AFIX Technical Workshop: Session 1

High-level overview of the Internet

Contents

Baby steps: Internet prehistory........1
Childhood: Growth and slow commercialisation........2
Adolescence: The world wide web is born........3
Network Access Points and Internet Exchanges........4
All about Exchange Points........4
Internet Exchanges: Increasing the power of peering........7
The African Picture........8
Glossary........12
Sources........12

Overview

This session provides a brief look at the history and current status of the internet and its architecture, including the role of ISPs and the relationships between them. We introduce (or refresh) key concepts including those of transit and peering, as well as considering various options for peering architectures.

Baby steps: Internet prehistory

The internet as we know it began as ARPANET, a project of the Advanced Projects Research Agency (ARPA) in the United States Department of Defence. In the late 1960s ARPA researchers started working on basic concepts of networking, using dialup telephone lines to transmit information packets. The first logon to a remote computer was made in 1969 by researchers at UCLA, who logged on to a computer hosted at the nearby Stanford Research Institute.

By 1971 the ARPANET linked 23 host computers, most of them on the west and east coasts of the US. In 1972 the now renamed US Defense Advanced Research Projects Agency (DARPA) began investigating ways to link packet networks together. They wanted to develop communication protocols which would allow computers to communicate across multiple, linked packet networks. This was called the Internetting project and the system of networks which emerged was known as the "Internet."

The key moment in this process was the development of two protocols which are still the foundation of internet architecture today: Transmission Control Protocol (TCP) and Internet Protocol (IP). The crucial decision was made to implement an open architecture with the following characteristics:

- Each network should be able to work on its own, requiring no modification to participate in the Internet.
- Each network would have a 'gateway' linking it with other networks, which would have the necessary software to transmit and redirect data packets.
- This gateway would retain no information about the traffic passing through. The intention was to reduce workload and speed up traffic, but it also removed a possible means of censorship and control.

- Packets would be routed through the fastest available route. If one computer was blocked or slow, the packets would be rerouted through the network until they eventually reached their destination.
- The gateways between the networks would always be open, and they would route the traffic without discrimination.
- The operating principles of the internet would be freely available to all the networks.

These principles are the foundation of the Internet as we know it today. Without the agreement to use TCP/IP-based protocols there would be no Internet at all: just a large collection of independent networks unable to talk to each other.

The ongoing development of the TCP/IP Protocol suite is managed by the Internet Engineering Task Force (IETF).

Childhood: Growth and slow commercialisation

Throughout the 1970s and early 1980s the Internet was still a network of large mainframe computers, dominated by universities and government institutions. However, throughout this period small-scale experiments were beginning to make it available to a wider public:

- In 1972 ARPA scientists developed a program to allow the sending of person-to-person messages – the beginning of e-mail.
- In 1974 Stanford University opened up Telenet, the first openly accessible public 'packet data service' (a commercial version of ARPANET).
- In 1976 a Unix-to-Unix protocol was developed by AT&T Bell laboratories and freely distributed to all Unix computer users. Since Unix was the main operating system at universities, this opened up networking to the broader academic community.
- In 1979 the Usenet news group system, still active today, was established.
- In 1981 City University New York developed Bitnet ("Because it's Time") to link university scientists using IBM computers, regardless of discipline, in the eastern US.
- CSNet, funded by the US National Science Foundation, was established to facilitate communication for computer scientists in universities, industry and government.
- In 1982 a European version of the Unix network, Eunet, was established, linking networks in the UK, Scandinavia and the Netherlands.
- This was followed in 1984 by a European version of Bitnet, known as EARN (European Academic and Research Network).

Throughout this period ARPANET was still the backbone to the entire system; but rapidly improving computer technology and the introduction of fibre-optic cables enabled the system to expand far more quickly than its founders had ever imagined. Not only did the number of hosts linked to the system grow rapidly (it reached 1,000 in 1984), but the amazing success of email meant that the volume of traffic per host was much larger than had been planned for.

Two important developments helped the internet to survive and grow through this period:

1. The Domain Name System (DNS) was introduced in 1984. Until then there was a single integrated list of host names and IP addresses that could easily be consulted – but this rapidly became unwieldy as the number of hosts grew. The DNS replaced this static host file database with a system which was easier to configure and update. The DNS uses a tree structure which divides the internet into a hierarchical structure of domains and subdomains. Top-level domains (TLDs) include .com, .edu and .org; every country also has its own country code TLD (ccTLD), such as .uk, .za, .gh or .ke. (For a full list of African ccTLDs see www.afridns.org, or for global ccTLDs see www.iana.org). Administrators of each TLD can create as many subdomains as they wish.

   An international network of DNS servers, which communicate with each constantly, maintains information about which domain name goes with each IP address; changes can be made on any one server and are rapidly propagated through the network.

2. The second development was the decision by two national governments to encourage the use of the internet throughout the higher education system, regardless of discipline. In 1984 the British government announced JANET (Joint Academic Network); in 1985 the US National Science Foundation established NSFNet. This involved a number of decisions that were crucial for the further development of the Internet:

   - All participants had to use TCP/IP.
   - Government agencies would share the cost of establishing common infrastructures (such as trans-oceanic connections) and support the gateways.
   - NSFNet signed shared infrastructure 'no-metered-cost' agreements with other scientific networks (including ARPANET), which formed the model for all subsequent agreements.
   - The NSF encouraged international cooperation in further research.
   - Finally, NSFNet agreed to provide the backbone for the US Internet service, and provided five supercomputers to service the envisaged traffic. There was one proviso: this facility excluded "purposes not in support of research and education".

The introduction of NSFNet had dramatic consequences. In the first place, it broke the capacity bottleneck in the system. Secondly, it encouraged a surge in Internet use. Where it had taken a decade for the number of hosts to reach 1,000, by 1986 the number of hosts had reached 5,000 and a year later the figure had climbed to 28,000.

Thirdly, the exclusion of commercial users from the backbone had the (intended) consequence of encouraging the development of private Internet providers. In 1985 the Internet Activities Board organised the first workshop specifically targeting the private sector, to discuss the potential and limitations of TCP/IP protocols. This began dialogue not only between scientists and the private sector, but also among entrepreneurs, who were thus able to ensure the interoperability of their products from the beginning. In 1987 the first subscription based commercial internet company, UUNET, was founded.

Adolescence: The world wide web is born

The net continued to grow through the late 1980s and early 1990s, still largely led by governments and the academic community. But it was becoming ever more international and its capabilities were expanding:

- The number of hosts climbed to 100,000 in 1989 and 300,000 in 1990.
- Researchers at McGill University in Canada developed Archie, the world's first internet search engine, in 1990.
- In 1991 the NSF removed its restriction on private access to its backbone computers.

Also in 1991, the world wide web was released to the public for the first time. It was developed by Tim Berners-Lee and scientists at CERN, the European Centre for High Energy Physics, who initially wanted to make it easier to read research documents online. They developed the Hypertext Transfer protocol (HTTP), which enabled users to find documents by name, and the Hypertext Markup Language (HTML), which enabled users to link documents to each other.

Tim Berners-Lee explains the difference between the internet and the world wide web as follows:

> "The Internet is a network of networks. Basically it is made from computers and cables… [that can be used] to send around little "packets" of information… The Web is an abstract (imaginary) space of information. On the Net, you find computers -- on the Web, you find document, sounds, videos,.... information. On the Net, the connections are cables between computers; on the Web, connections are hypertext links."

With the advent of the web, commercial and private use of the internet began to take off, and an increasing number of private companies began to provide and sell bandwidth.

By 1994, the NSFNet was no longer the internet's primary backbone, as other competing commercial providers created their own backbones and interconnections.

Network Access Points and Internet Exchanges

The terms “Network Access Point” (NAP), “Internet Exchange Point” (IXP) and “Internet Peering Point” (IPP) are used interchangeably by many people and there are no fixed, agreed definitions of these. We will follow the broad consensus by using “IXP” to refer to exchange points broadly defined, and “NAP” to refer only to the original four NAPs which provided access to the NSFNET backbone.

These four NAPS were all located in the US, in Chicago, New Jersey, Virginia and California. Only government and those receiving government research funding (mostly universities) were allowed to access the NSFNET, but the NAPS encouraged the development of a tier of regional access providers who could aggregrate traffic from different universities. These regional access providers, essentially an intermediate tier of networks between the universities and the NSFNET, eventually began to interconnect directly. This laid the groundwork for the development of a commercial Internet independent of the NSFNET backbone, which was eventually dismantled in the mid-1990s. Today the internet is not dependent on any single backbone – one of the reasons for its robustness.

The original four NAPs were soon joined by other peering and exchange points, and today there are several hundred IXPs around the world.

All about Exchange Points

So much for the history: how do IXPs fit into the Internet as it exists today?

One of the key points to remember is that, just as there would be no Internet without the agreement to use TCP/IP-based protocols, so there would be no Internet without the exchange of traffic between networks. The Internet is not a single entity, but a large group of independent networks who agree to share traffic with each other’s customers. Without these agreements it would be impossible, for example, for customers of two different ISPs to send each email.

To put it another way, Geoff Huston, author of The ISP Survival Guide, defines the internet as:

> “A collection of interconnected component networks that share
>
> - A common addressing structure
> - A common view of routing and traffic flow and
> - A common view of a naming system”.

One of the consequences of this interconnectedness is that ISPs who compete directly with each other for market share – whether in a single region, an entire country or internationally – still have to co-operate on a technical level so that they can provide the connectivity and end-to-end service their customers expect.

This co-operation takes two major forms:

1. Transit is “a business relationship whereby one ISP provides (usually sells) access to all destinations in its routing table” [1].

In the diagram below, taken from the work of William B Norton, EastNet is a transit customer of the upstream transit provider. This relationship gives it access to every other network known to the upstream provider. Any EastNet customer can communicate with any other customer of any other ISP in the world.

1. [1] William B Norton, “Internet Service Providers and Peering”.

Figure 1: The Transit Relationship

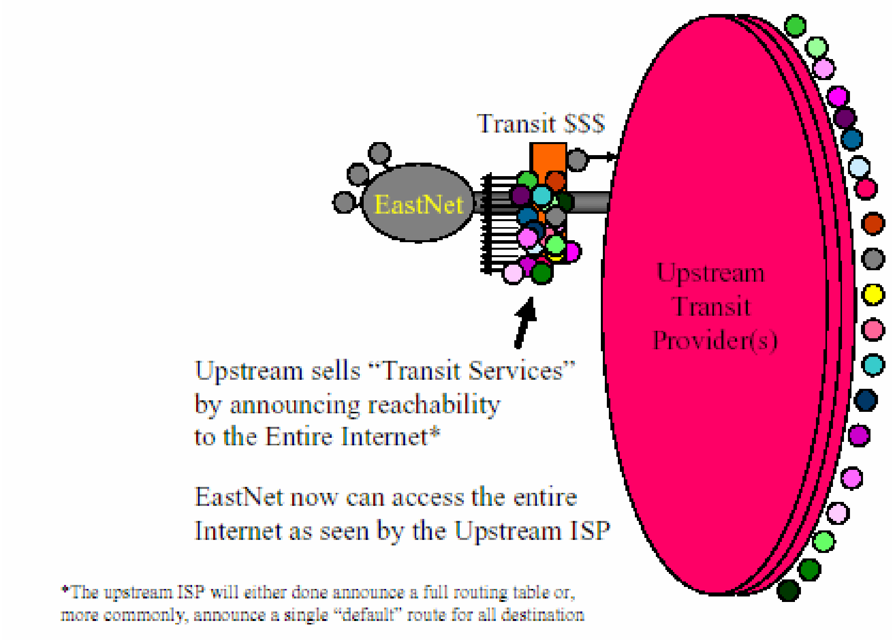

Source: William B Norton

Andrew McLaughlin of the Berkman Center for Internet & Society provides a similar definition and emphasises the commercial, for-payment nature of the relationship:

> "A *transit* agreement is … a bilateral business and technical arrangement, but one in which the transit provider agrees to carry traffic from the customer to third parties, and from third parties to the customer. The customer ISP is thus regarded as an end point for the traffic; the transit provider serves as a conduit to the global Internet. Generally, the transit provider will undertake to carry traffic not only to/from its other customers but to/from every destination on the Internet. Transit agreements typically involve a defined price for access to the entire Internet. **For virtually all developing country ISPs, the only option for connectivity to the global Internet is a transit agreement."** (emphasis added)

2. A simpler form of co-operation between ISPs is **peering** – defined as "the business relationship whereby ISPs reciprocally provide access to each other's customers". This is often, but not always, done on a barter rather than a payment basis.

In the figure below, both EastNet and WestNet have peering relationships with USNet: their customers can reach all of USNet's customers. USNet customers can reach all the customers of both its peers. However, customers of EastNet cannot reach customers of WestNet, and vice versa, unless they have a separate peering agreement.

Figure 2: The Peering Relationship

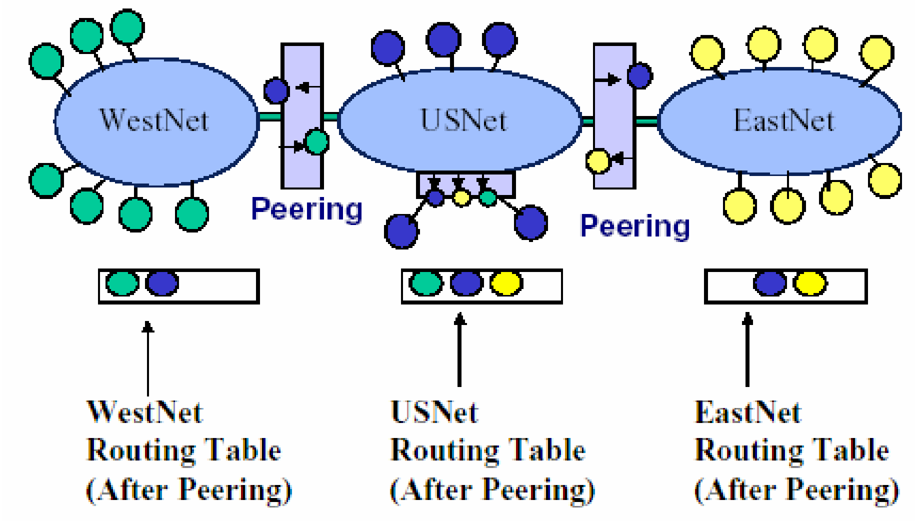

Source: William B Norton

McLaughlin's definition:

> "A *peering* agreement is a bilateral business and technical arrangement in which two connectivity providers agree to accept traffic from one another (and from one another's customers, and their customers' customers). In a peering agreement, there is no obligation for the peer to carry traffic to third parties. There are no cash payments involved – rather, it is more like barter, with each ISP trading direct connectivity to its customers in exchange for connectivity to the ISP's customers."

These are some of the pros and cons of peering:

Reasons to peer	**Reasons not to peer**
1. Lower Transit Costs Transit fees are one of the largest costs for most ISPs: a peering relationship at no (or lower) cost can make a big difference to the bottom line.	1. Traffic and Investment Asymmetry If one member of a peering relationship is a net provider of content to the other (for example if it hosts a number of web sites), the second provider could end up carrying a great deal more traffic than the first, and thus more expense.
2. Lower latency Customers of peered ISPs experience lower latency than if they had to pass through several relays of transit providers.	2. Potential transit sales An ISP may not want to peer with another which is a potential transit customer.
3. Increased usage revenues Lower latency encourages greater bandwidth use. For ISPs who bill according to usage, this means greater revenues.	3. Resource use Peering uses ISP resources such as time, router interface slots, etc – but does not necessarily bring in any.
4. More and better services Peering effectively makes more bandwidth available for the same cost – this encourages ISPs to introduce new value-	4. Competition on service quality When ISP A agrees to peer with ISP B, it effectively improves ISP B's performance and thus makes it a stronger competitor. (But ISP

added services.	A benefits at the same time!)
5. Increased customer satisfaction Because of all the above!	5. No service guarantees A successful peering relationship creates incentives for both parties to solve problems quickly – but there are no contractual guarantees to ensure this happens, as there would be in a conventional provider-customer relationship.

Most ISPs find it in their interests to peer where they can, so long as both parties see more or less equal benefit from the relationship. Geoff Huston concludes: “The structure of the Internet is one where there is a strong business pressure to create a rich mesh of interconnection at various levels”.[2]

In much of Africa, however, the structure of the internet creates few incentives to peering. On the whole, no overseas ISP would consider peering with an African ISP: most African ISPs have such a small customer base that international providers have no business incentive to enter into shared-cost peering agreements – it makes far more sense, from a commercial point of view, to take an African ISP on as a transit customer.

We will deal with the African case in more detail below; in the meantime, it is worth reviewing the basics of internet exchanges.

Internet Exchanges: Increasing the power of peering

Figure 3 below shows the worst-case scenario of an Internet without peering: in this case, traffic from one local ISP to another must pass up and down through several tiers of transit providers. In many cases, particularly in developing nations, the ultimate transit providers are located in a different country from the two local ISPs: the result is that traffic between ISPs in the same city in Africa or Asia might be routed via servers in Europe or the US. This is expensive and reduces the efficiency of the network.

2. [2] Geoff Huston, “Interconnection, Peering and Settlements”. 1999.

Figure 3: The Internet without peering

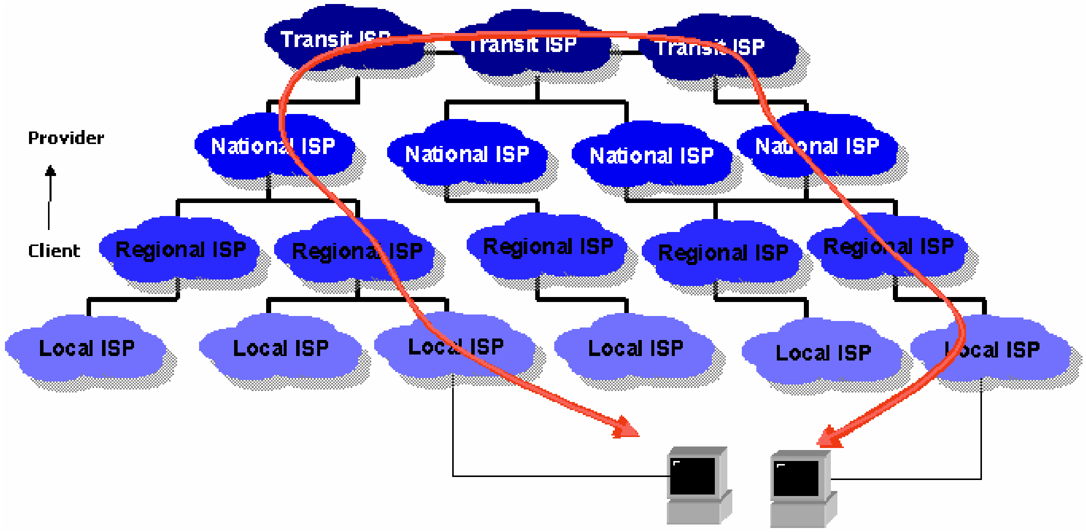

Source: Geoff Huston

The solution is for local ISPs to peer: the shorter the distance a packet must travel, the cheaper, faster and more efficient its journey.

Where there are more than two local ISPs in a market, it makes sense for all of them to get together for peering at an exchange point, rather than concluding a series of separate peering agreements.

McLaughlin provides the following definition of an Exchange Point:

> "An Internet Exchange Point is a facility operated by a single entity to facilitate the exchange of Internet traffic between three or more ISPs. (In technical terms, the IXP is a Layer 2 physical network facility.) An IXP is characterized by neutrality among all user/subscriber ISPs: often, it will be administered by a non-profit ISP association.
>
> "Typically, the IXP owns and operates the switching platforms used to interconnect the various users/subscribers. That is, the IXP consists of a shared switch fabric, where users arrange peering via bi-lateral agreements and then establish sessions between their routers to exchange routes and traffic."

The major benefits of peering, again, are:

1. Lower costs: the shorter the distance a packet must travel, the fewer resources are consumed and the lower the cost. The greatest cost savings come from eliminating the need to use costly international satellite or fibre links.

2. Improved efficiency and quality of service: Again, the shorter the distance a packet must travel, the faster it will reach its destination and the lower the chance of packet loss.

The African Picture

African ISPs have been slower than those in the US and Europe to develop exchange points, with the result that a large proportion of intra-African traffic still passes through servers overseas (see map overleaf).

The most recent continent-wide information was compiled by African internet specialist Mike Jensen in 2002. He summarised the situation at that stage as follows:

> "With the exception of some ISPs in Southern Africa, most of the international Internet circuits in Africa connect to the USA and Canada, with some to Belgium, Germany, the Netherlands, United Kingdom, Italy, and France. Of the 1.5Gbps of outgoing bandwidth, approximately

1Gbps lands in the US, 375 Mbps in Europe, 2 Mbps in Asia and only 13Mbps is intra-African….

“In Southern Africa, however, ISPs in countries with borders shared with South Africa benefit from the low tariff policies instituted by the South African telecom operator for international links to neighbouring countries. As a result South Africa acts as a hub for some of its neighbouring countries - Lesotho, Namibia, and Swaziland. Aside from this, there is only one other intra-regional Internet link between neighbouring countries - a 4Mbps connection between Gambia and Senegal, operated by the two national telecom operators, which is also used for VOIP. The main reason there are so few such intra-African links is that the high international tariffs charged by telecom operators discourage private ISPs from establishing multiple international links. As a result ISPs are forced to consolidate all of their traffic over their high cost international circuits.

This is also the reason behind the common practice amongst popular African Internet sites to be hosted on servers that are in Europe or the U.S. This is especially necessary for the many countries where ISPs operate their own independent international links without local interconnections (peering), such as in Tanzania and Nigeria, which means that traffic between the subscribers of two ISPs in the same city must travel to the US or Europe and back. This makes it more efficient to host outside-country, and is also being encouraged because web hosting costs can be very high, while there are even a number of free hosting sites in the US and Europe.

Local peering problems are now being addressed in some countries through the establishment of national Internet Exchange Points (IXPs) where all of the ISPs transfer local traffic. These have been set up by national ISP associations in Kenya, Mozambique and South Africa. Plans are at an advanced stage to establish similar facilities in Ghana and Uganda. Although local traffic is only 15-25% of total traffic, this can still result in significant savings on international bandwidth and improves performance for the user.”

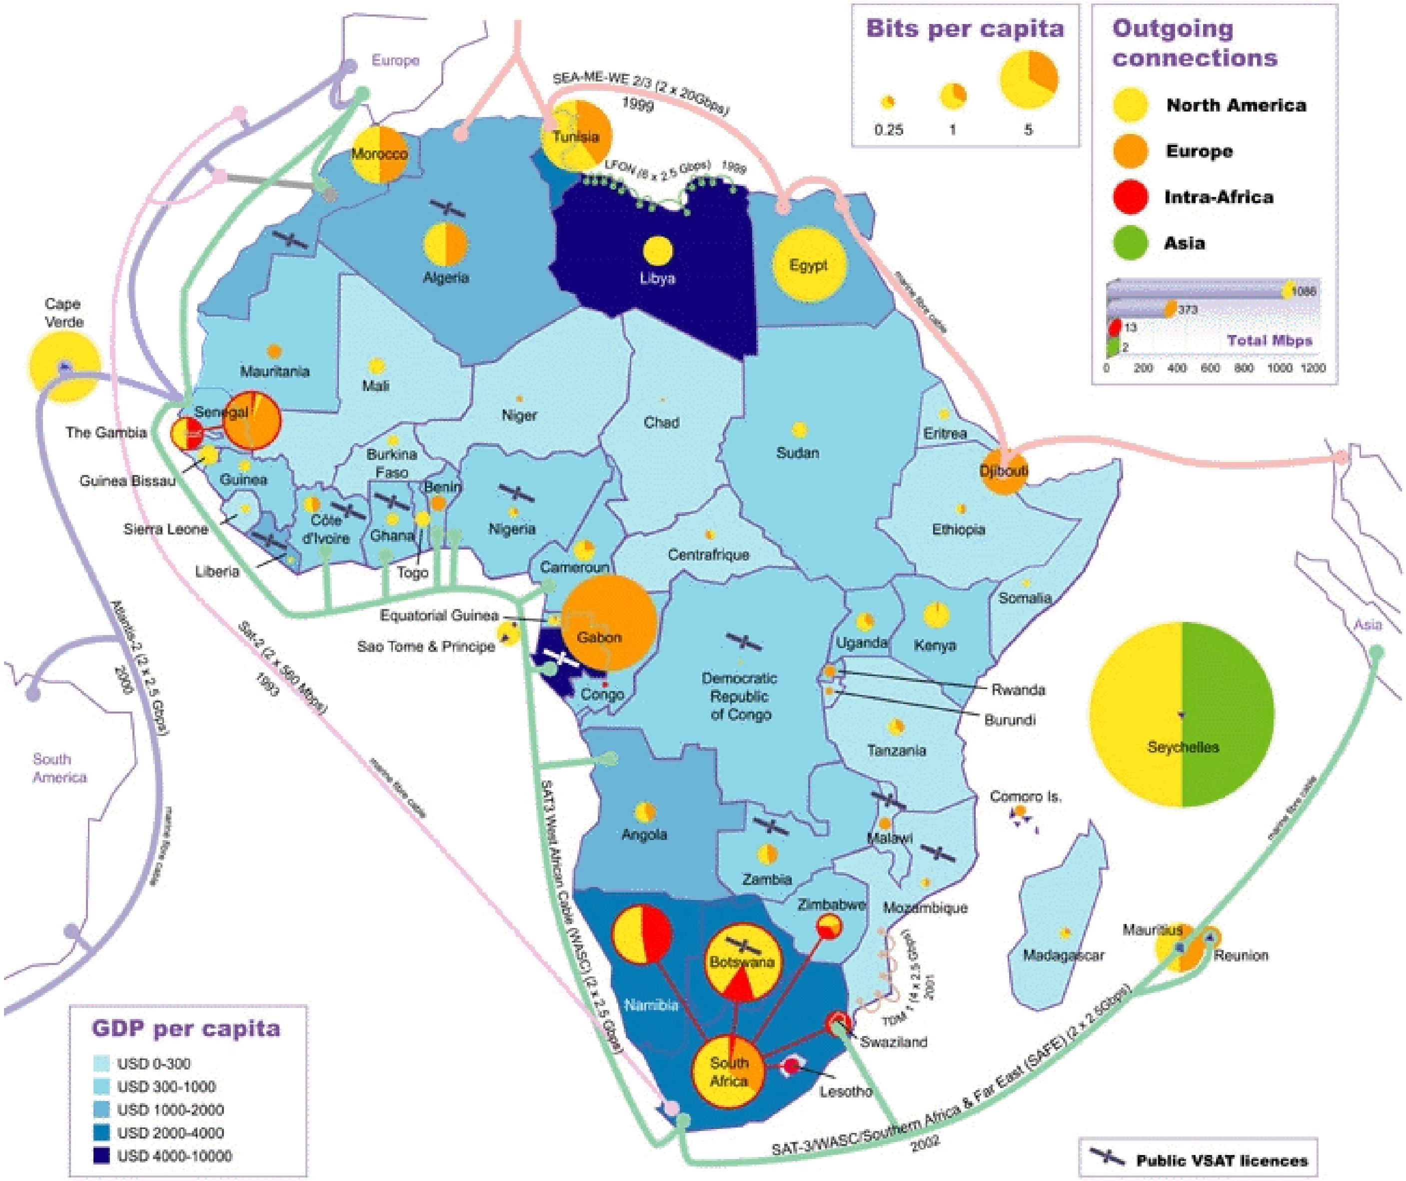

Figure 4: Internet Traffic out of Africa.

As this map demonstrates, the vast majority of outgoing Internet traffic from African countries is destined for North America and Europe. Intra-African traffic occurs only in small pockets in the far south and in Senegal.

Source: IDRC / Acacia Initiative

In October 2002 the African Internet Service Providers Association (AFRISPA) produced a "Halfway Proposition" that listed some of the problems confronting internet development in Africa, and suggested solutions.

The major issues are:

- The price of international bandwidth is one of the biggest contributors to the high cost of internet access in Africa.
- One effect of African ISPs' reliance on international transit providers is that African Internet users end up effectively subsidising users in other countries – there is a net transfer of capital from the developing to the developed world. This is especially nonsensical when one remembers that most of this traffic is destined for Africa itself!
- African ISPs rely heavily on satellite connectivity to the rest of the world, partly because of the lack of fibre optic links and partly because the few links that do exist are administered by monopoly telcos and prices are high.
- Satellite connections have much higher latency than fibre optic connections. As a result, internet users in Africa experience very slow connection speeds, for which they pay high prices. In addition, it is very difficult for internet-based businesses to develop and succeed because can't offer the same speed and quality of service as overseas competitors. For example, many African websites are hosted in the US or Europe.

These problems could be largely solved by implementing IXPs not only in every African country, but also between countries (a great deal of traffic is regional, between neighbours and trading partners).

IXPs already exist, or are in the process of being established, in the following countries in Africa:

Country	**IXP**	**Date Established**
South Africa	JINX	1999
Zimbabwe	ZIX	1999
Kenya	KIXP	2002
Mozambique	MOZ-IX	2002
Egypt	EG-IX	2002
Kinshasa / DRC	KINIX	2002
Uganda	UIXP	2003
Tanzania	TIXP	2003
Nigeria	IBIX	2003
Nigeria	Lagos IX	2003

As of 2003,then, only 10 out of 53 countries in Africa had IXPs. There are, however, some major obstacles to increasing the number of IXPs:

- The current providers of international cable and satellite links fear the introduction of domestic peering as this would reduce their revenues from international traffic.
- Many of these current providers are incumbent telco operators or others with strong links to governments and policy makers. Governments also benefit from the extra revenues earned by these telcos. As a result regulators and governments are often indifferent or even hostile to the idea of peering.

- Even without incumbent influence, African regulatory regimes are often extremely closed and make it difficult to implement new initiatives. In many cases regulators are not initially familiar with all the technical and economic aspects of the internet that provide the best rationale for establishing exchange points.
- In some cases ISPs themselves are unwilling to co-operate because they have what Andrew McLaughlin describes as "an exaggerated sense of competitiveness" and are unwilling to do anything that may benefit their competitors – even if they benefit at the same time.

AFRISPA suggests the following four-step plan to grow the internet in Africa:

1. Develop national IXs throughout Africa.
2. Develop regional carriers to interconnect these IXs.
3. Encourage massive investment in new fibre optic links between Africa and the rest of the world (donors and G8 governments may have an important role to play here).
4. Invest in infrastructure to link major cities within countries, and neighbouring countries.

Glossary

AFRISPA: African Internet Service Providers Association

ARPA[NET]: Advanced Projects Research Agency [Network]

DARPA: Defense Advanced Projects Research Agency

IETF: Internet Engineering Task Force

IXP: Internet Exchange Point

ICANN: Internet Corporation for Assigned Names and Numbers

NSFNET: National Science Foundation Network

PSTN: Public Switched Telecommunications Network

TCP/IP: Transmission Control Protocol / Internet Protocol

Sources

African Internet Service Providers Association: The Halfway Proposition: Background Paper on Reverse Subsidy of G8 Countries by African ISPs. Draft 4, October 2002, available at www.afrispa.org.

Huston, Geoff: Interconnection, Peering and Settlements. Internet Society, www.isoc.org/inet99/proceedings/1e/1e_1.htm.

Huston, Geoff: Where's the Money? Internet Interconnection and Financial Settlements. http://www.potaroo.net/ispcol/2005-01-isp.htm

Longwe, Brian: Intra-African Connectivity: Bridges to a continental backbone. Presented at Internet Week, 2003.

Norton, William B: Internet Service Providers and Peering. Draft 2.4 available at www-2.cs.cmu.edu/~srini/15-744/F04/readings/Nor01.pdf

McLaughlin, Andrew: Internet Exchange Points: Their Importance to Development of the Internet and Strategies for their Deployment – The African Example. Available from the Global Internet Policy Initiative at www.gigiproject.org/practices/ixp.pdf.

McLaughlin, Andrew and Ethan Zuckerman: Introduction to Internet Architecture and Institutions. Berkman Center for Internet and Society, cyber.law.harvard.edu/digitaldemocracy/internetarchitecture.html#Notes