International Charging for Internet Services, Module 1

Chapter Three

DOMESTIC AND INTERNATIONAL CHARGING ISSUES

3.1 Executive Summary

Internet Service Providers ("ISPs") interconnect facilities using different terms and conditions. A single model or description does not suffice, and preexisting telecommunication models require significant adjustment and qualification to provide much insight on the current situation. Complexity in interconnection and charging stems from the ability of ISPs to provide consumers with "seamless" access to most of the individual networks that comprise what we call the Internet. The packet-switched nature of the Internet coupled with switching and routing protocols, provides robust and diverse network access. At a macro-level the Internet constitutes a "network of networks," because individual ISP facilities so easily blend together.

Technoligical ease in switching and routing Internet traffic contrasts with the difficulty ISPs currently experience when they attempt to charge other ISPs for using their network. ISPs increasingly differ in terms of the reach and capacity of their networks, as well as the number of direct subscribers. Despite increasing differentiation among ISPs, the "network of networks" orientation, coupled with a general desire to promote ubiquitous access, has challenged initiatives to revamp pricing policies. During a period of early promotion, government sponsors and private ventures alike have tended to emphasize network connectivity at the expense of profit maximization. As the Internet matures government financial sponsorship and other incubation efforts have declined, or ended.

Likewise ISPs have begun to rethink their promotional pricing efforts. This examination focuses on what an ISP charges customers for access to the Internet and for creating a World Wide Web site. It also reconsiders what charge, if any, an ISP should charge other ISPs who routinely seek to use the ISPs' facilities for access to its subscribers and for transit services taking messages and data requests onward to the final destination.

Chapter Three addresses the changes in ISP philosophies, strategies and pricing rationales relating to facilities interconnection. In a broad sense we examine a shift in philosophy from one promoting access and use of the Internet, without a finely calibrated pricing mechanism, to one concerned about making the Internet more efficient, reliable and of higher quality. To achieve these goals, ISPs have to revise their pricing policies so that they can recoup existing and additional investments. Their pricing strategy increasingly focuses on determining who has triggered an expense or the need for an infrastructure investment and how can the ISP properly and fairly recover such expenses and costs

3.2 Overview

As discussed in the previous section, the Internet provides global service over diverse routes thanks to the interconnection and integration of numerous discrete networks. The Internet accrues what economists call positive networking externalities, because its value and utility to users grow as a function of the number of points served. The proliferation of Web sites enhances consumer welfare in terms of economic theory and personal experience. Internet users benefit when they have access to more content sources, points of communications and potential e-mail addresees even if we visit a tiny fraction of what is available.

Absent network congestion and other impediments to accommodating demand growth and Internet proliferation, the value of the World Wide Web rises with the number of subscribers linked via the Internet. Heretofore, Internet service providers have conferred expanding benefits to consumers without tying such enhancements to a schedule for raising rates. On the other hand, many telecommunications carriers historically have used expanding value accruing from broadened service opportunities as the justification for increased rates. For example, a domestic regulatory authority might require flat-rated, usage insensitive telephone service rates, but accept the carrier=s view that Avalue of service@ ratemaking justifies an automatic rate increase once the carrier reaches specific benchmarks, e.g., a specific number of subscribers, or the substitution of main lines for shared (Aparty@) lines.

Some Internet and telecommunication subscribers have grown to expect fixed, AAll You Can Eat@ (AAYCE@) usage-insensitive rates. Even though consumers pay on a usage-sensitive, metered basis for some services, e.g., local exchange telephone service, they nevertheless may expect to pay a monthly flate rate for Internet access. Such users assume that the ISP incurs no additional expense in furnishing another minute of access. Internet subscribers may go one step further in their assumptions that increasing usage imposes no additional costs by expecting their initial monthly subscription rate for Internet service to remain the same for the foreseeable future. This expectation becomes increasingly unrealistic as most ISPs have made substantial additional investments to support access by a larger number of subscribers to an increasingly diversified array of Internet services, some of which require comparatively larger, "broadband" capacity. For example, most Internet consumers previously used the Internet primarily for email, a narrowband application that need not be transmitted and received instantaneously. Now many users have requirements that require immediate "streaming" of video and other broadband applications. In both the Internet and telecommunications worlds service providers incur real costs when adding long haul trunk lines to new servers and routers, as well as short haul loops to end users. [1] But in the Internet arena how costs are recovered might not ensure that the persons and companies responsible for such costs fully and fairly compensate the service provider.

ISPs may unfairly underwrite infrastructure development if they carry more traffic generated by other ISPs than they originate for carriage by other ISPs. Similarly an ISP also might disproportionately bear network operation costs if it accommodates growing demand outside its network and subscriber base to transit or terminate such traffic on its network. A greater potential exists among ISPs for asymmetry between cost causer and cost reimburser, because of four major factors:

$ Abest efforts@ traffic routing achieved through TCP/IP networking provides the Aseamlessness@ needed for fast exchange of traffic among different ISPs and to provide the kind of routing diversity and reach envisioned for a Global Information Infrastructure;

$ most ISPs initially agreed to interconnect disparate networks without regard to traffic flows and without a financial settlement;

$ ISPs historically have not metered traffic flows and accordingly have not erected a pricing mechanism based on traffic; and

$ Taxpayers in some nations, including the United States, incurred a large percentage of initial development costs for the Internet.

Because of these four factors Internet pricing arrangements historically have lacked the precision reached in telecommunication interconnection and access arrangements. While both systems may fail to

---

1 "Considering that the end client is paying the local ISP for comprehensive Internet connectivity, when a clients's packet is passed from one ISP to another at an interconnection point, where is the revenue for the packet? Is the [appropriate] revenue model one where the packet sender pays or one the packet receiver pays?" Geoff Huston, Interconnection, Peering and Settlements, Sec. 4, Settlement Models for the Internet, available at http://www.telstra.net/gih/peerdocs/peer.html (viewed March 18, 1999). The author answers this question: "the retail base of the Internet is not an end-to-end tariff base. The sender of traffic does not fund the first hop ISP for the total costs of carriage through the Internet to the traffic's destination, nor does the ultimate receiver pay the last hope ISP for these costs. The ISP retail pricing structure reflects an implict [rough justice] division of cost between the two parties, and there is no consequent structural requirement for inter-provider financial balancing between the originating ISP and the terminating ISP." Id. At sec. 6.

achieve maximum efficiency, i.e., Apareto optimality,@ the Internet pricing model has farther to go. Put another way, the current Internet pricing model creates the greater potential for financial Afree riders,@ and unintended cross-subsidy beneficiaries that do not pay their fair share, as well as involuntary underwriters which pay more than their fair share. Cost allocation and recovery is more complex for the Internet relative to conventional telecommunication models, because of preexisting arrangements for interconnection, and the seamless integration of all networks into a universal Internet.

Beneficiaries of the status quo will balk at a new and fairer deal that reduces or eliminates their free rider or subsidy opportunities. Advocates for change may have difficulty identifying ISP free riders from an ISP that has fairly borne its financial burdens. Simply refering to an assemetry in traffic stream may not sufficiently make the case for a transfer payment or access charge. Other factors may matter including the value to consumers of what content an ISP and its customers provide to others and what kinds of access to other ISPs and other locations the ISP offers. For example, a U.S.-based ISP may generate substantially more outbound traffic as a result of external file transfer requests and World Wide Web visits to sites located on the ISP's network, or accessible via the U.S. ISP. Relying solely on traffic streams to determine who should pay for access might support the argument that the U.S. ISP should compensate other ISPs who terminate or transit the traffic "originated" by the U.S. ISP. But a more complete examination of the complete transaction would show that the U.S. ISP has responded to an initial request for service originated via another ISP. Under traditional telecommunication access pricing, the point of traffic origination is dispositive for purposes of determining which carrier pays for terminating transiting traffic, even if, in the case of call-back services, a carrier deliberately "re-originates" the traffic stream to exploit an accounting rate arbitrage opportunity. [2] In an Internet-mediated transaction, value of service pricing and other factors might warrant deviation from the telecommunication model.

As could be expected, reasonable people can disagree as to what constitutes Afair share.@ But in the case of the Internet as currently constituted, ISPs have lacked the resources to track traffic and cost causation and until now have had little incentive to devise metering and tracking systems. [3]

---

2 For more information on call-back see Rob Frieden, "The Impact of Call-Back and Arbitrage on the Accounting Rate Regime," 21 Telecommunications Policy, No. 9/10, 819-827 (1997); see also Rob Frieden, "Falling Through the Cracks: International Accounting Rate Reform at the ITU and WTO." 22 Telecommunications Policy No. 11 (December 1998).

3 "[N]ot so long ago . . . many of today's 'problems' provided a viable solution within the limited parameters of the non-commercial Internet--with the added attraction of being easy to implement." Bernadette Jew and Rob Nichols, Internet Connectivity: Open Competition in the Face of Commercial Expansion," p. 6, paper presented at the Pacific Telecommunications Conference, Honolulu, Hawaii (January, 1999) available at http://www.gtlaw.com.au/gt/bin/frameup.cgi/gt/pubs/opencompetition.html (viewed March 18, 1999).

3.3 The Differences Between Telecommunication and Internet Pricing Systems

The proliferation of ISPs, varying in terms of capacity, subscriber numbers, and geographical reach has triggered greater attention to pricing issues. Such attention has led to a clustering of similarly situated ISPs for purposes of determining who is entitled to hand off traffic at no charge and who must pay for the privilege. The opportunity for cost-free interconnection--what one might call true peering--has become much more limited and now is available primarily to large ISPs who provide Internet backbone facilities having the widest bandwidth. ISPs with less bandwidth and geographical reach now typically do not qualify for true peering , and instead must pay for the privilege primarily by leasing lines from one of the major backbone operators.

The development of a dichotomy between ISPs qualifying for true peering and lesser ISPs has resulted in a more hierarchical industrial structure for the Internet. ISPs that must pay for access are not peers of the major, so-called Tier-1 carriers (UUNet, a subsidiary of MCI Worldcom, Cable & Wireless, having acquired MCI=s pre-Worldcom merger Internet business, BBN, a GTE subsidiary and Sprint) that control between 85% and 95% of the total backbone traffic in the United States [4] They are customers and in effect resellers of Tier-1 ISP facilities and services. It should come as no surprise that most Tier-1 ISPs are subsidiaries or affiliated of the major facilities-based telecommunication carriers. If Tier-1 ISP management receives pricing guidance from their telecommunication carrier parents, the likelihood exists for the migration to charging arrangements analogous to the financial settlements currently used by telecommunication carriers. [5]

However, until such time as conditions favor financial settlements and cost-based pricing systems the Internet largely differs from the conventional telecommunication model. Bear in mind that despite momentous changes in a very short period of time, the Internet originated in a largely noncommercial, or promotional environment with governments serving as both incubator and user. During this time, ISPs seemed more keen on building networks and a customer base than on maximizing profits. In a test, development and promotional environment with mostly government operators and substantial government subsidies the pricing system lacked the sophistication and precision that more commercial and established systems would erect. In this environment operators emphasized

---

4 Kenneth Neil Cukier, APeering and Fearing: ISP Interconnection and Regulatory Issues, n.17-18 and accompanying text; available at http://www.ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html (viewed February 7, 1999)[hereinafter cited as Peering and Fearing].

3 We also note that while ISPs appear to move toward a carrier financial settlement system, some carriers migrate to whole circuit provisioning for routes where facilities-based and international simple resale opportunities exist. In telecommunication self-provisioning may achieve greater efficiency while also enabling the carrier to avoid having to pay an access charge to a foreign correspondent based on substantially above cost accounting rates.

connectivity over profits, cared little whether traffic flows were symmetrical, and usually had ample capacity available for the mostly text-based applications. This environment juxtaposes with the current conditions characterized by Athe for-profit nature of the Internet today; (2) asymmetry of global interconnections; (3) the increasing mix of Internet-based applications; and (4) the congestion and brown-outs on the Internet.@ 6

3.4 Internet Daisy-Chains and Telecommunication Correspondents

The Internet and telecommunication networks originate from a different baseline in terms of interconnection and coordination. Internet packet switching, the TCP/IP protocols and shared interest in global access have supported what one could term an Aall for one and one for all@ service orientation. An ISP need not network engineer a complete pathway for all conceivable routings that its customers might require. It can access the networks installed by other ISPs without having established a direct operating agreement with all other ISPs and without having agreed to an exchange of traffic. Until such time as the Internet became partitioned into self-contained Aintranets@ and Aextranets@ with parts of the transmission capacity available exclusively to a subset of all ISPs and Internet users, theoretically any single ISP could access all other ISPs and any part of the networks collectively constituting the Internet.

The Anetwork of networks@ slogan refers to the daisy-chaining of different networks into the appearance of one, integrated and seamless network. By way of an analogy to aviation consider the Internet a Afree fly zone,@ i.e., an environment where users can traverse all of cyberspace without significant restrictions. In the physical world of airspace nations must secure AFlights of Freedom@ to operate aircraft within the sovereign airspace of nations, to land, to take on and deliver passengers, and to transport passengers from one location to another within a nation. In the virtual world of the Internet, users face little if any such impediments. Better yet, for Internet travelers the transporter through foreign lands is none other than the networks owned and operated by foreign ISPs. The Internet traveler perceives that he or she need not secure and pay for transport beyond the Afirst and last@ kilometer of the circuit-switched public telephone network and the single network operated by the ISP providing Internet access. An Internet user might acknowledge that costs are incurred when the user requires the

---

6 See Organisation for Economic Co-Operation and Development, Committee for Information, Computer and Communications Policy, New Technologies and Their Impact on the Accounting Rate System, OECD/GD(97)14 (Paris, 1997); Organisation for Economic Co-Operation and Development, Committee for Information, Computer and Communications Policy, Refile and Alternative Calling Procedures: Their Impact on Accounting Rates and Collection Charges, OECD/GD(95)19 (Paris, 1995); Robert M. Frieden, "Falling Through the Cracks: International Accounting Rate Reform at the ITU and WTO." 22 Telecommunications Policy No. 11 (December 1998); Robert M. Frieden "The Impact of Call-Back and Arbitrage on the Accounting Rate Regime," 21 Telecommunications Policy, No. 9/10, 819-827 (1997).

services of other ISPs to route traffic beyond the reach of the ISP the user paid for service. But such a user might simply conclude that his or her ISP provides reciprocal and roughly equivalent switching and routing services for other ISPs thereby obviating the need to meter traffic flows.

Telecommunication networks have never had such porousness and ambiguity in routing arrangements, nor do telecommunication carriers expect traffic flows to even out. At least in theory a telecommunication carrier knows the composition of all traffic it carries in terms of which carrier Acorrespondent@ originated the traffic and which carrier will bear financial responsibility for such carriage. The telecommunication pricing model works to eliminate free riders even though economically speaking the routing of either Internet and telecommunication traffic may trigger little if any incremental cost under circumstances where no congestion exists. For Internet applications, near zero carriage costs favoured the absence of a pricing system, i.e., a zero cost arrangement for the termination of traffic at an intended destination, or the intermediary transiting of traffic onward toward its final destination. For telecommunication applications, near zero carriage costs favoured creative pricing arrangements like usage insensitive private lines and circuitous routing which takes advantage of distance insensitivity, i.e., the technical capability of routing traffic at the speed of light coupled with the financial capability of loading traffic at near zero cost where time of day and other factors mean that plenty of transmission capacity remains available.

Even if a particular routing might trigger no incremental cost, telecommunication carriers have devised a pricing arrangement that attributes a cost for each call, or at least for the assumed volume generated by a particular link. Such specificity means that telecommunication carriers bear the burden of acquiring or leasing all the transport facilities their customers need. Likewise, these carriers establish operating agreements with each other for direct service, and for indirect, transiting service using the transmission facilities of another carrier. Accordingly, a telecommunication carrier knows what facilities it owns and to which it has access.

ISPs historically have lacked such privity of contract and routing specificity. In exchange for a general offer of routing reciprocity a single ISP has access to all other ISPs= network resources. ISP A, having an interconnection arrangement with ISP B, can tap into both ISP B's network and whatever interconnection and transiting opportunities ISP B has secured from ISP C and others. Such reciprocity works well and efficiently when all ISPs generally handle the same volume of traffic, operate the same type of transmission facilities and metering would impose significant financial, administrative and technical costs. Satisfaction with a Arough justice@ unmetered exchange of traffic based on an assumption that traffic pattern were at parity even without the ability to measure traffic, made it possible for ISPs to forego the expense and efficiency loss in having to erect systems for tracking traffic volumes. A larger percentage of each packet could handle actual traffic rather than contain billing and other information needed if all ISPs had to meter traffic.

Such reciprocity does not work well when ISPs are not equals, an increasingly likely scenario as governments reduce or eliminate financial subsidies and as an increasing number of ISPs enter the marketplace. With the privatization of the Internet a variety of commercial ventures have begun providing service with different subscriber numbers, inventory of transmission capacity, interconnection agreements with other ISPs, number of locations where the venture can exchange traffic and locations where the venture does business.

Telecommunication carriers have established pricing arrangements that take into consideration the likelihood that the carriers will differ in one or more of the factors identified above. Telecommunications carriers agree to become correspondents for the purpose of securing traffic transiting and terminating service, but they do not expect each other to be true peers. In telecommunications one carrier will operate in a country with more people, income and competition typically resulting in a comparatively higher volume of outbound calls than what a correspondent would generate in a smaller and poorer country. Under the anticipated and likely condition of asymmetric traffic patterns, the carriers have established a financial settlement arrangement whereby the carrier generating more traffic pays the terminating carrier for the comparatively greater use of that carrier=s facilities. A similar arrangement might apply for Internet correspondents in the future. [7]

3.4 The Similarities Between Telecommunication and Evolving Internet Pricing Systems

As discussed above, a maturing Internet will encompass an increasingly diversified and heterogenous set of ISPs. A unitary peering system for network interconnection cannot work under such circumstances, particularly when an increasingly diversified set of carriers generate different traffic volumes and congestion becomes a worrisome factor. Already alternative pricing systems have come into place that create a hierarchy of ISPs based on such factors as traffic volume, number of interconnection points accessible to the ISP, and the bandwidth of lines used by the ISP. It appears that a declining number of ISPs will consent to, and be satisfied with, zero cost, "all for one and one for all" interconnection. Such public peering has resulted in significant congestion resulting in degraded service. In response to the combination of congestion and the free-rider/subsidized ISP, Tier-1 ISPs have established private peering points and have largely exited preexisting public peering points. Free access to these private peering points is limited to similarly situated, Tier-1 ISPs. Other ISPs can secure

7. For an analysis of this scenario see Robert M.Frieden, AWithout Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization.@ 3 Virginia Journal of Law & Technology 8 (Fall, 1998) available at http://vjolt.student.virginia.edu/.

access to these new interconnection points for a price. The terms and conditions under which lesser ISPs secure access to Tier-1 ISP facilities has begun to parallel the contractual, private correspondent telecommunications model: lesser-ISPs now must lease facilities from a Tier-1 ISP. Possibly in the future a Tier-1 ISP might require a traffic-based financial settlement.

The diversification and proliferation of ISPs occurs at the same time as ISPs have grown more conscious of Internet traffic volumes and patterns. This subject has become more important to ISPs, because demand has grown so quickly and dramatically that operators have had difficulty in making timely system upgrades. Necessary transmission and switching capacity enhancements to keep up with demand force ISPs collectively to make substantial additional investments lest the Internet become bogged down and inoperable. ISPs, which conscientiously make infrastructure upgrades, rightly object to a cost-free interconnection arrangement available to other ISPs who have failed to make commensurate investments.

A single ISP=s failure to invest in infrastructure improvements may occur due to a lack of financial resources, comparatively lower concerns about customer service, or perhaps even the absence of an immediate need for upgrades to serve its direct subscribers. ISPs that do not invest in network upgrades may not constitute a weak link in the global sense as not much traffic might transit their facilities. But to the extent these non-upgrading ISPs nevertheless experience the same kind of growth in subscribership, the traffic demands of such growing subscriber ranks aggregated with growing demand everywhere, exacerbate the potential for congestion throughout the Internet.

The growing ranks of Internet subscribers threaten to convert the World Wide Web into the AWorld Wide Wait@ absent substantial and expedited upgrades in bandwidth, modem banks, routers and the variety of facilities that make up the Internet. The Tier-1 ISPs and those operators most conscientious about making necessary service improvement may consider a financial settlement mechanism essential for maintaining the Internet=s ability to handle vastly more subscribers, traffic and Web sites. These ISPs have an additional justification for changing the terms and conditions for interconnection, viz., the pattern of Internet traffic. The potential for congestion and degraded service over the Internet results from both aggregate increase in demand and specific traffic routing and switching requirements on particular routes and for specific routers and servers.

An instructive example results from the mixed results and unfulfilled opportunities in an Internet promotion recently attempted by Victoria=s Secret a women=s intimate apparel and fashion chain. During the 1999 National Football League championship game, commonly referred to as the Suberbowl, the company spent approximately 1.5 million U.S. dollars for thirty seconds of advertising time. The company attempted to capitalize on the vastly expanded number of visits to its World Wide Web site and the upcoming Valentines Day holiday by organizing an Internet-mediated fashion show.

The company spent 5 million U.S. dollars to promote its AWebcast@ [8] and prepared to serve as many as 500,000 viewers. While exploiting the Internet during a peak time for mens= purchases of intimate apparel and their comparatively greater use of and comfort with the Internet for electronic commerce, Victoria=s Secret nevertheless managed to disappoint and frustrate over one million prospective participants who could not access the program, or who experienced technical difficulties. The onslaught of visitors resulted in a bottleneck that either blocked any viewership, or so degraded the streaming of video and sound as to diminish the experience and presumably the mood to buy products from the company.

The mixed results in the electronic commerce promotion by Victoria=s Secret and other ventures including, real time stock market transactional services exemplify the risks and rewards in using the Internet for peak demand, mission critical applications. But perhaps most importantly for our purposes they demonstrate that the current Internet pricing system cannot:

$ prioritize traffic based on willingness of sender or recipient to pay a premium price;

$ reserve capacity; and

$ offer more reliable service than Abest efforts@ routing for a price higher than the zero cost/true peering, SKA or capacity leases of smaller ISPs. [9]

In view of the stakes involved and the negative fallout that did occur companies like Victoria=s Secret gladly would have paid for temporarily more reliable service. But a fundamental pricing problem is that the company had no way to secure such premium service. No clearing house, broker or facilitator exists for accommodating Internet users having particular premium Internet application needs. Under current conditions Victoria=s Secret would have had to erect its own Internet network accessible from the general World Wide Web, but still controlled by it. While the company could construct an intranet or extranet over time, it could not have done so temporarily for the broadcast of its fashion show.

3.5 Traffic Patterns Matter in Both Telecommunication and Internet Applications

---

8. Edward Rothstein, A A Sex Metaphore, by Victoria=s Secret,@ The New York Times, Section E, Part 1, p. 28 (February 5, 1999).

9. Extensive efforts are underway to address these limitations. For example a Resource Reservation Protocol (ARSVP@) would provide a means to provide and charge a premium price for dedicated links.

Traffic patterns have always mattered in telecommunications and a pricing mechanism was implemented at the outset to account for differences in flow. Traffic patterns similarly matter in the Internet, but an appropriately calibrated pricing mechanism did not exist at the outset. Accordingly, implementing one now will require time, difficult commercial negotiations and perhaps a public policy/public relations campaign to convince the larger set of stakeholders that the overall health of the Internet depends, in large part, on the establishment of a new, fairer and more rational pricing regime. Of course beneficiaries of the status quo have every incentive to delay the date when they lose their significant financial windfalls.

A 1997 study conducted by the International Telecommunication Union succinctly stated that A[t]raffic patterns on the Internet are seriously distorted . . . [and not easily remedied because] there is considerable inertia . . ..@ [10] The study reported that in 1994 more than 80 percent of the total global Internet traffic traversed the United States largely due to the number of desirable World Wide Web sites located in that country. Under a telecommunications settlement scheme, service providers generating more outbound traffic than inbound would make transfer payments to the carriers transiting or terminating the comparatively more traffic.

Despite the importance of traffic patterns in both Internet-mediated and telecommunication environments, the Internet pricing system historically has operated from a "rough justice," Avalue of service@ model. Because so many Internet subscribers seek access to U.S.-based Web sites and thereby derive comparatively greater utility and value from such access, the ISPs providing such access must largely bear the expense in engineering complete links to U.S. sites. The balance of utility has skewed heavily to the United States, because of the attractiveness of what U.S. sites have to offer and because the ISPs seeking access historically have had less to offer by way of both networking facilities and number of Web sites. [11]

Absent a change in the flow of traffic toward parity, equity, utility and demand inelasticity factors support the requirement of U.S. ISPs that interconnection take place in the U.S as opposed to some equidistant location. The rationale for requiring off-shore ISPs to self-provision the lines achieving such

---

10. International Telecommunication Union, Challenges to the Network Telecoms and the Internet at p. 25 (Geneva: 1997).

11. AWhen foreign countries wish to connect to a U.S. backbone network they are obligated to cover the full costs of the international private line (both half circuits). There is some logic in this in that the countries connecting to the U.S. backbone network are not providing the U.S. backbone with much more than one network and a limited number of routes, all of which may be in just one country. The connecting network benefits greatly from connecting to the U.S. backbone; whereas, the U.S. backbone networks benefit much less from connecting with the outside network.@ Id.

access all the way to a U.S. location lies in the assumption that such access enhances the value and commercial appeal of the off-shore ISP. It can deliver to its subscribers U.S.-based content, access to the robust set of U.S. networks and as well the use of U.S. networks for accessing the facilities and services of other networks beyond. Arguably an off-shore ISP, which does not establish peering arrangements throughout the world, could secure subsidized or under-priced access to U.S. ISP networking arrangements for global access.

Just as some Asia-Pacific ISPs and national government object to self-provisioning in the face of migrating traffic patterns, U.S. ISPs object to what has been termed the Ahot potato@ [12] routing: the strategy of some ISPs to find a relatively nearby peering site for cheap off-loading of traffic onto another ISPs= network that will transit the traffic over a long haul to the final destination. In the worst case scenario, an ISP in the Asia-Pacific region might secure access to the rest of the world courtesy one or more U.S. ISPs and via one peering point the Asia-Pacific ISP accessed via self-provisioned international private lines. While seemingly burdensome to force self-provisioning, the ability to access the rest of the world via other ISP facilities presents quite a bargain, particularly in light of comparatively cheap private line rate for access to the United States relative to other points.

Such transiting opportunities do exist as evidenced by route tracing performed by the publication *TeleGeography*. [13] Intra-regional traffic within Asia-Pacific may transit the United States to take advantage of several combined factors: distance insensitivity in telecommunication routing, comparatively cheaper private line rates to the United States than between Asia-Pacific nations and the robust network access opportunities accruing from having secured even just one network access point in the United States.

3.6 Traditional Internet Pricing Arrangements and Reasons for Closer Scrutiny

For our purposes we can consider the traditional Internet pricing model as one involving public peering with no financial settlement. In an environment dominated by government agencies, the military and academic organizations, the goal of maximizing positive networking externalities predominated. Likewise even with the exit of government sponsorship and the onset of privatization, a promotional phase still supported affirmative efforts to spread the reach and accessibility of the Internet even if operators failed to maximize profitability. One working definition of this public peering environment identifies the landscape as an:

---

12. ABecause of >shortest exit routing,= the standard means of exchanging traffic among networks, the data goes onto the receiver=s network at the earliest point. All networks have an interest in passing off the data like a hot potato, to the other network at the earliest point, since it saves the sender-ISP valuable bandwidth.@ APeering and Fearing.

13. See Kenneth Neil Cukier, AThe Global Internet: A Primer, Fig. 8, Internet Traceroute Maps, July 1998, From Hong Kong to Japan via Silicon Valley, California, p, 121 in Gregory C. Staple, Ed. TeleGeography (1999).

> interconnection of two public networks that provide connectivity to hosts whose routes are advertised on the global Internet, on a settlement-free basis that allows customers of one network to exchange traffic to customers directly on the second ISP=s network. [14]

Set out below are several Internet network interconnection and pricing models.

Sender Keep All

The SKA model allows ISPs to retain all subscriber payments without having to settle accounts with other ISPs who participate in routing and delivering traffic. This model promotes the daisy-chaining of unaffiliated networks and delivers global access to sources of information, commerce and entertainment. This model has served as the primary template for Internet traffic routing, because of its administrative convenience and the willingness of ISPs to promote network connectivity regardless of whether traffic flows are symmetrical. SKA involves network interconnection without a metering mechanism either because the parties do not care whether traffic symmetry exists, assume that such symmetry exists, or believe that metering and the settlement of financial accounts trigger more cost and inconvenience than a Arough justice@ agreement to accept and route onward each others= traffic.

The SKA model promotes positive network externalities and universal service, because smaller and rural ISPs typically enjoy opportunities to generate more outbound traffic for free carriage by other ISPs than they receive from other ISPs for carriage onward to another ISP or for terminating traffic. The opportunity to avoid paying a penalty for being comparatively less necessary and operationally more expensive than larger ISPs translates into an opportunity for users in rural and high cost service areas to access the Internet on terms and conditions similar to what urban subscribers pay.

Peer-to-Peer Bilateral

This Internet-specific model adopts SKA, but with the expectation of traffic symmetry. Two unaffiliated ventures agree to use this model, which requires a direct and meterless connection, if and only if they have Athe same, size, experience, technology and customer base.@ [10] This model may eliminate the opportunity for Internet users in rural and high cost areas to pay less than full service costs,

---

14. Peering and Fearing.

if ISPs in such areas must resort to more expensive transit arrangements with bigger ISPs resulting in a one-way transfer payment from the smaller ISP to larger ones.

Hierarchical Bilateral

Even before the threat of network balkanization, a hierarchy of ISPs has developed based on geographical scope of service, available bandwidth, traffic volume and subscribership. The hierarchical bilateral model applies when different types of ISPs agree to interconnect their networks. The terms and conditions of this two-party contract reflect unequal bargaining strength in the sense that a smaller ISP, denied SKA and other cost-free interconnection opportunities, now must persuade a larger ISP to handle its traffic. A negotiation in this context establish a customer-provider, resale relationship rather than a carrier-to-carrier or ISP peering arrangement. Increasingly this model predominates as most Tier-1 ISPs treat regional and local ISPs as Aclients.@ Accordingly, the smaller regional and local ISP typically has to transfer funds to the Tier-1 ISP, because the bigger operator has incurred a greater infrastructure investment burden and has the capacity and wherewithal to route the smaller ISP=s traffic onward to another network, or to the final destination.

While transfer payments occur in this model, it is important to note that no incentive exists for a likely transfer payment recipient, i.e., the Tier-1 ISP, to discriminate or to deny interconnection with a smaller ISP who will pay for access. If this model continued to dominate the Internet, then the free ridership and unintended subsidy problem could abate without any single, financially qualified network operator facing denied access to other networks. The smaller network operator would simply have to agree to make the necessary transfer payments and thereby have transit [11] access to servers and e-mail recipients, etc. via the switching and routing facilities of other, larger and geographically diverse networks.

Third Party Administrator

The Third Party Administrator model involves a neutral Aclearinghouse@ function managed by a paid administrator that might not even operate a network. Before relinquishing all Internet management responsibilities, the National Science Foundation operated Network Access Points that served as public peering sites for the exchange of traffic between networks. Now commercial ventures, like the Commercial Internet Exchange Association (ACIX@), perform the same function. These businesses place greater emphasis on generating a profit from administrative fees, and reflect less of a quasi-common carrier orientation, i.e., agreeing to nondiscrimination and open access to any ISP on a rational,

---

13. For background on the manner by which international carriers agree to switch and route each others' traffic <u>see</u> Rob Frieden, International Telecommunications Handbook (Boston: Artech, 1996).

14.

traffic volume-based price structure.

The Third Party Administrator model works well when the administrator has the financial wherewithal to expand capacity and routing functions to meet demand and to maintain an adequate level of service. Third Party Administrators must enforce requirements that ISPs maintain bandwidth and traffic processing capabilities commensurate with vastly expanding traffic growth. Currently some traffic exchange locations have become so bogged down with traffic that packets of information must be resent, or are lost altogether.

ISPs, particularly ones with the largest traffic volume and available transmission capacity, are offloading some or all of their traffic onto Aprivate peering@ locations, because public peering points have become congested bottlenecks. This migration has the most adverse effect on smaller ISPs who lack the facilities investment to interconnect individually with one or more of the Tier-1 ISPs at several locations, typically an individual ISP's APoint of Presence.@ A small ISP could expect to provide its subscribers with access to just about any other ISP's network simply by interconnecting with the large number of peers at a Third Party Administrator=s public peering site. When Tier-1 ISPs boycott such sites and agree to handle traffic of lesser ISPs only upon payment for capacity leases, the smaller ISP must scramble to find substitute ways to access the major carrier's disparate networks, typically at several different locations and a higher cost.

3.7 Pricing Responses to Changed Circumstances

Concerns about network congestion, burgeoning demand, diversification of Internet applications, including high capacity video, free-riders, and hot potato routing collectively have caused ISPs to rethink the peering and settlement system. Private peering has become the most recent interconnection model and the one most likely to involve some degree of discrimination, entrance requirements, or transfer payment. This model involves the overlay of quasi-private Internets unavailable to every ISP or Internet user, or available at a price. Private peering users purposefully deem their networks inaccessible to outsiders ostensibly to preserve Anetwork integrity@ and minimal quality of service levels. However, the migration to private peering also results from the real or perceived need to safeguard a sizeable and expanding investment from the congestive effects of free riders.

The onset of private peering does not eliminate public peering opportunities, nor does it necessarily reduce the Internet=s robustness, seamlessness and integration. However, it does make the Internet more hierarchical in the sense that the Tier-1 ISPs have reduced the number of ISPs with which they peer on an SKA, cost-free basis. [13] It also makes the pricing structure arguably more rational and costly to smaller (Alower-tier@) ISPs as they must pay more for the privilege of accessing Tier-1 carrier networks.

Currently lower-tier ISPs can maintain tier-1 ISP access by leasing the tier-1 ISP-specified number of private lines, priced on a flat, usage-insensitive basis. Economists assert that an even more efficient and equitable pricing systemBparticularly during network congestion-would involve settlements on a per connection basis, and perhaps even on a next-hop basis, i.e., routerBby-router settlement for all links between the content requester and the content provider, presumably for routes both to and from the server providing the content. [14] Such precision in network management would require more extensive monitoring and perhaps more overhead and less Apaying@ content in packets. On the other hand it would impose market-based discipline in the switching and routing of Internet traffic. ISPs would face appropriate incentives and burdens with fewer opportunities to hand off traffic without incurring attributable costs. Likewise subscribers would face more exacting price signals and payment obligations with the mixed result of perhaps having to pay for services previously available on a flat-rated, AYCE basis, but also having the opportunity to pay for and receive premium (faster, better, more reliable) service.[15]

3.8 Emerging Internet Charging Issues

The diffuse and largely unregulated nature of the Internet will largely foreclose any attempt to manage its evolution and to enforce specific policy outcomes. The migration from government-sponsorship to private enterprise dilutes an already limited degree of control.
Notwithstanding these substantial limitations, governments do need to monitor changes in the Internet, particularly as they relate to public policy matters for which government has a legitimate interest. Put another way, governments cannot shape the manner by which the Internet matures, nor can it mandate a particular outcome. But government should understand the nature and scope of the Internet=s influence on markets and individuals.

Internet traffic patterns, coupled with emerging financial settlement mechanisms will have a

---

14. For an extensive collection of Internet access and pricing papers and articles see Hal R. Varian, The Information Economy, The Economics of the Internet, Information Goods, Intellectual Property and Related Issues, available at http://www.sims.berkeley.edu/resources/infoecon/ (viewed Feb. 9, 1999).

15. See David D. Clark, ACombining Sender and Receiver Payments in the Internet,@ in Gregory L. Rosston and David Waterman, eds., Interconnection and the Internet, 95-112 (Lawrence Erlbaum Assoc. 1997).

profound impact on who pays for needed infrastructure deployment and upgrades, and how one pays. Stakeholders in the Agolden age@ of carefree, zero cost settlements will balk at new pricing regimes that divest them of opportunities to ride the Internet without paying the full fare. Likewise advocates for new pricing regimes may have efficiency and equity arguments on their side, but may fail to secure immediate relief given the decentralized and largely ungoverned nature of the Internet. Relief or discomfort for either side probably will evolve on an incremental, often bilateral basis as ISPs negotiate and renegotiate the terms and conditions for interconnection.

Internet forecasting is a particularly risky undertaking. Indeed, follow-on modules in this endeavour may provide the kind of empirical and statistical data needed to assess what areas, routes and stakeholders have a rightful justifications for revised interconnection terms and conditions. It is beyond the scope of this Issues Paper to scrutinize or second-guess the provisions contained in bilateral and multilateral ISP interconnection agreements. What this Issues Paper can do is identify areas where APEC members may have a legitimate interest or concern relating to Internet pricing and interconnection. Additionally in the next section we will provide suggestions on what data collection could assist APEC members in understanding and addressing the nature of problems and controversies. In the last part of this section we set out our preliminary and non-exhaustive list of questions for additional examination and analysis:

1) Have Internet traffic patterns within Asia-Pacific and between Asia-Pacific and other locations changed recently to such a significant degree that existing peering arrangements place a disproportionately high financial burden on one party as compared to the allocation contemplated by parties initially in their voluntary, commercial transaction?

2) Do Internet traffic patterns between and among APEC member nations continue to support the requirement that ISPs operating outside the United States provision at their own expense the transmission capacity needed for peering within the United States?

3) What constitutes the best measure, or set of measurements for tracking Internet traffic? To what extent do traditional telecommunication measures, e.g., erlangs, bandwidth, minutes of use, etc. provide a proper gauge?

4) Do government have a role in monitoring, managing, approving, or regulating Internet peering, interconnection and pricing arrangements?

5) Can stakeholders in Internet peering, interconnection and pricing arrangements adequately self-regulate and resolve disputes, particularly relating to changed circumstances and the desire by one party to renegotiate and revise the terms and conditions of commercial contracts?

6) What deviation from absolute parity of Internet traffic streams might an ISP tolerate under a Sender Keep All pricing regime?

7) Presumably U.S. ISPs have been able to require home-based peering (instead of peering at an intermediary location, or a split in the cost of transmission costs) on the assumption that a large majority of Internet traffic between a U.S. and off-shore point will involve U.S. outbound traffic. What if anything might trigger a change in this assumption? What degree of change are ISPs likely to tolerate before one or both parties in a bilateral transaction seek to change the terms and conditions of a peering agreement based on such changed circumstances?

8) Where and on what routes might a significant change in Internet traffic patterns occur in the near term?

9) Are there any reasons that an ISP might have a different peering policy, coupled with different financial terms, for Internet traffic which transits rather than terminates an ISP=s network? What if any different peering terms and conditions might exist when a U.S.-based ISP provides transiting services for traffic terminating within Asia-Pacific or beyond the U.S., e.g., Hong Kong to/from Japan via the U.S. and Honk Kong to/from the United Kingdom via the U.S.?

10) What statistical index or compilation might be used to gauge the value of an Internet routing? Could an event, e.g., the 2000 Olympics in Sydney, or the proliferation of Internet servers/Web sites in a particular country trigger changes in peering policies, terms and conditions?