# Latin America Case Study

### Introduction

This Case Study of Latin American Internet performance was performed by Shahryar Khan and Les Cottrell of the SLAC IEPM-PingER group in anticipation of the ICFA Digital Divide Workshop: Mexico City October 24-27, 2007.

After indicating the deployment and coverage of the measurements made by PingER, it presents PingER measurements of round trip time (minimum and average), losses, jitter, and derived results such as throughput and Mean Opinion Score. The information is presented to show how Latin America looks as a region to and from the rest of the world and compares Latin America with other regions. Then it presents the status of the Research & Education network backbone in Latin America. Following this we show the historical performance of the Internet seen from SLAC for each individual country. This is followed with the current states of the various National Research and Education Networks (NRENs) for the various countries (this information is mainly from a presentation by Alberto Santoro at the April APS meeting 2007). We then look at the routing of Internet traffic within Latin America as seen from Brazil and Bolivia. Finally we compare the derived throughputs (a reasonable measure the quality of the network connection) with various Human and Economic indices.

## PingER Coverage

The sites from which there is PingER monitoring are shown in **red** below, sites which are monitored by all monitoring hosts (these are referred to as **beacon** sites and also include monitor sites) are shown in **blue**, and other monitored remote sites are shown in **green**. From SLAC PingER monitors about 460 sites, CERN used to monitor about 132 sites, however after re-installing in December 2006) they now only monitor beacon sites (56). ICTP Trieste monitors 105 sites but only has data going back to September 2006. To the right of the PingER map we also show a map which defines how we are assigning countries to regions.



Latin America includes South America, Caribbean and Mexico. We have three monitoring hosts in Brazil (two in Sao Polo and one in Rio De Janeiro), one in Mexico (Cuidad Juarez) and one in Bolivia. The map below shows the location of the sites.



Latin America Compared to the rest of the World

Before we start to compare L. America to the rest of the world it is useful to look at a World Map of Internet Users. This shows that for most the developed world (US and Canada, W. Europe, Japan, Taiwan, S. Korea) typically 40% or more of the people have Internet connectivity while for Latin America and the Caribbean that contains about 8% of the world's population, the penetration is about 17%, i.e. typically a factor of 2-3 less.

### Minimum RTT and Packet Loss from US to Latin American Countries



The above figure shows the Min RTT and packet loss to Latin American countries from SLAC. Mexico seems to have the least RTT and acceptable loss due to its direct connectivity with the US. Argentina, Brazil and Chile are also performing reasonably well. Bolivia, Ecuador and Paraguay are catching up. However Cuba is still stuck to satellite and the performance is dreadful.

From SLAC: 1. Packet Loss, 2. Min RTT to World Regions, 3. Unreachability, 4. Jitter



The left hand figure shows the packet loss to various regions of the world as seen from N. America. Since losses are fairly distance independent no attempt has been made to normalize the data. For typical bursty Internet loss patterns, losses of >1% will result in annoying interactive VoIP and video conferencing. It is seen that the world divides up into two major super-regions: N. America, Europe, E. Asia and Oceania with losses below 0.1%, and Latin America, C. Asia, Russia, S.E. Asia, S. Asia and Africa with losses > 0.1% and as high as as a few per-cent. All countries are improving exponentially, but Africa is falling further behind most regions.

The left middle figure shows the drop in the MinimumRTTs of > 400ms result in poor real-time interactivity. The large step for S. Asia in 2003 was due to the change over from satellite to fibre. Central Asia (also Afghanistan) has hardly moved in its minimum RTT since it continues to use geostationary satellites.

The right middle graph shows the unreachability of world regions seen from the US. A host is deemed unreachable if all pings of a set fail to respond. It shows the fragility of the links and is mainly distance independent (the reasons for fragility are usually in the last mile, the end site or host). Again the developed regions US and Canada, E. Asia, and Oceania have the lowest unreachability (< 0.3%) while the other regions have unreachability from 0.7% to 2%, and again Africa is not improving, with Latin America having the fourth worst unreachability (after Africa, S. Asia and the Middle East).

The right hand graph shows the jitter or variability of world regions seen from the US. The jitter is defined as the Inter Quartile Range (IQR) of the Inter Packet Delay Variability (*IPDV*<sub>*i*</sub> = *RTT*<sub>*i*-1</sub>). The Jitter is relatively distance independent, it measures congestion, and has little impact on the Web and email. It decides the length of VoIP codec buffers and impacts streaming. For real-time multimedia (H.323) Performance Measurement and Analysis of H.323 Traffic gives for one way: jitter = 0-20ms = Good, jitter = 20-50ms = acceptable, > 50ms = poor. We measure round-trip jitter which is roughly two times the one way jitter. We see the usual division into developed versus developing regions.

### MOS for various regions

The telecommunications industry uses the Mean Opinion Score (MOS) as a voice quality metric. The values of the MOS are: 1= bad; 2=poor; 3=fair; 4=good; 5=excellent. A typical range for Voice over IP is 3.5 to 4.2 (see VoIPTroubleShooter). In reality, even a perfect connection is impacted by the compression algorithms of the codec, so the highest score most codecs can achieve is in the 4.2 to 4.4 range.

There are three factors that significantly impact call quality: latency, packet loss, and jitter. We calculate the jitter using the Inter Packet Delay Variability (IPDV), see the tutorial for further details.

Most tool-based solutions calculate what is called an "R" value and then apply a formula to convert that to an MOS score. Then the R to MOS calculation is relatively standard. The R value score is from 0 to 100, where a higher number is better. To convert latency, loss, and jitter to MOS we follow Nessoft's method. The graphs below shows the Exponentially Weighted Moving Average (using *EWMI* <sub>*i*</sub> = *alpha* \* *EWMI* <sub>*i*-1</sub> + (1 - *alpha*) \* *Obs* <sub>*i*</sub> where *alpha* = 0.7 and *EWMI* <sub>*i*</sub> = *Obs* <sub>1</sub>) for the MOS as seen from the W. Coast of America (SLAC) and Switzerland (CERN). N.B. MOS values of one are reported for heavy loss (loss > 40 %).



CERN monitors fewer remote hosts than SLAC (56 versus over 400) so the data is not as complete in the CERN case. Comparing the two graphs, it can be seen that, as expected, the MOS is better for the shorter distances (i.e. CERN to Europe is better than SLAC to Europe, SLAC to N. America is better than CERN to N. America.).

It is also seen in both graphs that the Balkans, Russia and Latin America improved dramatically in 2000-2002. Much of Latin America and Russia moved from satellite to land lines in this period. It can be seen from the above plot that VoIP ought to be successful between SLAC and the US, Europe, E. Asia, Russia, Latin America and the Mid East (all above MOS = 3.5). S. E. Asia is marginal, S. Asia people will have to be very tolerant of one another, and C. Asia and Africa are pretty much out of the question in general. In general the CERN graphs looks similar to the SLAC graph to the various regions, except that S. E. Asia is worse for CERN than SLAC as is S. Asia

The third graph shows the Mean Opinion Score (MOS) from Europe two various regions. We have five monitoring sites in Europe (one at CERN, one at ICTP, one in Germany, and two at UK. The improvement in Latin America and Russia in 2002 is the result of shift from satellite to fiber. The drop for Russia in Sep, 2006 is because we installed a new version of PingER and it started monitoring 9 hosts in Russia whereas previously it was monitoring 20 hosts. For Central Asia the number of sites went up from 3 to 15 in Sep 2006, so the latter results are a better indication of the overall performance of Central Asia. For Sub Sahara Africa the coverage improved in Sep, 2006 (increased from 8 to 39 sites). So for sub Sahara Africa the results after Sep 2006 presents a better picture. South Asia as seen from Europe is performing better than as seen from US because MOS is derived from average RTT which is distance dependent.

### TCP throughput from CERN & SLAC to World Regions



The graphs above show the derived TCP throughput using the Mathis formula. The macroscopic behavior of the TCP congestion avoidance algorithm by Mathis, Semke, Mahdavi & Ott in Computer Communication Review, 27(3), July 1997, provides a short and useful formula for the upper bound on the transfer rate:

Rate <= (MSS/RTT)\*(1 / sqrt(p))

where:

Rate: is the TCP transfer rate or throughputd

MSS: is the maximum segment size (fixed for each Internet path, typically 1460 bytes)

RTT: is the round trip time (as measured by TCP)

p: is the packet loss rate.

In the left hand two graphs the data points (average throughput per month) are fitted to exponential functions and for simplicity the trend lines only are shown. These lines enable us to see that Russia and Latin America are 6 years behind Europe, the Mid-East and SE Asia are 7 years behind, and S. Asia, C. Asia and Africa are 10 years or more behind. What is even more concerning is that Africa (in particular), South and Central Asia are not catching up.

The right hand most figure shows the yearly average derived TCP throughput normalized by the minimum RTT for the region (to reduce the proximity effects). It is seen that the throughputs are not simply exponential straigh lines, but typically change in steps as major changes are made in the routing and circuits.

### ALICE and RedCLARA

Latin America has benefited greatly by connection in 2001 to the AMPATH Internet Exchange Point in Miami that brought connectivity to North America, and the ALICE/RedCLARA project started in 2004 to connect Latin American NRENs to Europe. The launching of these projects has helped bridge the "Digital Divide" both within the region and with the rest of the world. The projects have enabled Latin American researchers to become key players in the global research community.



The Alice/RedCLARA project seen above began to provide direct connectivity with 155 Mbps, in "a ring" topology, linking the NRENs of Argentina, Brazil, Chile, Panama and Mexico, and connecting them to GÉANT at 622 Mbps throughout a connection between São Paulo, Brazil, and Madrid, Spain. RedCLARA now connects 14 countries across Central and South America, allowing 738 universities and research institutes to communicate with their European and global peers directly.

The performance of Research and Educational Networks of Argentina, Brazil, Chile, Mexico, Guatemala and Uruguay has significantly improved due to the AMPATH and ALICE/RedCLARA projects. See the graphs below for the country details, which we now summarize:

- In December 2001, Argentina's (RETINA) connected to AMPATH, so the increase in performance to Argentina can be seen from US. Around
  June 2006 the TCP throughput of Argentina increased significantly. The average RTT reduced from 600 ms to 200 ms. The minimum RTT
  reduced from 250 ms to 200 ms. The packet loss has also decreased significantly.
- From January 2001 July 2001 the loss from US to Brazil decreased from 8 % to 1 %. The Average RTT decreased from 375 ms to 200 ms. The TCP throughput improved but it has not stabilized (still lots of variations).
- In June 2001, Chile's REUNA becomes the first Latin American NREN (National Research and Educational Network) to get connected to the AMPATH POP in Miami. As a result the TCP throughput increased significantly. Minimum RTT reduced from 230 ms to 200 ms and average RTT reduced from 725 ms to 250ms. Packet loss reduced from 10 % to 1 %.
- In July 2003 Mexico's connectivity improved. Minimum RTT became stable from 100 ms to 60 ms and average RTT also became stable. Since Uruguay joined in June, 2005 its TCP throughput has improved. Average RTT reduced from about 325 ms to 225 ms. Minimum RTT reduced from about 270 ms to 220 ms The loss situation has not seen much improvement.
- In the beginning of 2006 Paraguay moved from satellite to fiber but Cuba is still connected through satellite. Average and min RTT of Cuba are
  greater than 800ms and the packet loss is around 2-4 % which is dreadful. RedClara is also planning for a connection to Cuba in the near future,
  so that it can also play role in global research initiatives.

### Status of Latin American Countries

A very interesting presentation of the status of networking in Latin America can be found in "Closing the Digital Divide in Latin America" by Alberto Santoro or UERJ, Brazil. Below we show the measured performance of connections from N. America to Latin America.

#### Argentina

In December 2001, Argentina's (RETINA) connected to AMPATH, so the increase in performance to Argentina can be seen from SLAC. Around June 2006 the TCP throughput of Argentina increased significantly. The average RTT reduced from 600 ms to 200 ms. The minimum RTT reduced from 250 ms to





#### Brazil

Around Janruary 2001 - July 2001 the loss decreased from 8 % to 1 %. The Average RTT decreased from 375 ms to 200 ms The TCP throughput improved but it hasn't stabilized (still lots of variation)



#### Chile

In June 2001, Chile's REUNA becomes the first LA NREN(National Research and Educational Network) to get connected to the Ampath POP in Miami. As a result the TCP throughput increased significantly. The minimun RTT reduced from 230 ms to 200 ms and Average RTT reduced from 725 ms to 250 ms Packet loss reduced from 10 % to 1 %.



### Cuba

Poor performance due to satellite link. Average and min RTT are about 800ms. Packet loss is around 2-4 % in 2006 (The only hope for Cuba is to connect to RedClara) RedClara plans to connect to Cuba in the near future



Mexico





#### Guatemala

Janruary 2005 packet loss becomes stabilizing >1% April 2003 minimum and average RTT improved



#### Paraguay Shift from satellite to fibre at the start 2006.



Uruguay Since Uruguay joined in June, 2005 its TCP throughput has improved. Average RTT reduced from about 325 ms to 225 ms Minimum RTT reduced from about 270 ms to 220 ms The loss situation hasn't seen much improvement



Present Situation of NRENs

Country	Organization	NREN?	National connections	External Capacity	Sites	HEP.
Argetin	RETINA	yu	130 Mops	99 Miles CLARA	38	grid
Boliria	BOLart	yu	64 - 128 3Dpr	Unknown.	18	1
Bearl	3107	316	backhone 10 Ohys	282 +1280 +155 RC Mbge	369	1
citale	REUNA	316	90 Mips	20 Mbge Clara	18	2
Colombia	RENATA	уы	11 Mips	10 Migo Clara	6	ō
Carta Rána	CROM	y10	40 Miga	10 Migs Class	34	
Ceba	RedDate	Univ. Net.	19-2 30opt- 2MRps	uskoren	20	5
Brunke	CEDIA	905	45,11,4 Mbps	If Maps Class	20	5
ElSalvator	RAICES	340	witneys-	10 Mbge Clara		
Gustemals	RAGIE	ун	900 Mbps	30 Milejo Chern	В	S I
Hoedene	RHUTA	311	unknown.	<b>animen</b>		
Mexica	0.06	yes	295 Maps	45 Maps Clara		2
Hounges	RBMAA.	yes	200 Mbps	10 Mbps - Clara	,	SE
Panens	ReCYT	yu	2 Mbps	10 Mbps -Class	11	2
Paraguay	Anindu	No	unknown.	unicoren		<b>D</b>
Peru	3445	Yes	2Mops	10 Mbps-CLara		
Ungay	RAU	yes	6432ope to 1348pe	10 Mbge Clara	-86	1
Tenerro la	REACCIUN	yes	299 Mbps & 34 Mbps	45 Mbps-ClaraAmpath	78	28

### Argentina

Connected to RedCLARA (90 Mpbs) 42 universities, 8 National Research Laboratories, 5 Government Agencies



### Bolivia

Not connected to RedCLARA, no topology map available.

### Chile



### Brazil



### Costa Rica

One ministry, 2 non governmental institutions, 3 Universities , 3 Scientific-technological Institutes or Academies



#### 45 Mbps 10 Mbps

### Cuba

22 Universities and Research centers



### Ecuador

20 Universities , schools and Research and Development Institutions



### **El Salvador**



### Guatemala

6 Universities, 2 non governmental institutions ans 1 research center (backbone) 100 Mbps



Honduras



### Mexico

4 4 Non Government Institutions, 19 Research centers, 33 Universities, 11 Institutes, 1 Faculty Independent, 7 Superior Schools 1 government institution.



### Nicaragua



### Panama

7 Universities and 3 Governmental Institutions



### Peru



## Uruguay



### Venezuela



Routing

### **Traceroute from Brazil to Latin American Countries**



### **Routes from Brazil to other Latin American Countries**

Looking at the traceroutes from Brazilian sites to other Latin American countries, it can be seen that links with other Brazilian sites are direct while to other Latin American countries they go via the US and or Europe.

Brazil to Brazil: Direct Brazil to Ecuador: Brazil -> Spain(Madrid) -> Ecuador OR Brazil -> US (Kansas) -> Ecuador Brazil to Cuba: Brazil -> US(Miami) -> (Satellite Provider) -> Cuba Brazil to Venezuela: Brazil -> Europe -> Venezuela OR Brazil -> US -> Venezuela OR Brazil -> Spain(Madrid)-> Venezuela Brazil to Costa Rica: Brazil -> US -> Costa Rica Brazil to Argentina: Brazil -> US -> Uruguay -> Argentina Brazil to Mexico: Brazil -> US(California) -> Mexico Brazil to Chile: Brazil -> US -> Uruguay -> Chile Brazil to Uruguay: Brazil -> US (California) -> Uruguay Brazil to Bolivia: Brazil -> US (Kansas) -> Italy (Rome) -> Bolivia Brazil to Paraguay: Brazil -> Spain (Madrid) -> Paraguay Brazil to Panduras: Brazil -> US(Kansas) -> Guatemala -> Honduras Brazil to Peru: Brazil -> US(California) -> Uruguay -> Peru Brazil to Colombia: Brazil -> US (Kansas) -> Guatemala Brazil to Guatemala: Brazil -> US (Kansas) -> Colombia Brazil to Guatemala: Brazil -> US -> Gutemala Brazil to El Salvador: Brazil -> US -> Gutemala Brazil to Panama: Brazil -> US -> Italy(Rome) -> Panama!tracerouter\_Brazil\_.png|thumbnail! Routing as Seen from Brazil to Latin American Countries



### Routes from Bolivia to other Latin American Countries

Looking at traceroutes from Bolivia to other Latin American countries, the routes to Bolivian sites are direct, while to other Latin American countries they go via the U.S. and/or Europe.

Bolivia to Brazil: Bolivia -> US -> Brazil

Bolivia to El Salvador: Bolivia -> Italy -> US(Miami) -> El Salvador Bolivia to Venezuela: Bolivia -> Italy -> Venezuela Bolivia to Mexico: Bolivia -> US(Kansas) -> Mexico Bolivia to Argentina: Bolivia -> Italy -> US -> Argentina Bolivia to Guatemala: Bolivia -> US -> Guatemala

Bolivia to Colombia: Bolivia -> US(Kansas) -> Colombia

### Comparisons with "Development" Indices

The size of the Internet infrastructure is a good indication of a country's progress to an information based economy. However measuring the number of users is not easy in developing countries because many people share accounts, use corporate and academic networks, or visit the rapidly growing number of cyber cafés, telecentres and business services. Furthermore, the number of users does not take into account the extent of use, from those who just write a couple of emails a week, to people who spend many hours a day on the net browsing, transacting, streaming, or downloading. New measures of Internet activity are needed to take these factors into account. Most of the Internet traffic in a developing country is international (75-90%). We measure international Internet performance which is an interesting (good?) indicator. To see how well this correlates with "development" indices we scatter plot the Mathis derived throughput from PingER against various development indices. If it correlates well then we may be able to make much quicker snap-shots of a country's/regions performance without subjective biases.

There are many "development" indices today:

- ITU Digital Access Index (2003) and the Digital Opportunity Index (2006), both 180 countries
- World Economic Forum's Network Readiness Index (2004, 2005, 2006-2007: 122 countries)
  - Harvard University Network Readiness Index (2002, 75 countries)

Typically these indices use some combination of GDP/capita, knowledge (e.g. tertiary education enrolloment), life expectancy, network (hosts/capita, access, policy, usage, afford ability, users/capita); technology (patents, royalties, exports, phones/capita)

In the scatter plots below of the derived throughputs vs. the "development" index, the US, Canada and Mexico are typically excluded since the distance from the measuring point (the US) RTT is small so the derived throughput from the Mathis formula will be artificially high. Hosts in well connected countries such as Finland, Sweden, Japan also have their losses poorly measureed since only 14,400 packets are sent to a host in a month, so measuring losses of < .01% is inaccurate.

### Comparison of TCP Throughput with Digital Access Index



In 2003, the ITU's Market, Economics and Finance Unit launched the Digital Access Index (DAI), a new index, which measures the overall ability of individuals in a country to access and use new ICTs. The DAI is built around four fundamental vectors that impact a country's ability to access ICTs: infrastructure, affordability, knowledge and quality and actual usage of ICTs. The DAI has been calculated for ~180 economies where European countries were among the highest ranked. The DAI allows countries to see how they compare to peers and their relative strengths and weaknesses. The DAI also provides a transparent and globally measurable way of tracking progress towards improving access to ICTs." from http://www.itu.int/ITU-D/ict/ai/. In 2005 the ITU launched the Digital Opportunity Index(DOI) The DoI evaluates the opportunity, infrastructure and utilization of Information and Communication Technologies (ICTs) for 180 economies worldwide. The Index monitors the mobile communications that promise to bridge the digital divide in many parts of the world, as well as more recent technologies such as broadband and mobile Internet access. The right hand map above shows the DOI coverage and values worldwide.

### Comparison of TCP Throughput with the Network Readiness Index (NRI) 2006-2007

The Network Readiness Index (NRI) comes from the "The Global Information Technology Report 2006-2007" of the World Economic Forum. NRI measures the degree of preparation of a nation or community to participate in and benefit from ICT developments. The NRI is composed of three component indexes which assess:

- environment for ICT offered by a country or community
- readiness of the community's key stakeholders (individuals, business and governments)
- usage of ICT among these stakeholders.

A map of the NRI for the 122 countries of the 2006-2007 NRI are shown in the map below.



### **UseFul Links**

- The PingER Project http://www-iepm.slac.stanford.edu/pinger/
- http://www.telegeography.com/
- http://www.internetworldstats.com/stats.htm
- http://www.apdip.net/projects/dig-rev/info/
- http://www.nationmaster.com/cat/int-internet