



---

## Interconnection, Peering and Settlements - Part I

March 1999

Geoff Huston

Technology and business models share a common evolution within the Internet. To enable deployment of the technology within a service environment, a robust and stable business model also needs to be created. 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 an interaction at a level of technology, in terms of routing signalling and traffic flows, and also an interaction of business models, in terms of a negotiation of benefit and cost in undertaking the interconnection. This article examines 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 that 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 ISP network must interoperate with neighbouring 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 offering public Internet services, and therefore, every ISP not only must coexist with other ISPs but also must operate in cooperation with other ISPs.

This article examines both the technical and business aspects that surround this ISP interaction, commonly referred to as "interconnection, peering, and settlements." It examines the business motivation for interconnection structures, and then the technical architectures of such environments. The second part looks at the business relationships that arise between ISPs in the public Internet space, and then examines numerous broader issues that will shape the near-term future of this environment.

### Interconnection: Retailing, Reselling, and Wholesaling

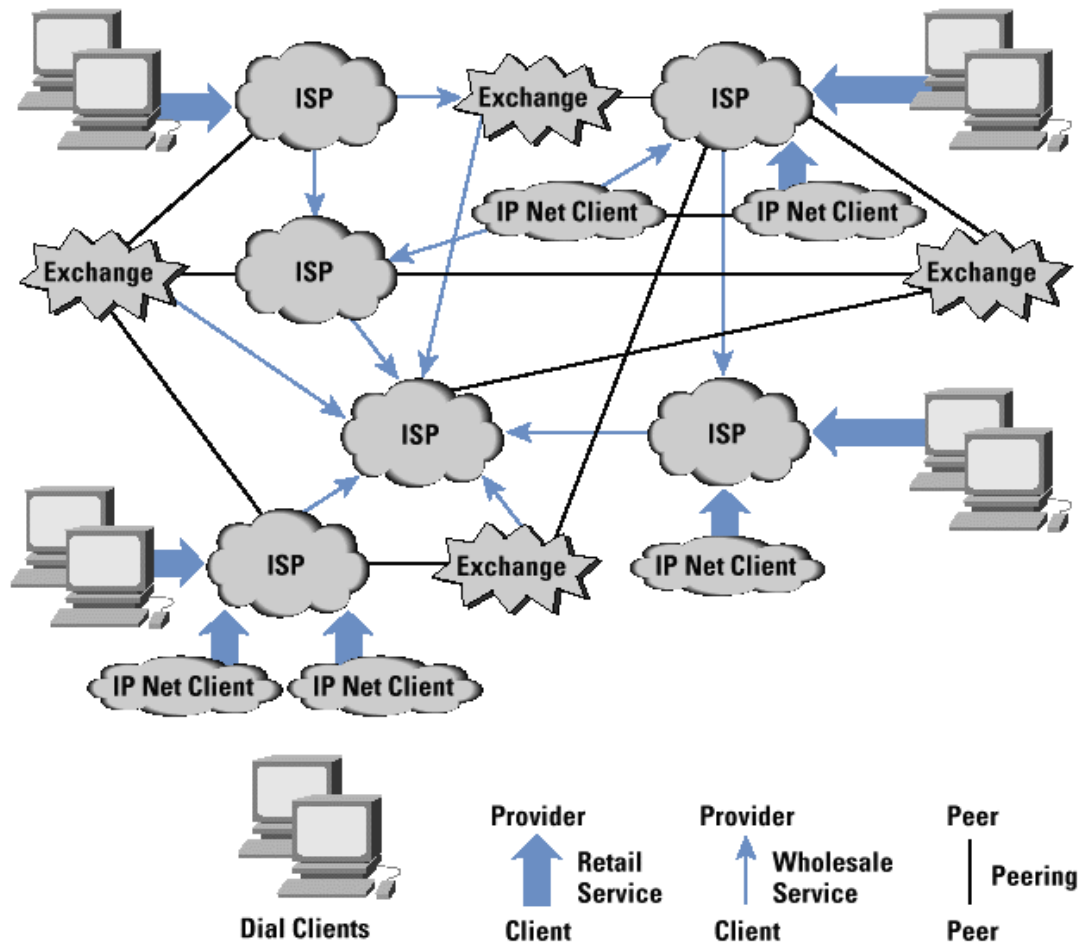
To provide some motivation for this issue of ISP interconnection, it is first appropriate to look at the nature of the environment. The regulatory framework that defined the traditional structure of other communications enterprises such as telephony or postal services was largely absent in the evolution 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. This section examines 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, because 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.

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 relatively 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 competitively discriminate the wholesale product effectively, 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 clients and as providers. 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.

Figure 1: ISP Roles and Relationships



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 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. Because 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 in which 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.

## Peer or Client?

One of the significant issues that arises here is: Can an objective determination be made of whether an ISP is a peer to, or a client of, another ISP? This is a critical question, because 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 becomes 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 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 1990s 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 a 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 arrangement 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 attached to the exchange structure. This model is inherently more flexible, as the bilateral exchange structure enables each represented ISP to make its own determination of whether to agree to a peer relationship or not with any other colocated 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 its 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 viewing 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 in which 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. Physical 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?

## 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 is strictly adhered to, traffic between two local ISPs is forced to transit a national ISP, and traffic between two national ISPs transits a global ISP even if both national ISPs operate within the same country. In the worst case, traffic between two local ISPs needs to transit a national ISP, 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.


Figure 2: A Purely Hierarchical Structure for the Internet



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 a 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.

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




### 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 becomes 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 the base level of functionality that is 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.

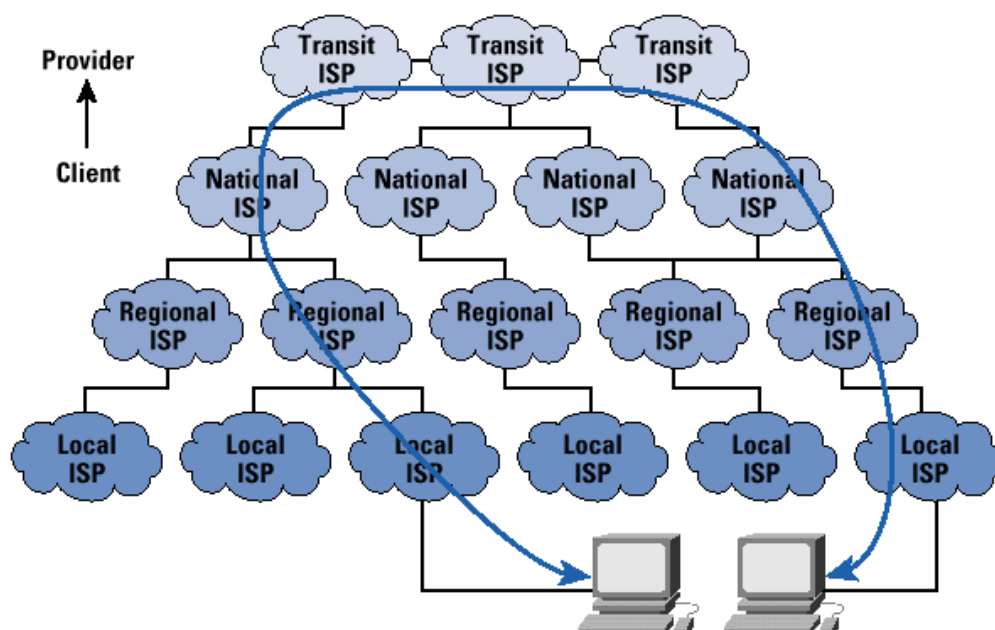
### The Exchange Switch

The modification to the interprovider 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, as shown 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.

Early interprovider 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 became congested.



Figure 5: An Exchange LAN



Subsequent refinements to the model have included the use of Ethernet switches as a higher capacity LAN, and the use of Fiber Distributed Data Interface (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 collocation accommodates a model of diversity of access media, in which the provider's colocated 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 centre. This structure may include multiple power utility connections, uninterruptible power supplies, multiple trunk fibre 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. Thus, 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 carries 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.

## 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 also have been deployed using LAN Emulation (LANE) to mimic the Layer 2 exchange switch functionality. Distributed exchange models attempt to address the significant cost of operating a single colocation 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 signalling, in which the switching element must inform the ingress element of switch contention. To increase the throughput of the switch, the latency of this signalling must be reduced. The dictates of increased switching speed have the corollary of requiring the switch to exist within the confines of a single location, if exchange performance is a paramount concern.

In addition to speed, the cost shift must be considered. 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 providers 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 colocation 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.

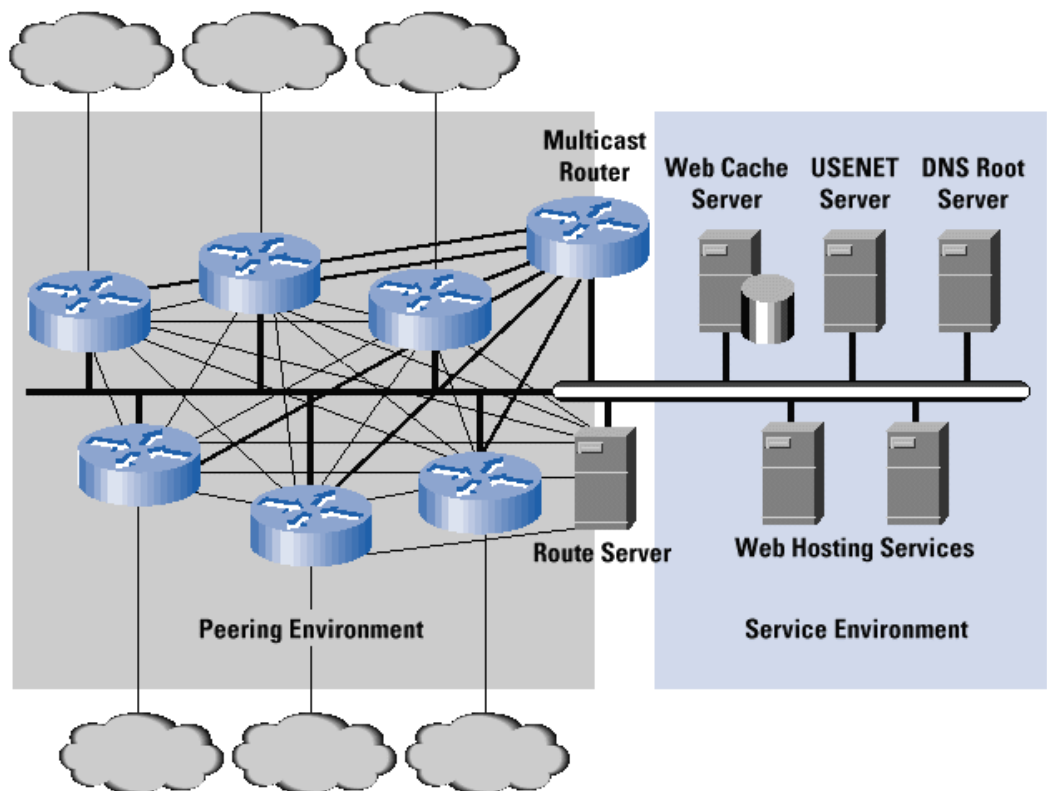
### Other Exchange-Located Services

The colocation 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 nonlocal 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, because the location is 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 colocation 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



## Network Access Points

The role of the exchange was broadened with the introduction of the Network Access Point (NAP) in the architecture proposed by the National Science Foundation (NSF) in 1995 when the NSFNET backbone was being phased out.

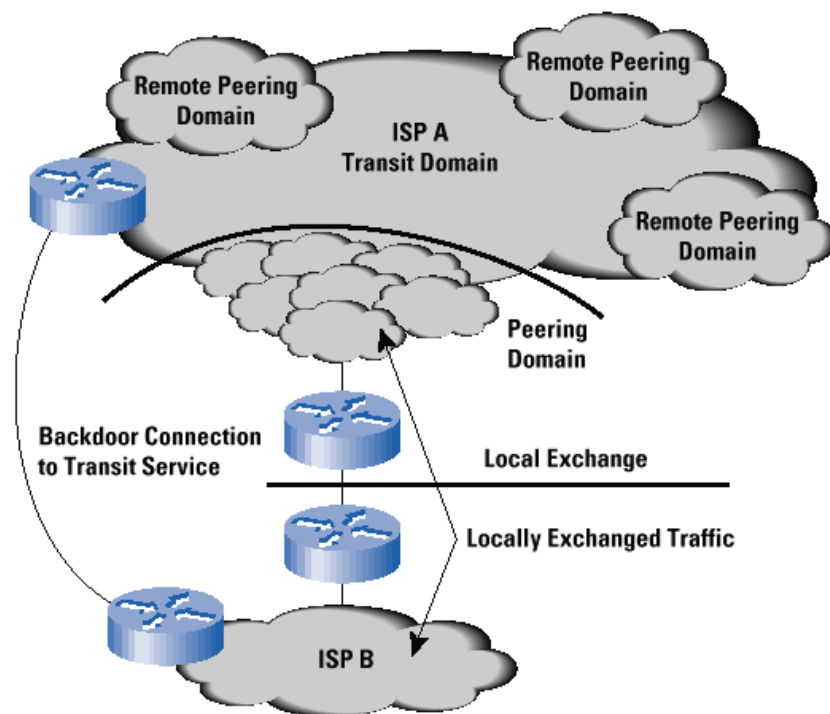
The NAP was seen to undertake two roles: the role of an exchange provider between regional ISPs who want to execute bilateral peering arrangements and the role of a transit purchase venue, in which regional ISPs could execute purchase agreements with one or more of a set of trunk carriage ISPs also connected 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 band-width 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?

The issue that the quality of the purchased transit service is coloured by the quality of the service provided by the NAP operator should also be considered. 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 colocation 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 role. Such a configuration is illustrated in Figure 8.

Figure 8: Peering and Transit Purchase



## Exchange Business Models

For the ISP industry, many 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
- Operates 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 colocate 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, because 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 among numerous ISPs. In these models, the exchange is a nonprofit common asset of the cooperative body. Although widely used, these models are 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.

## **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 interprovider connectivity environments, a situation that 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-Autonomous System (AS) routing policy tools that we use today, this structure must be a source of 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 fragility 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.

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 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.

---

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

---

[This article is based in part on material in *The ISP Survival Guide*, by Geoff Huston, ISBN 201-3-45567-9, published by Wiley. Used with permission.]