SLIDE NOTES for 0711 IHY-Africa_Cottrell.ppt

SLIDE 1

This talk will say why and how we measure Internet Performance for the Digital Divide. It will illustrate the overall Internet performance for the world, the trends for the last decade and how Internet performance correlates with International human and development indices. It will then go on to focus on the situation in Africa, and finish up with the geographical correlations between PingER and IHY GPS and magnetometer sites.

SLIDE 3

"it is not necessary for a large number of people in a village to each have his/her individual computer to make progress. People have done interesting things in terms of marketing materials made in their village using shared PC resources and shopping software. Access to networking here is the critical point." Paul Avery University of Florida

"networks [are an] enabling infrastructures that permit large numbers of endusers to take part in global commerce" Paul Avery University of Florida

See http://www.news.com/8301-10784_3-9782444-7.html

"Part of the dilemma in the entire Digital Divide debate is the issue of equality. Society cannot provide equality; it can provide equal opportunity (EO). EO is not driven by a computer, but by the opportunity to learn, followed by an inevitable selection process based on abilities and/or commitment of the individual." Harvey Newman, Caltech

"The paradigm shift that is coming, and this is relevant to the digital divide, is the critical mass of people who have access to high performance networks. This is no longer a privilege of the few, but is becoming an indispensible commodity. One of the opportunities this brings, in my opinion, is the transformation of software from a product to a service. And this, is important as it allows a much more flexible business model on how software is provided and how the costs are reflected to the end user.

The use of remote display (e.g Ajax), and management technologies allows the remote management or running of applications from compute farms with virtualized services. I know this sounds like going backwards, but it makes sense from an economic perspective. Large compute installations can be remote where economies of scale are possible and where electricity and cooling are cheap. Conversely an end user can conceivably have a device that is cheap to produce yet produced the same level of interaction with applications as the traditional approach of locally installed PC's. However, data management, application installation, virus and a combination of other headaches can be eliminated.

This idea is not fundamentally new, it stems from the old proposals of the "thin client" which can be cheap to produce. However, the economic value is only realized when there is a critical mass of users capable of having such a device. This in turn requires a large deployment of high performance networks.

For the developing regions therefore, it is vitally important that networking is the fundamental investment. Without this, the economic value of cheap managed end user systems which can offer software services at low cost cannot be realized."
David Foster, CERN

"Good networking to a critical mass of end users is now a fundamental component of emerging business models without which it will be impossible for the businesses to emerge that will help close the digital divide. So, in effect, it will widen. Whilst non DD countries will leverage increasingly cost effective access to technology and services, DD countries will fall further behind." David Foster, CERN

SLIDE 6

The size of the Internet infrastructure is a good indication of a country's progress towards an information-based economy. Africa's Internet infrastructure is the least developed in the world, with on average less than 1 in 100 people having access. However averages obscure the great diversity of the African continent, which is reflected in wide variations in levels of Internet-use.

But measuring the numbers 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, simply measuring 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. As a result, new measures of Internet activity are needed to take these factors into account.

One indicator that is becoming increasingly popular is to measure the amount of international Internet bandwidth used by a country - the 'size of the pipe', most often measured in Kilobits per second (Kbps), or Megabits per second (Mbps). Most of the Internet traffic in a developing country is international (75-90%), so the size of its international traffic compared to population size provides a ready indication of the extent of Internet activity in a country. In Africa some of these international links may only be as big as the circuit used by a small or medium sized business, or even a broadband home user in a developed country - about 128Kbps, or about 3-4 times standard modem dialup speeds. In most cases these are confined to very small and poor African countries, but there are many other regulatory, historic and social factors that also influence the extent of Internet use.

Credits:
Research & coordination - Mike Jensen mikej@sn.apc.org
Conceptualisation, management and refinement - Richard Fuchs rfuchs@idrc.ca & Heloise Emdon hemdon@idrc.ca
Design, DTP and Layout: Adam Martin lee@wildcoast.com
Background research and liason: Lee Martin & Rochelle Martin lee@wildcoast.com
Printing: The Printshop, Durban, South Africa.

With the above in mind PingER uses the Internet ping facility to measure the international Internet performance in terms of RTT, Loss, and derived throughput to characterize a country's performance. Coverage missing in Central Africa, Myanmar, Cambodia, Arabian Peninsula, parts of W. Africa, Guyana, Suriname, French Guiana.

SLIDE 8

Many regions of the world are connected by geostationary satellite links. The minimum RTT between two hosts connected by geostationary satellites is close to 600ms, which is about twice the minimum time needed to propagate a round trip signal over land lines. The minimum RTT therefore provides an effective signature for whether a geostationary satellite is used in the end-to-end path. Fig. shows the minimum RTTs measured from SLAC to the countries of the world in Jan 2000 and July 2006. In 2006 the main regions using satellite links are East Africa, the Central Asian Republics (currently accessed through the Silk Road project[1] with the major ground-station in Hamburg Germany), East Africa (which plans to upgrade to terrestrial lines via the EASSy project[2]), plus Bangladesh, Nepal, Nigeria, Niger and Ghana. Comparing with July 2005, it is apparent that where possible countries are moving from satellite link to land lines.

[1] "Silk Project Description", available at http:www.**silkproject**.org/**project**.htm
[2] "East Africa Submarine cable System", see http://eassy.org/

Coverage missing in Central Africa, Myanmar, Cambodia, parts of Arabian Peninsula, Somalia, parts of W. Africa, Guyana, Suriname, French Guiana.

SLIDE 14

Using the Mathis formula we can derive an estimate of the TCP throughput from our loss and RTT measurements. These estimates are only rough since the losses experienced by TCP[1] are different from those measured by ping, also PingER only sends about 14,400 pings a month between a monitoring host / remote host pair so one cannot see monthly losses of < 0.1% such as are often experienced on today's high quality paths. In addition the RTTs on high quality paths are approaching the limits of the speed of light in a fiber, so further improvement is difficult. None-the-less, especially for poorer quality paths, combining loss and RTT into a single metric is very useful. Fig. 8 shows the derived TCP throughputs measured from SLAC to the world's major regions, in some cases going back for the last 11 years. Similar plots (not shown here) are seen for the data measured from CERN in Geneva, Switzerland thus indicating that the effect is not just an anomaly associated with the measurements being from the U.S. The data for several of the developing countries only extends back for about five years and can vary greatly from month to month, so some care must be taken in interpreting the long term trends. With this caveat, it can be seen that links between the more developed regions including the U.S. and Canada, E. Asia and Europe are much better than elsewhere (3 - 10 times more throughput achievable). Regions such as Russia, S.E. Asia, S.E. Europe and Latin America appear to be 3-6 years behind. Russia and S.E. Asia are catching up slowly. However, Africa, S. Asia and C. Asia are 8-10 years behind and even worse appear to be falling further behind due to slow growth. Sites in many countries have less bandwidth than a residence in developed countries (typical residential DSL or cable bandwidths are of the order of a few hundred megabits/sec). Looking forward ten years to 2015, if the current rates of progress continue, then performance from N. America to Africa will be 1000 time worse than to Europe, to S. Asia and C. Asia will be 100 times worse than to Europe.

[1] TCP deliberately provokes loss as part of its congestion detection algorithm.

SLIDE 15

Information Technology is the engine of growth in today's Information Age.
Classifying countries by their development is difficult. One has to determine what to measure that is related to development and then measure it. There are then costs and practicality concerning: what can be measured, how useful it is, how pervasive it is, how well defined it is, how it changes over time, whether one is measuring the same thing for each country, and the cost of measuring. Various organizations such as the ITU, UNDP, CIA, World Bank etc. have come up with Indices based on measured items such as life expectancy, GDP, literacy, phone lines, Internet penetration etc. (see below). These measures take time to gather and so are often dated and only available, if at all, at widely separated intervals.
Another approach is to look at measures of the Internet that can be gathered automatically without running surveys, which in turn can be ambiguous. Today bandwidth is the life-blood of the information age

SLIDE 16

Human Development Index (HDI) composed of indicators like GDP, life expectancy, tertiary education enrollment.

SLIDE 21

There is only one intercontinental fibre link to Sub-Saharan Africa (SAT-3) which provides connections to Europe and the Far East for eight countries along the West Coast of the Continent. Except for some onward links from South Africa to its neighbours, and from Sudan to Egypt and from Senegal to Mali, the remaining 33 African countries are unconnected to the global optical backbones, and depend on the much more limited and high-cost bandwidth from satellite links. Even the few countries that have access to international fibre through SAT-3 are not seeing the benefits because it is operated as a consortium where connections are charged at monopoly prices8 by the stateowned operators which still predominate in most of Africa, and in many other developing regions.

The first large-scale international fibre project in sub-Saharan Africa, SAT-3/WASC's first segment connects Portugal to the Cape in South Africa reaching eight coastal countries along the way: Senegal, Ivory Coast, Ghana, Benin, Nigeria, Cameroon, Gabon and Angola. A second segment, in the Indian Ocean, connects South Africa to Malaysia while passing through Mauritius and India (SAFE). Jointly funded by 36 members10 and spearheaded by South African Telkom which invested US$85 million for a 13 per cent stake, the project cost about US$650 million dollars. The cable was expected to lead to much reduced international bandwidth costs, but so far this has not occurred due to the business models used to develop the project.

Landlocked African operators who have tried to purchase international fibre capacity directly from one of the consortium's international members have found themselves being charged as much to reach the SAT-3 landing point as they were charged to get from the landing station to Portugal. Sadly, the high costs have made it cheaper to send the traffic directly by satellite, even for SAT-3 shareholders such as Telecom
Namibia, which has no landing point of its own.

The only large-scale international fibre link in Africa (SAT-3/WASC/SAFE) connects eight countries on the west coast of the continent to Europe and the Far East. Operating as a cartel of monopoly state-owned telecommunication providers, prices have barely come down since it began operating in 2002.
Mike Jensen

Paul Hamilton, an independent consultant specialising in African telecommunication markets, is an associate of Balancing Act. Formerly the Telecoms Research Manager at World Markets Research Centre (WMRC), he has undertaken a range of research, analysis and consulting assignments for operators, vendors, NGOs and regulators. He continues to write for WMRC (now Global Insight) as the African telecom analyst, and other key publications.

SLIDE 23

Senegal now (Oct 02) has 1.5Gbps international bandwidth (45Mbps in '02)!

SLIDE 24

In 2002 "there is almost no intra-African Internet connectivity and the vast majority of international bandwidth lands in the G8 countries - principally North America followed by Europe (Belgium, France, Germany, Italy, Netherlands, Norway, Portugal, and the UK). High intra-regional telecom prices have limited the establishment of links between neighbouring countries to just 5 - Gambia-Senegal, and South Africa's links to Namibia, Lesotho, Swaziland and Botswana. As a result increasing amounts of intra-African traffic must be transitted through high cost cross-continental links."

Credits:
Research & coordination - Mike Jensen mikej@sn.apc.org
Conceptualisation, management and refinement - Richard Fuchs rfuchs@idrc.ca & Heloise Emdon hemdon@idrc.ca
Design, DTP and Layout: Adam Martin lee@wildcoast.com
Background research and liason: Lee Martin & Rochelle Martin lee@wildcoast.com
Printing: The Printshop, Durban, South Africa.

SLIDE 27

Institutions in African countries pay thousands of dollars a month for internet connections which a home broadband user in North America would pay US$20 a month for. Aside from the general dampening effect this has had on uptake, unaffordable bandwidth has actually excluded African scientists from gaining access to the services of global research networks which now expect their member countries to have at least 1Gbps on international connections in order to access the advanced services and petabit data sets they now provide.
Telemedicine and genetic research in particular require high bandwidths for the transfer of images, video and large data sets, other examples include high definition video, super-computing, and physics, and remote sensing data. The ATICS survey of 84 leading tertiary institutions in Africa found 850,000 students and staff with access to a total of only 100Mbps international

bandwidth (www.atics.info). By contrast, Australia's tertiary community of 250,000 share 6Gbps of international bandwidth (although even this is still insufficient to meet their needs).
Currently prices on SAT-3 are up to US$15 000 / Mbps/month, while it is estimated to cost the consortium only about US$300/Mbps/month

[From a posting by Dewayne Hendricks on Dave Farber's Iper list -- BSA]

The biggest cause [of lack of bandwidth] is the high cost of international connections to the global telecommunication backbones. This is mainly the result of the lack of international optic fibre infrastructure, which is necessary to deliver sufficient volumes of low-cost bandwidth, and the consequent dependency on much more expensive satellite bandwidth. Less than 20 of the 54 African countries have international optic fibre cable connections, and these are currently controlled by inefficient state-owned operators which charge monopoly prices while neglecting to build the national backbones needed to carry local and international traffic. As a result, circuits from Africa to the US or Europe usually cost more than US$5000 /month1, while cross-Atlantic links between North America and Europe can now be obtained for US$2.5/Mbps/month and for US$16–30/Mpbs/month on international routes in Asia2.
Mike Jensen

Only in Egypt, Libya, Mauritius is the annual cost of Internet access < 10% of the annual income
Market Research conducted by Paul Budde Communications

OECD Broadband Statistics to June 2006

This page is directly accessible at <http://www.oecd.org/sti/ict/broadband>
Over the past year, the number of broadband subscribers in the OECD increased 33% from 136 million in June 2005 to 181 million in June 2006. This growth increased broadband penetration rates in the OECD from 11.7 in June 2005 to 15.5 subscriptions per 100 inhabitants one year later. The main highlights for the first half of 2006 are:
Northern European countries have continued their advance with high broadband penetration rates. In June 2006, six countries (Denmark, the Netherlands, Iceland, Korea, Switzerland and Finland) led the OECD in broadband penetration, each with at least 25 subscribers per 100 inhabitants.

Denmark now leads the OECD with a broadband penetration rate of 29.3 subscribers per 100 inhabitants.

The strongest per-capita subscriber growth comes from Denmark, Australia, Norway, the Netherlands, Finland, Luxembourg, Sweden and the United Kingdom. Each country added more than 6 subscribers per 100 inhabitants during the past year.

Fibre to the home is becoming increasingly important for broadband access, particularly in countries with high broadband penetration. In Denmark, Danish power companies are rolling out fibre to consumers as they work to bury overhead power lines. Municipal broadband projects are also expanding in many northern European countries and throughout the OECD. Telecommunciation operators in several OECD countries have also begun or announced large fibre-to-the-premises rollouts.

Japan leads the OECD in fibre-to-the-premises (FTTP) with 6.3 million fibre subscribers in June 2006. Fibre subscribers alone in Japan outnumber total broadband subscribers in 22 of the 30 OECD countries.

The total number of ADSL subscriptions in Korea and Japan have continued to decline as more users upgrade to fibre-based connections.

DSL continues to be the leading platform in 28 OECD countries. Cable modem subscribers outnumber DSL in Canada and the United States.

The United States has the largest total number of broadband subscribers in the OECD at 57 million. US broadband subscribers now represent 36% of all broadband connections in the OECD, up from 31% in December 2005.

Canada continues to lead the G7 group of industrialized countries in broadband penetration.

The breakdown of broadband technologies in June 2006 is as follows:

- o DSL: 63%
- o Cable modem: 29%
- o Other technologies (e.g. satellite, fibre and fixed wireless) : 8%

Released: 13 October 2006

•African universities pay on average 50 times more for bandwidth than, for example, U.S. universities ($5.46/ Kbps/month vs. $0.12/Kbps/month). W. Africa $8/Kbps/month, N Africa $0.52/Kbps/month.
Promoting African Research & Education Networking (PAREN), A study sponsored by IDRC, Jan 2005

SLIDE 30

The availability of low-cost VSAT licenses is expected to have a major impact on the Internet infrastructure in countries that allow it, now that consumer and small-business oriented services have become available over the whole continent via Ku-band satellite footprints from operators such as Panamsat, Intelsat and News Skies. Satellite hubs can even be built at the new marine cable landing points to more economically provide onward satellite connectivity to regions without terrestrial telecommunication infrastructure. In addition associations of Internet service providers in Africa are planning to interlink their national Internet exchange points via fibre and satellite to reduce the amount of internal traffic that must flow off-continent.

Credits:
Research & coordination - Mike Jensen mikej@sn.apc.org
Conceptualisation, management and refinement - Richard Fuchs rfuchs@idrc.ca & Heloise Emdon hemdon@idrc.ca
Design, DTP and Layout: Adam Martin lee@wildcoast.com
Background research and liason: Lee Martin & Rochelle Martin lee@wildcoast.com
Printing: The Printshop, Durban, South Africa.

Internet growth see http://www.internetworldstats.com/stats.htm

How does one spur IT Development (from “Digital Prosperity” by Robert D. Atkinson and Andrew S. McKay of the Information Technology & Innovation Foundation, see www.innovationpolicy.org
Actively encourage Digital Innovation and Transformation in the Economic Sectors (e.g. RFIDs, GIS, smart cards, broadband, mobile commerce, e-government)
Use the tax code to spur investments (avoid taxing broadband access, IT imports, rapid depreciation of IT investments, spur inestments)
Encourage universal Digital Literacy and Digital Technology Adoption (paying bills online, home computers & linkupd
Do no Harm (don’t over-regulate RFIDs, VoIP, Internet video, encourage on-line competition to powerful off-line incumbents (e.g. banks, realtors)

“Indeed one reason why IT can help developing nations is that unlike technologies from the old economy that required massive investment in fixed plant capacity and high levels of technical skills, IT capital equipment is relatively cheap and east to use (Steinmueller 2001). For example, intruction manuals, user guides, and other assistance on IT products and services are usually

online." 'Digital Prosperity' Robert D. Atkinson & Andrew S. McKay The information Technology & Innovation Foundation, see www.innovationpolicy.org.
Steinmueller, W. Edward. "ICT's and the Possibilities for Leapfrogging by Developing Countries." Interntaional Labor Review 140.2 (2001): 194-210.

SLIDE 35

The unreachability metric (a host is deemed unreachable if all ten pings in a set are lost) like loss is useful in that it identifies the fragility of the network paths or end hosts and is also largely distance independent (as opposed to say RTT or derived throughput). Fig. illustrates the Unreachability for the regions of the world seen from SLAC in July 2006. As expected the developed regions of U.S., Canada, Europe, E. Asia and Oceania lead the way. Most regions are improving at a rate of approximately a factor of 10 in 8 years. The exception is Africa which is almost flat. The fragility of the networks in Africa and South Asia stands out, and the poor improvement seen for Africa is a cause for concern.

SLIDE 38

Using the Mathis formula we can derive an estimate of the TCP throughput from our loss and RTT measurements. These estimates are only rough since the losses experienced by TCP[1] are different from those measured by ping, also PingER only sends about 14,400 pings a month between a monitoring host / remote host pair so one cannot see monthly losses of < 0.1% such as are often experienced on today's high quality paths. In addition the RTTs on high quality paths are approaching the limits of the speed of light in a fiber, so further improvement is difficult. None-the-less, especially for poorer quality paths, combining loss and RTT into a single metric is very useful. Fig. 8 shows the derived TCP throughputs measured from SLAC to the world's major regions, in some cases going back for the last 11 years. Similar plots (not shown here) are seen for the data measured from CERN in Geneva, Switzerland thus indicating that the effect is not just an anomaly associated with the measurements being from the U.S. The data for several of the developing countries only extends back for about five years and can vary greatly from month to month, so some care must be taken in interpreting the long term trends. With this caveat, it can be seen that links between the more developed regions including the U.S. and Canada, E. Asia and Europe are much better than elsewhere (3 - 10 times more throughput achievable). Regions such as Russia, S.E. Asia, S.E. Europe and Latin America appear to be 3-6 years behind. Russia and S.E. Asia are catching up slowly. However, Africa, S. Asia and C. Asia are 8-10 years behind and even worse appear to be falling further behind due to slow growth. Sites in many countries have less bandwidth than a residence in developed countries (typical residential DSL or cable bandwidths are of the order of a few hundred megabits/sec). Looking forward ten years to 2015, if the current rates of progress continue, then performance from N. America to Africa will be 1000 time worse than to Europe, to S. Asia and C. Asia will be 100 times worse than to Europe.

[1] TCP deliberately provokes loss as part of its congestion detection algorithm.

SLIDE 40

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.com). In reality, even a perfect connection is impacted by the compression algorythms 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, jitter. Most tool-based solutions calculate what is called an "R" value and then apply a formula to convert that to an MOS score. We do the same. This 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, jitter to MOS we follow Nessoft's method. They use (in pseudo code):

```
#Take the average latency, add jitter, but double the impact to latency
#then add 10 for protocol latencies EffectiveLatency = (AverageLatency+Jitter*2+10)
#Implement a basic curve - deduct 4 for the R value at 160ms of latency
#(round trip). Anything over that gets a much more agressive deduction
```

if EffectiveLatency < 160 then R = 93.2 - (EffectiveLatency / 40)
else R = 93.2 - (EffectiveLatency - 120) / 10
#Now, let's deduct 2.5 R values per percentage of packet loss R = R - (PacketLoss * 2.5)
#Convert the R into an MOS value.(this is a known formula)
MOS = 1 + (0.035) * R + (.000007) * R * (R-60) * (100-R)

The graphs are smoothed using ewmi(i)=alpha*ewmi(i-1)+(1-alpha)*obs(i) whnere alpha =0.7 and ewmi(10=obs(1).

We are holding video conferences with Palestine and Pakistan. The latter voice is OK if one is careful (echo cancel), speak clearly and slowly. Palestine is mkuch better. Probably mainly an RTT thing.

SLIDE 43

In India we have four PingER monitoring sites: CDAC sites in Pune and Mumbai, VSNL in Mumbai and ERnet in Bangalore. In Pakistan we have five working monitoring sites: two at NIIT/NUST Rawalpindi (one on the Pakistan Educational and Research Network (PERN), the other on a Micronet DSL link), one at the National Center for Physics (NCP) at the Quaid-e-Azam university (QAU) Islamabad, one at COMSATS university Islamabad and one at PERN itself. In addition we have 3 remote (monitored) sites in Afghanistan, 3 in Bangladesh, 2 in Bhutan, 9 in India, 2 in the Maldives, 3 in Nepal, 16 in Pakistan and 6 in Sri lanka.

- end -