, using perfSONAR active measurement tools BWCTL (throughput) and OWAMP (one-way latency). This service would be provided on user demand, and based on site source and destination as designated by the user. A historical view of ongoing active measurements already being collected would be presented, if available. On-demand measurements will be made where necessary or requested, consistent with local policy. Similarly, longer term persistent end-to-end performance measurements would be set up, if requested by the user and consistent with local policy.
3. Hop-by-hop link data: On-demand collection of SNMP interface counters provided by remote SNMP Measurement Archives for network devices that comprise a specified network path. If feasible, on-demand active measurements across individual links would also be supported as part of this service.

We envision these services providing a set of hierarchical network perspectives to the user, allowing him to select the views of greatest value to his circumstance.

2.3.1.2 perfSONAR-based Services – Forecast Network Conditions:

Network resources are inherently topographically distributed, and constraints on user access to these resources are time-varying over different time scales. The amount of raw data available, or potentially available, from the network about its state is too great to be directly comprehended by the user. Data management to reduce this flood of network data and comprehensibly present its content elements of user interest constitutes a significant statistical task. Also, the complexity of the network's evolving state makes it impossible to deterministically calculate the quality of service (QoS) potentially available for any given user-specified process. Forecasting the availability of desired network resources is therefore necessarily statistical in nature. Both of these statistical challenges must be met for successful network weather reporting.

Vast amounts of status and performance data are readily available from the network through services based on perfSONAR [15,16], pingER [17] and similar web services-based infrastructures. The network measurements delivered by these services vary at all times scales and present an effectively noisy characteristic. We propose a two-step statistical methodology of sampling and filtering to manage these data. The choice of sampling rate sets the desired time scale for data presentation. Sampling is followed by filtering to extract trends at the desired time scale. We propose generalized median filtering [18] for this purpose. Generalized median filtering is, by dint of its flexibility and absence of parametric modeling assumptions, attractive for trend capture in network status and performance data. This approach to data reduction allows trends to be extracted at any time scale—minutes, hours, days, or weeks—and can support summarization of past performance, current status, and forecasting into the future. This can support a time-scale "zoom" capability for the visualization element of the Network Weather and Performance Service E-Center.

The statistics that form the substance of network forecasts can be developed in three general ways: from parametric reliability models of network components, from time series models, and from data mining. Parametric component reliability models have been successful in a range of network applications, both for optimization and performance prediction [19,20,21,22]. Parametric reliability models, though, derive from, and depend to some degree on, knowledge of one or more layers of the network. Time series models, including wavelet-based models [23], are generic models with adaptable parameters that require no attendant knowledge of network structure. Data mining is the process of extracting hidden patterns from large sets of data. In predictive data mining these patterns can be codified into association rules or artificial neural networks for forecasting. Predictive data mining has been given some consideration for network traffic analysis [24,25]. These considerations suggest that it can be successful in the present application. It is significant for our approach that each of these three prediction methodologies yields estimates of RMS prediction error.

We propose an inferential architecture that can exploit the advantages of each of parametric component reliability modeling, time series analysis, and data mining, as well as other prediction methodologies that may be identified. We propose to structure each predictive tool as an independent module in a suite. Some modules would rely on routinely collected and readily available network data. Others could incorporate, and even initiate, active network performance measurements. Each forecast request to a predictive module would carry a specified time scale, and the module would respond using sampled/filtered data for that time scale. Attached to the module's forecast would be an RMS error estimate.

The actual forecast presented to the user would be based on the predictive module with the best RMS error performance. This architecture for identifying the best forecast has been shown to perform as well as, or slightly better than, the best forecasting method in the suite [26]. Additionally, this inferential architecture allows other statistical forecasting techniques to be incorporated as they become available or as the aims of reporting change. Finally, network usage and resource availability dynamics have distinct diurnal and weekly components. Different predictive methods may prove better for predictions over different time scales: one minute, one hour, one day, and one month into the future. Our inferential architecture inherently addresses this prediction requirement.

We propose to offer the user an intuitive, interactive visualization interface based around a network schematic with an optional geographic underlay powered by Google Map and Google Earth. This network presentation will initially support path highlighting that yields 1) end-to-end performance measurements for latency, throughput, and packet drops and 2) hop-by-hop link statistics across the network path. Though not proposed for this project, path highlighting can be extended to tree highlighting for visualizing two-way, multi-user, and broadcast processes.

This interactive presentation will also have architecture and time zoom capabilities. To comprehend wide-area data movement, the user must be able visualize the broad network across the area of interest, see into each domain within that area, see the local network environment, and even look into individual host systems. We propose to introduce a zoom slider to allow the user to easily zoom up and down this architecture scale as the ability to see at each scale becomes available. We have described time scale-dependent statistics for current state, past performance, and forecasting for the network and any highlighted path at each architecture scale. We propose to also include a zoom slider in the network presentation to allow the user to easily visualize past trends and future extrapolations at different time scales. At each time and architecture scale we propose user-optional balloon tags containing selectable statistics for each network/path component to create images that can be user-commented and saved. These images can be used to support and drive user conversations and indicate and document network modifications and associated performance improvements.

We recognize benefits of statistical conceptualization to different development areas within the proposed Network Weather and Performance Service E-Center. Statistical analysis of current and archived network performance measurement data are required to facilitate forecasting of network conditions, both on a short-term and long-term basis. We propose to develop a predictive network service with robust statistical data management, an intuitive user visualization interface, and a flexible, robust forecasting architecture for the E-Center following the conceptual directions described above.

2.3.2 Web-based Services:

We will design our Web Services Portal as a single Point Of Service (POS) for the users. We define users as multi-faceted set of network administrators, end-site users or managers and general public. The proposed design of the portal should allow multiple views, user's added multimedia content in form of blogs or wiki pages, interactive mashups where users will determine what level of details is required and what set of graphical artifacts is sufficient.

2.3.2.1 WebServices Graphical View

The Global View is a default "Main View" for any user who logged in first time on the portal. As was described above, this level will be based on the Google maps. Later, user will be able to customize this view by choosing only those physical locations of the interest or we will provide an opportunity to see network weather map for some predefined set of the locations determined by the known DOE sponsored projects or research Virtual Organizations. We propose to implement 4 different layers of graphical views. Next after the Global View we introduce a Group View (or Organization/Domain View) with more detailed network status graphs for particular organization. Also, there will be Network Path View (or Subnet View) followed by the Device View.
The further detailing will be allowed only for the authorized network administrators because some data might be sensitive for the general public and may expose security risks. The Device View will be represented as clickable map of the network interfaces with network utilization graphs.

2.3.2.2 Web Enabled Services

We will provide the set of web enabled services for the users. Depending on the role acquired through the authorization process any visitor will be able to see multiple layers of the graphical representation of the network monitoring information. We will design and implemented a Network Troubleshooting Interface based on the Network Performance troubleshooting Methodology [11]. The authorized user will be able to walk through the steps of Network Performance Assessment process by entering all available networking information. The underlying E-Center Control Module will try to identify if extra information exists for the network path entered by user and will provide some conclusive answer to the user or will redirect the original request to the available network expert. We will provide the capability of built-in web chat service for the users in the US business hours in order to provide some level of expert assistance for the urgent requests.

2.3.2.3 User's Generated Content

We will provide a capability for users to add content interactively. We will utilize existing blogging and wiki platforms to allow logged in and authenticated users to post their comments on the network performance or desired improvements. This will be part of our user's experience assessment. By utilizing roles based authorization model we will allow identified network experts to post answers on these questions. Therefore we will create an indexable and searchable knowledge database. We anticipate that this growing knowledge database will be a first layer of expertise for most of the users.

2.3.2.4 Web Services Client API

We will design and provide web services API for most advanced users who will be able to utilize all capabilities of the E-Center at full extent. The use case scenario we envision will be presented by some third-party networked application which will require some assessment of the particular network path been done. Another use case will be the same E-Center deployed at another DOE lab and requesting some external network monitoring information which is not available from the remotely deployed perfSONAR-PS network monitoring services due some authentication or authorization barriers.

2.3.3. Legacy Services:

There are several legacy network performance monitoring services that are not perfSONAR-based. These include the pingER round trip time and packet loss service, based on the ICMP ping tool, and the Internet Performance Monitoring; Bandwidth to the World (IEPM/BW) service that provides regular active throughput measurements. While significant effort has been expended at Fermilab to adapt pingER to the perfSONAR NPToolKit package, neither service is being actively supported at the present time. However, both services still have significant historical, and potentially useful performance monitoring data. The E-Center's legacy services will add pingER-generated and IEPM/BW-generated performance monitoring data as available user services, under the common E-Center user interface.

III Long-term Strategic Vision:

3.1 Federation of Services

The target customer base for the project is the DOE/SC research community, and we will focus the E-Center deployment on the DOE network infrastructure. However, we recognize our customer base frequently participates in large, globally-distributed collaborations, and that our initial coverage will not provide complete coverage over their network paths of interest. We will design this project to be the initial stage of a three-phase deployment. Figure 2 depicts the coverage expansion that we envision it over time. In the first phase we intend to involve mainly DOE Labs, and the ESnet backbones. ESnet has a widely deployed perfSONAR infrastructure that will be used to provide monitoring data. We expect to have a prototype of the E-Center operating in about one year to allow users start using benefits of this project as well as developers to better understand user's needs and technologies that need to be evolved, deployed and developed.

In the second phase, we plan to extend the scope of this project to incorporate the wider US R&E community. Our focus will be to extend monitoring & measurement coverage to universities connected via Internet2. We will design the E-Center to support a federated model for exchanging perfSONAR-based monitoring & measurement data. Our strategy will be to get Internet2 to deploy a comparable E-Center functionality, and get the two E-Centers seamlessly interoperating. We anticipate this will be present significant challenges, requiring investigation and development efforts. However, since Fermilab and ESnet already work closely with Internet2 staff on perfSONAR development and deployment, the working relationships to overcome those challenges is already in place.

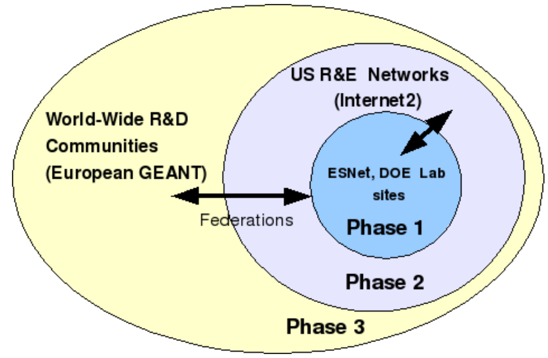

Figure 2 - Federated E-Centers

In the third phase we plan to extend scope of this project to European GEANT2 community and other global R&E network entities. In the case of GEANT, Fermilab is already participating in the GEANT perfSONAR deployment effort through the Multi-Domain Monitoring (MDM) project to monitor the Large Hadron Collider Optical Private Network (LHCOPN).

3.2 End-to-End Scope:

Users have an end-to-end perspective. The scope of E-Center monitoring & measurement coverage in this proposal is the wide-area network path between sites. This would include the ESnet backbone, and potentially site-supported perfSONAR-based data collection on their boundaries. Our strategic vision is to extend the perfSONAR coverage not only into the site networks, but also into host systems themselves (Figure 3). Our vision is to be able to collect performance data through perfSONAR measurement points (MPs) on host operating system and hardware parameters that can affect end-to-end performance. With performance measurement data available from the wide-area network path, the local area network infrastructures, and the host systems involved in the data transfer, a true performance analysis capability could be deployed within an E-Center. Collection of host performance data through perfSONAR services is beyond the scope of this proposal. However, we will design the E-Center with a vision of this capability being developed at some point in the future. If that occurs, we would envision a future version of the E-Center deploying an automated version of the Structured Performance Analysis services that Fermilab supports today.

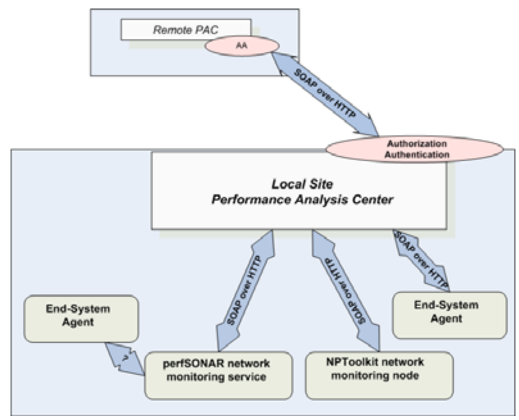

Figure 3 – perfSONAR Host Data Collection

IV. Relevant Work & Research:

A critical area of relevant work is Fermilab & LBL's participation the US-based perfSONAR development effort. Both staffs actively contributed to the development of the perfSONAR-PS product, the package that provides basic perfSONAR measurement point (MP) and measurement archive (MA) services. In addition, both facilities have been active contributors to the NPToolKit, the perfSONAR active performance monitoring distribution package. ESnet has deployed perfSONAR-based services at all its backbone hub locations. It is also worth noting that Fermilab has deployed and supports an NPToolKit-based perfSONAR server to actively measure performance across it's data paths to US sites involved in the LHC Compact Muon Solenoid (CMS) experiment. One final related area of perfSONAR-related work is the GEANT's MDM network appliance effort, another perfSONAR-based performance measurement tool. Fermilab staff has contributed to the development of the MDM appliance, and will be supporting an appliance at the Laboratory in the near future.

Fermilab has been a co-developer, with SLAC, of the pingER path monitoring tool, based on ICMP ping packet round-trip measurements. Fermilab staff adapted pingER support into the NPToolKit (perfSONAR).

The Lambda Station alternate network path selection service project [12], a DOE-funded R&D effort, uses many of the same technologies as the E-Center would use. Lambda Station is based on a Service-Oriented Architecture (SOA), deployed with Web 2.0 network services,

Finally, Fermilab has also developed an end-to-end network performance analysis. The E-Center would build on the experiences and knowledge gained with deployment of that service. In conjunction with that effort, the Fermilab staff has developed considerable expertise in analysis of wide-area data movement, an important element in providing performance expertise to E-Center users.

We are aware of a number of performance monitoring and analysis efforts, including the MonALISA [13] project, developed at Caltech, and the Web100 [14] effort, developed by Lawrence Berkeley Laboratory. We will investigate the adaptability and usefulness of these products into the E-Center.

V. Collaboration:

For success of this project we will extend the existing collaborations and establish new connections in the following areas:

- IP Backbone providers (Internet2, European GEANT3, other international R&E networks)
- Developers of monitoring frameworks, perfSONAR (Internet2, Dante)
- Users scientific communities as customers of the E-Center such as OSG (Open Science Grid), US LHC Communities such as US CMS Tier1 Facility at Fermilab, US Atlas Tier1 at BNL, LHC Tier2 and Tier3 facilities at various universities.

5.1 perfSONAR collaboration

Previously noted, both LBL and Fermilab actively participate in perfSONAR development efforts. Both facilities plan to continue their current level of involvement in the perfSONAR project, focusing on developing missing functionalities, researching new network monitoring services, and investigating new network monitoring data exchange protocols. Similarly, Fermilab will continue to participate in the European MDM project team. We anticipate higher level of synergy between perfSONAR services and proposed E-Center and possibility of developing E-Center web services API based on perfSONAR set of interoperable protocols.

5.2 Open Science Grid (OSG)

The OSG consortium is based at Fermilab. Optimizing wide-area data movement is obviously of high interest to OSG members. We will coordinate closely with OSG on our development efforts, and solicit OSG participation in seeking useful data transfer candidates to try our service. We will work with OSG to get our service advertised to their general user community. Finally, we recognize that our service may be useful for OSG middleware developments, and will be receptive to OSG suggestions on adapting our product to their services.

VI. Program of Work & Project Management:

6.1 Program of work breakdowns and milestones

Our deliverable will be an operational E-Center facility that provides the services detailed in section 2.3. A secondary deliverable is a software package that can be deployed elsewhere to create other E-Centers. Our proposed product is expected to be completed within the three year time frame of the project. The outline of our program of work to is shown in the table(s) below. We have organized our work effort in a manner to develop a working prototype system by the end of year 1, a deployable test system by the end of year 2, and a fully functional product by the end of year 3. The principal investigator will coordinate the overall integration effort, and align resources to ensure adequate development effort within each component.

Year 1:

E-Center Control	Prototype Control Module	FNAL
	Design common database	FNAL, ESnet
	Design framework for authentication & authorization services	FNAL, ESnet
perfSONAR Services	Prototype perfSONAR Services Interface	FNAL
	Design and prototype Data Collection Module	ESnet
	Design and prototype Data Fusion Module	FNAL
	Functional basic data collection from ESnet backbone (scheduled)	ESnet
	Prototype network weather map service	ESnet
	Develop modular predictive services inferential architecture	Bucknell
	Prototype Tier-1 forecasting service, based on parametric reliability	Bucknell
Portal Services	Prototype Portal Services Interface	FNAL
Web Services	Prototype Web Services Interface	FNAL
	Create templates and hierarchy for customizable user views	FNAL
	Create Weather map views based on Google Maps API	FNAL
Legacy Services	Prototype Legacy Services Interface	FNAL
Integrated Product Deliverables	Prototype with basic functions working by end of year	FNAL, ESnet, Bucknell

Year 2:

E-Center Control	Develop first release version of Control Module	FNAL
	Add Data Archive capability to database	FNAL, ESnet
perfSONAR Services	Functional on-demand data collection from ESnet backbone	ESnet
	Data archiving of collected perfSONAR data	ESnet
	Prototype E2E measurement service	FNAL
	Prototype hop-by-hop path data service	FNAL
	Develop data collection capability from R&D (100GE) test bed	ESnet
	Prototype Tier-1 forecasting service, with integrated data mining	Bucknell
	Prototype time services analysis module for Tier-1 forecasting	Bucknell
Portal Services	Deploy Internal Documents database	FNAL
	Develop Local "Wizards" Interface	FNAL
Web Services	Deploy initial release version of Web Services Interface	FNAL
	Add interface views for Portal and Legacy services	FNAL
	Prototype user-customizable Weather map views	FNAL
	Implement E-blog, E-chat, and Wiki services	FNAL
Legacy Services	Port pingER service	FNAL
Integrated Product Deliverables	First release user documentation	FNAL, ESnet, Bucknell
	First integrated release version (conduct user trials and collect feedback)	FNAL, ESnet, Bucknell

Year 3:

perfSONAR Services	Functional on-demand perfSONAR performance measurements	ESnet
	Roles-based AA for perfSONAR data collection	ESnet
	perfSONAR data collection from non-ESnet sites (Labs)	ESnet
	Develop Tier-2 forecasting service, with active network performance measurements	Bucknell
Portal Services	Develop Remote "Wizards" Interface	FNAL, ESnet
	External Document links	FNAL, ESnet
Web Services	Roles-based AA for Web Services interface	FNAL
	Develop user-customizable E2E, hop-by-hop, & predictive services views	FNAL
	Develop & test prototype API for federation of services	FNAL, ESnet
Legacy Services	Port IEMP service	FNAL
Integrated Product Deliverables	Final user documentation	FNAL, ESnet, Bucknell
	Production version for deployment	FNAL, ESnet, Bucknell

6.2 Project management structure

We will adhere to standard Project Management Body of Knowledge (PMBOK) methodology, consistent with the scale and scope of the project. Our first step in the project will be to draft a project definition document. Based on that document, we will draft a detailed project execution plan. The execution plan will specify deliverables, milestones, and reporting mechanisms.

We will utilize a web-based content management system for all relevant information, documentation, and developed code associated with the project. Software versions will be tracked at the component level, but all testing will be done using a current integrated release.

6.3 Resources Required

Personnel expenditures are estimated at approximately 2.2 FTEs/year over three years. This effort will be spread across the project partners at levels of ~1.4FTEs at Fermilab, ~0.4FTE at ESnet, and ~0.4 FTE at Bucknell. The program of work reflects this relative allocation of effort.

Material resources include 2 servers for basic interface testing, a developer workstation and licenses for visualization software. Since some of this work involves interface with perfSONAR and ESnet developers, travel expenses for 2 trips per year are included for each project partner.

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. Breakdown of project costs, including indirect overheads, are estimated below for each site.

FNAL:

Budget Item	Year 1	Year 2	Year 3	Total
PI	$0K	$0K	$0K	**$0K**
Developers	$114K	$120K	$120K	**$354K**
Fringe	$64K	$68K	$68K	**$200K**
M & S	$25K	$10K	$10K	**$45K**
Travel	$5K	$5K	$5K	**$15K**
Indirect Costs	$142K	$147K	$147K	**$436K**
Total	**$350K**	**$350K**	**$350K**	**$1050K**

LBL:

Budget Item	Year 1	Year 2	Year 3	Total
PI	$22K	$20K	$18K	**$60K**
Developers	$29K	$30K	$31K	**$90K**
Fringe	$0K	$0K	$0K	**$0K**
M & S	$0K	$0K	$0K	**$0K**
Travel	$5K	$5K	$5K	**$15K**
Indirect Costs	$94K	$95K	$96K	**$285K**
Total	**$150K**	**$150K**	**$150K**	**$450K**

Bucknell:

Budget Item	Year 1	Year 2	Year 3	Total
PI	$34K	$36K	$37K	**$107K**
Developers	$9K	$9K	$10K	**$28K**
Fringe	$13K	$13K	$14K	**$40K**
M & S	$14K	$15K	$15K	**$44K**
Travel	$15K	$16K	$17K	**$48K**
Indirect Costs	$28K	$29K	$31K	**$88K**
Total	**$113K**	**$118K**	**$124K**	**$355K**

VII Personnel Expertise and Skill Sets:

The PIs form an ideal team to lead the proposed research. PI DeMar is a network architect at Fermi National Accelerator Laboratory, with more than 20 years experience in network management and operation. Co-PI Dr. Crawford is an experienced network protocol designer and head of the Data Movement and Storage department at Fermi National Accelerator Laboratory, author of ten RFCs, and works in the areas of large storage and high-speed networking. Co-PI Tierney is a Staff Scientist and group leader of the Advanced Network Technologies Group in ESnet. Prior to this he was group leader of the Data Intensive Distributed Computing Group in the Distributed Systems Department at LBNL. He has been the PI for several DOE research projects in network and Grid monitoring systems for data intensive distributed computing.

The PIs will provide strategic decision for the project, lead the design effort in product development and deployment, and provide oversight on the project's progress.

Developers expertise that will be brought to the project:

Dr. Wenji Wu: Dr. Wu will serve a leadership role in the technical development 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 will also have primary responsibility for development of the customizable user interface..

Maxim Grigoriev: Contributing developer for perfSONAR services and projects, including integration of the pingER monitoring tool into the base perfSONAR software package. Mr. Grigoriev will have primary responsibility for development of the perfSONAR Services, Legacy Services, and Web-services components of the project. He will bring his extensive perfSONAR expertise, both in terms of software development and practical deployment experience, to ensure tight integration with existing perfSONAR infrastructure.

Andrey Bobyshev: Project leader and principal developer of the Lambda Station project. Mr. Bobyshev bring expertise to the development of the Portal services.

VIII Facilities Description:

The researchers, developers, and supporting staff for the proposed project will have access to a wide spectrum of advanced computational and networking facilities:

8.1 Fermilab

The computational facilities at Fermilab include large scale computing clusters, mass storage systems of petabyte scale, and distributed disk cache systems, also of petabyte scale. The Fermilab Tier 1 Center for the Large Hadron Collider's (LHC) Compact Moun Solenoid (CMS) experiment, for example, consists of approximately 2000 systems, with 12 petabytes of tape storage and 3 petabytes of disk cache storage. The systems are connected to high performance network switches, interconnected at 40Gb/s. The CMS Tier-1 Center generates very large volumes of wide-area network traffic, and makes extensive use of data circuit technology to move that data. It serves as an excellent user facility to fully exercise advanced wide-area network technologies and services.

Fermilab's off-site network access infrastructure is similarly capacious and extensive. The Laboratory has a total of six 10Gb/s links to ESnet for off-site data movement, including four dedicated to circuit-based data movement across ESnet's Science Data Network (SDN). An additional two 10Gb/s channels run down to the StarLight Optical Network Exchange in Chicago, and are available for network R&D projects.

The Laboratory also has an extensive test bed environment dedicated to computing and networking R&D efforts. This environment includes a small computing cluster (14 systems), several high performance 10Gb/sconnected test systems, and 20Gb/s of off-site network infrastructure completely separate from production network facilities.

8.2 ESnet

ESnet is a high-speed production network that directly connects to approximately 42 sites nation wide. The ESnet backbone is comprised of 16 hub sites interconnected by 10Gb/s rings, and is logically separated into two networks, the ESnet IP core network, and the ESnet Science Data Network (SDN). The ESnet IP core network is architected primarily to transport IP packets, and the ESnet SDN is engineered to support dynamic virtual circuits. At each of the hub sites, there is a combination of core routers, peering routers, and performance monitoring systems strategically deployed based on the functional requirements of the location. In addition, there are ESnet owned and managed routers at all primary customer sites. Overall ESnet manages about 270 routers and systems at the Lawrence Berkeley National Laboratory (LBNL), remote Network Operation Centers (NOCs) and sites. ESnet is headquartered at the Lawrence Berkeley National Laboratory (LBNL, CA), but also maintains four remote NOCs (Livermore, CA; Ames, IA; Brookhaven, NY, and Seattle, WA) and two remote data centers (Livermore, CA, and AOA, NY). Critical services are replicated to remote data centers, and engineering staff at the remote NOCs can maintain network services in the event of disruption to its headquarters site.

ESnet operates a network test lab that has various specialized equipment representative of the production routers, switches and testers in use in ESnet in order to test and reproduce network scenarios.

ESnet operates two science services: Federated Trust and PKI systems, and video, audio, and data collaboration systems. The Authentication and Trust Federation (ATF) group has deployed a set of Certification Authorities supporting SciDAC and other DOE science collaborations, as well as ESnet operations.

ESnet support infrastructure consists of dozens of systems that support email, engineering databases, Web servers, etc.

IX. References

[1] Wenji Wu et al., 'End-to-End Network/Application Performance Troubleshooting Methodology," Proceedings of Computing in High Energy Physics (CHEP) 2007, Vitoria, Canada.

[2] http://www.perfsonar.net/

[3] http://www-iepm.slac.stanford.edu/bw/

[4] http://www-iepm.slac.stanford.edu/pinger/

[5] http://maps.google.com/

[6] http://earth.google.com/

[7] A. Bell, Exploring web 2.0: Second Generation Interactive Tools – Blogs, Podcasts, Wikis, Networking, Virtual Words, and More, CreateSpace Publisher, ISBN – 1441449868.

[8] http://www.internet2.edu/performance/pS-NPToolkit

[9] http://e2epi.internet2.edu

[10] http://e2epi.internet2.edu/ndt/

[11] https://plone3.fnal.gov/P0/WAN/netperf/methodology/

[12] http://www.lambdastation.org/

[13] http://monalisa.caltech.edu/monalisa.htm

[14] http://www.web100.org/

[15] J. Boote *et al*., "Instantiating a Global Network Measurement Framework, *LBNL Technical Report,* LBNL-1452E, Jan. 2009.

[16] http://www.perfsonar.net/

[17] http://www-iepm.slac.stanford.edu/pinger/

[18] Y. Lee and S. Kassam, "Generalized Median Filtering and Related Nonlinear Filtering Techniques," *IEEE Transactions on Acoustics, Speech and Signal Processing*, 33(3), 672–683, Jun. 1985.

[19] M. Frey, "Scaling Factors for Estimating the Reliability of Series-Parallel Composition Systems," Hawaii International Conference on Statistics, Honolulu, Hawaii, June 5–9, 2002.

[20] M. Frey, "Responsive buffers—properties in the fluid flow limit," Journal of The Franklin Institute, 343, 137–151, 2006.

[21] M. Frey et al., "A Distributed, Near-Optimal Call Set-Up Protocol for All-Optical Wavelength Routing Networks," Journal of Interconnection Networks, 5(2), 165–180, June 2004.

[22] M. Frey et al., “Wavelength Conversion and Call Connection Probability in WDM Networks,” IEEE Trans. on Communications, 49(10), October 2001.
[23] R. Ogden, Essential Wavelets for Statistical Applications and Data Analysis, Birkhauser, 1997.
[24] M. Baldi et al., “Data Mining Techniques for Effective and Scalable Traffic Analysis,” Proceedings of the 9th IFIP/IEEE International Symposium on Integrated Network Management, 105–118, May 2005.
[25] G. Iannaccone, “Fast Prototyping of Network Data Mining Applications,” Passive and Active Measurement 2006, Adelaide, Australia, 2006.
[26] R. Wolski, “Dynamically Forecasting Network Performance using the Network Weather Service,” Cluster Computing, 1, 119–132, 1998.

Matt Crawford

Computing Division
Fermi National Accelerator Laboratory +1 630 840 3461
P.O. Box 500, MS-120 +1 630 840 3109 FAX
Batavia, IL, 60510 crawford@fnal.gov

Education

Doctor of Philosophy in Physics, University of Chicago, 1985
Thesis: "Statistical Mechanics in a Covariant Gauge," advisor: David N. Schramm
Bachelor of Science (with Honors) in Applied Mathematics and Physics, California Institute of Technology, 1978

Professional Appointments

Fermi National Accelerator Laboratory: Department Head, Data Movement & Storage, 2006-present. Group Leader, Wide Area Systems (2005-2006); CPPM/Computer Security Coordinator (1998-2005); Network Analyst (1992-1997).

University of Chicago: Senior Research Associate, Department of Astronomy and Astrophysics, Physical Sciences Division, and Office of the Provost (1987-1992); Research Associate, Department of Astronomy and Astrophysics (1985-1987).

HONORS AND AWARDS

[1] SuperComputing 2005 Bandwidth Challenge Award for "Distributed TeraByte Particle Physics Data Sample Analysis," with L. Cottrell (SLAC), H. Newman (Caltech), D. Petravick (FNAL).
[2] Fermilab Employee Performance Recognition Award, 2002, for leading computer security technical program.
[3] University of Chicago's Valentine Telegdi Prize, 1978, for doctoral candidacy exam in the Department of Physics.

Publications Related to Proposed Project

[1] W. Wu, M. Crawford, "Potential Performance Bottleneck in Linux TCP," International Journal of Communication Systems (Wiley) 20, Issue 11, pp. 1263-1283 (2007).
[2] W. Wu, M. Crawford, "Performance Analysis of Linux Networking - Packet Receiving," Computer Communications (Elsevier) 30, Issue 5, pp. 1044-1057 (2007).
[3] W. Wu, M. Crawford, "Interactivity vs. Fairness in Networked Linux Systems," Computer Networks (Elsevier) 51, Issue 14, pp. 4050-4069 (2007).
[4] W. Wu, M. Crawford, "End-to-End Network/Application Performance Troubleshooting Methodology," Proceedings of Computing in High Energy Physics (CHEP) 2007, Victoria, BC, Canada, 2007.
[5] W. Wu, M. Crawford, "The Performance Analysis of Linux Networking–Packet Receiving," Proceedings of Computing in High Energy Physics (CHEP) 2006, Mumbai, India, 2006.
[6] M. Crawford, Building Global HEP Systems on Kerberos, Proceedings of Computing in High Energy Physics (CHEP) 2004, Interlaken, Switzerland, 2004.

[7] M. Crawford, Protecting an Open-Network Environment, Computer and Network Security Workshop, sponsored by Institute of Electrical and Electronic Engineers (IEEE), Fox Valley Subsection, the Illinois Institute of Technology (IIT), and the Federal Bureau of Investigation (FBI), Wheaton, Illinois, 2002.
[8] M. Crawford, B. Haberman, Ed. "IPv6 Node Information Queries," RFC 4620, 2006.
[9] M. Crawford. "Router Renumbering for IPv6," RFC 2894, 2000.
[10] M. Crawford, C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering," RFC 2874, 2000.
[11] M. Crawford, "Binary Labels in the Domain Name System," RFC 2673, 1999.
[12] M. Crawford, "Non-Terminal DNS Name Redirection," RFC 2672, 1999.
[13] M. Crawford, "Transmission of IPv6 Packets over Ethernet Networks," RFC 2464 (1998).
[14] M. Crawford, "Transmission of IPv6 Packets over FDDI Networks," RFC 2467 (1998).
[15] M. Crawford, T. Narten, S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks," RFC 2470,1998.

Synergistic Activities

[1] Member of the IPv6 Forum's Technical Directorate.
[2] Member of the Global Grid Forum.
[3] Invited reviewer of the Globus Toolkit, version 4 (gt4) design.
[4] Program co-chair for the Internet2 Joint Techs meetings in January 2008, July 2008, and February 2009.
[5] Invited participant in the NICT-DOE-NSF Optical Network Testbeds Workshop 3, Tokyo, September, 2006.
[6] Invited participant in the Large Scale Networking Coordination Group's (LSN CG) Networking Research Challenges Workshop, Seattle, September, 2008.

List of Collaborators and Co-editors

A. Bobshev, Fermilab
B. Haberman, JHU APL
C. Huitema, Microsoft
F. Wuerthwein, University of California at San Diego
H. Newman, California Institute of Technology
M. Bowden, Fermilab
M. Grigoriev, Fermilab
M. Livny, University of Wisconsin
P. DeMar, Fermilab
W. Wu, Fermilab

Graduate and Postdoctoral Advisors and Advisees

David N. Schramm, University of Chicago

Phil DeMar

Computing Division

Fermi National Accelerator Laboratory +1 630 840 3678
P.O. Box 500, MS-120 +1 630 840 3109 FAX
Batavia, IL, 60510 demar@fnal.gov

Education

Master, June 1984, Computer Science, DePaul University, Chicago, USA

Professional Appointments

Fermi National Accelerator Laboratory

Group Leader, Wide-Area Network and Systems Group, 2006 – Present

Section Head, Networking, 2002 - 2006

Group Leader, Data Communication Group, 1998 – 2002,

Network Analyst, 1985 – 1998

Publications Related to Proposed Project

"End-to-End Network/Application Performance Troubleshooting Methodology," W. Wu, A. Bobyshev, M. Bowden, M. Crawford, P. DeMar, *Proceedings of Computing in High Energy Physics (CHEP) 2007*, Vitoria, Canada.

"Use of Alternate Path WAN Circuit at Fermilab," P. DeMar, A. Bobyshev, M. Crawford, *Proceedings of Computing in High Energy Physics (CHEP) 2007*, Vitoria, Canada.

"Lambda Station: Alternate network path forwarding for production SciDAC applications," A. Bobyshev, M. Crawford, P. DeMar, *Proceedings of Computing in High Energy Physics (CHEP) 2007*, Vitoria, Canada.

"Effect of Dynamic ACL (Access Control List) Loading on Performance of Cisco Routers," A. Bobyshev, P. DeMar, *Proceedings of Computing in High Energy Physics (CHEP) 2006*, Mumbai, Inida.

"Lambda Station; On-demand flow based routing for data intensive grid applications over multi-topology networks," A. Bobyshev, M. Crawford, P. DeMar, *IEEE proceedings of the Third International Conference on Broadband Communications, Networks and Systems*, California, USA, October 1-2, 2006.

"Lambda Station: A forwarding and admission control service to interface production network facilities with advanced research network paths," Phil Demar, Don Petravick, *Proceedings of Computing in High Energy Physics (CHEP) 2004*, Switzerland.

Synergistic Activities

Chair, ESnet Site Coordination Committee; August, 2008 – Present

Co-Chair, Internet2/ESCC Joint Techs Workshops, January, 2008 – Present

Program Co-Chair, Internet2/ESCC Joint Techs Workshops, August 2007 – Present

Member, LHC Optical Private Network Operation Group, January 2007 – Present

Review Panel for DOE Early Career Principle Investigator for Network Research and DOE Small Business Innovative Research (SBIR)

List of Collaborators and Co-editors

Harvey Newman, California Institute of Technology

Dantong Yu, Brookhaven National Laboratory

Shawn McKee, University of Michigan

List of Graduate and Postdoctoral Advisors and Advisees

Not applicable

<u>Brian Tierney</u>

Staff Scientist, ESnet
Lawrence Berkeley National Laboratory (http://acs.lbl.gov/~tierney)

Education

M.S. Computer Science, San Francisco State University, 1990
B.A. Physics, Minor in Computer Science, The University of Iowa, 1985

Research Interests

High-performance networking and network protocols; distributed system performance monitoring and analysis; network tuning issues; data intensive computing and high performance storage systems; cyber security and intrusion detection systems; the application of distributed computing to problems in science and engineering.

Narrative

Brian L. Tierney is a Staff Scientist and group leader of the Advanced Network Technologies Group in ESnet. Prior to this he was group leader of the Data Intensive Distributed Computing Group in the Distributed Systems Department at LBNL. He has been the PI for several DOE research projects in network and Grid monitoring systems for data intensive distributed computing.

Selected Publications

“A User Driven Dynamic Circuit Network Implementation”, Chin Guok, David Robertson, Evangelos Chaniotakis, Mary Thompson, William Johnston, Brian Tierney, Proceedings of the Distributed Autonomous Network Management Systems Workshop (DANMS 2008), November 2008.

“Essential Grid Workflow Monitoring Elements”, Daniel K. Gunter, Keith R. Jackson, David E. Konerding, Jason R. Lee and Brian L. Tierney, The 2005 International Conference on Grid Computing and Applications (GCA’05).

“Instrumentation and Monitoring”, B. Tierney and J. Hollingsworth, Chapter 20, The Grid Blueprint for a New Computing Infrastructure, 2nd Edition, Elsevier, 2003

“On-Demand Grid Application Tuning and Debugging with the NetLogger Activation Service”, D. Gunter, B. Tierney, C. E. Tull, V. Virmani, 4th International Workshop on Grid Computing (Grid2003).

“Enabling Network Measurement Portability Through a Hierarchy of Characteristics”, B. Lowekamp, B. Tierney, Les Cottrell, R. Hughes-Jones, T. Kielmann, M. Swany, 4th International Workshop on Grid Computing (Grid2003).

“An Infrastructure for Passive Network Monitoring of Application Data Streams”, D. Agarwal, J. M. González, G. Jin, B. Tierney, Passive and Active Measurement Workshop, April, 2003.

“Dynamic Monitoring of High-Performance Distributed Applications”, D. Gunter, B. Tierney, K. Jackson, J. Lee, M. Stoufer, Proceedings of the 11th IEEE Symposium on High Performance Distributed Computing, July 2002.

“Monitoring Data Archives for Grid Environments”, J. Lee, D. Gunter, M. Stoufer, B. Tierney, Proceeding of IEEE Supercomputing 2002 Conference, Nov. 2002.

“A TCP Tuning Daemon”, T. Dunigan, M. Mathis and B. Tierney, Proceeding of IEEE Supercomputing 2002 Conference, Nov. 2002.

“Enabling Network-Aware Applications”, B. Tierney, D. Gunter, J. Lee, M. Stoufer, Proceedings of the 10th IEEE Symposium on High Performance Distributed Computing, August 2001.

“TCP Tuning Guide for Distributed Application on Wide Area Networks”, Tierney, B, Usenix ;login, Feb. 2001.

“Using High-Speed WANs and Network Data Caches to Enable Remote and Distributed Visualization”, Bethel, W., Tierney, B., Lee, J., Gunter, D., Lau, S., Proceeding of the IEEE Supercomputing 2000 Conference, Nov. 2000.

“The NetLogger Methodology for High Performance Distributed Systems Performance Analysis”, B. Tierney, W. Johnston, B. Crowley, G. Hoo, C. Brooks, D. Gunter, Proceeding of IEEE High Performance Distributed Computing conference, July 1998.

Collaborators: E. Boyd (Internet2), A. Chervenak (ISI), E. Deelman (ISI), I. Foster (ANL), K. Keahey (ANL), C. Kesselman (ISI), R. Kettimuthu (ANL), M. Livny (UWisc), R. Madduri (ANL), L. Pearlman (ISI), V. Paxson (ICIR), J. Schopf (NSF), R. Sommer (ICIR), M. Swany (UDel), M. Zekauskas (Internet2)

Graduate Advisor: William Johnston, LBNL

Wenji Wu

Computing Division

Fermi National Accelerator Laboratory | +1 630 840 4541
P.O. Box 500, MS-120 | +1 630 840 3109 FAX
Batavia, IL, 60510 | wenji@fnal.gov

Education

Ph.D., Dec. 2003, Computer Engineering, University of Arizona, Tucson, USA
Master, May 2001, Industrial Engineering, University of Arizona, Tucson, USA
Master, May 1997, System Engineering, Zhejiang University, Hang Zhou, China

Professional Appointments

June 2005 – Present, Network Researcher, Fermi National Accelerator Laboratory
September 2004 – June 2005, Research Assistant Professor, ECE dept., Univ. of Arizona, Tucson
September 2003 – September 2004, Research Engineer, ECE dept., Univ. of Arizona, Tucson

Publications Related to Proposed Project

Journals

[1] Wenji Wu, Matt Crawford, "Interactivity vs. Fairness in Networked Linux Systems," *Computer Networks* (Elsevier), Volume 51, Issue 14, pp. 4050 – 4069, 2007.
[2] Wenji Wu, Matt Crawford, "Performance Analysis of Linux Networking – Packet Receiving," *Computer Communications* (Elsevier), Volume 30, Issue 5, pp. 1044 – 1057, 2007.
[3] Wenji Wu, Matt Crawford, "Potential Performance Bottleneck in Linux TCP," *International Journal of Communication Systems* (Wiley), Volume 20, Issue 11, pp. 1263 – 1283, 2007.
[4] Wenji Wu, Natalia Gaviria, Kevin M. McNeill, "Two-layer Hierarchical Wavelength Routing for Islands of Transparency Optical Networks," *Computer Communications* (Elsevier), Volume 29, Issue 15, pp. 2952-2963, 2006.
[5] Wenji Wu, Ralph Martinez, Peng Choop, "A Modeling Process and Analysis of GMPLS-based Optical Switching Routers," *Journal of Photonic Network Communications*, Volume 8, Issue 1, Jun 2004.

Conferences:

[1] Wenji Wu et al., 'End-to-End Network/Application Performance Troubleshooting Methodology," *Proceedings of Computing in High Energy Physics* (CHEP) 2007, Vitoria, Canada.
[2] Wenji Wu and Matt Crawford, "The Performance Analysis of Linux Networking–Packet Receiving," *Proceedings of Computing in High Energy Physics* (CHEP) 2006, Mumbai, India, 2006.
[3] Wenji Wu, Ralph Martinez, Peng choop, "Simulation-Based GMPLS Photonic Router using the OPNET MPLS Module," *OPNETWORKS2002*, Aug. 2002, Washington. (Best paper award)
[4] Wenji Wu, Ralph Martinez, Peng choop, "Constraint-based Routing for Islands of Transparency Optical Networks," *OPNETWORK2003*, Aug. 2003, Washington D.C, 2003

Grants

(1) Wenji Wu, Co-Principal Investigator, “Adaptive Voice Quality Enhancement Mechanisms for VoIP”, NSF Connection One Grant ($84,000), July 2004.
(2) Wenji Wu, Co-Principal Investigator, “BAE Connection One Non-Core Research”, supported by BAE SYSTEMS ($120,000), September 2004.

Professional Society

(1)IEEE Member
(2)IEEE Communications Society Member
(3)LHC Optical Networking Group

Reviewing

(1) DOE SBIR/STTR review panel, 2006
- Reviewing proposal “Bandwidth Aware Network Interface Card”

List of Collaborators and Co-editors

Dantong Yu, Ph.D., Brookhaven National Laboratory
Kevin

McNeill, Ph.D., BAE Systems

Mark Bowden, Fermi National Accelerator
Laboratory
Matt Crawford, Ph.D., Fermi National Accelerator Laboratory
Mingkuan Liu, Ph.D., University of Arizona
Natalia Gaviria, Ph.D., University of Arizona
Phil DeMar, Fermi National Accelerator Laboratory
Ralph Martinez, Ph.D., BAE Systems
Xian-He Sun, Ph.D., Illinois Institute of Technology

Graduate and Postdoctoral Advisors and Advisees

Kevin McNeill, Ph.D., BAE Systems
Pitu Mirchandani, Ph.D., University of Arizona
Ralph Martinez, Ph.D., BAE Systems

Maxim Grigoriev

Com
puting Division
Fermi National Accelerator Laboratory +1 630 840 6024
P.O. Box 500, MS-120 +1 630 840 3109 FAX
Batavia, IL, 60510 maxim@fnal.gov

Education

Master Diploma, June 1993, Mathematics and Computer Science,
Moscow State Lomonosov University, Moscow, Russia

Professional Appointments

Fermi National Accelerator Laboratory:
Systems Analyst II, Wide Area Network and Research Group, 2001-present

Publications Related to Proposed Project

1 M. Crawford, P. DeMar, M. Grigoriev et al., "End-to-End Network/Application Performance Troubleshooting Methodology," *Proceedings of Computing in High Energy Physics (CHEP) 2007*, Victoria, Canada.

2 M. Crawford, P. DeMar, M. Grigoriev et al., "Deploying perfSONAR-based End2-End monitoring for production US CMS networking", *Proceedings of Computing in High Energy Physics (CHEP) 2007*, Victoria, Canada.

3 A. Bobyshev, P. DeMar, V. Grigaliunas, M. Grigoriev, "Use of Flow Data for Traffic Analysis and Network performance characterization", *Proceedings of Computing in High Energy Physics (CHEP) 2007*, Victoria, Canada.

4 A. Bobyshev, M. Crawford, P. DeMar, M. Grigoriev et al., "Lambda Station: Production Applications Exploiting Networks in Data Intensive High Energy Physics", *Proceedings of Computing in High Energy Physics (CHEP) 2006*, Mumbai, India.

5 M. Grigoriev, E. Berman, P. DeMar et al., "Network Information and Monitoring Infrastructure (NIMI"), Proceedings of CHEP2006, Mumbai, India

6 M. Grigoriev, R. Les Cottrell, C. Logg, M. Chhaparia et al., "Evaluation of Techniques to Detect Significant Network Performance Problems using End-to-End Active Network Measurements", Proceedings of IEEE NOMS2006, Vancouver, Canada

7 A. Bobyshev, M. Crawford, P. DeMar, "Lambda Station; On-demand flow based routing for data intensive grid applications over multi-topology networks," *IEEE*

Synergistic Activities

Member, perfSONAR-PS project, December 2006 – Present Member,
LambdaStation project, 2005 – Present
Member, PingER project, 2001 – Present

Andrey Bobyshev

Computing Division
Fermi National Accelerator Laboratory +1 630 840 2499
P.O. Box 500, MS-120 +1 630 840 3109 FAX
Batavia, IL, 60510 bobyshev@fnal.gov

Education

Master of Science Degree in Computer Science and Applied Mathematics, February 1982, Moscow State Engineering and Physics Technical University

Professional Appointments

Fermi National Accelerator Laboratory:
Network Architect I, Wide Area Network and Research Group, 2000-present

Research Interests

High-performance networking, analysis of flow data and traffic characterizations, web-based services for networking

Narrative

Andrey Bobyshev is a network architect of the Wide Area Network Group in Fermilab. He leads design and operational support for networking at US CMS Tier 1 facility. He played a leading role in DOE funded Lambda Station project. Prior to this he was working for Deutsches Elektronen-Synchrotron (Information Technology Division), Hamburg, Germany, and leaded or actively participated in several International projects on establishing Internet connectivity for High Energy Physics community.

Publication related to th proposed project:

1. A.Bobyshev, M.Crawford, P. DeMar, V. Grigaliunas, M. Grigoriev, A.Moibenko, D. Petravick, R. Rechenmacher,H. Newman, J. Bunn, F. Van Lingen, D. Nae, S. Ravot, C. Steenberg, X. Su, M. Thomas, Y. Xia, ***Lambda Station: Production applications exploiting advanced networks in data intensive high energy physics***, Proceedings of CHEP06, TIFR, Mumbai, India, 13-17 February 2006.
2. A.Bobyshev, M.Crawford, V.Grigalinus, M.Grigoriev, R. Rechenmacher, ***Investigating the behavior of network aware applications with flow-based path selection***, Proceedings of CHEP06, TIFR, Mumbai, India, 13-17 February 2006.
3. A.Bobyshev, Effect of dynamic ACL (access control list) on performance of Cisco routers, Proceedings of CHEP06, TIFR, Mumbai, India, 13-17 February 2006.
4. A. Bobyshev , M. Crawford , P. DeMar, V. Grigaliunas, M. Grigoriev, A.Moibenko, D. Petravick, R. Rechenmacher, H. Newman, J. Bunn, F. Van Lingen, D. Nae, S. Ravot, C. Steenberg, X. Su, M. Thomas, Y. Xia, ***Lambda Station: On-demand flow based routing for data intensive grid applications over multitopology networks***, IEEE proceedings of the Third International Conference on Broadband Communications, Networks and Systems, publication ID 116, San Jose, California, USA, Oct 1-2, 2006
5. Maxim Grigoriev,Andrey Bobyshev, Matt Crawford, Phil Demar, Vyto Grigaliunas, Alexander Moibenko, Don Petravick, Harvey Newman, Conrad Steenberg, Michael Thomas . ***Lambda Station: Alternate Network Path Forwarding for Production SciDAC Applications***,

Proceedings of CHEP07, Victoria BC, Canada, 2-4 September 2007

6. CHEP04, oral presentation, Interlaken, Switzerland, 27th September - 1st October 2004.
7. A.Bobyshev, M.Grigoriev, Methodologies and techniques for analysis of network flow data, CHEP04, paper and oral presentation, Interlaken, Switzerland, 27th September - 1st October 2004.
8. A.Bobyshev, R.Rechenmacher, P.Demar, M.Ernst, WAN Emulation Development and Testing at Fermilab, CHEP04, paper and oral presentation, Interlaken, Switzerland, 27th September - 1st October 2004.
9. A.Bobyshev,P.Demar,D.Lamore, AutoBlocker: A system for detecting and blocking of network scanning based on analysis of netflow data, CHEP04, paper and oral presentation, Interlaken, Switzerland, 27th September - 1st October 2004.
10. CHEP06, oral presentation, poster,CHEP06, TIFR, Mumbai, India, 13-17 February 2006
11. GridNets2006, oral presentation, San Jose, California, USA, Oct 1-2, 2006
12. JointTech/ESCC, LambdaStation, oral presentation, Fermilab, July 15-19, 2007
13. CHEP07, Lambda Station, oral presentation,Victoria BC, Canada, 2-4 September 200