A Business Case for ISP Peering

William B. Norton <wbn@equinix.com>

Abstract

During the recent economic turmoil, Internet Service Providers are looking for ways to reduce the costs associated with providing Internet services. Chief among their costs is telecommunications costs, the cost of getting traffic to the networks that make up the Internet.

Discussions with ISP Peering Coordinators identified Internet Service Provider (ISP) "Peering" as one of the most effective methods of reducing telecommunications costs for ISPs. The ISP Peering Coordinator's job is to establish and effectively manage this interconnection between ISPs. Goals include maximizing efficiency and minimizing the costs of interconnection. As demonstrated in the "Internet Service Providers and Peering " earlier work, this job requires a rare combination of technical and business acumen with good people and negotiating skills.

This paper introduces the ISP Peering Coordinators terminology and the tools and analysis typically used. We demonstrate these tools with a specific implementation, and then generalize them in the form of a "Peering Break Even Analysis Graph." The financial models are included in the appendix so the reader can adjust the cost components to match their environment. This paper presents a business case for Internet Service Provider Peering based on current practices and market prices.

Introduction and Definitions

Over two hundred ISP Peering Coordinators were interviewed to determine the processes of and motivations for peering[1], and chief among the motivations was reducing the cost of "**transit**". To describe this motivation in brief let's first introduce a couple of definitions.

First, the most basic definition of the Internet Service Provider:

Definition: An Internet Service Provider (ISP) is an organization that sells access to the Internet.

By definition therefore, ISPs must somehow themselves connect to the Internet. For most ISPs this means purchasing a service called "**transit"** from an ISP that is already attached to the Internet.

Definition: A Transit Relationship is a business arrangement whereby an ISP provides (typically sells) access to the Global Internet[2].

To illustrate, consider figure 1 in which ISP A has customer attachments to the left shown as small gray circles and a transit service connection to the Internet to the right.

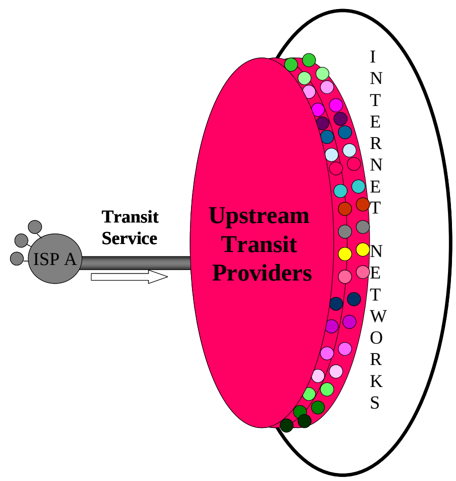

Figure 1 - Transit Service

In this picture, ISP A "*purchases transit*" from an "Upstream Transit Provider" who has the responsibility of providing access to the Global Internet. After the transit relationship is in place, ISP A customers can access the Internet, and the Internet can access ISP A customers. Put simply, transit is a plug in the wall that says "Internet This Way".

Cost of Traffic Exchange in a Transit Relationship

The cost of transit varies widely but is typically metered and charged based upon a 95th percentile traffic sampling technique[3]. Traffic flow is typically

[1] Norton, William B., "Internet Service Providers and Peering", available from the author via e-mail to wbn@equinix.com.

[2] According to the Tony Bates CIDR Report, the Global Internet includes approximately 100,000 network entries in the routing table. See http://www.employees.org/~tbates/cidr-report.html for details.

[3] The peak rate is typically measured using 5-minute samples over a months' time, using the 95th percentile number to determine the billing rate.

measured in Mega-bits-per-second (Mbps) and prices range between $100/Mbps and $1200/Mbps[4]. There is typically an initial startup cost[5], and the price per Mbps generally decreases slightly as more traffic is exchanged. Figure 2 below shows a sample tiered transit fee pricing structure[6].

Transit Costs

Mbps	$ Per Megabit-per-second
1-15 Mbps	$425
16-30	$395
31-44	$365
45+	$325

Figure 2 – Cost function for Transit Services

This cost function for transit is graphed below. On the Y-axis we see the Unit cost for transit on a $-per-Mbps basis.

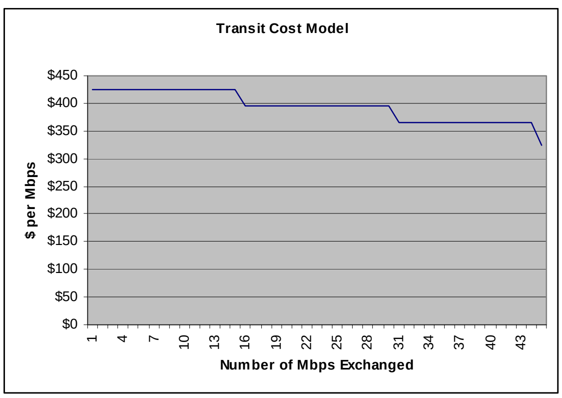

Figure 3 - Cost Function for Transit Service

The volume of Internet traffic has historically increased[7] and all indications are that this trend continues today. Even though the price of transit has declined by an average of about 30% per year[8], the peak traffic rate has typically increased at least 3 times that fast[9]. As customers expand use of innovative and high bandwidth services such as multimedia streaming of radio, video broadcasts, large volume music exchange services and live non-cacheable event casting across the Internet, ISPs carry much more traffic and realize dramatically increased transit fees. To manage this, some[10] Internet Service Providers measure their transit traffic flows to determine *where* their transit traffic is ultimately delivered[11]. These destinations are associated with ISPs as shown pictorially in figure 4 below.

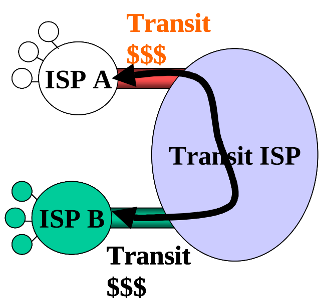

Figure 4 - Aggregate Traffic Flow Measurement

Once the top traffic destinations are identified and associated with specific ISPs, these ISPs are targeted for potential peering relationship discussions. Below in figure 5 is a sample top destinations list from a large global content-heavy ISP[12]. These ISPs

4 Based on conversations with ISPs in the 2001 calendar year.

5 For simplicity we will ignore these one-time startup costs in our pricing models.

6 This price point is on the lower end of the transit cost spectrum. (Courtesy of Wolfgang Tremmel, Director of Peering and Network Planning for Via Net.works Inc.)

7 Some quote Internet traffic doubling as often as every 180 days.

8 Gnanasekaran Swaminathan (Savvis)

9 Michael Hrybyk, General Manager BCNet and David Prior (PBIMedia) – Internet Traffic Growth rate estimated at 96% compounded annual growth rate from 1996-2005.

10 Based on my conversations, less than 5% of ISPs perform this detailed analysis.

11 Cisco NetFlow and Juniper equivalents are tool of the trade. There are issues dealing with these tools as the volume of data, processing and analysis, access to staff expertise, and the impact on routers were all cited as challenges. David Prior mentioned CAIDA's CoralReef software as a solid real-time analysis tool addressing these issues.

12 From left to right, we see the assigned ISP Autonomous System (AS) Number and the average number of Mbps destined to that ISP's customer. This AS number is then mapped to the ISP Name and Contact information from

are targeted as ideal candidates for a "**peering**" relationship. The intent is jointly yield a lower cost and more direct traffic exchange in a "peering relationship".

Internet Service Provider A			
AS Number	Mbps	Destination ISP	Contact
6172	24.35	HOME-NET-1	[HOME-NOC-ARIN]
701	8.90	ALTERNET-AS	[IE8-ARIN]
1668	8.14	AOL-PRIMEHOST	[AOL-NOC-ARIN]
4766	7.08	APNIC-AS-BLOCK	[SA90-ARIN]
3320	5.12	RIPE-ASNBLOCK4	[RIPE-NCC-ARIN]
577	4.24	BACOM	[EQ-ARIN]
6327	3.90	SHAWFIBER	[IAS-ARIN]
1	3.89	BBNPLANET	[CS15-ARIN]
7018	3.66	ATT-INTERNET4	[JB3310-ARIN]
9318	3.13	APNIC-AS-3-BLOCK	[SA90-ARIN]
5769	2.67	VIDEOTRON	[NAV1-ARIN]
6830	2.30	HCSNET-ASNBLK	[MD205-ARIN]
9277	2.22	APNIC-AS-3-BLOCK	[SA90-ARIN]
10994	2.08	TAMPA2-TWC-5	[JD6-ARIN]
1239	2.05	SprintLink	[SPRINT-NOC-ARIN]

Figure 5 - Sample list of Peering Candidates sorted by Traffic Volume[13]

<u>Definition:</u> *<u>Peering</u>* is the business relationship whereby ISPs reciprocally provide access to each others' customers[14].

It is important to note that peering is not a substitute for transit. *Transit* provides access to the <u>entire</u> Internet routing table for a fee, while *peering* is typically a no-cost arrangement providing access only to each others' customers.

To illustrate this point, consider Figure 6 below. ISP A has entered into a peering relationship with ISP B. ISP A sends all traffic destined to ISP B directly to ISP B, and ISP B reciprocally sends all traffic destined to ISP A directly to ISP A. In this example, both ISP A and ISP B continue to purchase transit from an upstream transit provider to access the rest of the Internet. Both ISPs reduce their transit costs proportionately to the amount of traffic they exchange with each other in this peering relationship.

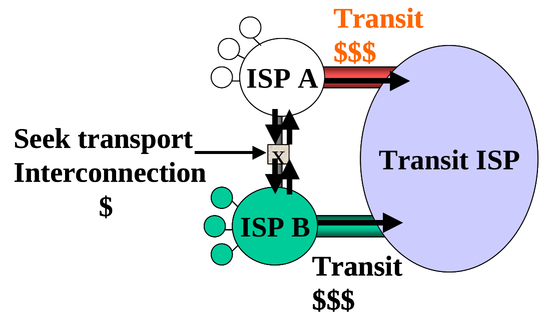

Figure 6 - ISP Peering Moves Transit Traffic to a Lower Cost (and More Direct) Peering Path

There are several methods of implementing peering interconnections. Most common is "Public Peering" which refers to an ISP interconnection across a shared fabric[15]. This is typically done in a location where ISPs collocate routers. The other type of peering is called "Private Peering" and refers to the direct point-to-point interconnection. Private peering is increasingly popular at exchange points where it implemented using fiber or copper cross connects[16]. Private peering is also accomplished using point-to-point leased circuits[17]. Both of these last two private peering models are described and modeled in "Interconnection Strategies for ISPs[18]". For this paper we will model the most common peering approach: public peering at an exchange point using a shared switch fabric[19].

The Cost of Traffic Exchange in a Peering Relationship

the appropriate assigning authority.

13 Some of these destinations represent ASes that are not specifically listed in the ARIN as.txt file so are shown as aggregates contained in aggregates such as RIPE-ASNBLOCK4.

14 While the definition refers to ISPs peering, large volume content providers are starting to aggressively pursue peering as well. Yahoo! for example is very active in the peering arena, reducing their millions of dollars in transit costs by aggressively pursuing peering relationships with ISPs. Jeffrey Pappen (Yahoo!) presented by far the most extensive traffic analysis process the author has seen at the Equinix Gigabit Peering Forum III in Dallas on July 17th, 2001.

15 Such as a Gigabit Ethernet or ATM.

16 Some exchange points do not allow private peering or require the purchase of a port on the public peering fabric in order to then purchase a private cross connect.

17 Over SONET services for example.

18 Available from the author. Send e-mail to wbn@equinix.com with "Interconnection Strategies for ISPs" in the subject line.

19 It is important to note that, due to the wide variety of network equipment architectures and configurations, we ignore the cost of network equipment for both the transit and peering models.

The cost of peering at an exchange point typically includes three cost components:

1. Transport[20] into the exchange point,
2. A port on the exchange point shared fabric[21], and
3. Collocation space at the exchange point[22].

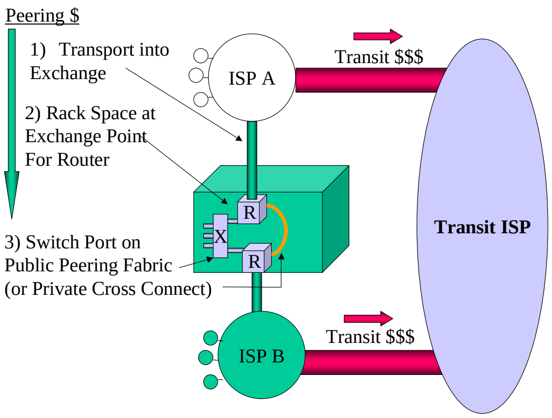

Figure 7 - Components of the Public Peering Cost Model

Unlike transit service, traffic exchange in peering relationships is not metered[23]. ISPs can send as much traffic as can fit across the transport circuit and peering fabric for the cost of the interconnection[24].

Question: When does it make sense to Peer? When is it less expensive to send traffic over a Peering Interconnection versus simply sending all traffic to Upstream ISP(s) in a Transit relationship(s)?

In order to compare Peering and Transit we need to describe peering costs in the same terms as transit costs. We need to compare both Peering and Transit on a per-megabit-per-second basis.

Example: Consider a recent pricing snapshot of peering costs. A large telecommunications company[25] offers a DS3 (45Mbps) transport circuit into the Equinix Internet business Exchange for $1000/month. At the Equinix Internet Business Exchange, the Exchange Pak I product includes a 100Mbps port and a half rack for $1000/month as pictured in figure 8 below[26].

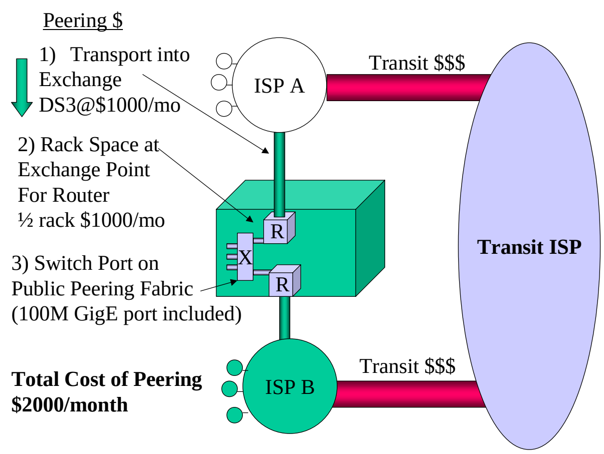

Figure 8 – Sample Public Peering Costs[27]

The monthly costs of *peering* in this example are fixed at $2000/month. The average cost per Mbps of traffic exchanged will vary based upon how many Mbps are exchanged at the peering point.

For example, if ISP A exchanges only 1 Mbps of traffic with the population at the exchange, the cost per Mbps is $2000/Mbps. If the ISP A exchanges 2 Mbps with the population at the exchange, the cost

---

20 Definition: *Transport* refers to a physical/data link layer media interconnection (e.g. circuits, gigE switching fabric, gigE over fiber cross connects).

21 Popular exchange point fabrics include gigabit Ethernet, ATM, and FDDI.

22 This is true for exchange points that support or require collocation of routers with the switch gear.

23 They are typically free and therefore not metered for billing purposes. Peering connections are sometimes monitored and measured for engineering purposes, and to ensure that traffic flow ratios are within the range agreed upon between the ISPs.

24 In some cases ISPs may have a for-free peering relationship up to the point when a traffic ratio is not exceeded, and then fee-based beyond that point. According to David Prior (PBIMedia), Telia has published a 2:1 imbalance ratio as the largest acceptable imbalance before financial compensation is required. This is more the exception than the rule today.

25 This was promotional pricing for DS3 from midtown San Jose to the Equinix San Jose Internet Business Exchange (IBX). Note that pricing is mileage based and highly variable – this is for demonstration purposes only.

26 Once again, promotional pricing for demonstration purposes only.

27 Source: Equinix 2001 Gigabit Peering Forum in San Jose Promotional Pricing

per Mbps is $1000 per Mbp exchanged, and so on. The cost per Mbps declines as the number of Mbps exchange increases as shown in the table below.

Mbps Exchanged	Peering Cost Per Mbps
1	$2,000
2	$1,000
3	$667
4	$500
5	$400
6	$333
7	$286
8	$250
9	$222
10	$200

Figure 9 - Peering Costs allocated over Traffic Volume

Assume further that the cost of transit is $400/Mbps. From a strictly financial position, *peering* makes sense when the unit cost of peering is less than the cost of transit, that is, when more than 5 Mbps will be sent to the Exchange Point. This we call the "**Peering Break Even Point**."

Generalize the Peering vs. Transit Tradeoff. If we continue plotting the cost per Mbps across the size of the peering bandwidth[28] we get the graph below. The conclusion is clear - if an ISP can peer, exchanging more traffic than indicated at the "Peering Breakeven Point", then the ISP should prefer to peer instead of *solely* purchasing transit from an upstream ISP.

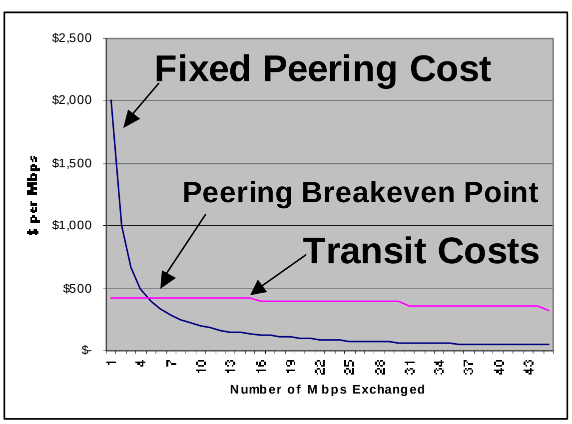

Figure 10 - Example of Peering Breakeven analysis

[28] Recall in this example we are using a DS3 (45 Mbps) capacity transport circuit into the exchange.

To the right of the breakeven point (see figure 10), ISPs completely cover their peering costs with reduction in transit fees. This savings is proportional to the amount of traffic exchanged with the peering population.

To the left of the Peering Breakeven Point is the "Peering Risk", the possibility of not realizing sufficient peering traffic volume to offset the cost of peering.

The complete financial model for this example is included in Appendix A as an Excel spreadsheet. Readers are encouraged to enter their own transport and exchange point costs to determine their total cost of peering and Peering Breakeven Point. Appendix A also includes one additional column that shows the unit cost savings for sending traffic over a peering link rather than through an upstream transit provider.

What is the Maximum Reduction in Transit Fees? So far this analysis of peering scales up to what I call the "**Effective Peering Bandwidth**", or the smaller of the transport and port speed. For example, in this last case where the transport into the exchange is a DS3 (45Mbps) and the exchange fabric uses a 100Mbps port, the Effective Peering Bandwidth is 45 Mbps since that is the maximum amount of traffic that can be used for peering traffic exchange[29]. The cost per Mbps at that 45Mbps turns out to be about $45/Mbps. Note that this is significantly less than transit at $400/Mbps. The cost of traffic exchange at this extreme allows us to quantify the minimum unit cost for traffic exchange. Beyond this Effective Peering Bandwidth point additional transport and/or capacity must be provisioned.

Here is a generalization of the calculation of this minimum cost per Mbps:

$$MinCostOfPeering = \frac{CostofPeering}{EffectivePeeringBandwidth}$$

$$= \frac{Transport + RackFee + PortFees}{\min(DS3BW, 100MFastE)}$$

$$= \frac{\$2000}{45Mbps} = \$45\, perMbps$$

Note that peering scales very well. The "Interconnection Strategies for ISPs" study showed

[29] Given the transport of 45 Mbps and the port speed of 100Mbps, one can fill the 45 Mbps pipe and send no more traffic even though the exchange point port can handle up to 100Mbps.

that an ISP can typically acquire four times the bandwidth for a twofold increase in transport cost. Therefore, large scale peering can result in very low unit cost for traffic exchanged.

Example of Large Scale Peering: Consider a large scale peering ISP purchasing an OC-12 (622 Mbps) transport circuit into the Equinix IBX using Exchange Pak II which consists of a half rack and a gigabit (1000Mbps) Ethernet port. Pricing studies[30] have recently shown metro OC-12 prices around $10,000 per month. For this example, assume that the Exchange Pak II package is priced at $2000 per month. What is the minimum cost of traffic exchange across this peering configuration?

$$MinCostOfPeering = \frac{CostofPeering}{EffectivePeeringBandwidth}$$

$$= \frac{Transport + RackFee + PortFees}{\min(OC12BW, 1000MGigE)}$$

$$= \frac{\$12{,}000}{622Mbps} = \$19\, perMbps$$

Again, at this minimum cost of traffic exchange, $19 is substantially lower than the $400/Mbps transit fee[31].

There are also some innovative Metropolitan Area Network providers[32] that provide Ethernet-based metered transport services. With these business models, the transport part of peering is no longer fixed but proportional to the amount of bandwidth actually used. This shifts the Peering Breakeven Point to the left, reducing the "Peering Risk", or the risk of installing at an exchange point and not realizing the traffic volumes necessary to reduce traffic exchange costs.

Additional Motivations for Peering

Finally, it is important to highlight some technical motivations for peering that lead to less easily quantifiable motivations for peering:

1. Peering provides the *lowest latency path* between ISP customers. Peering has been found to improve performance by as much as 40-50 milliseconds[33].
2. Peering gives ISPs *more control* over routing, and have more flexibility to route around congested paths that could cause packet loss.
3. Peering provides *redundancy*. If peering sessions fail, the transit services provide backup connectivity to the peer networks. If the transit connectivity fails, the peering connectivity is unaffected.
4. Some exchange points support both Peering and Transit traffic exchange, allowing for the aggregation of transit traffic and peering traffic back to the ISP network. This reduces local loop costs for access to transit services[34].

These can have a significant financial impact, but there are also factors that are more difficult to model and quantify. For example, packet loss causes data transfers[35] to timeout, and retransmissions cause the data transfer window size to decrease, ultimately resulting in lower aggregate data transfer rate. This means that customers not only have a degraded experience, but the resulting decrease in data transferred, results in the ISP not make as much money either! To maximize revenue therefore, ISPs should seek to minimize packet loss and latency. Peering gives the ISP lower latency (direct path) and lower loss (assuming at least one ISP is motivated) with greater control over the routing.

Conversations with ISPs highlighted a few other subtle advantages to peering as well.

1. By peering, ISPs build and maintain a relationship with other ISPs and as a side effect get a better sense of the competitive environment in which they operate.

30 Gigabit Peering Forum II in Reston Virginia Nov 2000

31 Naturally at this volume of traffic, transit fees would be substantially lower than $400/Mbps, but still much higher than $19/Mbps!

32 Telseon, Yipes, etc.

33 Data point from Jalil Sanad Halim, Program Manager/Network Planner for 9 Telecom.

34 In some cases local loop costs can represent a significant cost. Purchasing transit at an exchange point can reduce local loop costs with cross connect fees several orders of magnitude less expensive.

35 We are assuming TCP-based (perhaps web-based) transfers here.

2. Peering was cited as substantial marketing collateral, particularly for content heavy or hosting ISPs. The implied assumption is that more peering means more redundancy, and larger capacity peering with the largest ISPs yields greater performance.
3. ISPs can improve their network reliability by peering at multiple points. Internet routers are very good at detecting failed paths and automatically rerouting traffic around breaks. Wide scale peering decreases the affect of any one failed network component.

There are some challenges with peering that are highlighted more thoroughly in “Internet Service Providers and Peering”. For example:

1. Peering requires greater network expertise than simply purchasing from a single upstream transit provider.
2. There are administrative startup costs associated with peering. Peering often requires contracts and negotiations iterations between ISP legal departments.
3. Peering is not always granted and is sometimes impossible to obtain. Beyond the difficulty in finding the contact person to initiate peering discussions, some ISPs have unpublished peering prerequisites that prevent all but the largest ISPs from peering with them.
4. There is greater operational overhead with peering than if one managed only a few transit relationships. Some ISPs cited transit (customer) outage trouble tickets get addressed more quickly than peering tickets. A couple ISPs mentioned that they prefer the “teeth” of a customer-based contract over the softer peering assurance that both ISPs will work diligently to fix peering-related issues.
5. The process of peering is slow. It may take months to get peering up and operationally passing traffic[36].

<See Appendix B for “The Top Reasons for NOT Peering>

All in all, peering can offer substantial benefits for ISPs and large-scale content players that exchange a lot of traffic. The larger the traffic volume and the greater the difference between transit costs and peering costs, the greater the motivation to explore peering as a cost savings strategy.

Summary: Generalized Peering Breakeven Analysis Graph

Let’s summarize in the form of a generalization of the examples we have seen so far in a “Peering Breakeven Analysis” graph (figure 10).

The graph below (figure 11) generalizes the cost of traffic exchange in a transit relationship against the cost of traffic exchange in peering relationships across an exchange point.

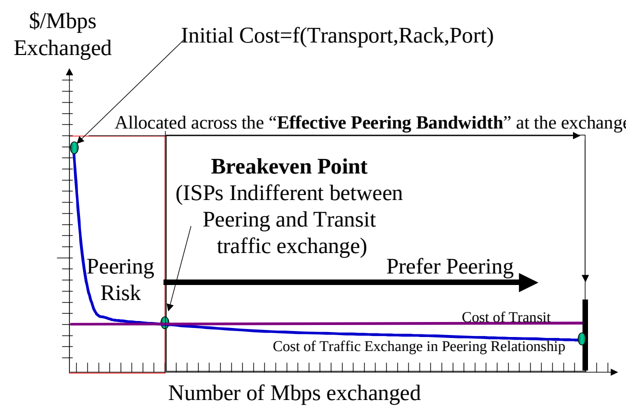

Figure 11 - Generalized Peering Breakeven Analysis Graph

The unit cost of traffic exchange is on the Y-axis in cost-per-Mbps. The X-axis shows the volume of traffic exchanged in Mbps. The cost of transit is shown as a relatively flat unit cost line[37]. The sloped line shows the cost of traffic exchange in a peering relationship.

Peering costs are fixed and include the cost of transport into an exchange, the cost of a partial rack for routing equipment, and the cost of a port on a switch for peering with the exchange population. To compare Peering and Transit, this peering cost is allocated across the amount of traffic exchanged between the ISP and the population of ISPs that are peering with the ISP at the exchange point.

The more traffic exchanged at the peering point, the lower the unit cost of traffic exchanged.

36 Jon Castle (Comdiso)

37 This line is actually stepped but shown as a flat line for simplicity.

There is a "Peering Breakeven Point" where ISPs are financially indifferent between peering and simply sending all traffic through its upstream ISP[38]. Once traffic volume between the ISP and the peering population reaches the breakeven point, ISPs start saving money by peering. The amount of money saved is proportional to the amount of traffic sent to the population of ISPs at the exchange point. The "Peering Risk" is the range of traffic exchange where an ISP fails to exchange enough traffic with the other ISPs at the exchange point to offset the cost of peering.

The amount of traffic sent to the exchange is capped by the minimum of the transport bandwidth and the port bandwidth, termed the "Effective Peering Bandwidth".

The minimum cost of traffic exchange can be calculated to be the cost of peering divided by the Effective Peering Bandwidth.

Acknowledgement

The author wishes to express appreciation for the following individuals who assisted and provided valuable feedback during the creation of this document: Jeffrey Pappen (Yahoo!), Mitchell Rose (Digital Island), Dennis L. Shreve (FACET Networks), Jon Castle (Comdisco), Jeff Swinton (Conxion), Jalil Sanad Halim (9 Telecom), Glen Larwill, Nick Monjo (PSINet Australia), Tim Griffin (AT&T Research), Rene Hatem (CA*net3), Todd Braning (Level3), Joe Klein (Cogent), Ren Nowlin (Carrier1), Joe Provo (RCN), Vince Fuller (MFN), Mike Leber (Hurricane Electric), Brokaw Price (Yahoo! Broadcast), Waqar Khan (Qwest), John Osman (rigozsaurus.com), Mike Hughes (LINX), Ted Hardie (Equinix), Lane Patterson (Equinix), Jason Starr (Equinix), Nic Lewis (AduroNet), Wolfgang Tremmel (via Net.Works), John Zweibel (Procket networks), James Rice (BBC Internet Services).

About the Author

Mr. Norton's title is Co-Founder and Chief Technical Liaison for Equinix. In his current role, Mr. Norton focuses on research on large-scale interconnection and peering research, and in particular scaling Internet operations using optical networking. He has published and presented his research white papers ("Interconnections Strategies for ISPs", "Internet Service Providers and Peering", "A Business Case for Peering") in a variety of international operations and research forums.

From October 1987 to September 1998, Mr. Norton served in a variety of staff and managerial roles at Merit Network, Inc., including directing national and international network research and operations activities, and chairing the North American Network Operators Group (NANOG) Internet industry forum. Mr. Norton received a B.A. in computer Science and an M.B.A. from the Michigan Business School, and has been an active member of the Internet Engineering Task Force for the past 15 years.

[38] At this point, all peering costs are covered by the cost savings of free traffic exchange with peers at the exchange.

Appendix A – Peering Financial Model

Mbps	Transit Costs $ Per Megabit-per-second
1-15 Mbps	$425
16-30	$395
31-44	$365
45+	$325

DS3 Circuit Cost	$	1,000
100M Port Cost	$	1,000

Mbps Exchanged	Peering Cost Per Mbps	Transit Cost per Mbps	Savings Per Mbps
1	$ 2,000	$425	$ (1,575)
2	$ 1,000	$425	$ (575)
3	$ 667	$425	$ (242)
4	$ 500	$425	$ (75)
5	$ 400	$425	$ 25
6	$ 333	$425	$ 92
7	$ 286	$425	$ 139
8	$ 250	$425	$ 175
9	$ 222	$425	$ 203
10	$ 200	$425	$ 225
11	$ 182	$425	$ 243
12	$ 167	$425	$ 258
13	$ 154	$425	$ 271
14	$ 143	$425	$ 282
15	$ 133	$425	$ 292
16	$ 125	$395	$ 270
17	$ 118	$395	$ 277
18	$ 111	$395	$ 284
19	$ 105	$395	$ 290
20	$ 100	$395	$ 295
21	$ 95	$395	$ 300
22	$ 91	$395	$ 304
23	$ 87	$395	$ 308
24	$ 83	$395	$ 312
25	$ 80	$395	$ 315
26	$ 77	$395	$ 318
27	$ 74	$395	$ 321
28	$ 71	$395	$ 324
29	$ 69	$395	$ 326
30	$ 67	$395	$ 328
31	$ 65	$365	$ 300
32	$ 63	$365	$ 303
33	$ 61	$365	$ 304
34	$ 59	$365	$ 306
35	$ 57	$365	$ 308
36	$ 56	$365	$ 309
37	$ 54	$365	$ 311
38	$ 53	$365	$ 312
39	$ 51	$365	$ 314
40	$ 50	$365	$ 315
41	$ 49	$365	$ 316
42	$ 48	$365	$ 317
43	$ 47	$365	$ 318
44	$ 45	$365	$ 320
45	$ 44	$325	$ 281

Appendix B – Top Reasons NOT to Peer

During the conversations with ISP and Content Player Peering Coordinators, several reasons for not peering were uncovered. The following represents the most common reasons that ISPs have given for not peering:

1) "We already get the traffic 'For Free' through existing peering relationships."

In this example, Content Player Y is purchasing transit from ISP E and has targeted ISP A as a high volume destination and therefore as a target for peering.

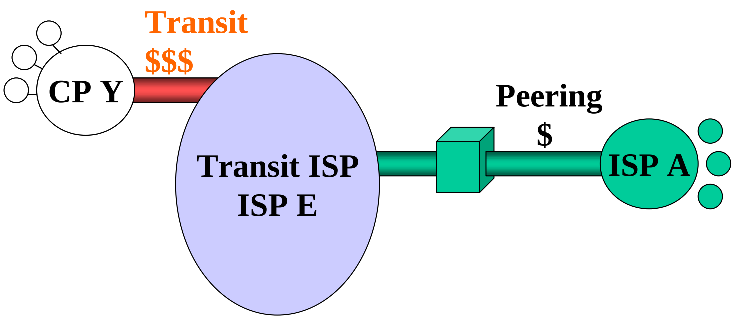

However, ISP A already receives this traffic via a peering arrangement with ISP E, and therefore has no financial incentive to peer directly with Content Provider Y.

As stated in this document, there is a *side effect* financial benefit in that the more direct traffic flow across a direct peering relationship will improve the performance, but in this example, it was insufficient to warrant peering.

2) "We are not true peers."

The reasoning here is that the benefits of peering to the two ISPs are disproportional; this ISP will benefit far less by peering than the potential peer. In the example below,

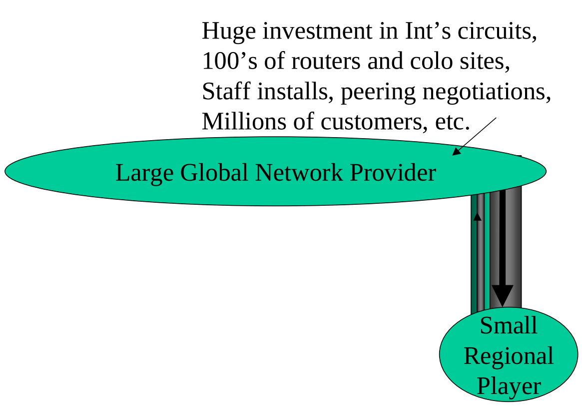

the small regional player seeks peering with a large global ISP. The small regional network offers 100 dial-up customers worth of traffic and routes while the large global player can offer 10,000 customers around the world that required years of expensive deployment and infrastructure. Clearly the balance of value is asymmetric, and for this reason the Large Global ISP chooses not to peer with the small regional player. The phrase often heard associated with this reason is "I don't want to haul your traffic around the globe for free.[39]"

3) "Lack of Technical Competence."

Since peering is of mutual benefit to the peers, there is a mutual dependence on the reliability of the peering components (transport, exchange equipment, etc.). When problems arise such as configuration errors, precious resources can be squandered debugging problems that the peer should have been able to solve independently. This lack of technical expertise causes a drain on the operations resources, and unless the value of peering is high, the ISP will prefer not to peer with this substandard ISP.

4) "Transit Sales Preferred."

This argument suggests that the potential peer is also a potential customer for the ISP. Since there is revenue associated with a sale and none typically associated with peering, peering requests sometimes get funneled to sales[40].

5) "BGP is Tough."

There are conceptual hurdles associated with the required network engineering for peering.

---

[39] John Curran, XO Communications and formerly CTO of Genuity.

[40] Conversation with Martin Levy (formerly with Concentric and XO Communications) indicated that peering@concentric.net was automatically forwarded to sales@concentric.net !

Specifically, there are three conceptual hurdles associated with peering.

First, purchasing transit requires little of no network engineering expertise. All traffic is simply forwarded to the upstream ISP.

There is a conceptual hurdle as the ISP purchases transit from a primary and a backup ISP. Setting this up requires configuration of a router using BGP to select the proper path when the primary is operating properly. To ensure the backup transit service works the network engineers must test the failover case, and make sure the primary service is reselected when it returns to service.

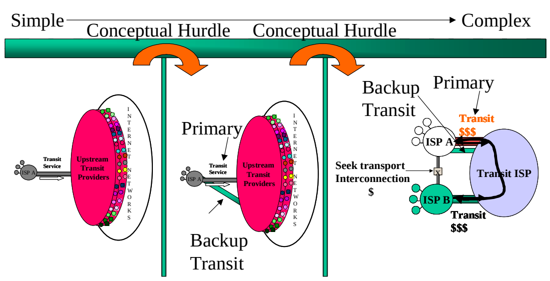

Finally, there is a further conceptual hurdle as the ISP explores migrating from a primary with backup transit service to configuring BGP to peer with potentially many ISPs at an exchange point.

These represent a great conceptual hurdle for some ISPs and Content Providers, but one that is easily overcome with networking staff.

6) Personality Clashes

Interestingly, there are sometime personality clashes that prevent two ISPs from peering. Further, peering discussions and negotiations are sometimes contentious and troublesome, amplifying any personality conflicts that may already exist. These have led to a surprising number of failed peering negotiations.

<wbn@equinix.com>