International Internet Connectivity and its Impact on Australia

Final Report on an Investigation for the

Department of Communication Information Technology and the Arts,

Canberra, Australia

By

John Hibbard,

John de Ridder ,

Dr George R. Barker and

Professor Rob Frieden

May 31, 2004

LECG
ECONOMICS
FINANCE

www.lecg.com

TABLE OF CONTENTS

The views in this report are those of the Consultant and should not be held to be necessarily those of DCITA or the Australian Government.

The Consultants wish to thank those industry players who contributed to make this report possible in a short time. The research assistance of Nick Berger, Kieran Boesenberg and Dr Wen Li Cheng is also gratefully acknowledged.

Executive Summary

The international internet connectivity market has evolved over the past 5 years due to both technology and market drivers. While these developments actually reversed the trend in settlements that were a legitimate cause of concern 5 years ago, we are now at another cusp where we expect to see out payments increasing again driven by migration from dial-up to broadband. But, this is not cause for alarm as most sites that are rich in content can now be accessed at public peering points and unit costs of transmission capacity will continue to fall. The current cost of international connectivity is less than 1 cent a Megabyte (MB) and will continue to remain small.

While some dimensions of the market and technology have changed significantly, one thing that has not changed is that the internet is and will remain dominated by the US in the short to medium term, currently accounting for more than 80% of Australia's international traffic. The structure and operation of the international internet market will be dictated by the US for the next 5 years.

When the internet was commercialized in the early 1990s, non-US participants had to connect to the USA for access to content and for other international delivery. The burgeoning of the internet saw the costs for international connectivity soar to the stage where they represented up to 70% of all the costs of an Australian ISP.

International connectivity costs comprise the transmission link across the Pacific and the cost of access within the US. Under the internet charging arrangements, the non-US entity paid 100% of the transmission link costs to the US because the (peering or transit) agreements applied at the exchange point in the US. This seemed increasingly unfair as the balance of traffic shifted from 10:1 in favour of the US to 70: 30[1] and a heavy impost on non-US ISPs[2] . The price of capacity from Australia was over US$100,000 per Mbps per month around 1993.

In addition, transit access within the US had to be paid at a common price offered by the few operators who could offer satisfactory access anywhere in the US (and globally). Foreign ISPs were required to purchase transit at US$1,000 per Mbps per month or more. While US ISPs had settlement free peering amongst themselves, no overseas ISP was able to obtain peering until NTT acquired Verio.

The emergence of much higher capacity fibre optic cables has greatly reduced unit transmission costs. Since 2000, with healthy competition between SX and AJC coupled

---

[1] Communications Research Unit 1998, *Internet Congestion and Pricing: A Preliminary Study*, Department of Communications, Information Technology and the Arts, Canberra.

[2] In the Telstra/C&W litigation in the U.S. (Cable & Wireless v. Federal Communications Commission, 166 F.3d 1224 (D.C. Cir. 1999) the court explicitly refused to address the Internet charging issue asserting that the case involved only voice telephony settlements.

with discounts for greater volumes, prices for international transmission capacity have moved down. More recently, a number of the initial capacity owners in these new cables have offloaded their unneeded capacity at heavily discounted prices, in some cases as low as US$200 per Mbps per month for short term leases.

In parallel, competition for the delivery of Australian traffic into the US has increased. A major factor has been the movement of content providers and overseas entities to public peering points where they can peer with each other. As a result of using such opportunities, transit now accounts for less than 50% of an Australian ISP's traffic in the USA. This public peering induced competition for the supply of transit between the Tier 1 carriers such that the price of transit fell from US$1,000 to the current figure of around US$70 per Mbps per month.

As a consequence of these changes, the share of costs for international connectivity for Australian ISPs has fallen from around 70% to about 10%. We estimate the cost of international connectivity has fallen from A$700M in 2000 to A$150M in 2003 equating to 0.5 cents per MB downloaded from overseas. Our projected estimates indicate that the $A150M is likely to double by 2006 due to growth in volume driven by broadband, although the unit cost of transmission capacity will continue to fall. We estimate that if charging arrangements were to move to one based on traffic flows, there may be up to a A$130 million improvement in GDP

In addition, assuming a very simplistic approach, our estimates indicated that the global benefit derived by the US, in requiring other countries to pay the transmission and transit costs into the US, was US$1.3 billion in 2003. This figure is extrapolated from the Australian situation and does not reflect the circumstances applying to any other countries. It is based on the Australian share of global internet traffic with the US and extrapolated for the world. Also extrapolating it forward on the same basis, the US global benefit is projected to rise to US$2.7 billion in 2006.

1 Structure and Operation of the International Internet Market

History: The internet started in the US. The US still has the major hosts for content. While the distribution of email is quite global (like telephony), it represents less than 10% of internet traffic. The vast majority of traffic (around 80%) is file transfer and web browsing/downloads. The last element is driven by the need to access content and this is why the US remains the centre of the internet. When the internet was commercialized in the early 1990s, non-US participants had to connect to the US for access to content and for all other international delivery. With the reduced prices for transit in the US it is still more cost effective to deliver traffic to other countries via the US.

There are major content sites outside the US. These draw substantial access particularly where there are language factors. The countries with pictographic alphabets (China, Korea) seem to attract a bigger share of content access.

International connectivity costs comprise the transmission link across the Pacific and the cost of delivery within the US. Initially, these were bundled and sold by entities such as MCI and Sprint. Bundling produced inflexibility in shopping for the best value connectivity, and led to the first initiatives to obtain more competitive pricing. Subsequently, transmission and US delivery components were unbundled and could be acquired from different parties. In the late 1990s, whole circuits replaced half circuits as the method of capacity ownership so that the ease of selecting alternative US delivery partners was increased. However, the non-US entity paid 100% of the transmission link costs to the US representing a heavy impost on the non-US ISPs.

At the same time, delivery within the US remained expensive. This transit cost was held up by a common price offered by the few operators who could offer satisfactory delivery anywhere in the US. While these parties had settlement free peering amongst themselves, peering was precluded from ISPs from overseas places like Australia.

Private Peering: Private peering with Tier1 operators has been the goal for most non-US parties to reduce delivery costs. Few have achieved it as the criteria of multiple and geographically widespread North American interconnection point locations, accessed by very high capacity connections posed too high an economic cost. Tier 1 operators prefer to sell transit. Since they have a bigger customer/content base, they have negotiating leverage. Even if you get peering, the non-transitive nature of peering presents some access limitations.

Where the thresholds were reached, private peering with non-US parties only applied at the connection of a "foreigner's" US node and not in the latter's home country or at a notional mid-point. With these few exceptions, foreign ISPs such as from those Australia were required to purchase transit at US$1,000 per Mbps per month or more.

One interviewee complained of the difficulty of achieving private peering in the USA saying that the criteria were unreasonable for operators in Asia and Australia. The requirement to have several US east coast connections as well as a west coast connection seemed a tactic to preclude private peering with the Tier 1 operators. However this issue is less of a concern with the availability of public peering and lower transit rates.[3]

Public Peering: A milestone occurred with the re-emergence of public peering which had gone into disfavour because the commercialization of NSFNET was followed by severe declines in the quality of service. The establishment of new public peering points where Tier2 content providers and overseas entities could peer with each other

[3] Where an ISP has established multiple U.S. POPs it may obtain a comparative advantage over its Australian and Asia-Pacific competitors. Multiple POPs would also enhance its leverage in negotiations for private peering.

provided a very competitive alternative to transit. For the price of about US$1,500 per 100mbps port plus the cost of a router, you can peer with other operators at no transaction charge. This represents an effective price of about US$25 per month per Mbps.

As a result of public peering, transit now represents less than 50% of delivery into the US[4]. The re-emergence of public peering points as a real force has also made private peering less important. This peering alternative induced competition for transit between the Tier 1 carriers such that the price of transit fell from the US$1000 to the current figure of around US$70 per Mbps per month.

US Transit: Competition for the provision of access in the US has increased. This is not only because public peering has induced competition for the remaining transit business, but also because the group of Tier 1s has expanded. The initial group of 3 major internet connectivity providers (MCI, Sprint and BBN) colloquially called "Tier 1s" has grown to 5 and some say 8 (MCI, Sprint, AT&T, Level 3, Savvis (formerly C&W), Qwest, Global Crossing and NTTVerio). This competition has seen a reduction in transit prices from over $1000 per Mbps per month to around US$70 (we found a range of US$50 to US$100 for those who bought directly) which has reduced the desire and breakeven point for pursuing private peering.

Our survey indicated that the proportion of traffic which was delivered via transit is now only between 20% and 50%. Of the balance, most (in some cases all) traffic was delivered via public peering.

Global Transit: An Australian ISP might find task of negotiating US transit, plus private or public peering, plus establishing a node in the US plus acquiring transpacific capacity too daunting. An alternative is to purchase global delivery at the wholesale level from an operator in Australia. Currently such transit is available at a price of US$200 - US$300 per Mbps per month which includes both transpacific capacity and the US (and global) delivery charges.

Market Power: We have considered whether international internet connection agreements may be influenced by the exercise of market power, with specific reference to particular competition bottlenecks, and whether there has been market intervention internationally.

The relevant market to assess market power can be considered as three distinct submarkets based on separate geographical aspects. These are the:

- Australian domestic market for wholesale delivered transit

[4] But an Australian content provider like, say, the ABC would not contemplate establishing a site in the US. First, Australian ISPs charge for downloads, not uploads (i.e. the ABC does not face hefty port charges Google faced as a transit customer). Second, it is not at a competitive disadvantage with, say, CNN based in the US as Americans do not pay for downloading content from the ABC under the current arrangements.

- International market for long haul carriage services
- Foreign (E.g. US) domestic markets for wholesale delivered transit.

The first submarket is the subject of a current investigation by the Australian Competition and Consumer Commission[5] and is beyond the scope of this inquiry. On the second, we find below that the market for international transmission capacity is competitive. And, as discussed above, there is now more competition in the market for transit; partly as a result of increased public peering.

2 Australia's International Internet Connectivity

Market: Our findings are that for Australia, the proportion of downloads from international sources varies between 47% and 90% with a median value of 70%. One party indicated that while for consumers, it was 70%+, for corporate customers it was as low as 30%.

US Traffic: The proportion of international traffic which is to/from the US varies between 80% and 95% of total. But, all Australian operators interviewed send some traffic via the US for other countries, and some send all their international traffic via the US.

Most ISPs do not analyse packet headers continuously to determine the ultimate destination. It is not worth it. Samples are sometimes used to get an indicative picture. The incremental cost of using the high capacity links to the US, coupled with the lower transit fees paid to US ISPs for global connectivity is in most cases less than the cost of establishing a thin route to, say, Hong Kong or Singapore[6]. Notwithstanding this, the growth rate of capacity to Asia is higher than to the US.

Few of those interviewed measure traffic volumes on a regular basis. The resources required to do so are considerable and the analysis is time consuming. In general, traffic volumes are derived from capacity measurements. All monitored their capacity usage and used tools such as the 95th percentile of 5 minute periods[7] as the prime method for assessing fill. The average fill is around 50% (the range reported figures varied from 43% to 53%) although this will vary from operator to operator depending on customer base and QOS objectives. From this we get the rule of thumb that to carry 1 Terabyte per month requires 6Mbps of capacity.

[5] See http://www.accc.gov.au/content/index.phtml/itemId/269309

[6] Direct routes to Asia result from international affiliations (e.g. SingTel and Reach) or cable routes (i.e. AJC)

[7] It is interesting to note that Korea has become a major centre for content partly because it charges on the basis of total traffic rather than peak traffic. This gives content providers greater cost certainty.

Both the cumulative sum of advised traffic and the result inferred from capacity using the rule of thumb suggest that the traffic from the US is about 2,000 TB/month (2,700 TB per month from all sources). The growth of capacity to the US is shown in the diagram below.

The flow to the US was significantly less. The respondents indicated traffic balances between 60/40 and 80/20 with a median value of 70/30. Thus the level of traffic to the US was 42% of the traffic level from the US or approximately 800TB per month. Most significantly 70/30 was also the ratio in 1997 and it has been surprisingly stable during the intervening period notwithstanding such developments as mirror sites, caching, and the general increase in content housed in Australia.

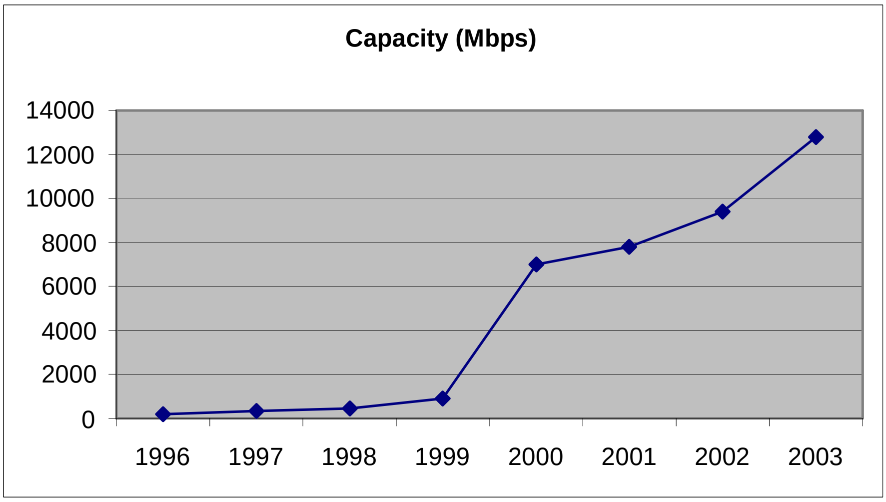

Transmission Capacity: The emergence of much higher capacity fibre optic cables using DWDM heralded a major step change in the cost structure. Firstly Jasuraus, but more significantly Southern Cross (SX) and then Australia Japan cable (AJC) and SEA-ME-WE 3 (SMW3) have provided an abundance of capacity [see Table 1].

Table 1 Potential International Cable Capacity—Australia

International cables out of Australia	Year began service	Channel capacity	Mbps	Length (km)
Tasman2	1992	15,120	2 x 560Mbps	Sydney to Auckland—2,200
PacRimEast	1993	7,560	2 x 560Mbps	7,800

PacRimWest	1995	7,560	2 x 560Mbps	Sydney to Guam—7,000
Jasuraus	1997	60,480	5Gbps	Perth to Indonesia—2,800
Southern Cross	2000	120,000	Up to 480Gps (protected)	Sydney to US via NZ
SeaMeWe 3	1999	120,000	Up to 40Gbps	48,000
Australia Japan Cable	2001		Up to 320 Gbps (protected)	Sydney to Guam via Japan

Source: The authors and the Communications Research Unit 1998, *Internet Congestion and Pricing: A Preliminary Study*, Department of Communications, Information Technology and the

The total activated bandwidth out of Australia for internet is estimated to be 16.3Gbps at December 2003 comprising 12.8 Gbps to the US and 3.5Gbps to NZ (predominantly for uploads to NZ) and Asia. There is no significant direct capacity to Europe for internet services.

Transmission Prices: Traditionally, cable capacity has been sold as IRUs which involve payment upfront (although time scheduled payment schemes are available). Initially prices for an STM1 out of Australia cost US$38M up front, but this rapidly reduced under price protection clauses and competition between SX and AJC to a figure around US$10M for a fully protected (i.e. an alternative fibre pair with automatic switchover) STM1 circuit. Further competition coupled with increased volumes over diverse routes has seen interest in unprotected capacity. This has resulted in the following typical IRU prices for cable capacity.

TABLE 2 TYPICAL CAPACITY (UNPROTECTED) PRICES

	IRU US$m	Lease Equivalent US$ per Mbps per month
STM1	3.8	700
STM4	14	650
STM16	45	500
STM64	100	300

For ease of comparison, these IRU prices have been converted into approximate lease equivalents, however such prices are clearly dependent on market and supplier influences.

This reduction in prices has been assisted by the emergence of additional sources of supply. Traditionally the only source has been the cable company. Resale if permitted was a small market as demand exceeded supply. However the new forms of financing associated with private cables often required substantial pre-sales tempting sponsor carriers to over-commit in the short term. The dot-com bubble produced many ISPs with unrealistic aspirations who had subscribed for capacity beyond their means. As they went into receivership, administrators put their capacity onto the market.

Recently, unneeded capacity has been leased at heavily discounted prices, apparently t lower than US$200 to cover the O&M expense (approx US$70). Given that the capacity cost as an IRU is more than US$1,000 and possibly as high as US$4,000, this leased capacity is extraordinarily cheap. However the availability of this leased capacity is very limited and the word on the market is that the short term leases will not be renewed. If so, the spot market price will rise to close to the numbers in Table 2 above.

Some capacity holders own IRUs with/without price protection, others own leases of different durations and others own a mix of some or all of these. Some buy STM1s, others larger blocks with implicit volume discounts. With IRUs, different accounting policies are adopted. Hence establishing a typical price each year is difficult and a median value of the gathered information has been necessary to set a figure for each year. The adopted rates are in the following graph.

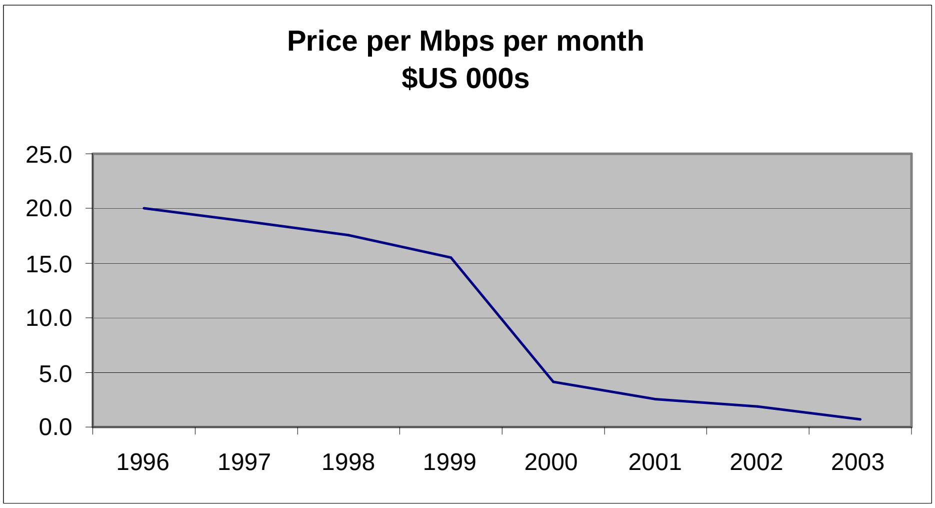

Cost Structures: At the outset of the commercial internet, international connectivity costs dominated the economics of Australian ISPs. The international component represented 50-70% of the costs in the early 1990s.

With the reduction in both transit costs and capacity costs, together with public peering, the international connectivity component of costs has fallen significantly. On the other hand, the growth of the internet and the rollout of broadband have seen other cost elements become increasingly significant in the array of costs. Information gathered indicates that international connectivity now represents about 5 – 15% of total wholesale costs, with one response giving a figure of 2%.

The value of international internet contracts to Australia reflects a combination of the growth in demand (and hence capacity) and the change in prices. The aggregate value of these contracts peaked in 2000 when there was a surge in capacity and supply was no longer constrained. Since then prices have fallen faster than capacity has increased.

3 Impact on Australian Economy and Consumers.

The macroeconomic impact of current charging arrangements is small.

Impacts: Our analysis suggest that the Australian Internet industry imported (i.e. paid) around A$150 million to achieve connectivity with the United States in 2003. These charges comprise A$144 million in transmission fees; A$4 million in transit fees; and A$2 million in peering fees.

Consumers: The A$150m cost of international connectivity in 2003 translates into a cost of 0.5 cents per MB[8] compared with marginal retail prices of 15 cents per MB.

Economy: LECG estimate that if charging arrangements were to move to one based on traffic flows, there may be up to a A$130 million improvement in GDP. This is comprised of two effects:

- a A$76 million increase in consumption as a result of lower retail prices for internet usage to households and businesses; and
- a A$53 million increase in net exports due to a fall in import payments, and a rise in export payments associated with international internet connection arrangements.

4 Future Trends

The introduction of high capacity fibre optic cables has seen the capacity ex Australia change from being supply constrained to one of abundance. There is now potential capacity in excess of 1 Tbps if all wavelengths are put in service. Currently about 200Gbbps (or 20%) of fibre is lit. Of this about 20Gbps (2% of potential capacity) is activated.

Over the period from 1997 to 2003, annual traffic growth was about 70% pa. However this demand was driven by dial-up internet. The migration to broadband access has

[8] Based on interviews and capacity data, we estimate about 2,700 TB per month is downloaded from overseas currently. (1TB is 1 million MB).

stimulated higher growth. No respondent indicated a growth of less than 100% currently with a number suggesting significantly higher numbers.

Two interviewees reported that on the basis of their measurements, the download volume of international traffic was 10 times higher for a broadband customer than for a dial-up customer. ABS statistics[9] on combined international and domestic downloads suggest corresponding multipliers of 12 and 18 for residential and business customers.

As a result of migration to DSL, volume increases will now overtake price decreases such that net out-payments on international charging arrangements will start to increase again after declining for three years. In addition, the availability of cheap short term leases will cease and this will temper the decline in unit capacity rates.

The projections below stop at 2006 due to the uncertainties about the rollout of FTTH and other factors. No new trans-oceanic submarine cables ex Australia are expected before 2008.

It is projected that traffic to non-US locations will increase faster than traffic with the US but it still seems unlikely that less than 75% of traffic will involve the US before 2007. This may eventually change with China but, given its relatively low base, it will be many years before it could, if ever, dominate Australia's internet traffic volumes.

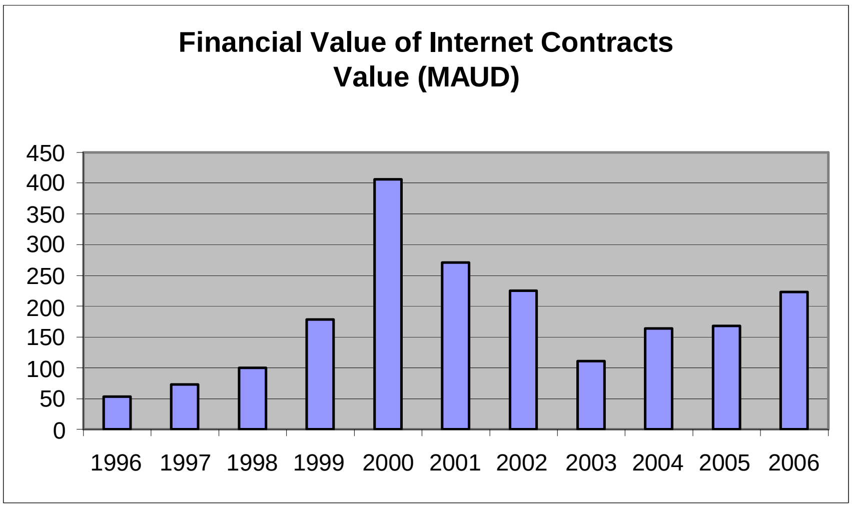

Charging: None of the respondents believed that traffic balance alone should be used as the metric for determining the relative contributions to the cost of the connectivity

[9] Australian Bureau of Statistics, *Internet Activity*, 8153.0, September 2003

between two countries. All acknowledged that it was subject to abuse or misinterpretation, and that the resources involved were considerable. Precise measurement was difficult and results subject to debate. However the majority did indicate that with a simplified measurement approach (eg a sample per month) that traffic balance did represent one important input to the assignment of benefit. In this context, the innovative work being done at ETRI looks interesting but is still at an early stage of development.

There was general acknowledgement that the determination of benefit has to be more broadly based than traffic balance. Benefit depends not only on traffic but also on customer base, quality of onward connectivity, etc. It needs to be something that can be applied multi-laterally. Examination of headers is resource heavy and even then is not necessarily accurate in attributing benefit.

5 Glossary

ABC	Australian Broadcasting Corporation
ABS	Australian Bureau of Statistics
ADSL	Asymmetric digital subscriber line
AJC	Australia-Japan Cable
DCITA	Dept of Communications, Information Technology and the Arts
DSL	Digital subscriber loop (of which ADSL is the main type)
DWDM	Dense wave division multiplexing
ETRI	Electronics and Telecommunications Research Institute (Korea)
FTTH	Fibre to the home
Gbps	Gigabits per second (1,000 Mbps)
GDP	Gross Domestic Product (measures size of the economy)
IRU	Indefeasible right of use (long term capacity)
ISP	Internet service provider
Kbps	Kilobits (or thousand bits) per second
MAUD	Millions of Australian dollars
Mbps	Megabits (or million bits) per second (pipe capacity)
MB	Megabyte (measure of volume delivered)
NSFNET	The original academic US internet backbone network
O&M	Operations and maintenance
Peer	A party that has agreed to exchange traffic settlement-free
Peering Criteria	Sometimes published terms of peering
POP	Point of presence
Private Peering	Bi-lateral traffic exchange on settlement free basis (SKA)
Public Peering	Bi-lateral and/or multi-lateral traffic exchange settlement free

QOS	Quality of service
SKA	Sender Keeps All (no charging)
STM1	155Mbps
STM4	4 x STM1
STM64	64 STM1s (often called 10Gbps since so close to this number)
SX	Southern Cross (a trans-oceanic cable network)
Tbps	One thousand Gbps (measure of pipe capacity)
Terabyte	One TB is 1M MB (measure of volume delivered)
Tier 1	Tier 1 operators do not buy transit in the US
Transit	Service provided by a Peer to a Customer ISP
US	United States of America
WDM	Wave division multiplexing