

**PingER: Internet Performance Monitoring - How No Collisions
Make Better Physics.**

W. Matthews, L. Cottrell

Presented at International Conference on Computing in High Energy Physics and
Nuclear Physics (CHEP 2000), 2/7/2000—2/11/2000, Padova, Italy

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309

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

PingER: Internet Performance Monitoring

How No Collisions Make Better Physics.

W. Matthews, L. Cottrell

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

Abstract

Internet connectivity is critical infrastructure for modern high energy nuclear and particle physics experiments at laboratories around the world. Achieving the ambitious computing goals is dependent on reliable and fast connections between collaborators in geographically separate regions.

Keywords: Network, Performance, Monitoring

1 Introduction

The ambitious computing goals of experiments such as *BaBar*, *RHIC* and the *LHC* place internet connectivity in a highly critical position. More over the ability to monitor performance and identify weak points for upgrades has become pivotal to recruiting collaborators not only overseas from the location of the experiment but in locations previously considered remote.

The methodology of the Internet End-to-End Performance Monitoring (IEPM) project and long term trends in regional and trans-oceanic performance measured by the *PingER* tools has previously been described [1]. The project has grown significantly and now 593 nodes at 424 sites in 72 countries are monitored by 28 monitoring sites in 15 countries. A total of 2138 end-to-end pairs are monitored, making *PingER* probably the largest performance monitoring project in the world. Recently particular effort has been made to extend the monitoring of locations in East Europe and the former USSR and to Central and South America and the Middle East, reflecting the increasing reach of high energy nuclear and particle physics research.

2 Inter-Regional Connectivity

In this paper, the performance from SLAC to East Europe, the Middle East, Central and South America is studied in detail. Not all the nodes or even countries studied are collaborators on HENP experiments at SLAC or elsewhere, but there are universities in these countries listed in the particle physics data group book and network connectivity will be understood ready for when they do wish to collaborate.

In most cases, a national research network provides connectivity to all the universities in each country. So it is assumed the monitored nodes are typical of the performance to nodes on the same network throughout the country. This is not always a valid assumption, but care is taken when selecting a node. The node is expected to be an end node, (not an intermediate router ¹) that it is lightly loaded and always available.

¹a router treats ICMP packets sent to it very differently to transit traffic, which may make performance appear much poorer than it is in reality

Performance measured by *pingER* is used in conjunction with *traceroute* and information gathered from websites to explain the performance. Qualitative sounding measures of performance such as "very poor" and "unusable" are defined in [1].

2.1 Performance to the Former Eastern Block, Middle East and India

Performance between North America and East Europe is typically considered very poor, and even performance between West Europe and East Europe is typically classified poor.

Performance between SLAC and nodes in the Baltic region (Estonia, Latvia and Lithuania) is poor to unusable. However there is considerable difference in packet loss and round trip time to each node. This variation may be caused by differing bandwidth, utilization and especially how much headroom is available above average utilization. Also peering arrangements and the resulting route taken may be different, and is often the cause of large differences in RTT between geographically local nodes. *ping* packet loss and RTT to Latvia and Lithuania is particularly high, but to Estonia is significantly lower. This is most likely related to congested links, but also Estonia is connected to the high performance NorduNet network, which is likely to be less congested than commercial trans-Atlantic links.

The national research network of Croatia (CARnet) has an impressive OC12 (622Mbps) backbone, but only a 4Mbps connection to the USA. Most likely this is a bottle neck and consequently the cause of the high packet loss between SLAC and the monitored node on CARnet. Similarly *ping* RTT is around twice as long as the *ping* RTT to the node monitored in Moldova, which has connectivity through Moldova Telecom to New York. Belarus, Bulgaria, Turkey and Ukraine have separate trans-Atlantic connections to the New York Area with commercial providers. NorduNet has a link to Ukraine, but the connection from SLAC to the node in the Ukraine does not route via NorduNet.

Macedonia has a connection to Deutsches Telekom. The link between SLAC and the node in Macedonia has one of the lowest *ping* RTT, but the highest *ping* packet loss of the sites studied in East Europe. Albania and Romania have satellite links, with commercial providers. The *ping* RTT between SLAC and the node in Albania is almost twice as much as any other *ping* RTT to this area, perhaps because geosynchronous satellite links add a RTT delay of over 500ms. The Romanian link is 2Mbps and also has a long *ping* RTT but packet loss is amongst the lowest in the region.

Performance between North America and the Middle East is also typically considered very poor. The region is geographically opposite to the US, and much of the infrastructure has traditionally pointed towards their large neighbor, Russia.

Performance between SLAC and nodes in the Caucasus region (Armenia, Azerbaijan and Georgia) is acceptable to very poor. *ping* packet loss and RTT to Armenia and Georgia are better than to Azerbaijan. In this case Armenia and Georgia are connected by satellite to DESY and the German DFN network, whereas Azerbaijan has a commercial provider. It is not necessarily the case that academic networks are better than commercial ones, but academic and research networks are likely to peer together and the acceptable use policy (AUP) of a research network often results in lower transit traffic and traffic to non-research related sites.

Kazakhstan has a satellite connection to DFN, and the node in Kazakhstan has the lowest packet loss in the region. RTT is comparable to all nodes.

Packet loss between SLAC and India has been steadily increasing since October 1999, but the round trip time dropped suddenly at that time. It is possible the node monitored in India changed provider. Packet loss to Pakistan and Iran are lower, but the node in Pakistan was unreachable for a significant amount of time.

2.2 Performance to Central and South America

Performance between North America and South America is consistently considered "poor". Performance between West Europe and South America is also poor, however traffic is typically routed through the USA anyway.

By far the lowest packet loss of the sites studied was between SLAC and the University of Simon Bolivar (USB) in Venezuela which is routed to ESnet via a commercial provider peering in Washington. The link between SLAC and USB also has the lowest RTT, constantly around 200 ms. The RTT between SLAC and the nodes in Colombia and Brazil are around 300 ms and to Uruguay, Argentina, Chile and Peru are above 600 ms. Packet Loss to sites in Brazil is comparable and constant. The best connection between SLAC and Uruguay employs a satellite link.

SLAC to Central America is mostly very poor. In January 2000, the daily average RTT between SLAC and the National Autonomous University of Mexico (UNAM) ranged from around 150ms to nearly 300ms. In the same period it took around 650ms for the round trip from SLAC to the University of Costa Rica and as much as 950ms for the round trip from SLAC to the Center for Mathematics and Theoretical Physics in Havana, Cuba. In all three cases routing was via New York where a commercial provider was used to peer with ESnet. UNAM is particularly interesting because it is a collaborator on the D0 experiment at Fermi Lab. UNAM is also a member of the Mexican National research network (CUDI) and it is expected performance would increase dramatically if CUDI connected to Internet2 and traffic was routed through Abilene and MREN to Fermi Lab. A direct connection to ESnet's hub in Albuquerque would also significantly improve performance to Fermi Lab.

Plans for an Internet2 gigapop in Miami and peering arrangements with the South American crossing leg of the Global Crossing project may improve performance to South America. There are other plans to provide network connectivity for Astronomy research in Chile which will connect to U.S. networks in California.

3 Conclusion

The correct bandwidth and peering arrangements can provide good inter-regional network performance, but in many cases the economics are prohibitive. Universities are often the first site to gain access to the Internet and the nature of international collaborations in research is often the driving force behind improved infrastructure. As the needs of HENP and other research groups grow, so will the network infrastructure to all regions.

References

- 1 Warren Matthews, Les Cottrell and David Martin, "Internet Monitoring in the HEP Community", Poster and paper presented at CHEP'98.

Figure 1: The poster presented at CHEP2000 that this paper is based on. For your very own full size copy, send email to iepm@slac.stanford.edu

PingER: Internet End-to-end Performance Monitoring

How no collisions make better physics

Warren Matthews and Les Cottrell
Stanford Linear Accelerator Center

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

Introduction

The ambitious computing goals of High-Energy Nuclear and Particle Physics experiments such as BaBar, RHIC and the LHC place Internet connectivity in a highly critical position. Moreover the ability to monitor performance and identify weak points for upgrades has become pivotal to recruiting collaborators not only overseas from the location of the experiment but in locations previously considered remote.

Recently particular effort has been made to extend the monitoring of locations in East Europe and the former USSR and to Central and South America and the Middle East, reflecting the increasing reach of High Energy Nuclear and Particle Physics research.

The PingER Project

The PingER project has grown significantly since the methodology and some long-term trends in regional and trans-oceanic performance was described at CHEP98, now 593 nodes at 424 sites in 72 countries are monitored by 28 monitoring sites in 15 countries. A total of 2138 end-to-end pairs are monitored, making PingER probably the largest performance monitoring project in the world.

Performance to the Former Eastern Block, Middle East and India.



Performance between North America and East Europe is typically considered very poor, and even performance between West Europe and East Europe is typically classified poor.

Performance between SLAC and nodes in the Baltic region (Estonia, Latvia and Lithuania) is poor to unusable. However, there is considerable difference in packet loss and round trip time to each node. This variation may be caused by differing bandwidths, utilization and especially how much headroom is available above average utilization. Also peering arrangements and the resulting route taken may be different, and is often the cause of large differences in RTT between geographically local nodes. Ping packet loss and RTT to Latvia and Lithuania is particularly high, but to Estonia is significantly lower. This is most likely related to congested links, but also Estonia is connected to the high performance NorduNet network, which is likely to be less congested than commercial trans-Atlantic links.

The National Research Network of Croatia (CARnet) has an impressive OC12 (622Mbps) backbone, but only a 4Mbps connection to the USA. Most likely this is a bottle neck and consequently the cause of the high packet loss between SLAC and the monitored node on CARnet. Similarly, ping RTT is around twice as long as the ping RTT to the node monitored in Moldova which has connectivity through Moldova Telecom to New York. Belarus, Bulgaria, Turkey and Ukraine have separate trans-Atlantic connections to the New York area with commercial providers. NorduNet has a link to Ukraine, but the connection from SLAC to the node in the Ukraine does not route via NorduNet.

Macedonia has a connection to Deutsches telekom. The link between SLAC and the node in Macedonia has one of the lowest ping RTT, but the highest ping packet loss of the sites studied in East Europe. Albania and Romania have satellite links, with commercial providers. The ping RTT between SLAC and the node in Albania is almost twice as much as any other ping RTT to this area, perhaps because geosynchronous satellite links add a RTT delay of over 500ms. The Romanian link is 2Mbps and also has a long ping RTT but packet loss is amongst the lowest in the region.

Performance between North America and the Middle East is also typically considered very poor. The region is geographically opposite to the US and much of the infrastructure has traditionally pointed towards their large neighbor Russia. Performance between SLAC and nodes in the Caucasus region (Armenia, Azerbaijan and Georgia) is acceptable to very poor.

Ping packet loss and RTT to Armenia and Georgia is better than to Azerbaijan. In this case Armenia and Georgia are connected by satellite to DESY and the German DFN network, whereas Azerbaijan has a commercial provider. It is not necessarily the case that academic networks are better than commercial ones, but academic and research networks are likely to peer together and the acceptable use policy (AUP) of a research network often results in lower transit traffic to non-research related sites. Kazakhstan has a satellite connection to DFN, and the node in Kazakhstan has the lowest packet loss in the region. RTT is comparable to all nodes.

Packet Loss between SLAC and India has been steadily increasing since October 1999, but the round trip time dropped suddenly at that time. It is possible the node monitored in India changed provider. Packet loss to Pakistan and Iran are lower, but the node in Pakistan was unreachable for a significant amount of time.



Performance to Latin America.

Performance between North America and South America is consistently considered poor. Performance between West Europe and South America is also poor, however traffic is typically routed through the US anyway.

In January 2000, the daily average RTT between SLAC and the National Autonomous University of Mexico (UNAM) ranged from around 150ms to nearly 300ms. In the same period it took around 650ms for the round trip from SLAC to the University of Costa Rica and as much as 950ms for the for the round trip from SLAC to the Center for Mathematics and Theoretical Physics in Havana, Cuba. In all three cases, routing was via New York where a commercial provider was used to peer with ESnet. UNAM is particularly interesting because it is a collaborator on the D0 experiment at Fermi Lab. UNAM is also a member of the Mexican National research network (CUDI) and it is expected performance would increase dramatically if CUDI connected to Internet2 and traffic was routed through Abilene and MREN to Fermi Lab. A direct connection to ESnet's hub in Albuquerque would also significantly improve performance to Fermi Lab.

Plans for an Internet2 gigapop in Miami and peering arrangements with the "South American Crossing" leg of Global Crossing may improve performance to South America. There are other plans to provide network connectivity for Astronomy research in Chile which will connect to US networks in California.

By far the lowest packet loss of the sites studied was between SLAC and the University of Simon Bolivar (USB) in Venezuela which is routed to ESnet via a commercial provider peering in Washington. The link between SLAC and USB also has the lowest RTT, constantly around 200ms. The RTT between SLAC and the nodes in Colombia and Brazil are around 300ms and to Uruguay, Argentina, Chile and Peru above 600ms.

Packet loss to sites in Brazil is comparable and constant.

The best connection between SLAC and Uruguay employs a satellite link.



Conclusion

The correct bandwidth and peering arrangements can provide good Inter-regional network performance, but in many cases the economics are prohibitive. Universities are often the first site to gain access to the Internet and the nature of international collaborations in research is often the driving force behind improved infrastructure. As the needs of HENP and other research groups grow, so will the network infrastructure to all regions.

Map Legend:
■ Packet Loss
■ Round Trip Time
■ Unreachability

Monitoring Site
 Beacon Site

Countries shaded grey have a University or Institute collaborating on at least one of BaBar, Atlas, CMS, CDF, D0, RHIC or Zeus.

For Further Information, visit our website
<http://www-iepm.slac.stanford.edu>