Whitepaper: Peering Center Economics for ISPs

Henry Kilmer

Version 0.11, 14 April 2003

Abstract

This document provides an overview for the use and economic evaluations ISPs review for determining the efficiencies of building into a peering center.

Contents

1 Introduction

The economics of interconnection and the conditions under which interconnections are made by ISPs can present a market for a facilities provider to give rise to an exchange point. Understanding the requirements and economics involved in the ISPs interconnection design is the first step in understanding peering centers.

2 Peering Politics

A short description of peering and its politics can be useful in understanding "build" decisions made by ISPs with regard to peering centers. Not all ISPs are alike in their view on peering and their policies regarding interconnection options.

2.1 Peering versus Transit Defined

Interconnections between networks are described in terms of peering and transit. Peering is a business and technical arrangement whereby the providers agree to accept traffic from one another and from one another's customers. Peering does not include the obligation to carry traffic to third parties.

Transit is a business and technical arrangement whereby one provider agrees to carry traffic to third parties on behalf of the other entity. This encompasses the provider's network, their customers, AND their peers (ie: the traffic will *transit* the provider's network to reach all destinations possible).

2.2 ISPs Core Business

An ISP sells and attempts to deliver connectivity (transit service) to the *global* Internet. Since the Internet is a connected mesh of independently run networks, ISPs spend engineering time monitoring their routing table looking for significant changes indicating connectivity loss to a network (small changes are routine and are caused by end sites having connectivity issues or misconfigurations). ISPs achieve global connectivity through peering and transit services. The top backbones have no transit themselves and achieve global connectivity through peering relationships and their customer base. Tier-2 and tier-3 ISPs (intermediate ISPs) achieve global connectivity through transit service from one or more ISPs, where at least one of the ISPs is a top tier backbone. Some ISPs also have peering relationships but rely on the transit service to achieve global connectivity.

2.3 Tier-1 Status

The Tier-1 status label has gone through a few variations since it was first used to describe the hierarchy that existed among the ISPs. Today this label is generally used for a small set of ISPs whose networks cover multiple continents, carry a significant portion of the Internet traffic, and have negotiated amonst themselves the status of peer. It is important to note the last two parts of that definition as there can be multi-continental networks of like size (by some measures) but have not, for various reasons, been able to negotiate a peering relationship with one or more of the existing Tier-1 providers. In this case, it is common for the network provider to refer to themselves as a Tier-1 provider but still be at a disadvantage when negotiating peering or competing for sales versus the existing Tier-1 ISPs. Tier-1 status carries significant weight in competing for sales in the transit market and additionally carries weight for those companies that sell services to ISPs. For example, a router vendor that has not been able to sell to the Tier-1 ISPs will have a very difficult time getting into any backbone network, likewise an exchange point with one or more Tier-1 ISPs is much more likely to draw the Tier-2 and 3 ISPs than a competing exchange point without the Tier-1 networks.

2.4 A Peer is a Peer Unless its Mr. Ratio

The determination that one's network is a peer of another network is nebulous at best and can be very difficult to determine. Some ISPs have created guidelines for themselves in an attempt to evaluate each peer request under a standard, reviewable process to determine if peering would be of equal value to both parties. These guidelines usually indicate the general size of the network (a certain number of geographic regions connected by a certain amount of bandwidth) and the amount of traffic flowing between the two networks. Some ISPs add factors like the number of customers or the number of nodes on the network.

After all of these requirements are met, a network can still be viewed as not equal based on the ratio of traffic flowing between the two networks. This means that one network may be sending X times more traffic (some ISPs use a 2:1 ratio, others use a 4:1 ratio) to another potential peer's network than it is receiving from the same network, and if that difference is great enough it will be turned down for peering. The reason is that the peering would not be of equal value and the network sending more traffic would be getting more benefit from the peering.

There are technical ways to solve *some* of the arguements supporting this position but they are not elegant or scalable. Suffice it to say that this situation exists today and at least for the near future (perhaps until per-bit billing) and it puts content providers and access providers at odds when neither would have a service without the other. The end result is that one of the providers ends up buying transit or bartering service from a third party to maintain connectivity between the networks.

3 The Cost of Peering - It Really Isn't Free

It is often misunderstood that peering is free and transit costs money. This is simply untrue. In order to provide global connecitivity (an ISP needs to be able to deliver a packet to *any* destination on the Internet), an ISP either has to negotiate a significant amount of peering relationships, arrange to be connected to all of the peers networks through peering points and/or direct connections, or the ISP must purchase transit service from another ISP that is fully able to reach any destination on the Internet.

Peering costs can be broken down into a few categories: administrative (for the negotiations), interconnection costs and backbone costs.

3.1 Administrative Costs

Administrative costs for peering cover items listed below:

- Legal contract negotiation.
- Technical network review.
- Statistics collection and analysis.
- Router configuration and updates.
- Operational interactions (NOC contact information).
- Business case review.

For a small network, these can be small costs but for a large network, these costs become more significant. Some Tier-1 and large Tier-2 ISPs have a number of staff members solely dedicated to this process.

3.2 Interconnection Costs

Interconnection costs vary based upon the engineering of the interconnections. When ISPs negotiate peering it usually happens in one of three ways.

The first way is when both parties are connected to an existing service like a peering center's exchange. The ISPs agree and configure connectivity between the two networks via the exchange point(s). The costs are the peering center's space/power/connection with no additional costs to the ISPs.

The second way is when both parties are co-located within the same building and can get interconnections between their routers through the landlord or host. This could be a peering center but the providers do not wish to use the exchange service, or it could be any telecommunications data center or shared facility. The costs are the shared facility's space/power/connection plus an additional router port for the individual connection.

The third way involves the scenario when neither ISP is in a shared facility where the other ISP is also located. In this scenario, telecommunications services (options include a circuit with a service like ATM on it, a direct peer-to-peer circuit and direct fiber) need to be aquired to interconnect the two networks. The costs are the connection service of choice (usually a circuit or fiber) plus a router port to terminate the connection.

The first two interconnection scenarios utilize existing shared costs within an ISPs network (the peering/shared facility's base costs) and the first scenario carries that further by also utilizing a shared port on the router. Each option has different scaling properties and the large ISPs end up using a mixture of all three designs.

3.3 Backbone Costs

Backbone costs are mentioned here since it is important to note that any peering decision effects the traffic flows across an ISPs network. Adding additional peers in one location may necessitate upgrading the backbone capacity into that location. The decision of how, where and when to peer with another ISP ties directly to backbone design, current capacity constraints and costs.

4 The Cost of Transit - It Really Isn't Free Either

This may seem like an obvious statement but the purchasing of transit carries a cost associated with it. The not-quite-as-obvious part is the political aspect of purchasing transit. There are providers that hold to the belief that once a customer, always a customer. Some carry this beyond individual products as well, meaning buying *anything* from some companies (long distance for example) may prevent an ISP from obtaining peering well down the road.

There are benefits to paying for transit as opposed to a peering relationship. Some are as follows:

- Transit service often comes with a service level agreement (SLA) giving rebates if the service level is not met. Peering agreements are best effort.
- Transit service will not go away whereas peering agreements are commonly annual and subject to review against changing requirements.
- Transit service normally contains monitoring and repair.
- Transit service commonly offers traffic measuring reporting.
- Transit service commonly offers access to a security response center for assistance in any security related incidents.

There are benefits to a mixed peering and transit environment as well. Some are as follows:

- Peering can reduce transit dependency and costs.
- Peering can reduce latency and jitter.
- Peering can reduce downtime.

5 Peering Center Conditions and Costs

As we have indicated, a peering center can help ISPs gain efficiencies in their peering (network) design. A properly run center can reduce an ISPs interconnection costs, offer them the ability to sell services, and scale more gracefully than utilizing standard telecommunications circuits by offering more options at a lower cost.

5.1 Fundamental Requirements

In order to offer the service neccesary for the ISPs to efficiently utilize the center, a peering center must have the following:

- Rich telecommunications infrastructure built into the facility. This gives the ISPs reliable choices in covering their telecommunications needs.
- Reliable power. Power filters and uninterruptable backup generators are mandatory for a reliable service.
- Trained staff. Competent response when needed assure the ISPs that any issues will be resolved professionally and expediently minimizing risk or downtime, reducing length of downtime and generating reliable performance.

5.2 Peering Cost Controls

Having a reliable, quality service covers the performance aspects of the ISPs needs but does not necessary cover the cost aspects of their needs. If the costs of building their backbone into the center plus the costs of leasing the space and power are too great, an ISP will not build into the center. Obtaining a rich set of carriers in the center will help mediate the backbone build costs as much as possbile making that a significant cost control point. Additionally, offering enhanced services like remote-hands support, operating a shared peering exchange, or referral/reseller agreements with appropriate companies can help offset the other main cost control point: the space and power costs. The third main cost control point in a peering facility is the cost of the interconnection itself. This is the main product a peering facility is offering and the main reason for its existence.

5.3 "Public" Peering Service

Offering a peering exchange service allows ISPs to peer or sell transit over a shared connection (commonly an ethernet switch). Since the traffic runs over a shared router interface and switch, this form of peering is commonly referred to as "public" peering. Conversely, peering done over a direct, peer-to-peer connection (like a circuit) is referred to as a private peering connection. The decision for public or private peering with an ISPs network is both an economic and technical decision. Since most public peering LANs are ethernet based with Gigabit Ethernet being the largest connection offered, if there are 10 peers averaging 100 Mbps each, the connection is full. In this type of situation, ISPs normally move the top few peers onto private connections eliminating the congestion and allowing for more growth.

6 The Decision Point

In understanding what the ISPs are trying to do at peering centers, their costs and their reasoning, a cost/benefit analysis can be reached. If an ISP has sufficient costs associated in a region with telecommunications services, it is possible to reduce costs while offering the same services by building into a peering center within the region.

For example, taking only peering into account, let's assume an ISP has 10 peers in one region each connected via a T1 circuit costing $300/month average, equalling $3000/month. These 10 peers are also located in a local peering center. The cost for the space, power and a 100 Mbps ethernet connection for peering would cost the ISP $2000/month and the cost of two redundant DS3's to connect the peering center location into the ISPs backbone would cost $1000/month. On the surface, the ISP would pay $3000/month either way, but the growth properties of each scenario are very different. In the case of the point-to-point T1 circuits, if the ISP adds another peer, they add $300/month in cost whereas in the peering center, they would not add any cost since there is plenty of headroom on both the peering connection and the backbone circuits.

References

[1] Geoff Houston *"An ISP Survival Guide"* John Wiley, 1999

[2] Terrence P. McGarty *"Peering, Transit, Interconnection: Internet Access in Central Europe"*

`http://itc.mit.edu/itel/docs/2002/mcgarty_2002.pdf`

ITC Working Paper, January, 2002.

[3] Martin Weiss and SeungJae Shin *"Internet Interconnection Economic Model and its Analysis: Peering and Settlement"*

`http://itc.mit.edu/itel/docs/2002/wcc2002_comsys_weiss_shin.pdf`

17th World Computer Congress, Montreal, Canada, August 2002.

[4] David Gabel *"Interconnection Payments in Telecommunications: A Competitive Market Approach"*

http://itc.mit.edu/itel/docs/2002/gabel_interconn_pymnt.pdf