

Interconnection, Peering and Settlements

Geoff Huston

January, 1999

This article is based in part on material drawn from "ISP Survival Guide, by Geoff Huston, ISBN 201-3-45567-9, published by John Wiley & Sons, 1998. Used with permission.

Technology and business models share a common evolution within the Internet. To enable deployment of the technology within a service environment there is also the need to create a robust and stable business model. This tied destiny of technology and business factors is perhaps most apparent within the area of the interconnection of Internet Service Providers (ISPs). Here there is both an interaction at a level of technology, in terms of routing signaling and traffic flows, and also an interaction of business models, in terms of a negotiation of benefit and cost in undertaking the interconnection. In this article we will examine this environment in some detail, looking closely at the interaction between the capabilities of the technical protocols, their translation into engineering deployment and the consequent business imperatives such environments create.

It is necessary to commence this examination of the public Internet with the observation that the Internet is not, and never has been, a single network. The Internet is a collection of interconnected component networks that share a common addressing structure, a common view of routing and traffic flow, and a common view of a naming system. This interconnection environment spans a highly diverse set of more than 50,000 component networks, and this number continues, inexorably, to grow and grow. One of the significant aspects of this environment is the competitive Internet service industry, where many thousands of enterprises, both small and large, compete for market share at a regional, national and international level.

Underneath the veneer of a highly competitive Internet service market is a somewhat different environment, in which every Internet Service Provider (ISP) network must interoperate with neighboring Internet networks in order to produce a delivered service outcome of comprehensive connectivity and end- to-end service. No ISP can operate in complete isolation from others while still participating in offering public Internet services, and therefore, every ISP must not only coexist with other ISPs but also must operate in cooperation with other ISPs.

In this paper we will examine both the technical and business aspects which surround this ISP interaction, commonly referred to as "interconnection, peering and settlements". We will examine the business motivation for interconnection structures, then examine the technical architectures of such environments. In the second part of this paper we will commence with an examination of the business relationships that arise between ISPs in the public Internet space, and then examine a number of broader issues that will shape the near term future of this environment.

1. Interconnection: Retailing, Reselling, and Wholesaling

To provide some motivation for this issue of ISP interconnection it is first appropriate to examine the nature of the environment. The regulatory framework which defined the traditional structure of other communications enterprises such as telephony or postal services was largely absent in the evolution of the Internet service industry. The resultant service industry for the Internet is most accurately

characterized as an outcome of business and technology interaction, rather than a planned outcome of some regulatory process. In this section we will examine this interaction between business and technology within the ISP environment.

A natural outcome of the Internet model is that the effective control of the retail service environment rests with a network client of an access service rather than with the access service provider, as such a client of an ISP access service has the discretionary ability to resell the access service to third-party clients. In this environment, reselling and wholesaling are very natural developments within the ISP activity sector, with or without the explicit concurrence of the provider ISP. The provider ISP may see this reselling as an additional channel to market for its own Internet carriage services, and may adopt a positive stance by actively encouraging resellers into the market as a means of overall market stimulus, while tapping into the marketing, sales, and support resources of these reselling entities to continue to drive the volumes of the underlying Internet carriage service portfolio. The low barriers to entry to the wholesale market provide a means of increasing the scope of the operation, as, to lift business cash-flow levels, the business enters into wholesale agreements that effectively resell the carriage components of the operation without the bundling of other services normally associated with the retail operation. This process allows the ISP to gain higher volumes of carriage capacity, that in turn allow the ISP to gain access to lower unit costs of carriage.




Figure 1. ISP roles and relationships.

Given that a retail operation can readily become a wholesale provider to third party resellers at the effective discretion of the original retail client, is a wholesale transit ISP restricted from undertaking retail operations? Again, there is no such natural restriction from a technical or business perspective. An Internet carriage service is a commodity service that does not allow for a significant level of intrinsic product discrimination. The relative low level of value added by a wholesale service operation implies a low unit rate of financial return for that operation. This low unit rate of financial return, together with an inability to effectively competitively discriminate the wholesale product, induces a wholesale provider into

the retail sector as a means of improving the financial performance of the service operation. The overall result is that many ISPs operate both as a client and as a provider. Few, if any, reasonable technical-based characterizations draw a clear and unambiguous distinction between a client and service provider when access services to networks are considered. A campus network may be a client of one or more service providers, while the network is also a service provider to campus users. Indeed most networks in a similar situation take on the dual role of client and provider, and the ability to resell an access service can extend to almost arbitrary depths of the reselling hierarchy. From this technical perspective, very few natural divisions of the market support a stable segmentation into exclusively wholesale and exclusively retail market sectors. The overall structure of roles is indicated in Figure 1.

The resultant business environment is one characterized by a reasonable degree of fluidity, in which no clear delineation of relative roles or markets exists. The ISP market environment is, therefore, one of competitive market forces in which each ISP tends to create a retail market presence. However, no ISP can operate in isolation. Each client has the expectation of universal and comprehensive reachability, such that any client of any other Internet ISP can reach the client, and the client can reach a client of any other ISP. The client of an ISP is not undertaking a service contract that limits connectivity only to other clients of the same ISP. As no provider can claim ubiquity of access, every provider relies on every other provider to complete the user-provided picture of comprehensive connectivity. Because of this dependent relationship, an individual provider's effort to provide substantially superior service quality may have little overall impact on the totality of client-delivered service quality. In a best effort public Internet the service quality becomes something that can be impacted negatively by poor local engineering but cannot be uniformly improved beyond the quality provided by the network's peers, and their peers in turn. Internet wholesale carriage services in such an environment are constrained to be a commodity service, in which scant opportunity exists for service-based differentiation. In the absence of service quality as an effective service discriminator, the wholesale activity becomes a price-based service with low levels of added value, or in other words a commodity market.

The implication in terms of ISP positioning is that the retail operation, rather than the wholesale activity, is the major area where the ISP can provide discriminating service quality. Within the retail operation, the ISP can offer a wide variety of services with a set of associated service levels, and base a market positioning on factors other than commodity carriage pricing.

Accordingly, the environment of interconnection between ISPs does not break down into a well ordered model of a set of wholesale carriage providers and associated retail service providers. The environment currently is one with a wide diversity of retail-oriented providers, where each provider may operate both as a retail service operator, and a wholesale carriage provider to other retailers.

1.1 Peer or Client?

One of the significant issues that arises here is whether an objective determination can be made of whether an ISP is a peer to, or a client of, another ISP. This is a critical question, as, if a completely objective determination cannot be readily made, the question then becomes one of who is responsible for making a subjective determination, and on what basis.

This question is an inevitable outcome of the reselling environment, where the reseller starts to make multiple upstream service contracts, with a growing number of downstream clients of the reselling service. At this point, the business profile of the original reseller is little distinguished from that of the original provider. The original reseller sees no unique value being offered by the original upstream provider and may conclude that it is in fact adding value to the original upstream provider by offering the upstream provider high volume carriage and close access to the reseller's client base. From the perspective of the original reseller, the roles have changed, and the reseller now perceives itself as a peer ISP to the original upstream ISP provider.

This assertion of role reversal is perhaps most significant when the generic interconnection environment is one of zero sum financial settlement, in which the successful assertion by a client of a change from client to peer status results in the dropping of client service revenue without any net change in the cost

base of the provider's operation. The party making the successful assertion of peer interconnection sees the opposite, with an immediate drop in the cost of the ISP operation with no net revenue change.

The traditional public regulatory resolution of such matters has been through an administrative process of "licensed" communications service providers, who become peer entities through a process of administrative fiat. In this model, an ISP would become a licensed service provider through the payment of license fees to a communications regulatory body. The license then allows the service enterprise access to interconnection arrangements with other licensed providers. The determination of peer or client is now quite simple: a client is an entity that operates without a such a carrier license, and a peer is one that has been granted such an instrument. However, such regulated environments are quite artificial in their delineation of the entities that operate within a market, and this regulatory process often acts as a strong disincentive to large-scale private investment, thereby placing the burden of underwriting the funding of service industries into the public sector. The regulatory environment is changing worldwide to shift the burden of communications infrastructure investment from the public sector, or from a uniquely positioned small segment of the private sector, to an environment that encourages widespread private investment. The Internet industry is at the leading edge of this trend, and the ISP domain typically operates within a deregulated valued-added communications service provider regulatory environment. Individual licenses are replaced with generic class licenses or similar deregulated structures in which formal applications or payments of license fees to operate in this domain are unnecessary. In such deregulated environments no authoritative external entity makes the decision as to whether the relationship between two ISPs is that of a provider and client or that of peers.

If no public regulatory body wants to make such a determination, is there a comparable industry body that can undertake such a role? The early attempts of the Commercial Internet eXchange (CIX) arrangements in the United States in the early 1990's were based on a description of the infrastructure of each party, in which acknowledgments of peer capability were based on the operation of a national transit infrastructure of a minimum specified capability. This specification of peering within the CIX was subsequently modified so that CIX peer status for an ISP was simply based on payment of the CIX Association membership fee.

This CIX model was not one that intrinsically admitted bilateral peer relationships. The relationship was an multilateral one, in which each ISP executed a single agreement with the CIX Association and then effectively had the ability to peer with all other association member networks. The consequence of this multilateral arrangements is that the peering settlements can be regarded as an instance of zero sum financial settlement peering, using a single threshold pricing structure.

Other industry models use a functional peer specification. For example, if the ISP attaches to a nominated physical exchange structure, then the ISP is in a position to open bilateral negotiations with any other ISP also directly attaches to the exchange structure. This model is inherently more flexible, as the bilateral exchange structure enables each represented ISP to make their own determination of whether to agree to a peer relationship or not with any other co-located ISP. This model also enables each bilateral peer arrangement to be executed individually, admitting the possibility of a wider diversity of financial settlement arrangements.

The bottom line is that a true peer relationship is based on the supposition that either party can terminate the interconnection relationship and that the other party does not consider such an action a competitively hostile act. If one party has a high reliance on the interconnection arrangement and the other does not, then the most stable business outcome is that this reliance is expressed in terms of a service contract with the other party, and a provider/client relationship is established. If a balance of mutual requirement exists between both parties, then a stable basis for a peer interconnection relationship also exists. Such a statement has no intrinsic metrics that allow the requirements to be quantified. Peering in such an environment is best expressed as the balance of perceptions, in which each party perceives an acceptable approximation of equal benefit in the interconnection relationship in their own terms.

This conclusion leads to the various tiers of accepted peering that are evident in the Internet today. Local ISPs see a rationale to view local competing ISPs as peers, and they still admit the need to purchase

trunk transit services from one or more upstream ISPs under terms of a client contract with the trunk provider ISP. Trunk ISPs see an acceptable rationale in peering with ISPs with a similar role profile in trunk transit but perceive an inequality of relationship with local ISPs. The conclusion drawn here is that the structure of the Internet is one where there is a strong business pressure to create a rich mesh of interconnection at various levels, and the architecture of interconnection structures is an important feature of the overall architecture of the public Internet.

2. Interconnection Architectures: Exchanges and NAPs

One of the physical properties of electromagnetic propagation is that the power required to transmit an electromagnetic pulse over a distance varies in accordance with this distance. The shorter the distance between the transmitter and the receiver, the lower the transmission power budget required; *closer is cheaper*.

This statement holds true not only for electrical power budgets but also for data protocol efficiency. Minimizing the delay between the sender and receiver allows the protocol to operate faster and operate more efficiently as well; *closer is faster*, and *closer is more efficient*.

These observations imply that distinct and measurable advantages are gained by localizing data traffic, that is by ensuring that the physical path traversed by the packets passed between the sender and the receiver is kept as physically short as possible. These advantages are realizable in terms of service performance, efficiency, and service cost. How then are such considerations of locality factored into the structure of the Internet?

2.1 The Exchange Model

A strictly hierarchical model of Internet structure is one in which a small number of global ISP transit operators is at the "top"; a second tier is of national ISP operators; and a third tier consists of local ISPs. At each tier the ISPs are clients of the tier above, as shown in Figure 2. If this hierarchical model were strictly adhered to, traffic between two local ISPs would be forced to transit a national ISP, and traffic between two national ISPs would transit a global ISP, even if both national ISPs operated within the same country. In the worst case, traffic between two local ISPs would need to transit a national ISP, and then a global ISP from one hierarchy, then a second global ISP, and a second national ISP from an adjacent hierarchy in order to reach the other local ISP. If the two global providers interconnect at a remote location, the transit path of the traffic between these two local ISPs could be very long indeed.

As noted above, such extended paths are inefficient and costly, and such costs are ultimately part of the cost component of the price of Internet access. In an open competitive market, strong pressure always is applied to reduce costs. Within an hierarchical ISP environment, strong pressure is applied for the two national providers, who operate within the same market domain, to modify this strict hierarchy and directly interconnect their networks. Such a local interconnection allows the two networks to service their mutual connectivity requirements without payment of transit costs to their respective global transit ISP providers. At the local level is a similar incentive for the local ISPs to reduce their cost base, and a local interconnection with other local ISPs would allow local traffic to be exchanged without the payment of transit costs to the respective transit providers.




Figure 2. A purely hierarchical structure for the Internet.

Although constructing a general interconnection regime based on point-to-point bilateral connections is possible, this approach does not exhibit good scaling properties. Between N providers, who want to interconnect, the outcome of such a model of single interconnecting circuits is $(N^2 - N) / 2$ circuits and $(N^2 - N) / 2$ routing interconnections, as indicated in Figure 3. Given that interconnections exhibit the greatest leverage within geographical local situations, simplifying this picture within the structure of a local exchange is possible. In this scenario each provider draws a single circuit to the local exchange and then executes interconnections at this exchange location. Between N providers who want to interconnect, the same functionality of complete interconnection can be constructed using only N point-to-point circuits.




Figure 3. Fully meshed peering.

2.1.1 The Exchange Router

One model of an exchange is to build the exchange itself as a router, as indicated in Figure 4. Each provider's circuit terminates on the exchange router, and each provider's routing system peers with the routing process on the exchange router. This structure also simplifies the routing configuration, so that full interconnection of N providers is effected with N routing peer sessions. This simplification does allow greater levels of scaling in the interconnection architecture.

However, the exchange router model does become an active component of the interconnect peering policy environment. In effect, each provider must execute a multilateral interconnection peering with all of the other connected providers. Selectively interconnecting with a subset of the providers present at such a router-based exchange is not easily achieved. In addition, this type of exchange must execute its own routing policy. When two or more providers are advertising a route to the same destination, the exchange router must execute a policy decision as to which provider's route is loaded in the router's forwarding table, making a policy choice of transit provider on behalf of all other exchange-connected providers.

Because the exchange is now an active policy element in the interconnection environment, the exchange is no longer completely neutral to all participants. This imposition on the providers may be seen as unacceptable, in that some of their ability to devise and execute an external transit policy is usurped by the exchange operator's policies.




Figure 4. An exchange router.

Typically, providers have a higher expectation of flexibility of policy determination from exchange structures than this base level of functionality as provided by an exchange router. Providers want the flexibility to execute interconnections on a bilateral basis at the exchange, and to make policy decisions as to which provider to prefer when the same destination is advertised by multiple providers. They require the exchange to be neutral with respect to such individual routing policy decisions.

2.1.2 The Exchange Switch

The modification to the inter-provider exchange structure is to use a local layer 2 switch (or LAN) as the exchange element. In this model a participating provider draws a circuit to the exchange and locates a dedicated router on the exchange LAN. This structure is indicated in Figure 5. Each provider executes a bilateral peering agreement with another provider by initiating a router peering session with the other party's router. When the same network destination is advertised by multiple peers, the provider can execute a policy-based preference as to which peer's route will be loaded in the local forwarding table. Such a structure preserves the cost efficiency of using N circuits to effect interconnection at the N provider exchange, while admitting the important policy flexibility provided by up to $(N^2 - N) / 2$ potential routing peer sessions.




Figure 5. An exchange LAN.

Early inter-provider exchanges were based on an Ethernet LAN as the common interconnection element. This physical structure was simple, and not all that robust under the pressures of growth as the LAN become congested. Subsequent refinements to the model have included the use of Ethernet switches as a higher capacity LAN, and the use of FDDI rings, switched FDDI hubs, fast Ethernet hubs, and switched fast Ethernet hubs. Exchanges are very high traffic concentration points, and the desire to manage ever higher traffic volumes has led to the adoption of gigabit Ethernet switches as the current evolutionary technology step within such exchanges.

The model of the exchange co-location accommodates a model of diversity of access media, in which the provider's co-located router undertakes the media translation between the access link protocol and the common exchange protocol.

The local traffic exchange hub does represent a critical point of failure within the local Internet topology. Accordingly, the exchange should be engineered in the most resilient fashion possible, using standards associated with a premium quality data center. This structure may include multiple power utility connections, uninterruptible power supplies, multiple trunk fiber connections, and excellent site security measures.

The exchange should operate neutrally with respect to every participating ISP, with the interests of all the exchange clients in mind. Therefore, exchange facilities, which are operated by an entity that is not also a local or trunk ISP, enjoy higher levels of trust from the clients of the exchange.

There are also some drawbacks to an exchange, and a commonly cited example is that of imposed transit. If an exchange participant directs a default route to another exchange router, then, in the absence of defensive mechanisms the target router will carry the imposed transit traffic even when there is no routing peering or business agreement between the two ISPs. Exchange located routers do require careful configuration management to ensure that route peering and associated transit traffic matches the currently executed interconnection agreements.

2.1.3 Distributed Exchanges

Distributed exchange models also have been deployed in various locations. This deployment can be as simple as a metropolitan FDDI extension, in which the exchange comes to the provider's location rather than the reverse, as indicated in Figure 6. Other models that use an ATM-based switching fabric, using LAN Emulation (LANE) to mimic the layer 2 exchange switch functionality, also have been deployed. Distributed exchange models attempt to address the significant cost of operating a single co-location environment with a high degree of resilience and security, but do so at a cost of enforcing the use of a uniform access technology between every distributed exchange participant.




Figure 6. A distributed exchange.

However, the major challenge of such distributed models is that of switching speed. Switching requires some element of contention resolution, in which two ingress data elements that are addressed to a common egress path require the switch to detect the resource contention and then resolve it by serializing the egress. Switching, therefore, requires signaling, in which the switching element must inform the ingress element of switch contention. To increase the throughput of the switch, the latency of this signaling must be reduced. This dictates that increased switching speed has the corollary of requiring the switch to exist within the confines of a single location, if exchange performance is a paramount concern.

Besides speed, we must consider the cost shift. In a distributed exchange model, the exchange operator operates the set of access circuits that form the distributed exchange. This process increases costs to providers, while it prevents the provider from using a specific access technology that matches their business requirements of cost and supportable traffic volume. Not surprisingly, to date the most prevalent form of exchange remains the third-party hosted co-location model. This model admits a high degree of diversity in access technologies, while still providing the substrate of an interconnection environment that can operate at high speed and therefore manage high traffic volumes.

2.1.4 Other Exchange-Located Services

The co-location environment is often broadened to include other functions, in addition to a pure routing and traffic exchange role. For a high-volume content provider, the exchange location offers minimal transit distance to a large user population distributed across multiple local service providers, as well as allowing the content provider to exercise a choice in selecting a non-local transit provider.

The exchange operator can also add value to the exchange environment by providing additional functions and services, as well as terminating providers' routers and large-volume content services. The

exchange location within the overall network topology is an ideal location for hosting multicast services, as the location is quite optimal in terms of multicast carriage efficiency. Similarly, Usenet trunk feed systems can exploit the local hub created by the exchange. The overall architecture of a co-location environment that permits value-added services, which can productively use the unique environment created at an exchange, is indicated in Figure 7.




Figure 7. Exchange-located service platforms.

2.2 Network Access Points

The role of the exchange was broadened with the introduction of the Network Access Point (NAP) in the NSF-proposed post-NSFNET architecture of 1995.

The NAP was seen to undertake two roles: the role of an exchange provider between regional ISPs who want to execute bilateral peering arrangements and the role of a transit purchase venue, in which regional ISPs could execute purchase agreements with one or more of a set of trunk carriage ISPs also connected at the NAP. The access point concept was intended to describe access to the trunk transit service. This mixed role of both local exchange and transit operations leads to considerable operational complexity, in terms of the transit providers being able to execute a clear business agreement. What is the bandwidth of the purchased service in terms of requirements for trunk transit, versus the access requirements for exchange traffic? If a local ISP purchases a transit service at one of the NAPs, does that imply that the trunk provider is then obligated to present all the ISP's routes at remote NAPs as a peer? How can a trunk provider distinguish between traffic presented to it on behalf of a remote client versus traffic presented to it by a local service client?

We also should consider the issue that the quality of the purchased transit service is colored by the quality of the service provided by the NAP operator. Although the quality of the transit provider's network

may remain constant, and the quality of the local ISP's network and ISP's NAP access circuit may be acceptable, the quality of the transit service may be negatively impacted by the quality of the NAP transit itself.

One common solution is to use the NAP co-location facility to execute transit purchase agreements and then use so-called backdoor connections for the transit service provision role. This usage restricts the NAP exchange network to a theoretically simpler local exchange roles. Such a configuration is illustrated in Figure 8.




Figure 8. Peering and transit purchase.

2.3 Exchange Business Models

For the ISP industry, a number of attributes are considered highly desirable for an exchange facility. The common model of an Internet exchange includes many, if not all, of the following elements:

- Operated by a neutral party who is not an ISP (to ensure fairness and neutrality in the operation of the exchange)
- Constructed in a robust and secure fashion
- Located in areas of high density of Internet market space
- Able to scale in size
- Operate in a fiscally sound and stable business fashion

A continuing concern exists about the performance of exchanges and the consequent issue of quality of services that traverse the exchange. Many of these concerns stem from an exchange business model that may not be adequately robust under pressures of growth from participating ISPs.

The exchange business models typically are based on a flat-fee structure. The most basic model uses a fee structure based on the number of rack units used by the ISP to co-locate equipment at the exchange. When an exchange participant increases the amount of traffic presented over an access interface, under a flat-fee structure, this increased level of traffic is not accompanied by any increase in exchange fees. However, the greater traffic volumes do imply that the exchange itself is faced with a greater traffic load. This greater load places pressure on the exchange operator to deploy further equipment to augment the switching capacity, without any corresponding increase in revenue levels to the operator.

For an exchange operator to base tariffs on the access bandwidths is not altogether feasible, given that such access facilities are leased by the participating ISPs and the access bandwidth may not be known to the exchange operator. Nor is using a traffic-based funding model possible given that an exchange operator should refrain from monitoring individual ISP traffic across the exchange, given the unique position of the exchange operator. Accordingly, the exchange operator has to devise a fiscally prudent tariff structure at the outset that enables the exchange operator to accommodate large-scale traffic growth, while maintaining the highest possible traffic throughput levels.

Alternatively there are business models in which the exchange is structured as a cooperative entity between a number of ISPs. In these models the exchange is a nonprofit common asset of the cooperative body. This model is widely used, but also one that is prone to the economic condition of the Tragedy of the Commons. It is in everyone's interest to maximize their exploitation of the exchange, while no single member wants to underwrite the financial responsibility for ensuring that the quality of the exchange itself is maintained.

The conclusion that can be drawn is that the exchange is an important component of Internet infrastructure, and the quality of the exchange is of paramount importance if it is to be of any relevance to ISPs. Using an independent exchange operator whose income is derived from the utility of the exchange is one way of ensuring that the exchange is managed proficiently and that the service quality is maintained for the ISP clients of the exchange.

2.4 A Structure for Connectivity

Enhancing the Internet infrastructure is quantified by the following objectives:

- Extension of reachability.
- Enhancement of policy matching by ISPs.
- Localization of connectivity.
- Backup arrangements for reliability of operation.
- Increasing capacity of connectivity.
- Enhanced operational stability.
- Creation of a rational structure of the connection environment to allow scalable structuring of the address and routing space in order to accommodate orderly growth.

We have reached a critical point within the evolution of the Internet. The natural reaction of the various network service entities in response to the increasing number of ISPs will be to increase the complexity of the interconnection structure to preserve various direct connectivity requirements. Today, we are in the uncomfortable position of increasingly complex inter-provider connectivity environment which is stressing the capability of available technologies and equipment. The inability to reach stable cost distribution models in a transit arrangement creates an environment in which each ISP attempts to optimize its position by undertaking as many direct 1:1 connections with peer ISPs as it possibly can. Some of these connections are managed via the exchange structure. Many more are implemented as direct links between the two entities. Given the relative crudity of the inter-AS routing policy tools that we

use today, this structure must be a source of some considerable concern. The result of a combination of an increasingly complex mesh of inter-AS connections, together with very poor tools to manage the resultant routing space, is an increase in the overall instability of the Internet environment. In terms of meeting critical immediate objectives, however, such dire general predictions do not act as an effective deterrent to these actions.

The result is a situation in which the inter-AS space is the critical component of the Internet. This space can be viewed correctly as the demilitarized zone within the politics of today's ISP-based Internet. In the absence of any coherent policy, or even a commonly accepted set of practices, the lack of administration of this space is a source of paramount concern.

3. Interaction Financials: Peering and Settlements

We have 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. In continuing our examination of the technology and business considerations that are significant within the subject of Internet Service Provider (ISP) interconnection, we will now focus on the topic from a predominately business perspective.

Any large multi-provider 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 each provider can be compensated for their costs in an equitable structure of inter-provider 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 provider. The customer's original fare, paid to the original service provider, is distributed by the service provider 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 inter-provider 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 inter-provider cost distribution models and then look at the business models currently used in the inter-provider Internet environment. This area commonly is termed financial settlement, a term the Internet has borrowed from the telephony industry.

3.1 The Currency of Interconnection

What exactly is being exchanged between two ISP's 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 9, and a bilateral routing-mediated flow occurs only when routes are passed in both directions.




Figure 9. Routing and traffic flows.

Within the routing environment of an ISP there are a number of 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 BGP session with the client, or part of an ISP pool of addresses that are dynamically assigned to the client as part of the dial-up session.

Internal ISP routes fall into a number of 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, POP mail servers, and game servers. Some routes correspond to ISP-operated client services that require Internet-wide access, such as DNS forwarders and SMTP relay hosts. Lastly are internal services with no visibility outside the ISP network, such as 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 client 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 10.




Figure 10. External routing interaction.

The implicit outcome of this routing policy structure is that the ISP does not act in a transit role to peer ISPs and does not permit 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.

3.2 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 inter-provider 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, inter-provider peering takes place within one of three general models:

3.2.1 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, a number of 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 multi-party transit arrangements. The majority of 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, as 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 between the two parties. As 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 inter-provider 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.

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.

3.2.2 Sender Keep All

The second model, rarely used in telephony interconnection, is that of Sender Keep All (SKA), in which each service provider invoices their 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.

3.2.3 Transit fees

The third model is that of transit fees, in which the one party invoices the other party for services provided. For example, this arrangement is commonly used as the basis of the long-distance provider local access provider interconnection arrangements. Again, this 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.

3.2.4 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.

3.3 Internet Considerations

A number of critical differences exist between the telephony models of interconnection and the Internet environment, which 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 UDP, and various forms of tunneling 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 pre-determined.

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 honored, 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 also may 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.

4. 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 honor the service contract to the downstream ISP customer.

However, given two ISP's 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, the two animals 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.

4.1 Packet Cost Accounting

One potential accounting model is based on the observation that a packet incurs cost as it is passed 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 11.




Figure 11. Financial inter-provider settlement via packet cost accounting.

As with all strawman models, this one has a number of 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 only generates revenue to the ISP 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 inter-provider 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 only accept packets in which a specific route is maintained for the network. More critical is the issue that the mechanism is open to abuse. Packets, which 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 by using a tariff for domestic transit and a second for international transit.

4.2 TCP Session Accounting

These traffic-based metrics do exhibit some weaknesses because of their inability to resist abuse and likelihood of exacting an inter-provider 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 inter-provider 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 inter-provider boundary. The outcome of the routing configuration may be an asymmetric traffic path, so that a single inter-provider 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: Considering 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 where the packet sender pays or one 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.

5. 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 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 towards 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.

5.1 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 a number of network domains within the United States, and then passed back across the Pacific. Not only is this inefficient in terms of resource utilization, 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 such 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 due to the high cost of operating such 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, nor in the best interests of 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 act to ensure local interconnectivity between such ISPs. A consequent priority is to reach acceptable ISP peering arrangements.

5.2 Sender Keep All

Sender Keep 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 such cases, any SKA peering is likely to relate to only domestic traffic, and international transit would either be provided by a separate agreement or provided 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 criteria for a stable SKA peering structure is perceived equality in the peering relationship. This can be achieved in a number of 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, which is destined to peering ISPs, from upstream transit providers. The business model here is that each client of an ISP is 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 their customers at their 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 corresponded 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 12.




Figure 12. SKA peering using local cells.

The probable outcome of widespread use of SKA interconnections is a generalized ISP domain along the lines of Figure 13. 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, as 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 13. 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.

5.3 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, as 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 which 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. Where there is a rich mesh of connectivity, there is a strong financial incentive for each party to adjust the routing parameters to match the lowest financial expenditure by using restricted route advertisements with the greatest levels of revenue by using a local preference for received routes, with the highest preference for client-advertised routes and the next level of preference for financial settled peer advertised routes. Such settings of the routing system may not necessarily correspond to the optimal traffic path in network engineering terms, nor will these settings necessarily result in a highly stable routing and traffic configuration.

Of far 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 which 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, as they introduce significant financial risks to the ISP into the interconnection environment.

6. 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 widespread 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 sub-domain 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. Once 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 is indicated in Figure 14.




Figure 14. Partial path paired services.

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 inter-provider financial balancing between the originating ISP and the terminating ISP.

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, that 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 who 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 there 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, such as where 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 with 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 their 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, as 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 towards 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.

7. 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 multi-provider environment, the stability of any form of inter-provider 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 towards support of distinguished quality of service (QoS) mechanisms. Where a client signals the requirement for some level of preemption 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 inter-provider financial settlements within the Internet industry are highly dependant on the introduction of end-to-end service retail models. This in turn is dependant 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 inter-provider 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 inter-provider 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 signaling 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 to 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 inter-provider financial settlements. The alternative is the current uniform best effort environment with no logical role for inter-provider 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.

8. 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, or 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, once market saturation occurs and growth tapers off, these issues of financial settlement between the various Internet industry players will then shape the future of entire global ISP industry.

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc from the Australian National University. He has been closely involved within the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial deployment of the Internet within Australia as the program manager for the academic and research network. Geoff Huston is currently the Chief Technologist for Telstra's Internet Products area. He is also an active member of the Internet Engineering Task Force, and is currently Secretary for the Board of Trustees for the Internet Society. He is co-author of *Quality of Service; Delivering QoS on the Internet and in Corporate Networks*, published by John Wiley & Sons, ISBN 0-471-24358-2, a collaboration with Paul Ferguson. He is also the author of *ISP Survival Guide*, published by John Wiley & Sons, ISBN 201-3-45567-9. Email: gih@telstra.net