

Interconnection and Peering

Geoff Huston
October 2000

One of the topics that has had the most airplay within the ISP industry for many years, and is still a topic that generates all forms of excited reactions is the financial model of interconnection between ISPs. When you consider that the Internet is constructed of more than 8,500 distinct Internet networks, then interconnection is the glue that holds this together as one global network.

Technically, ISP interconnection is a relatively sophisticated affair, with individual providers being to define their import and export policies with some precisions. This allows each ISP network to connect to a number of other ISP's network and then express a preference as with which network will be used to source incoming traffic and which network will be used to pass outgoing data traffic. The protocol used to support this function is the Border Gateway Protocol (BGP), and it is the sum of these individual BGP-based pair wise interconnections defines the global flow of data across the Internet. However, in this column I don't want to dwell on the technology of IP interconnection, but instead I want to look at the business models that underpin ISP interconnection. For while the technology of Internet interconnection is extremely sophisticated, the business models are best described as being brutal in their crudity and inflexibility.

It is useful to introduce financial interconnection models by looking at two telephone service providers who interconnect their phone networks. The traditional telephone tariff model is that the originator of a phone call pays the local provider for the entire cost of the call, and the receiver of the call pays no fee to answer the call. What should happen when a call originates in one provider's network and terminates in the other provider's network? The originating provider collects revenue for the call, and the terminating provider collects no revenue, even though both networks have incurred some form of cost to support the call. To balance this inequity of cost and revenue, the telephone industry uses the concept of call accounting settlements. In this model the two providers agree on a standard unit and a standard rate for calls which pass across the interconnection. The unit typically used is that of a 'call minute'. Now when a call is made across the interconnection the call is logged, and the call originator uses a portion of the revenue from the call to pay the terminating provider by multiplying the duration of the call by the agreed rate. As calls will typically be made in both directions, the providers will normally keep a running total of these interconnection accounts, and at regular intervals they will settle the account by paying the balance. If more calls are made in one direction than the other, then the financial settlements will adjust to balance the costs of the two providers. The intent is that relative equity of cost and revenue between the two providers is restored through these arrangements.

Now lets look at the Internet interconnection environment. Unfortunately a number of useful concepts disappear from the picture straight away. Even in a model with only two ISPs interconnecting there is no such thing as a 'call' any longer, nor a clear concept of a 'caller' and a 'receiver'. Each provider may not want to inspect the address headers of every packet flowing over the interface in order to determine how the packet should be accounted. Even if the providers agree to ignore the fine detail of the packet header and simply total the volume of data sent in each direction across the interface, the task is apportioning relative value to each network remains unresolved. Packets may flow from a client of one provider's network to a client of the other because the receiver wanted the data (such as a web fetch), or because the sender wanted to send the data (such as an electronic mail message). It is also worth remembering that the Internet service is an unreliable data delivery service, so that packets may be dropped in

transit. If the interconnection regime has placed a real value on the packet, who bears the loss in value when the packet is dropped? We also need to remember that the cost of carrying an individual packet is extremely small, and the cost of accounting for each packet may well be greater than the cost of carrying the packet across the providers network in the first place!

All these questions point to a more basic aspect of the Internet. The issue is that with the Internet there is no longer a single application model for the network, and there is no longer a simple model of value transfer and cost apportionment, so that concepts of balanced financial settlements do not find a natural home within the Internet environment. Within the model of the Internet, services are part of the host-to-host interaction (or end-to-end as it is more commonly known). Within the packet switching fabric of IP transmission it is not possible to infer from packets passing through a particular interconnection where and how cost is incurred, where and how value is being added and where and how revenue is being injected into the network.

One answer to this conundrum has been to use a simpler model of interconnection, where one ISP is the supplier and the other ISP is the customer. In this model of interconnection the flow of money across the interconnection is in a single direction as the client pays the provider for an Internet connection service. Simple models that scale quickly are well suited to the internet, and we have seen quite widespread adoption of this supplier / customer interconnection model. The outcome of this is the creation of a number of levels of hierarchy within the Internet, where local ISPs are customers of larger regional ISPs, who in turn may be customers of larger national ISPs, who in turn may be customers of international transit ISPs. This structured view of ISP interconnection is shown in Figure 1. The structure is often referred to as a 'tiered' model, where the ISPs at the top level of the hierarchy are referred to as "Tier 1 ISPs", customer of these ISPs are referred to as "Tier 2 ISPs", and so on.

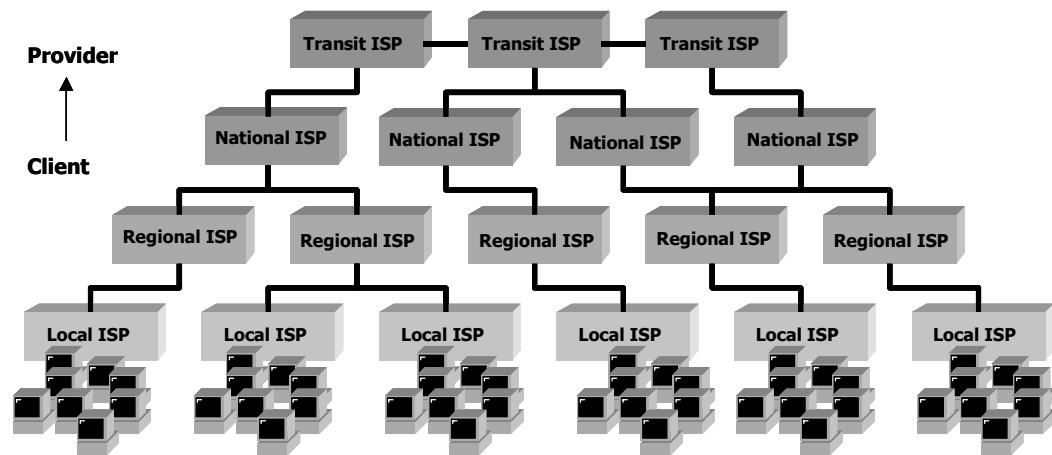


Figure 1 - the Hierarchy Model

This structured model breaks down when we look at the real world, and instead of a uniform collection of customer / supplier relationships we see a more unordered model where an ISP may use 2 or more suppliers and the strict order into wholesale and retail ISPs, and the tiering of wholesale ISPs is simply not present. However there is one further aspect of the interconnection environment which breaks away from the strict supplier / customer relationship, and that is the use of sender keep all peering (or "peering" as it is commonly referred to). In this form of interconnection both parties determine that the benefit of interconnection is equal, within some bounds of approximation, and there is no possibility of one provider assuming a supplier role to the other. Such peering occurs at all levels of the tiered hierarchy, and generally works

along the principle that peering is often a more cost attractive option to exchange traffic than the alternative of passing the traffic to a supplier and paying the incremental cost of so doing.

The resultant environment of the Internet today is one where there are three basic means of interconnection between two ISPs: firstly where there is no direct interconnection, in which case the mutually exchanged traffic uses the services of intermediaries to act as transit providers, secondly where one ISP acts as the supplier and is funded in this role by the other ISP assuming the role of customer, and, lastly, where the ISPs undertake a peering arrangement and no financial exchange takes place. The resultant structure of the Internet is a little more chaotic, as Figure 2 indicates.




Figure 2 - the Interconnection Model

None of these relationships are static. If an ISP is a customer of another ISP, then its business objective is to reduce the costs associated with that relationship by forming other peering relationships as a means of offloading traffic from this funded interconnection, and, ultimately, by converting the customer / supplier relationship to one of a peer. And within a peering relationship an ISP may perceive that it is adding greater value to the connection than its peer and the approximate equality of the relationship no longer holds. In this case the ISP will attempt to alter the relationship to that of a supplier to the other ISP. The general description of these changing interconnection relationships is that customers want to become peers and peers want to become suppliers.

While the Internet market continues to grow at explosive rates the business outcomes of interconnection arrangements pale in importance to the ISP's primary task of simply amassing adequate resources to fuel the Internet's growth pull. It is perhaps because of this that the relative crudity of the existing interconnection arrangements are enough to do the job and get packets through. When, or if, we get to the point where this phenomenal growth in the ISP market ever slackens off, then perhaps attention will turn back to interconnection business structures, and it is my guess that we will then see the ISP industry start to grapple with the somewhat difficult concept of the change in the value of a packet as it passes through the Internet.
