SLAC-PUB-8382
March 2000

International Network Connectivity and Performance, The Challenge from High Energy Physics *

Warren Matthews, Les Cottrell, Charles Granieri

Stanford Linear Accelerator Center,
Operated for the U.S. Department of Energy by Stanford University,
2575 Sand Hill Road, Menlo Park, CA, 94025, USA.

Abstract

International Connectivity between research networks is vitally important to the High Energy and Nuclear Physics (HENP) research community. In this paper, the challenge from the needs of HENP will be reviewed, and long term international performance of Internet2 International members and the future prospects for HENP will be discussed.

Talk presented to Internet2 International Task Force,
at the Internet2 Spring Meeting, Washington D.C., March 27, 2000.

*Work supported by Department of Energy contract DE–AC03–76SF00515.

1 Introduction

The requirements of the new generation of High Energy and Nuclear Physics (HENP) experiments such as the BaBar detector at the Stanford Linear Accelerator Center (SLAC), the Relativistic Heavy Ion Collider (RHIC) groups at the Brookhaven National Laboratory (BNL) and the LHC projects currently under development at the European Center for Particle Physics (CERN) are a huge challenge to networking.

In order to increase understanding, and to improve performance and connectivity by identifying bottlenecks and allocating resources, the HENP networking community has been actively monitoring the network for over five years.

2 The Special Needs of HENP

HENP researchers require the usual applications such as email, web browsing, FTP and SSH, and highly interactive applications such as voice-over-IP (VoIP) and video conferencing are also important and are playing an increasing role in collaborative research. Applications vary in their sensitivity to network performance but most will be severely affected by packet loss over a few percent and some will be unusable with even moderate packet loss. The rapid growth of commercial interests on the Internet has improved HENP research with better and faster-to-market technologies that have become essential tools but at the same time academic traffic often shares the same overloaded routers and congested links as the commercial traffic. Consequently reliance on networking, even in the prescence of the congestion they suffer, has made it essential that the network maintain a high level of performance. Dedicated research networks with an Acceptable Use Policy (AUP) that disallows traffic to commercial sites, or some sort of service quality such as dedicated bandwidth carved from a commodity link or preference for packets of a certain type or to a certain destination are available techniques and will become increasingly more utilized by HENP.

Figure 1 shows the reconstruction of a real event from the BaBar detector. Events such as this take around 60kbytes of disk space and occur approximately 10 times per second (the exact rate depends on how sensitive the physicists set the trigger). The detector is expected to run for 8-10 months a year and in all, the detector is expected to collect several petabytes (10^{15} bytes) of data during its lifetime. Such high statistics are required to obtained very accurate measurements. Much of this data will be distributed

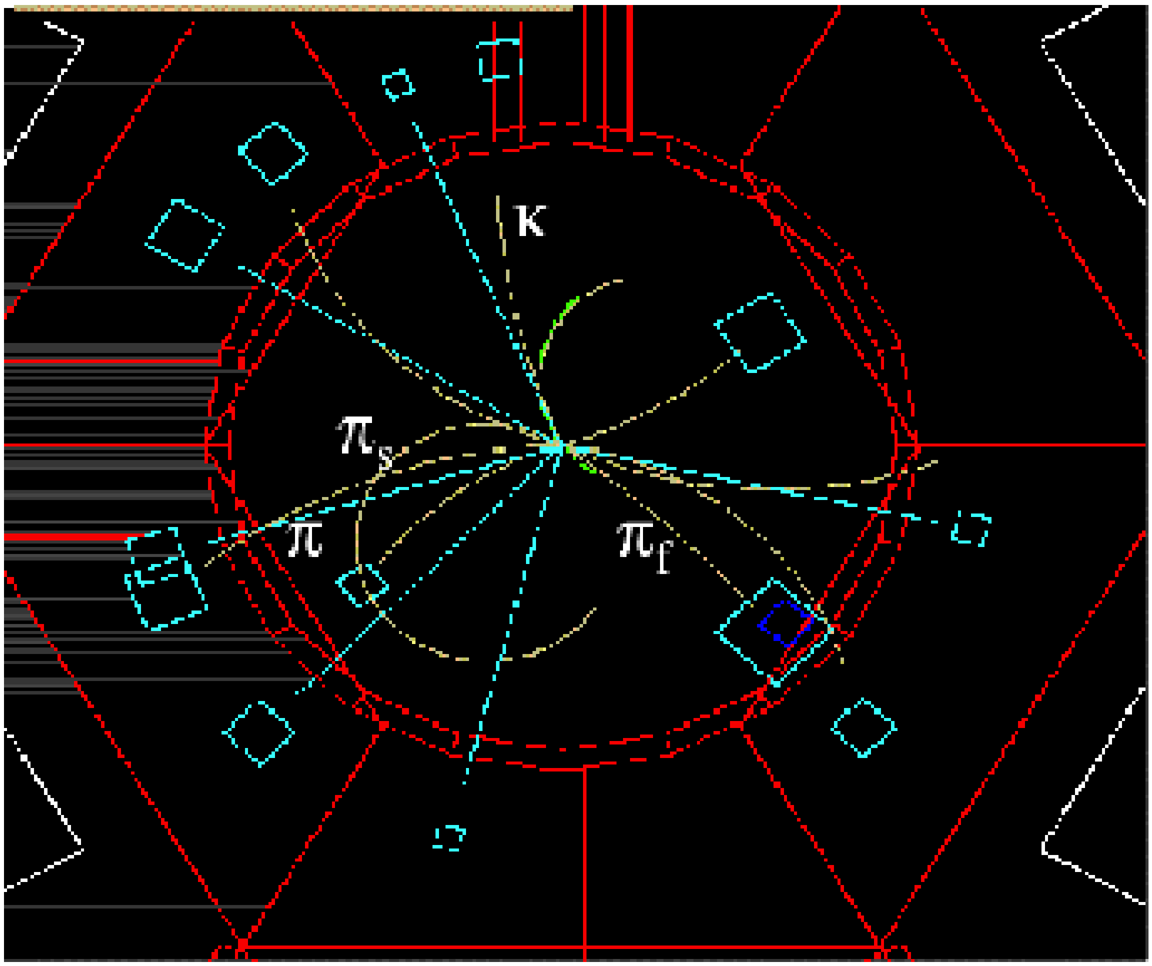

Figure 1: Reconstruction of a real event from the BaBar detector. This view is directly down the beam pipe. Electrons and positrons collide in the center and annihilate each other leaving only energy ($E=mc^2$). A very short time later the energy creates a set of new particles according to certain laws, that are thrown at high energy into the various subdetectors resulting in signals that are read out and the data is collected and written to disk.

to the experiment's collaborators at Universities and Institutes throughout the world. BaBar, for example, has collaborators in North America, West Europe, the Former Eastern Block and the Far East. Other experiments also include other regions. Such large international collaborations are typical in modern HENP. Infact, this large international collaborative nature of HENP research has meant that HENP has pioneered early Internet connectivity to many countries including China, Armenia, georgia, Kazakhastan, Ukraine and Belarus. Figure 2 shows the world wide distribution of HENP research and PingER monitoring. Countries shaded grey have a University collaborating on a large HENP experiment at CERN, SLAC, FNAL or BNL. In total 50 countries are involved in the experiments at these labs.

3 The PingER Project

The Mathematical, Informational and Computational Sciences (MICS) division of the US Department of Energy (DoE) supports an Internet End-to-end Performance Monitoring (IEPM) project which has developed a set of tools for Ping End-to-end Reporting, and has become know as `PingER`.

The emphasis for `PingER` monitoring is on sites involved in HENP re-

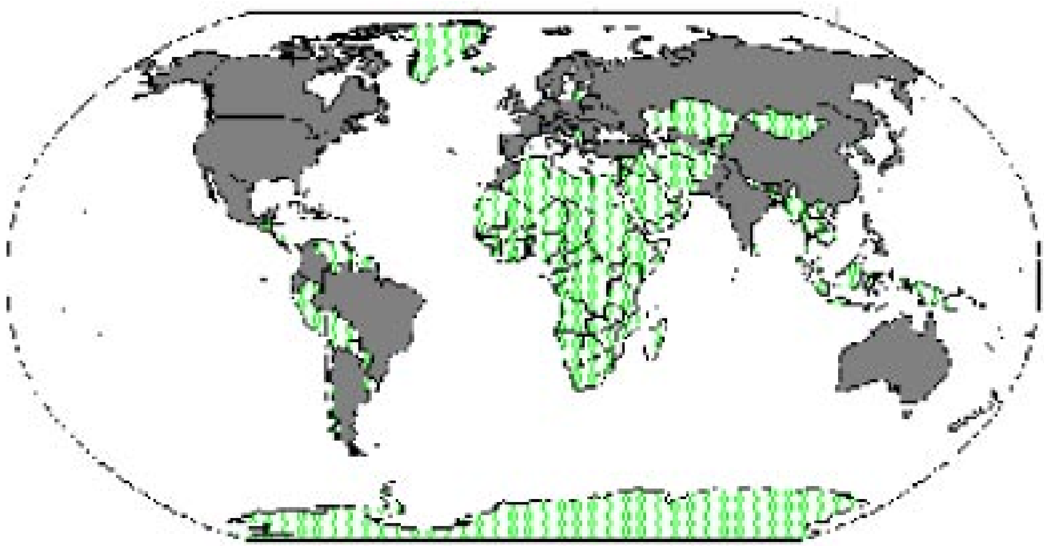

Figure 2: The reach of HENP research throughout the World. The darkly shaded countries have at least one University collaborating on a large HENP experiment. The experiments are at , SLAC, FNAL, and BNL. In total 50 countries are involved in the experiments at these labs.

search, but also Fusion, Basic Energy Sciences (BES), and commercial interests are reviewed.

At the time of writing, the `PingER` project consists of 30 monitoring sites[†] in 15 countries and nearly 600 remote sites in 72 countries. Typically, a monitoring site is a HENP laboratory and it monitors remote sites that are Universities that collaborate on experiments at the laboratory. Recently the `PingER` project coordinators have attempted to manage the monitoring to improve wider understanding by introducing a list of beacon sites that all monitoring sites are requested to include in their lists of remote sites.

`PingER` sends 11x100 byte `ping` packets followed by 10x1000 byte `ping` packets to each remote site every 30 minutes scheduled with cron. The `ping` packet loss and RTT are recorded. Further metrics are derived such as unreachability, quiescence and unpredictability. An estimate of throughput can also be derived.

Details of the PingER project and methodology are available in [1], comparisons between PingER and other monitoring experiments are available in [2].

[†]Several more are expected to come on in the next few months

4 Results

Across US research networks such as the Energy Sciences Network (ESnet), which provides connectivity for US labs, and the Internet2 backbone Abilene, average line utilization is low and consequently packet loss is typically less than 1%. It is desirable to operate at low average line utilization, or at least with significant head room because at times there will be a need to carve off significant bandwidth for a video conference or other intensive applications that would normally impair performance for other users. There are occassional glitches, such as outages caused by hardware or IOS upgrades, but the long term trend is mostly towards steady improvement, that is improvements in infrastructure have stayed one step ahead of bandwidth demands. This is a considerable achievement in engineering since traffic on ESnet for example, is doubling every year. There are continual further improvements planned and capacity across the ESnet backbone is expected to be to Terabit per second in 2003-2005.

4.1 International Connectivity

The International collaborative nature of HENP research described in section 2 and the associated movement of data, is the reason for the importance of network connectivity for HENP. It is so important in fact, that the International Committee for Future Accelerators (ICFA) created the Standing Committee for Inter-regional Connectivity (SCIC) specifically to oversee development of networks for HENP research.

Groups such as the ESnet International Working Group (ESI) and the Internet2 International Task Force (I2I) are a key part of this development effort.

The HENP community wants to improve connectivity between the labs and research Universities. ESI facilitates this by bringing together interested parties and co-ordinating with ESnet and peering with US and non-US networks, commercial and research networks to provide better connectivity between US labs and US and non-US Universities. Similarly, the I2I membership also facilitates this because they wish to peer with Abilene to improve performance to US Universities. CERN is a member of Internet2 and the researchers at 32 US Universities that are collaborating on the LHC Experiments that are planned to take data in 2005, will benefit directly. Also, there will be LHC Regional centers at BNL and Argonne National Lab (ANL).

In addition, there are 9 US Universities on the Zeus experiment at the Deutches Elektronen Synchrotron (DESY), which is connected to DFN.

In the future, the Particle Physics Data Grid (PPDG) and PPDG sites on Abilene such as the University of Wisconsin at Madison may improve performance by distributing data to locations in separate geographical and logical regions.

SLAC has its primary connectivity through ESnet but also has a connection to the California Regional Network (CALREN2) [‡], Fermi National Accelerator Center (FNAL) also has its primary connectivity through ESnet and also has a connection to MREN. Outside the US, academic and research (A&R) Institutions, including HENP Universities, are typically connected to a single National Research Network (NRN). The Canarie network in Canada, the CUDI network in Mexico, the DFN network in Germany, the Italian GARR network, HUNGARNET in Hungary, the Israeli IUCC, the RedIris network in Spain, TAnet in Taiwan and JAnet in the UK all connect Universities collaborating on Experiments at US labs.

Presently, many of these networks have both Abilene and ESnet peering. Some have one-or-the-other and a few have neither. In the near future, more will become connected to Abilene, and either plan to or have the opportunity

[‡]and since CALREN2 is a connector to Abilene and transits traffic, SLAC has a connection to Abilene

to peer with ESnet because Abilene peers with ESnet in many places, and the I2I member can peer with ESnet at the same location as their Abilene connection.

DFN closed its Perryman POP in mid-Feb and lost direct peering with ESnet. For several weeks packets were routed across a commercial provider until "direct" peering was re-established via Dante in New York in mid-March. Figure 3 shows packet loss between the DESY lab and FNAL during this time. Packet loss went from 0% to ranging from 3% to as much as 9%.

High packet loss is also observed between SLAC and the Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT) in Spain. The Spanish NRN, RedIris, connects with both ESnet and Internet2 via commercial networks. Figure 4 shows that the RTT between SLAC and CIEMAT and between Stanford and CIEMAT varies dramatically depending on the day, with significantly lower round trip times on the weekends compared to the working week. Such a pattern is typical of a congested link. It is intriguing to consider what level of performance will be achieved when RedIris peers with Internet2, or if routing to ESnet were via Dante as with DFN.

Most of the I2I members, who are already connected, peer with Abilene

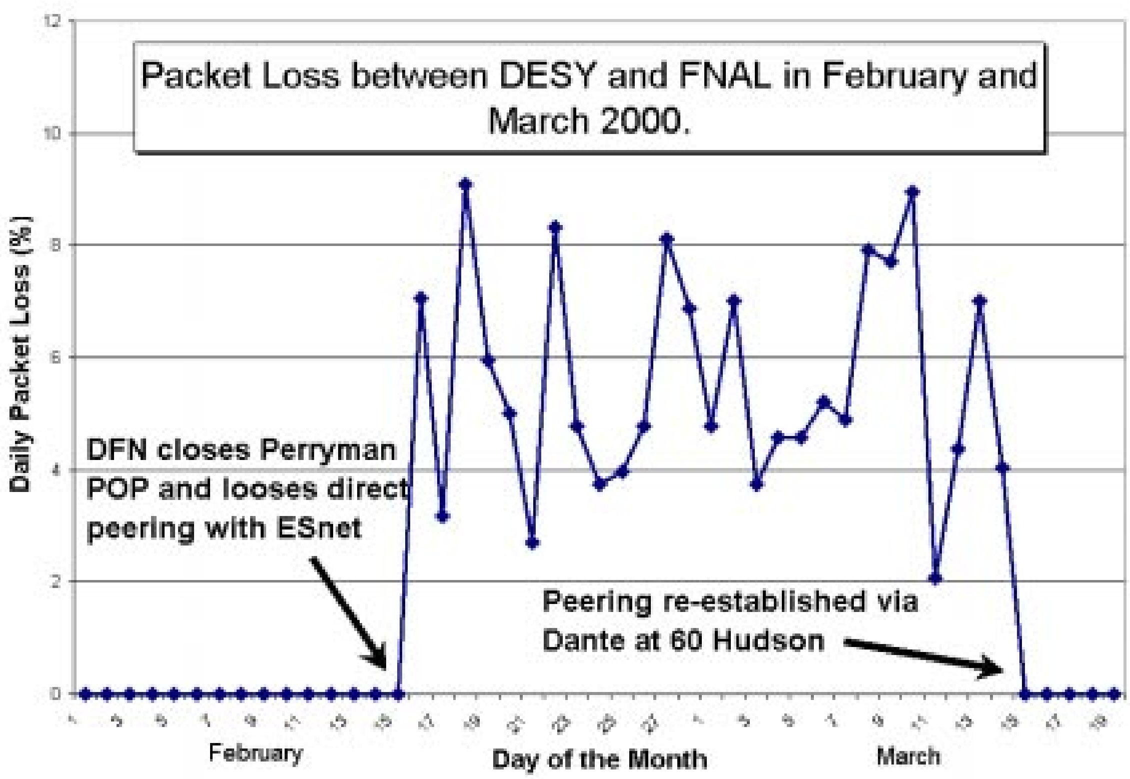

Figure 3: Packet Loss between DESY and FNAL in February and March 2000. DFN closed its POP in mid-February and packets were transitted across a commercial provider until peering was established via Dante in mid-March.

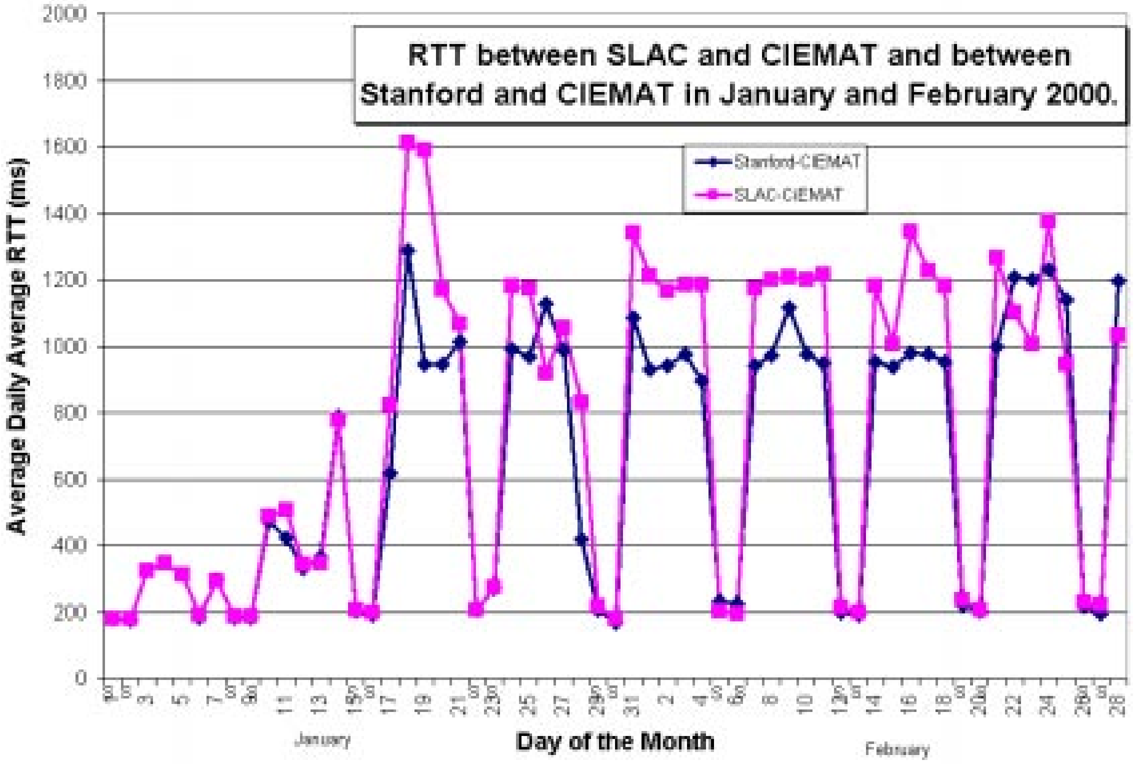

Figure 4: Round Trip Time between SLAC and CIEMAT and between Stanford and CIEMAT has a remarkable day-of-the-week effect that is typical of a congested link. Note the weekend days marked with S.

in New York or at the STAR TAP near Chicago, CUDI will connect with CALREN2 at San Diego, and CALREN2 will transit the traffic to Abilene.

Figure 5 shows the packet loss between National Autonomous University of Mexico (UNAM) and SLAC. There is a clear day-of-the-week effect, a stong indicator of congested line. During the working week, the packet loss is typically 10%, resulting in very poor performance for interactive applications.

UNAM is a collaborator on the D0 experiment at FNAL. Since performance across ESnet is good, the performance between UNAM and FNAL can be approximated by the performance between UNAM and SLAC.

Currently the route between SLAC and UNAM goes via ESnet to the Washington DC hub, and via commercial providers between Washington DC and Mexico. ESnet has a hub in San Diego at the General Atomic Corporation (GAC) where it peers with CALREN2. Performance to ESnet would be significantly improved if CUDI peered directly with ESnet here along with the peering to Abilene, and avoided the congested commercial Internet.

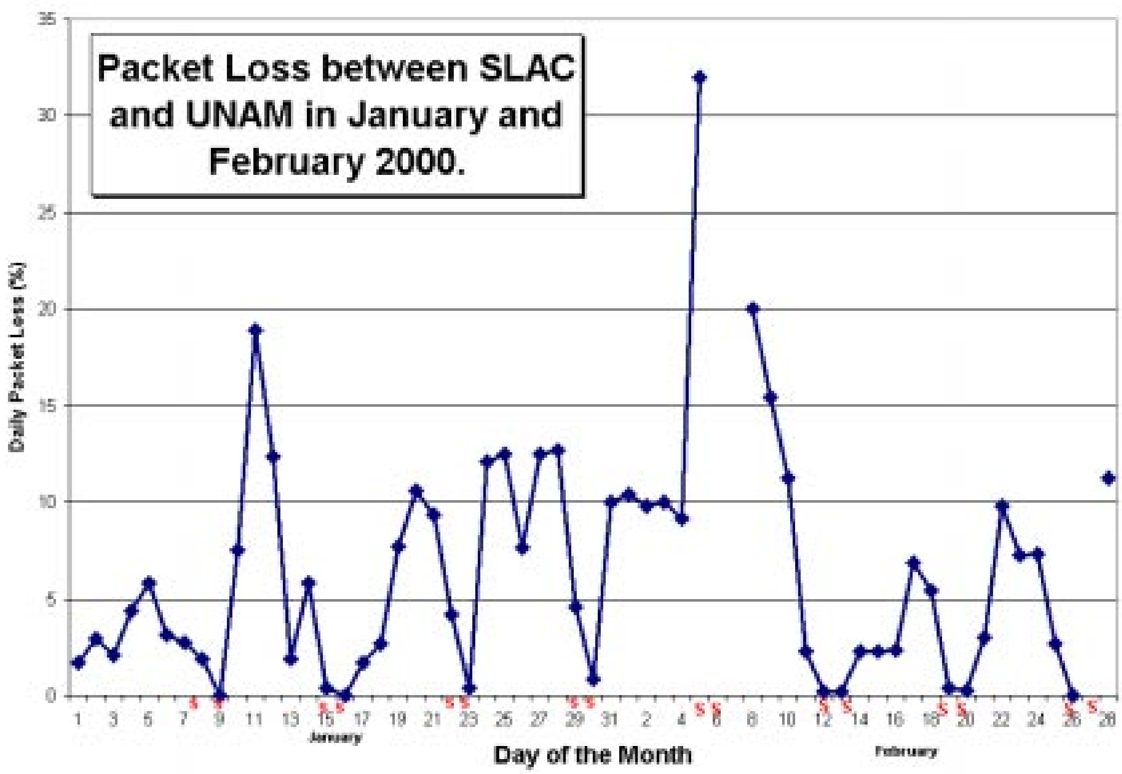

Figure 5: Packet Loss between SLAC and the National Autonomous University of Mexico (UNAM) in January and February 2000. Note the weekend days marked with S.

5 Prospects

Groups such as Internet2 International and ESnet International, and world-wide network initiatives such as Global Crossing and the PPDG are advancing high speed networking throughout the world and providing connectivity to an ever increasing diversity of scientific researchers.

The HENP community looks forward to the improved connectivity that the I2I members in the Former Eastern Block and South America will soon enjoy, and wishes to see the ESI and I2I membership include networks in Africa [§], Central Asia, and all remote regions.

Other I2I members such as Dante and NORDUnet may be able to improve performance by providing connectivity to some NRNs that otherwise would not connect to high speed networks.

Deployment of High Speed networks throughout the world, continuing improvements in bandwidth and peering between them, and optimising routing are all essential to the success of the new and future generation of HENP accelerators and experiments. In addition further use of Quality of Service techniques will be increasingly utilized.

Furthermore, closer cooperation between A&R networks to permit Inter-

[§] Morocco is on Atlas, and South Africa is an observer nation at CERN.

network Differentiated Services or relaxing of AUPs to allow transit traffic for specific cases may also increase performance for researchers.

In addition, the Arena project will assist network managers and engineers to keep track of the rapid improvment in services. The authors encourage everyone to contribute to the project. Visit the project webpages at

`http://www.internet2.edu/international`

Acknowledgements

This work could not be possible without the enthusiatic support of the myriad scientists and engineers who make high speed network for the HENP community a reality. We thank them all.

References

[1] Matthews, W., and Cottrell, L., "*The PingER Project: Active Internet Performance Monitoring for the HENP Community*," to be published in IEEE Communications Magazine on Network Traffic Measurements and Experiments.

[2] Matthews, W., and Cottrell, L., "*Internet End-to-end Performance Monitoring for the High Energy Nuclear and Particle Physics Community*," to be presented at the Passive & Active Measurement Workshop (PAM), Hamilton, New Zealand, April 2000.