



Interconnection, Peering and Settlements-Part II

July 1999

Geoff Huston

In Part I we examined the business drivers behind the adoption of the exchange model as the common basis of interconnection, and also examined the advantages and pitfalls associated with the operation of such exchanges within the public Internet. (See *The Internet Protocol Journal*, Volume 2, No. 1, March 1999.) In continuing our examination of the technology and business considerations that are significant within the subject of Internet Service Provider (ISP) interconnection, in this part we focus on the topic from a predominately business perspective.

Interaction Financials: Peering and Settlements

Any large multiprovider distributed service sector has to address the issue of cost distribution at some stage in its evolution. Cost distribution is the means by which various providers can participate in the delivery of a service to a customer who purchases a service from a single provider, and providers can each be compensated for their costs in an equitable structure of interprovider financial settlement.

As an example, when an airline ticket is purchased from one air service provider, various other providers and service enterprises may play a role in the delivery of the service. The customer does not separately pay the service fee of each airport baggage handler, caterer, or other form of service. The customer's original fare, paid to the airline, is distributed to other providers who incurred cost in providing components of the total service. These costs are incurred through sets of service contracts, and are the subject of various forms of interprovider financial settlements, all of which are invisible to the customer.

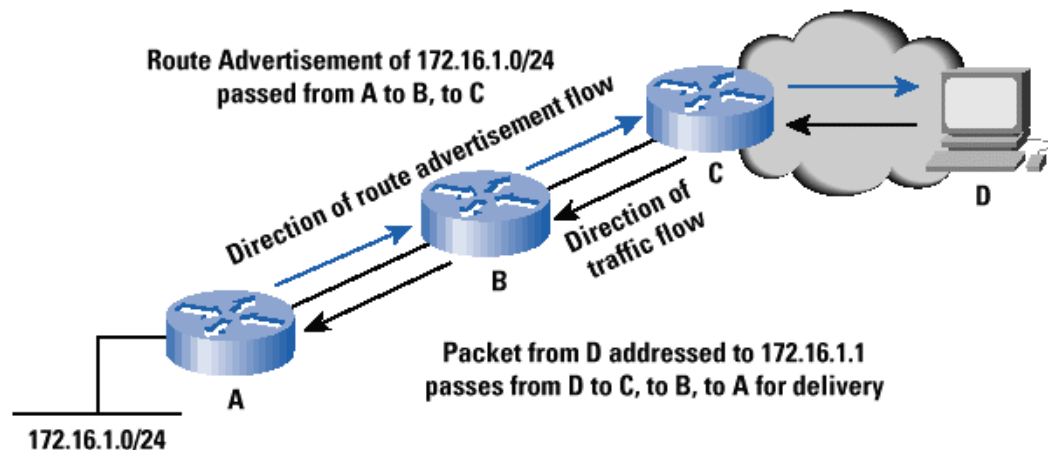
The Internet is in a very similar situation. Some 50,000 constituent networks must interconnect in one fashion or another to provide comprehensive end-to-end service to each client. In supporting a data transaction between two clients, the two parties often are not clients of the same network. Indeed, the two-client service networks often do not directly interconnect, and one or more additional networks must act in a transit provider role to service the transaction. Within the Internet environment, how do all the service parties to a transaction who incur cost in supporting the transaction receive compensation for their cost? What is the cost distribution model of the Internet?

Here, we examine the basis for Internet interprovider cost distribution models and then look at the business models currently used in the interprovider Internet environment. This area commonly is termed financial settlement, a term the Internet has borrowed from the telephony industry.

The Currency of Interconnection

What exactly is being exchanged between two ISPs who want to interconnect? In the sense of the meaning of currency as the circulating medium, the question is: What precisely is being circulated at the exchange and within the realm of interconnection? The technical answer to the question is: routing entries. When two parties exchange routing entries, the outcome is that traffic flows in response to the flow of routing entries. The route advertisement and traffic flows move in opposite directions, as indicated in Figure 1, and a bilateral routing-mediated flow occurs only when routes are passed in both directions.

Figure 1: Routing and Traffic Flows



Within the routing environment of an ISP there are many different classes of routes, with the classification based predominately on the way in which the route has been acquired by the ISP:

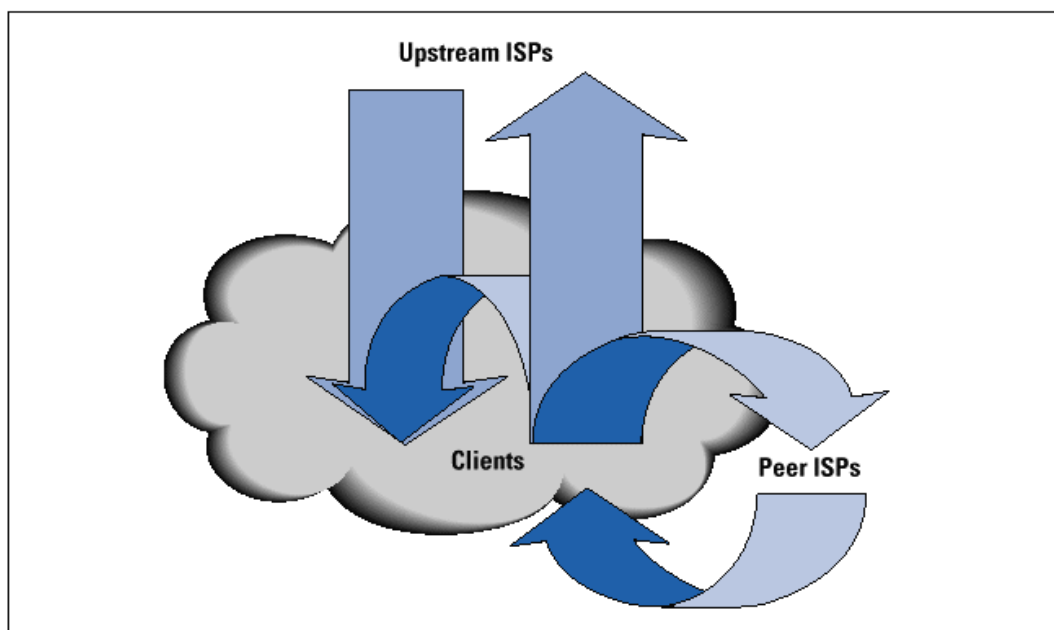
- Client routes are passed into the ISP's routing domain by virtue of a service contract with the client. The routes may be statically configured at the edge of the ISP's network, learned by a Border Gateway Protocol (BGP) session with the client, or they may constitute part of an ISP pool of addresses that are dynamically assigned to the client as part of the dialup session.
- Internal ISP routes fall into numerous additional categories. Some routes correspond to client services operated by the ISP, solely for access to the clients of the ISP, such as Web caches, Post Office Protocol (POP) mail servers, and game servers. Some routes correspond to ISP-operated client services that require Internet-wide access, such as Domain Name System (DNS) forwarders and Simple Mail Transfer Protocol (SMTP) relay hosts. Lastly are internal services with no visibility outside the ISP network, such as Simple Network Management Protocol (SNMP) network management platforms.
- Upstream routes are learned from upstream ISPs as part of a transit service contract the ISP has executed with the upstream provider.
- Peer routes are learned from exchanges or private interconnections, corresponding to routers exported from the interconnected ISP.

How then should the ISP export routes so that the inbound traffic flow matches the outbound flows implied by this route structure? The route export policy is generally structured along the following lines:

- Clients: All available routes in the preceding four categories, with the exception of internal ISP service functions, should be passed to clients, either in the form of a default route or as explicit route entries passed via a BGP session.

- Upstream providers: All client routes and all internal ISP routes corresponding to Internet-wide services should be passed to upstream providers. Some clients may want further restrictions placed on their routes being advertised in such a fashion. The ability for a client to specify such caveats on the routing structure, and the mechanism used by the ISP to allow this to happen, should be clearly indicated in the service contract.
- Peer ISPs: All client routes and all ISP routes corresponding to Internet-wide service should be passed to peer ISPs. Again the clients may want to place a restriction on such an advertisement of their routes as a qualification to the ISP's own route export policy.

This structure is shown in Figure 2.



The implicit outcome of this routing policy structure is that the ISP does not act in a transit role to peer ISPs and permits neither peer-to-peer transit nor peer-to-upstream transit. Peer ISPs have visibility only to clients of the ISP. From the service visibility perspective, client only services are not visible to peer ISPs or upstream ISPs, and, therefore, value-added client services are implicitly visible only to clients and only when they access the service through a client channel.

Settlement Options

Financial settlements have been a continual topic of discussion within the domain of Internet interconnection. To look at the Internet settlement environment, let's first look at the use of interprovider financial settlements within the international telephony service industry. Then, we will look at the application of these generic principles to the Internet environment.

Within the traditional telephony model, interprovider peering takes place within one of three general models:

Bilateral Settlements

The first, and highly prevalent, international peering model is that of bilateral settlements. A call-minute is the unit of settlement accounting. A call is originated by a local client, and the local client's service provider charges the client for the duration of the entire end-to-end call. The call may pass through, or transit, many providers, and then terminate within the network of

the remote client's local provider. The cost distribution mechanism of settlements is handled bilaterally. In the most general case of this settlement model, the originating provider pays the next hop provider to cover the costs of termination of the call. The next hop provider then either terminates the call within the local network, or undertakes a settlement with the next hop provider to terminate the call. The general telephony trunk model does not admit many multiparty transit arrangements. Most telephony settlements are associated with trunk calls that involve only two providers: the originating and terminating providers. Within this technology model, the bilateral settlement becomes easier, because the model simplifies to the case where the terminating provider charges the originating provider a per-call-minute cost within an accounting rate that has been bilaterally agreed upon between the two parties. Because both parties can charge each other using the same accounting currency, the ultimate financial settlement is based on the net outcome of the two sets of call-minute transactions with the two call-minute termination accounting rates applied to these calls. (There is no requirement for the termination rates for the two parties to be set at the same level.) Each provider invoices the originating end user for the entire call duration, and the financial settlements provide the accounting balance intended to ensure equity of cost distribution in supporting the costs of the calls made between the two providers. Where there is equity of call accounting rates between the two providers, the bilateral interprovider financial settlements are used in accordance with originating call-minute imbalance, in which the provider hosting the greater number of originating call-minutes pays the other party according to a bilaterally negotiated rate as the mechanism of cost distribution between the two providers.

As a side note, the Federal Communications Commission of the United States (FCC) asserts that U.S. telephone operators paid out some \$5.6 billion in settlement rates in 1996, and the FCC is voicing the view that accounting rates have now shifted into areas of non-cost-based settings, rather than working as a simple cost distribution mechanism. This accounting settlement issue is one of the drivers behind the increasing interest in voice-over-IP solutions, because typically no accounting rate settlement component exists in such solutions, and the call termination charges are cost-based, without bilateral price setting. In those cases

where accounting rates have come to dominate the provider's call costs, voice-over-IP is perceived as an effective lever to bypass the accounting rate structure and introduce a new price point for call termination in the market concerned.

Sender Keeps All

The second model, rarely used in telephony interconnection, is that of Sender Keeps All (SKA), in which each service provider invoices its originating client's user for the end-to-end services, but no financial settlement is made across the bilateral interconnection structure. Within the bilateral settlement model, SKA can be regarded as a boundary case of bilateral settlements, where both parties simply deem the outcome of the call accounting process to be absolutely equal, and consequently no financial settlement is payable by either party as an outcome of the interconnection.

Transit Fees

The third model is that of transit fees, in which one party invoices the other party for services provided. For example, this arrangement is commonly used as the basis of the long-distance/local access provider interconnection arrangements. Again, this case can be viewed as a boundary case of a general bilateral settlement model, where in this case the parties agree to apply call accounting in only one direction, rather than bilaterally.

Telephony Settlement Trends

The international telephony settlement model is by no means stable, and currently, significant pressure is being placed on the international accounting arrangements to move away from bilaterally negotiated uniform call accounting rates to rates separately negotiated for calls in each direction of a bilateral interconnection. Simultaneously, communications deregulation within many national environments is changing the transit fee model, as local providers extend their network into the long-distance area and commence interconnection arrangements with similar entities. Criticism also has been directed at the bilaterally negotiated settlement rates, because of the observation that in many cases the accounting rates are not cost-based rates but are based on a desire to create a revenue stream from accounting settlements.

Internet Considerations

Numerous critical differences exist between the telephony models of interconnection and the Internet environment; these differences have confounded all attempts to cleanly map telephony interconnection models into the Internet environment.

Internet Settlement Accounting by the Packet

Internet interconnection accounting is a packet-based accounting issue, because there is no "call-minute" in the Internet architecture. Therefore, the most visible difference between the two environments is the replacement of the call with the packet as the currency unit of interconnection.

Although we can argue that a TCP session has much in common with a call, this concept of an originating TCP call-minute is not always readily identified within the packet forwarding fabric, and accordingly it is not readily apparent that this is a workable settlement unit. Unlike a telephony call, no concept of state initiation exists to pass a call request through a network and lock down a network transit path in response to a call response. The network undergoes no state change in response to a TCP session, and therefore, no means is readily available to the operator to identify that a call has been initiated, and by which party. Of course the use of User Datagram Protocol (UDP), and various forms of tunnelling traffic, also confound any such TCP call-minute accounting mechanism.

Packets may be dropped

When a packet is passed across an interconnection from one provider to another, no firm guarantee is given by the second provider that the packet will definitely be delivered to the destination. The second provider, or subsequent providers in the transit path, may drop the packet for quite legitimate reasons, and will remain within the protocol specification in so doing. Indeed, the TCP protocol uses packet drop as a rate-control signal. For the efficient operation of the TCP protocol, some level of packet drop is a useful and anticipated event. However, if a packet is used as the accounting unit in a general cost distribution environment, should the provider who receives and subsequently drops the packet be able to claim an accounting credit within the interconnection? The logical response is that such accounting credits should apply only to successfully delivered packets, but such an accounting structure is highly challenging to implement accurately within the Internet environment.

Packet paths are not predetermined

Packet transit paths can be within the explicit control of the end user, not the provider. Users can exercise some significant level of control of the path a packet takes to transit the Internet if source routing is honoured, so that the relative packet flows between two providers can be arbitrarily manipulated by any client, if so desired.

Routing and traffic flow are not paired

Packet forwarding is not a verified operation. A provider may choose to forward a packet to a second provider without reference to the particular routes the second provider is advertising to the first party. A packet may also be forwarded to the second provider with a source address that is not being advertised to the second provider. Given that the generic Internet architecture strives for robustness under extreme conditions, attempts to forward a packet to its addressed destination are undertaken irrespective of how the packet may have arrived at this location in the first place, and irrespective of how a packet with reverse header IP addresses will transit the network.

Comprehensive routing information is not uniformly available. Complete information is not available to the Internet regarding the status and reachability of every possible Internet address. Only as a packet is forwarded closer to the addressed destination does more complete information regarding the status of the destination address become apparent to the provider. Accordingly, a packet may have incurred some cost of delivery before its ultimate undeliverability becomes evident. An intermediate transit provider can never be completely assured that a packet is deliverable.

Settlement Models for the Internet

Where a wholesale or retail service agreement is in place, one ISP is, in effect, a customer of the other ISP. In this relationship, the customer ISP (downstream ISP) is purchasing transit and connectivity services from the supplier ISP (upstream ISP). The downstream ISP resells this service to its clients. The upstream ISP must announce the downstream ISP's routes to all other customers and other egress points of the ISP's networks to honour the service contract to the downstream ISP customer.

However, given two ISPs who interconnect, the decision as to which party should assume the upstream provider role and which party should assume the downstream customer role is not always immediately obvious to either party, or even to an outside observer. Greater geographic coverage may be the discriminator here that allows the customer/provider determination. However, this factor is not the only possible one within the scope of the discussion. One ISP may host significant content and may observe that access to this content adds value to the other party's network, which may be used as an offset against a more uniform customer relationship. In a similar vein, an ISP with a very large client population within a limited geographic locality may see this large client base as an offset against a more uniform customer relationship with the other provider. In many ways, the outcome of these discussions can be likened to two animals meeting in the jungle at night. Each animal sees only the eyes of the other, and from this limited input, they must determine which animal should attempt to eat the other!

An objective and stable determination of which ISP should be the provider and which should be the client is not always possible. In many contexts, the question is inappropriate, given that for some traffic classes the respective roles of provider and client may swap over. The question often is rephrased along the lines of, "Can two providers interconnect without the implicit requirement to cast one as the provider and the other as the client?" Exploration of some concepts of how the question could possibly be answered is illustrative of the problem space here.

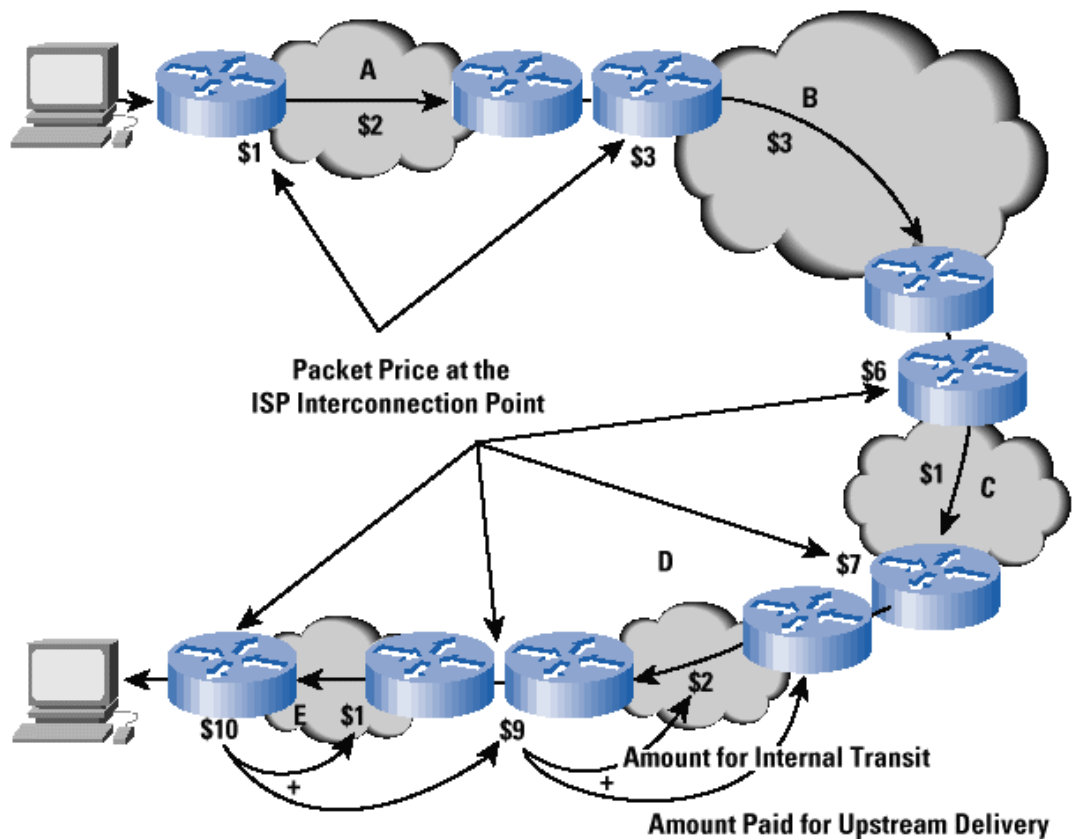
Packet Cost Accounting

One potential accounting model is based on the observation that a packet incurs cost as it passes through the network. For a small interval of time, the packet occupies the entire transmission capacity of each circuit over which it passes.

Similarly, for a brief interval of time, the packet is exclusively occupying the switching fabric of the router. The more routers the packet passes through, and the greater the number and distance of transmission hops the packet traverses, the greater the incurred cost in carrying the packet.

A potential settlement model could be constructed from this observation. The strawman model is that whenever a packet is passed across a network boundary, the packet is effectively sold to the next provider. The sale price increases as the packet transits through the network, accumulating value in direct proportion to the distance the packet traverses within the network. Each boundary packet sale price reflects the previous sale price, plus the value added in transiting the ISP's infrastructure. Ultimately, the packet is sold to the destination client. This model is indicated in Figure 3.

Figure 3: Financial Interprovider Settlement via Packet Cost Accounting



As with all strawman models, this one has numerous critical weaknesses, but let's look at the strengths first. An ISP gains revenue from a packet only when delivered on egress from the network, rather than in network ingress. Accordingly, a strong economic incentive exists to accept packets that will not be dropped in transit within the ISP, given that the transmission of the packet generates revenue to the ISP only on successful delivery of the packet to the next hop ISP or to the destination client. This factor places strong pressure on the ISP to maintain quality in the network, because dropped packets imply foregone revenue on local transmission. Because the packet was already purchased from the previous provider in the path, packet loss also implies financial loss. Strong pressure also is exerted to price the local transit function at a commodity price level, rather than attempt to undertake opportunistic pricing. If the chosen transit price is too great, the downstream provider has the opportunity to extend its network to reach the next upstream

provider in the path, resulting in bypassing the original upstream ISP and purchasing the packets directly from the next hop upstream source. Accordingly, this model of per-packet pricing, using a settlement model of egress packet accounting, and locally applied value increments to a cumulative per-packet price, based on incremental per-hop transmission costs, does allow for some level of reasonable stability and cost distribution in the interprovider settlement environment.

However, weaknesses of this potential model cannot be ignored. First, some level of packet drop is inevitable, irrespective of traffic load. Generally, the more remote the sender from the destination, the less able the sender is to ascertain that the destination address is a valid IP address, and the destination host is available. To minimize the liability from such potential packet loss, the ISP should maintain a relatively complete routing table and accept only packets in which a specific route is maintained for the network. More critical is the issue that the mechanism is open to abuse. Packets that are generated by the upstream ISP can be transmitted across the interface, which in turn results in revenue being generated for the ISP. Of course, per-packet accounting within the core of the network is a significant refinement of existing technology. Within a strict implementation of this model, packets require the concept of an attached value that ISPs augment on an ingress-to-egress basis, which could be simplified to a hop-by-hop value increment. Implementations feasibly can use a level of averaging to simplify this process by using a tariff for domestic transit and a second for international transit.

TCP Session Accounting

These traffic-based metrics do exhibit some weaknesses because of their inability to resist abuse and the likelihood of exacting an interprovider payment even when the traffic is not delivered to an ultimate destination. Of more concern is that this settlement regime has a strong implication in the retail pricing domain, where the method of payment on delivered volume and distance is then one of the more robust ways that a retail provider can ensure that there is an effective match between the interprovider payments and the retail revenue. Given that there is no intrinsic match of distance, and therefore cost, to any particular end-to-end network transaction, such a retail tariff mechanism would meet with strong consumer resistance.

Does an alternative settlement structure that can address these weaknesses exist? One approach is to perform significantly greater levels of analysis of the traffic as it transits a boundary between a client and the provider, or between two providers, and to adopt financial settlement measures that match the type of traffic being observed. As an example, the network boundary could detect the initial TCP SYN handshake, and all subsequent packets within the TCP session could be accounted against the session initiator, while UDP traffic could be accounted against the UDP source. Such detailed accounting of traffic passed across a provider boundary could allow for a potential settlement structure based on duration (call-minutes), or volume (call-volumes).

Although such settlement schemes are perhaps limited more by imagination in the abstract, very real technical considerations must be borne to bear on this speculation. For a client-facing access router to detect a TCP flow and correctly identify the TCP session initiator requires the router to correctly identify the initial SYN handshake, the opening packet, and then record all in-sequence subsequent packets within this TCP flow against this accounting element. This identification process may be completely impossible within the network at an interprovider boundary. The outcome of the routing configuration may be an asymmetric traffic path, so that a single interprovider boundary may see only traffic passing in a single direction.

However, the greatest problem with this, or any other traffic accounting settlement model, is the diversity of retail pricing structures that exist within the Internet today. Some ISPs use pricing based on received volume, some on sent volume, some on a mix of sent and received volume, and some use pricing based on the access capacity, irrespective of volume. This discussion leads to the critical question when considering financial settlements: Given that the

end client is paying the local ISP for comprehensive Internet connectivity, when a client's packet is passed from one ISP to another at an interconnection point, where is the revenue for the packet? Is the revenue model one in which the packet sender pays or one in which the packet receiver pays? The packet egress model described here assumes a uniform retail model in which the receiver pays for Internet packets. The TCP session model assumes the session initiator pays for the entire traffic flow. This uniformity of retail pricing is simply not mirrored within the retail environment of the Internet today. Although this session-based settlement model does attempt to promote a quality environment with fair carriage pricing, it cannot address the fundamental issue of financial settlements.

Internet Settlement Structures

For a financial settlement structure to be viable and stable, the settlement structure must be a uniform abstraction of a relatively uniform retail tariff structure. This conclusion is critically important to the entire Internet financial settlement debate.

The financial structure of interconnection must be an abstraction of the retail models used by the two ISPs. If the uniform retail model is used, the party originating the packet pays the first ISP a tariff to deliver the packet to its destination within the second ISP; then the first ISP is in a position to fund the second ISP to complete the delivery through an interconnection mechanism. If, on the other hand, the uniform retail model is used in which the receiver of the packet funds its carriage from the sender, then the second ISP funds the upstream ISP. If no uniform retail model is used, when a packet is passed from one provider to the other, no understanding exists about which party receives the revenue for the carriage of the packet and accordingly, which party settles with the other party for the cost incurred in transmission of the packet. The answer to these issues within the Internet environment has been to commonly adopt just two models of interaction. These models sit at the extreme ends of the business spectrum, where one is a customer/provider relationship, and the other is a peering relationship without any form of financial settlement, or SKA. These models approximately correspond to the second and third models described previously from traditional models of interconnection within the communications industry. However, an increasing trend has moved toward models of financial settlement in a bilaterally negotiated basis within the Internet, using non-cost-based financial accounting rates within the settlement structure. Observing the ISP industry repeat the same well-trodden path, complete with its byways into various unproductive areas and sometimes mistakes of the international telephony world, is somewhat interesting to say the least. Experiential learning is often observed to be a rare commodity in this area of Internet activity.

No Settlement and No Interconnection

Examining the option of complete autonomy of operation, without any form of interaction with other local or regional ISPs, is instructive within this examination of settlement options. One scenario for a group of ISPs is that a mutually acceptable peering relationship cannot be negotiated, and all ISPs operate disconnected network domains with dedicated upstream connections and no interconnection. The outcome of such a situation is that third-party connectivity would take place, with transit traffic flowing between the local ISPs being exchanged within the domain of a mutually connected third-party ISP (or via transit across a set of third-party ISPs). For example, for an Asian country, this situation would result in traffic between two local entities, both located within the same country, being passed across the Pacific, routed across numerous network domains within the United States, and then passed back across the Pacific. Not only is this scenario inefficient in terms of resource utilization, but this structure also adds a significant cost to the operation of the ISPs, a cost that ultimately is passed to the consumer in higher prices for Internet traffic.

Note that this situation is not entirely novel; the Internet has seen such arrangements appear in the past; and these situations are still apparent in today's Internet. Such arrangements have arisen, in general, as the outcome of an inability to negotiate a stable local peering structure.

However, such positions of no interconnection have proved to be relatively short-lived because of the high cost of operating international transit environments, the instability of the significantly lengthened interconnection paths, and the unwillingness of foreign third-party ISPs to act (often unwittingly) as agents for domestic interconnection in the longer term. As a result of these factors, such off-shore connectivity structures generally have been augmented with domestic peering structures.

The resultant general operating environment of the Internet is that effective isolation is not in the best interests of the ISP, nor is isolation in the interests of other ISPs or the consumers of the ISPs' services. In the interests of a common desire to undertake rational and cost-effective use of communications resources, each national (or regional) collection of ISPs acts to ensure local interconnectivity between such ISPs. A consequent priority is to reach acceptable ISP peering arrangements.

Sender Keeps All

Sender Keeps All (SKA) peering arrangements are those in which traffic is exchanged between two or more ISPs without mutual charge (an interconnection arrangement with no financial settlement). Within a national structure, typically the marginal cost of international traffic transfer to and from the rest of the Internet is significantly higher than domestic traffic transfer. In these cases, any SKA peering is likely to relate to only domestic traffic, and international transit would be provided either by a separate agreement or independently by each party.

This SKA peering model is most stable where the parties involved perceive equal benefit from the interconnection. This interconnection model generally is used in the context of interconnection or with providers with approximate equal dimension, as in peering regional providers with other regional providers, national providers with other national providers, and so on. Oddly enough, the parties themselves do not have to agree on what that value or dimension may be in absolute terms. Each party makes an independent assessment of the value of the interconnection, in terms of the perceived size and value of the ISP and the value of the other ISP. If both parties reach the conclusion that in their terms a net balance of value is achieved, then the interconnection is on a stable basis. If one party believes that it is larger than the other and SKA interconnection would result in leverage of its investment by the smaller party, then an SKA interconnection is unstable.

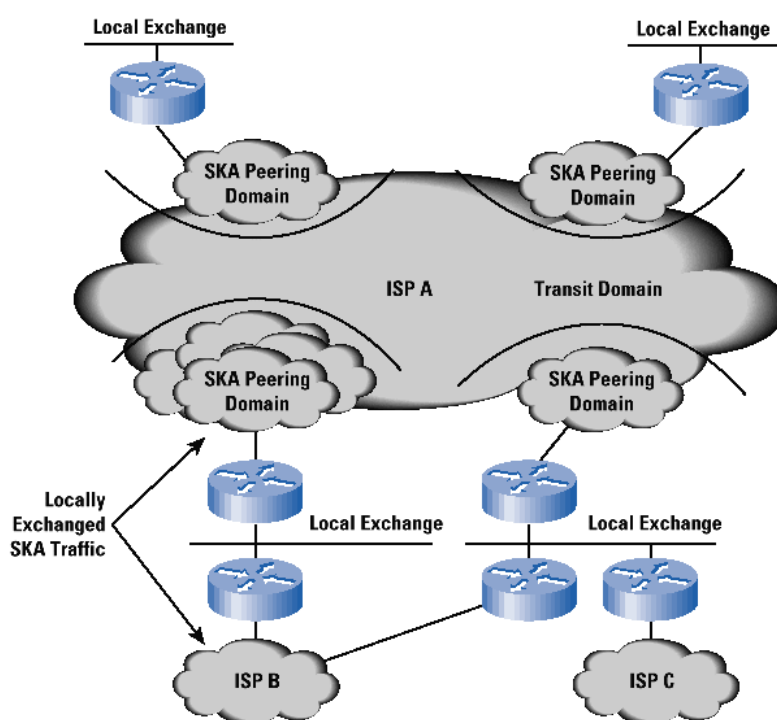
The essential criterion for a stable SKA peering structure is perceived equality in the peering relationship. This criterion can be achieved in many ways, including the use of entry threshold pricing into the peering environment or the use of peering criteria, such as the specification of ISP network infrastructure or network level of service and coverage areas as eligibility for peering.

A typical feature of the SKA peering environment is to define an SKA peering in terms of traffic peering at the client level only. This definition forces each peering ISP to be self-sufficient in the provision of transit services and ISP infrastructure services that would not be provided across a peering point. This process may not result in the most efficient or effective Internet infrastructure, but it does create a level of approximate parity and reduces the risks of leverage within the interconnection. In this model, each ISP presents at each interconnection or exchange only those routes associated with the ISP's customers and accepts only traffic from peering ISPs at the interconnection or exchange directed to such customers. The ISP does not accept transit traffic destined to other remote exchange locations, nor to upstream ISPs, nor traffic directed to the ISP's infrastructure services. Equally, the ISP does not accept traffic that is destined to peering ISPs, from upstream transit providers. The business model here is that

clients of an ISP are contracting the ISP to present their routes to all other customers of the ISP, to the upstream providers of the ISP, and to all exchange points where the ISP has a presence. The particular tariff model chosen by the ISP in servicing the customers is not material to this interconnection model. Traffic passed to a peer ISP at the exchange becomes the responsibility of the peer ISP to pass to its customers at its cost.

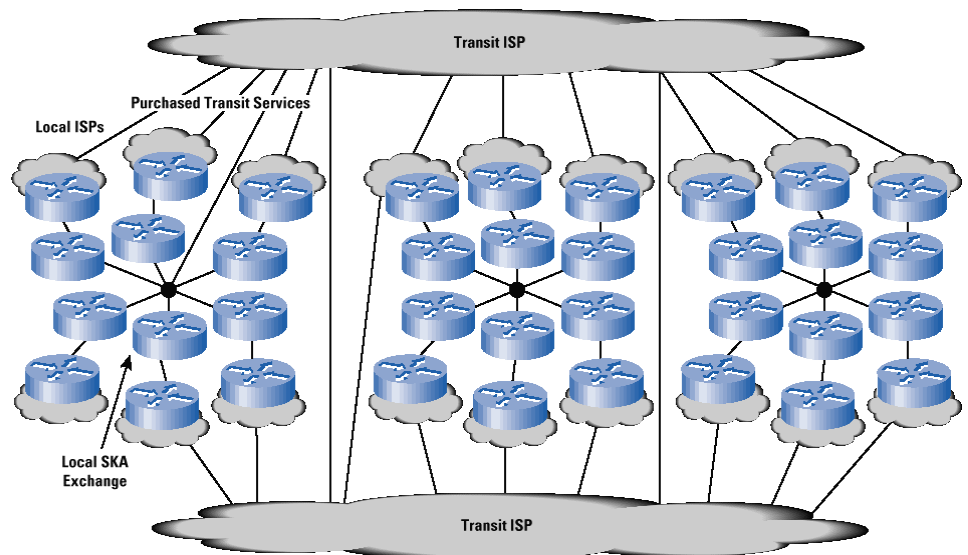
Another means of generating equity within an SKA peering is to peer only within the terms of a defined locality. In this model, an ISP would present routes to an SKA peer in which the routes correspond to customers located at a particular access POP, or a regional cluster of access POPs. The SKA peer's ability to leverage advantage from the greater level of investment (assuming that the other party is the smaller party) is now no longer a factor, because the smaller ISP sees only those parts of the larger ISP that sit within a well-defined local or regional zone. This form of peering is indicated in Figure 4.

Figure 4: SKA Peering Using Local Cells



The probable outcome of widespread use of SKA interconnections is a generalized ISP domain along the lines of Figure 5. Here, the topology is segregated into two domains consisting of a set of transit ISPs, whose predominate investment direction is in terms of high-capacity carriage infrastructure and high-capacity switching systems, and a collection of local ISPs, whose predominate investment direction is in service infrastructure supporting a string retail focus. Local ISPs participate at exchanges and announce local routes at the exchange on an SKA basis of interconnection with peer ISPs. Such ISPs are strongly motivated to prefer to use all routes presented at the exchange within such peering sessions, because the ISP is not charged any transit cost for the traffic under an SKA settlement structure. The exchange does not provide comprehensive connectivity to the ISP, and this connectivity needs to be complemented with a separate purchase of transit services. In this role, the local ISP becomes a client of one or more transit ISPs explicitly for the purpose of access to transit connectivity services.

Figure 5: ISP Structure of Local and Transit Operations



In this model, the transit ISP must have established a position of broad-ranging connectivity, with a well-established and significant market share of the wholesale transit business. A transit ISP also must be able to present customer routes at a carefully selected set of major exchange locations and have some ability to exchange traffic with all other transit ISPs. This latter requirement has typically been implemented using private interconnection structures, and the associated settlements often are negotiated bilaterally. These settlements possibly may include some element of financial settlement.

Negotiated Financial Settlement

The alternative to SKA and provider/client role selection is the adoption of a financial settlement structure. The settlement structure is based on both parties effectively selling services to each other across the interconnection point, with the financial settlement undertaking the task of balancing the relative sales amounts.

The simplest form of undertaking this settlement is to measure the volume of traffic being passed in each direction across the interconnection and to use a single accounting rate for all traffic. At the end of each accounting period, the two ISPs would financially settle based on the agreed accounting rate applied to the net traffic flow. Which way the money should flow in relationship to traffic flow is not immediately obvious. One model assumes that the originating provider should be funding the terminating provider to deliver the traffic, and therefore, money should flow in the same direction as traffic. The reverse model assumes that the overall majority of traffic, is traffic generated in response to an action of the receiver, such as web page retrieval or the downloading of software. Therefore, the total network cost should be imposed on the discretionary user, so that the terminating provider should fund the originating provider. This latter model has some degree of supportive evidence, in that a larger provider often provides more traffic to a smaller attached provider than it receives from that provider. Observation of bilateral traffic flow statistics tends to support this, indicating that traffic-received volumes typically coincide with the relative interconnection benefit to the two providers.

The accounting rate can be negotiated to be any amount. There is a caveat on this ability to set an arbitrary accounting rate, because where an accounting rate is not cost-based, business instability issues arise. For greater stability, the agreed settlement traffic unit accounting rate would have to match the average marginal cost of transit traffic in both ISP networks for the settlement to be attractive to both parties. Refinements to this approach can be introduced, although they are accompanied by significant expenditure on traffic monitoring and accounting systems.

The refinements are intended to address the somewhat arbitrary determination of financial settlement based on the receiver or the sender. One way is to undertake flow-based accounting, in which the cost accounting for the volume of all packets associated with a TCP flow is directed to the initiator of the TCP session. Here, the cost accounting for all packets of a UDP flow is directed to the UDP receiver. The session-based accounting is significantly more complex than simple volume accounting, and such operational complexity would be reflected in the cost of undertaking such a form of accounting. However, asymmetric paths are a common feature of the inter-AS environment, so that it may not always be possible to see both sides of a TCP conversation and perform an accurate determination of the session initiator.

Another refinement is to use a different rate for each provider, where the base rate is adjusted by some agreed size factor to ensure that the larger provider is not unduly financially exposed by the arrangement. The adjustment factor can be the number of Points of Presence, the range of the network, the volume carried on the network, the number of routes advertised to the peer, or any other metric related to the ISP's investment and market share profile. Alternatively, a relative adjustment factor can simply be a number, without any basis in a network metric, to which both parties agree.

Of course, such a relative traffic volume balance is not very robust either, and the metric is one that is vulnerable to abuse. The capability to adjust the relative traffic balance comes from the direct relationship between the routes advertised and the volume of traffic received. To reduce the amount of traffic received, the ISP reduces the number of routes advertised to the corresponding peer. Increasing the number of routes, and at the same time increasing the number of specific routes, increases the amount of received traffic. When there is a rich mesh of connectivity, the primary objective of routing policy is no longer that of supporting basic connectivity, but instead the primary objective is to maximize the financial return to the operator. If the ISP is paying for an "upstream" ISP service, the motivation is to minimize the cost of this contract, either by maximizing the amount of traffic covered under a fixed cost, or minimizing the cost by minimizing the traffic exchanged with the upstream ISP. Where there is a financially settled interconnection, the ISP will be motivated to configure its routing policies to maximize its revenue from such an arrangement. And of course an ISP will always prefer to use customer routes wherever possible, as a basic means of maximizing revenue into the operation.

Of greater concern is the ability to abuse the interconnection arrangements. One party can generate and then direct large volumes of traffic to the other party. Although overt abuse of the arrangements is often easy to detect, greed is a wonderful stimulant to ingenuity, and more subtle forms of abuse of this arrangement are always possible. To address this, both parties would typically indicate in an interconnection agreement their undertaking not to indulge in such forms of deliberate abuse. Notwithstanding such undertakings by the two providers, third parties can still abuse the interconnection in various ways. Loose source routing can generate traffic flows that pass across the interconnection in either direction. The ability to remotely trigger traffic flows through source address spoofing is possible, even where loose source routing is disabled. This window of financial vulnerability is far wider than many ISPs are comfortable with, because it opens the provider to a significant liability over which it has a limited ability to detect and control. Consequently, financial settlement structures based on traffic flow metrics are not a commonly deployed mechanism, because they introduce significant financial risks to the ISP interconnection environment.

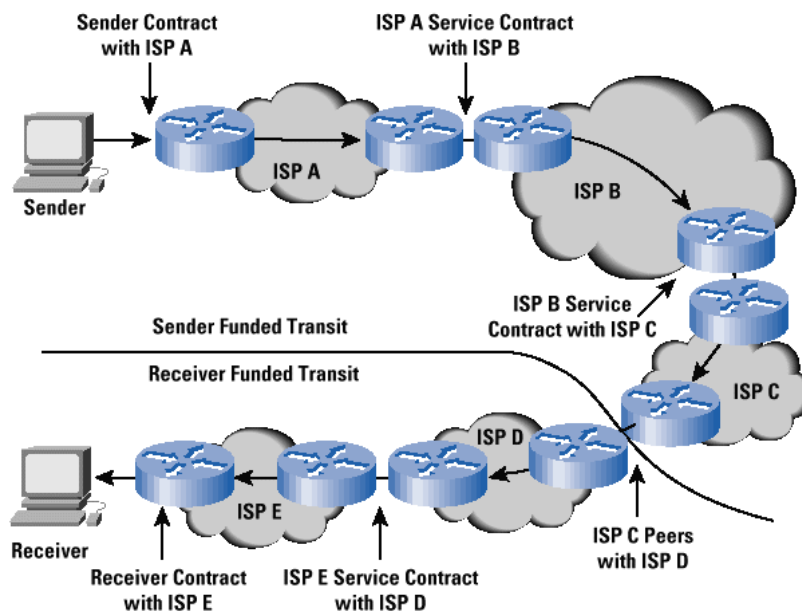
The Settlement Debate

The issue of Internet settlements, and associated financial models of settlement, has occupied the attention of a large number of ISPs, traditional communications carriers, public regulators, and many other interested bodies for many years now. Despite these concentrated levels of attention and analysis, the Internet interconnection environment remains one where there are no soundly based models of financial settlement in wide-spread use today.

It is useful to look further into this matter, and pose the question: "Why has the Internet managed to pose such a seemingly intractable challenge to the ISP industry?" The prime reason is likely to be found within the commonly adopted retail model of ISP services. The tariff for an ISP retail service does not implicitly cover the provision of an Internet transmission service from the client to all other Internet-connected hosts. In other words, the Internet service, as retailed to the client, is not a comprehensive end-to-end service.

In a simple model of the operation of the Internet, each ISP owns and operates some local network infrastructure, and may choose to purchase services from one or more upstream service providers. The service domain offered to the clients of this network specifically encompasses an Internet subdomain limited to the periphery of the ISP network together with the periphery of the contracted upstream provider's service domain. This is a recursive domain definition, in that the upstream provider in turn may have purchased services from an upstream provider at the next tier, and so on. After the client's traffic leaves this service domain, the ISP ceases to directly, or indirectly, fund the carriage of the client's traffic, and the funding burden passes over to a funding chain linked to the receiver's retail service. For example, when traffic is passed from an ISP client to a client of another provider, the ISP funds the traffic as it transits through the ISP and indirectly funds the cost of carriage through any upstream provider's network. When the traffic leaves the provider's network, to be passed to either a different client, another ISP, or to a peer provider, the sender's ISP ceases to fund the further carriage of the traffic. This scenario is indicated in Figure 6. In other words, these scenarios illustrate the common theme that the retail base of the Internet is not an end-to-end tariff base. The sender of the 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 hop ISP for these costs. The ISP retail pricing structure reflects an implicit division of cost between the two parties, and there is no consequent structural requirement for interprovider financial balancing between the originating ISP and the terminating ISP.

Figure 6: Partial-Path Paired Services



An initial reaction to this partial service model would be to wonder why the Internet works at all, given that no single party funds the carriage of traffic on the complete path from sender to receiver. Surely this would imply that once the traffic had passed beyond the sending ISP's service funded domain the traffic should be discarded as unfunded traffic? The reason why this is not the case is that the receiver implicitly assumes funding responsibility for the traffic at this

handover point, and the second part of the complete carriage path is funded by the receiver. In an abstract sense, the entire set of connectivity paths within the Internet can be viewed as a collection of bilaterally funded path pairs, where the sender funds the initial path component and the receiver funds the second terminating path component. This underscores the original observation that the generally adopted retail model of Internet services is not one of end-to-end service delivery, but instead one of partial path service, with no residual retail price component covering any form of complete path service.

Financial settlement models typically are derived from a different set of initial premises than those described here. The typical starting point is that the retail offering is a comprehensive end-to-end service, and that the originating service provider utilizes the services of other providers to complete the delivery of all components of the retailed service. The originating service provider then undertakes some form of financial settlement with those providers who have undertaken some form of an operational role in providing these service elements. This cost-distributed business structure allows both small and large providers to operate with some degree of financial stability, which in turn allows a competitive open service market to thrive. Through the operation of open competition, the consumer gains the ultimate price and service benefit of cost-efficient retail services.

The characteristics of the Internet environment tend to create a different business environment to that of a balanced cost distribution structure. Here there is a clear delineation between a customer/provider relationship and a peer relationship, with no stable middle ground of a financially settled inter-ISP bilateral relationship. An ISP customer is one that assumes the role of a customer of one or a number of upstream providers, with an associated flow of funding from the customer to the upstream provider, whereas an ISP upstream service provider views the downstream provider as a customer. An ISP peer relationship is where the two ISPs execute a peering arrangement, where traffic is exchanged between the two providers without any consequent financial settlement, and such peering interactions are only stable while both providers perceive some degree of parity in the arrangement; for example, when the two providers present to the peering point Internet domains of approximate equality in market coverage and market share. An ISP may have multiple simultaneous relationships, being a customer in some cases, an upstream provider in others, and a peer in others. In general, the relationships are unique within an ISP pairing, and efforts to support a paired relationship which encompasses elements of both peering and customer/provider pose significant technical and business challenges.

The most natural business outcome of any business environment is for each provider to attempt to optimize its business position. For an ISP, this optimization is not simply a case of a competitive impetus to achieve cost efficiency in the ISP's internal service operation, because the realization of cost efficiencies within the service provider's network does not result in any substantial change in the provider's financial position with respect to upstream costs or peering positioning. The ISP's path toward business optimization includes a strong component of increasing the size and scope of the service provider operation, so that the benefits of providing funded upstream services to customers can be maximized, and non-financially settled peering can be negotiated with other larger providers.

The conclusion drawn is that the most natural business outcome of today's Internet settlement environment is one of aggregation of providers, a factor quite evident in the Internet provider environment at present.

Quality of Service and Financial Settlements

Within today's ISP service model, strong pressure to change the technology base to accommodate more sophisticated settlement structures is not evident. The fundamental observation is that any financial settlement structure is robust only where a retail model exists that is relatively uniform in both its nature and deployment, and encompasses the provision of

services on an end-to-end basis. Where a broad diversity of partial-service retail mechanisms exists within a multiprovider environment, the stability of any form of interprovider financial settlement structure will always be dubious at best.

If paired partial path service models and SKA peering interconnection comfortably match the requirements of the ISP industry today, is this entire financial settlement issue one of simple academic interest? Perhaps the strongest factor driving change here is the shift towards an end-to-end service model associated with the current technology impetus toward support of distinguished Quality of Service (QoS) mechanisms. Where a client signals the requirement for some level of pre-emption or reservation of resources to support an Internet transaction or flow, the signal must be implemented on an end-to-end basis in order for the service request to have any meaning or value. The public Internet business model to support practical use of such QoS technologies will shift to that of the QoS signal initiator undertaking to bear the cost of the entire end-to-end traffic flow associated with the QoS signal.

This is a retail model where the application initiator undertakes to fund the entire cost of data transit associated with the application. This model is analogous to the end-to-end retail models of the telephony, postal, and freight industries. In such a model, the participating agents are compensated for the use of their services through a financial distribution of the original end-to-end revenue, and a logical base for inter-agent financial settlements is the outcome. It is, therefore, the case that meaningful interprovider financial settlements within the Internet industry are highly dependent on the introduction of end-to-end service retail models. These financial settlements are, in turn, dependent on a shift from universal deployment of a best effort service regime with partial path funding to the introduction of layered end-to-end service regimes that feature both end-to-end service-level undertakings and end-to-end tariffs applied to the initiating party.

The number of conditionals in this argument is not insignificant. If QoS technologies are developed that scale to the size of the public Internet, that provide sufficiently robust service models to allow the imposition of service level agreements with service clients, and are standardized such that the QoS service models are consistent across all vendor platforms, then this area of inter-provider settlements will need to change as a consequence. The pressure to change will be emerging market opportunities to introduce interprovider QoS interconnection mechanisms and the associated requirement to introduce end-to-end retail QoS services. The consequence is that there will be pressure to support this with interprovider financial settlements where the originating provider will apportion the revenue gathered from the QoS signal initiator with all other providers that are along the associated end-to-end QoS flow path.

Such an end-to-end QoS settlement model assumes significant proportions that may in themselves impact on the QoS signalling technologies. It is conceivable that each provider along a potential QoS path may need to signal not only their capability of supporting the QoS profile of the potential flow, but also the unit settlement cost that will apply to the flow. The end user may then use this cost feedback to determine whether to proceed with the flow given the indication of total transit costs, or request alternate viable paths in order to choose between alternative provider paths so as to optimize both the cost and the resultant QoS service profile. The technology and business challenges posed by such an end-to-end QoS deployment model are certainly an impressive quantum change from today's best effort Internet.

With this in mind, one potential future is that the public Internet environment will adopt a QoS mediated service model that is capable of supporting a diverse competitive industry through interprovider financial settlements. The alternative is the current uniform best effort environment with no logical role for interprovider settlements, with the associated strong pressures for provider aggregation. The reliance on Internet QoS technologies to achieve not only Internet service outcomes, but also to achieve desired public policy outcomes in terms of competitive pressures, is evident within this perspective. It is unclear whether the current state of emerging QoS technologies and QoS interconnection agreements will be able to mature and

be deployed in time to forge a new chapter in the story of the Internet interconnection environment. The prognosis for this is, however, not good.

Futures

Without the adoption of a settlement regime that supports some form of cost distribution among Internet providers, there are serious structural problems in supporting a diverse and well populated provider industry sector. These problems are exacerbated by the additional observation that the Internet transmission and retail markets both admit significant economies of scale of operation. The combination of these two factors leads to the economic conclusion that the Internet market is not a sustainable open competitive market. Under such circumstances, there is no natural market outcome other than aggregation of providers, leading to the establishment of monopoly positions in the Internet provider space. This aggregation is already well underway, and direction of the Internet market will be forged through the tension between this aggregation pressure and various national and international public policy objectives that relate to the Internet industry.

The problem stated here is not in the installation of transmission infrastructure, nor is it in the retailing of Internet services. The problem faced by the Internet industry is in ensuring that each provider of infrastructure is fairly paid when the infrastructure is used. In essence, the problem is how to distribute the revenue gained from the retail sale of Internet access and services to the providers of carriage infrastructure. While explosive growth has effectively masked these problems for the past decade, after market saturation occurs and growth tapers off, these issues of financial settlement between the various Internet industry players will then shape the future of the entire global ISP industry.

[This article is based in part on material in The ISP Survival Guide, by Geoff Huston, ISBN 0-471-31499-4, published by John Wiley & Sons in 1998. Used with permission.]

Annotated Reading List

The following articles and publications address various aspects of Internet interconnection and peering, and the underlying issues of the economics of Internet carriage.

- Huston, G., "Interconnection, Peering and Settlements-Part I," The Internet Protocol Journal, Volume 2, Number 1, March 1999. The first part of this article.
- Huston, G., ISP Survival Guide, ISBN 0-471-31499-4, John Wiley & Sons, November 1998. *A more comprehensive view of the technology, business and strategy behind the Internet service sector.*
- Halabi, B., Internet Routing Architectures, ISBN 1-56205-652-2, Cisco Press, April 1997. *An excellent information resource on how to configure BGP to express policies for interconnecting networks.*
- Frieden, R., "Without Public Peer: The potential Regulatory and Universal Service Consequences of Internet Balkanization," Virginia Journal of Law and Technology, ISSN 1522-1687, Vol. 3, Sept. 1998. http://vjolt.student.virginia.edu/graphics/vol3/vol3_art8.html . *A good briefing paper from an economic perspective on interconnection issues, with particular attention to the domestic situation in the United States.*
- Cukier, K., "Peering and Fearing: ISP Interconnection and Regulatory Issues," Presented paper at the Harvard Information Infrastructure Project Conference on the Impact of the Internet on Communication Policy, December 3-5 1997. Conference program is at: <http://ksgwww.harvard.edu/iip/iicompol/agenda.html> The Cukier paper is at: <http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html>
- Shapiro, C., Varian, H., Information Rules: A Strategic Guide to the Information Economy, ISBN 087584863X, Harvard Business School Press, November 1998. *A broader look at the Internet from an economic perspective, looking at both content and service provider economics.*
- Varian, H., "The Information Economy-The Economics of the Internet, Information Goods, Intellectual Property and Related Issues," <http://www.sims.berkeley.edu/resources/infoecon/> *This is a collection of references to other online resources, and is a useful starting point for further reading on this topic.*

GEOFF HUSTON holds a B.Sc and a M.Sc from the Australian National University. He has been closely involved with the development of the Internet for the past decade. He was responsible for the initial build of the Internet within the Australian academic and research sector. Huston is currently the Chief Technologist in the Internet area for Telstra. He is also an active member of the IETF, and is a member of the Internet Society Board of Trustees. He is author of *The ISP Survival Guide*, and coauthor of *Quality of Service: Delivering QoS on the Internet and in Corporate Networks*, a collaboration with Paul Ferguson. Both books are published by John Wiley & Sons. E-mail: gjh@telstra.net
