# Terapaths DWMI FY 2006 Accomplishments

This is the start of the FY2006 DoE Terapaths DWMI Progress Report due September 10, 2006

Submitted

## Report

**Terapaths: A QoS Collaborative Data Sharing Infrastructure for Petascale Computing Research: DWMI: Datagrid Wide Area Monitoring Infrastructure**
*Les Cottrell, Yee-Ting Li & Connie Logg, Stanford Linear Accelerator Laboratory (SLAC)*

**Summary**:
The main goal of the DWMI project is to build, deploy and effectively learn how to use an initially relatively small but rich, robust, sustainable, manageable network monitoring infrastructure focused on the needs of critical HEP experiments such as Atlas, BaBar, CMS, CDF and D0.

Today's data intensive sciences, such as High Energy Physics (HEP), need to share large amounts of data at high speeds. This in turn requires high-performance, reliable end-to-end network paths between the major collaborating sites. In addition network administrators need alerts of anomalous events and grid middleware and end-users need long and short-term forecasting for applications and network performance in order to plan, set expectations and to trouble-shoot problems. To enable this requires a network monitoring infrastructure between the major sites that can help to notify and identify potential problems.

**Active monitoring**: We have developed an active network monitoring toolkit (IEPM-BW) that provides network measurements, data archiving, analysis, reporting and visualization. This has been deployed and making regular measurements from the following major LHC related sites: CERN, BNL, Caltech, FNAL, SLAC, and Taiwan. We also have about 45 locations worldwide that are being monitored from these sites. We use a selection of probes based on the quality and interest in the path being measured utilizing metrics such as network routes, round trip time, one-way delays, available bandwidth and achievable throughput. We are working closely with the LHC-ATLAS physicists, the LHC Optical Private Network (LHCOPN) working group, BNL Terapaths, CERN and UltraLight to install the toolkit and ensure the data and presentations meet their needs.
With regards to cross-domain end-to-end MPLS circuits, we are currently developing mechanisms to automatically schedule active measurements using IEPM-BW, Terapaths and OSCARS to compare the performance of complete end-to-end QoS paths against normal production services.

Deployment of IEPM LHC monitoring, red dots show monitor sites, green show monitored sites

**Passive Monitoring**: We have studied and reported on limitations using current active end-to-end network measurement techniques in future high-speed networks. As a result of this we are exploring the effectiveness of using passive (e.g. Netflow) tools against active measurements.
In conjunction with BNL we are building a netflow monitoring toolkit using open source software to bring together quality tools to gather, store, process, analyze and visualize the performance information. The intent is to augment or possibly replace existing active monitoring mechanisms and to deploy it at LHC sites for both network performance monitoring and network security analysis.
In fact, much of our development is being steered by the requirements of the BNL site, specifically for the Terapaths project, where we have an early development version of the entire suite running collecting real netflow data from production network systems.

**Event Detection and Diagnosis**: With the expansion of network infrastructure and the increased networked applications in use, it is becoming increasingly impossible for network managers to manually review the large number of reports to detect, and more importantly diagnose problems. Thus, we are developing tools to automate this activity by forecasting network performance to detect anomalous events and reporting them.
We are currently field testing the Plateau, Holt-Winters and KS algorithms on production networks via IEPM-BW. As part of this, in the last year, we have also detected, reported (with in-depth case studies) and helped to manually diagnose major problems at sites such as BNL, Taiwan, SDSC, NRL, BINP, and CERN.
We are also in the process of building a framework from which these network alerts can be used to automatically diagnose and identify the cause of network problems. Utilizing heuristics analysis and an innovative scoring system to pin-point the cause of an event, we are actively working closely with network providers to field test, review and corroborate our design to identify, locate and verify network problems and their symptoms.

**Internet Measurement Confederation**: An important aspect of being able to both understand and diagnose network performance problems is the unification of reporting formats and the understanding of tool performance on the Internet.
We have recently started close collaboration with both Internet2 and ESnet to help develop and expand the functionalities of the international PerfSONAR project. PerfSONAR has gained much momentum over the last few months due to its open-source, open-community, open-standards based ethos of network monitoring. SLAC are delighted to help contribute our network analysis expertise and experience to deploy and apply PerfSONAR based technologies for projects such as the LHC.
We aim to apply much of our existing analysis frameworks and tools to benefit the PerfSONAR project - including that of data visualization, event detection and network problem diagnosis.

**High speed data transport**: Our world-leadership role in evaluating TCP transport algorithms in production networks for large scale science projects such as the LHC, has led Microsoft to request our help in evaluating their next generation TCP stack (CTCP). Given the extent of Windows deployment it is critical to ensure that CTCP performs well without a negative impact upon the Internet community.
As part of this we have identified and aided the testing of numerous added features to aid the performance of the delay-based congestion control algorithm used in CTCP. We have finalized our initial report into the deployment impact of using CTCP in production environments on both long and short distance high speed Internet paths, and are now writing a joint conference paper with Microsoft.

For further information contact:
Dr. Les. Cottrell,
Stanford Linear Accelerator Centre
Scientific Computing and Computing Services
Phone: 650-926-2523
cottrell@slac.stanford.edu