

# Economics of peering

*Steve Gibbard*

*[Packet Clearing House/Steve Gibbard Consulting](#)*

*October, 2004*

Economics of peering .....	1
Introduction.....	1
What is peering? .....	1
What is an exchange point? .....	1
Why do networks peer? .....	2
Why don't networks peer? .....	2
Considerations involved in deciding whether to peer.....	3
Economics of peering .....	3
History of peering perceptions.....	3
Bill Norton's paper .....	4
How to do your own calculations .....	5
Geographic factors.....	8
How does this affect big American cities? .....	9
Acknowledgements.....	10
Contact Information.....	10

## ***Introduction***

I've been asked to start this paper off with a brief introduction to peering. Those of you already familiar with that topic can safely skip the first section. I'm going to follow that with some information on the economics of peering, both in terms of how you can figure out if peering makes economic sense for your network, and how to fit this into a more global context.

## **What is peering?**

Peering is a relationship between ISPs or other Internet networks in which they exchange customer traffic, generally for free. This differs from the usual paid transit arrangements in that the two networks send each other traffic only going between their networks and their customers. Neither sends the other traffic intended to get out to the rest of the Internet. While peering can happen anywhere where two networks can connect to each other, it is usually done at an exchange point.

## **What is an exchange point?**

An exchange point is a facility where networks interconnect, such as the PAIX and Equinix facilities scattered throughout the US, the LINX in London, and many others. An exchange point is generally an ethernet switch that all the participants plug into and use to establish BGP sessions between their networks. Many exchange points, or the colocation providers who host them, also offer private cross connects, cables going directly between networks in their facilities that the networks can use for interconnections. Private cross connects are useful when two networks have a large amount of traffic going between them, and don't want to fill up the capacity of their exchange point switch ports.

## **Why do networks peer?**

Networks peer for a number of reasons. Sometimes they think it's cheaper, and sometimes it actually is. Sometimes they think it improves performance, and sometimes it does. Networks also peer to gain more control over their routing.

The arguments about lowering costs and improving performance are fairly simple. The price argument is that peering is free or close to it, and thus lowers costs over transit which has to be paid for. The performance argument is that data should get where it's going faster if it has to pass through fewer networks.

The control rationale is that the more paths you have, the more different ways you can send data. Your transit providers can fall over and die without impacting the traffic going between you and your peers. If you've got a problem with traffic following one path, you can divert the traffic to another path and hopefully make the problem go away.

## **Why don't networks peer?**

There are a number of reasons why networks avoid peering.

The most common is probably the "what's peering?" factor. When networks buy transit, they get a port that they can plug into, without having to think much about it. It's a standard purchase that small office network administrators and even home DSL users are familiar with, and even for fairly large network operators, it works. A lot of network operators either don't realize peering exists, don't realize it's available to them, or have no idea how to get started.

Peering can be expensive, so some network operators decide it doesn't make financial sense, or that they can't afford it. While peers don't generally charge each other for peering, that doesn't make it free. Exchange point operators often charge for use of the peering switch. If the traffic to be handed off via peering is originating somewhere other than the exchange point location, there are extra circuits and maybe extra routers to buy, and rack space to rent. Sometimes, the so-called free peering can end up being more expensive than the paid transit it would replace.

Peering can be a lot of work. Coordinating peering for a network means figuring out who to peer with and pestering them until they turn up the sessions, at which point the sessions need to be maintained. With transit, you've generally got a couple of connections that either work or they don't, and if they don't you can call your transit provider and yell at them. With peering, you've suddenly got a lot of other networks to deal with.

Additionally, if your transit provider is peering locally, performance gains from peering may be minimal, or they may be entirely fictional. If data is going from your network, to a router at the nearest exchange point, and then on to other exchange point participants, does it really matter from a performance perspective who owns the router at the exchange point? In this sense, you can think of transit as paying somebody to do your peering for you.

## **Considerations involved in deciding whether to peer**

When deciding whether to peer, there are a number of things to think about. There are the economic aspects of it, both in terms of the cost of getting to the exchange point, and your own labor cost to maintain it. There are the performance considerations: will peering improve your performance beyond what your transit provider is doing for you? And then there are some harder to quantify factors. How much is the extra control peering gives worth to you? How much are shorter AS paths worth, if they cause your own transit customers to send more traffic in your direction? Is there marketing benefit to it?

## ***Economics of peering***

### **History of peering perceptions**

Over my several years in this industry, not having been closely involved in peering until the last year or two, I've heard lots of perceptions about the economic effects of peering.

First, there was "Peering is free! Why pay for something when you could get it for free?" The thinking was that transit cost money, while peering was connectivity being given away. Who would pass that up?

The converse was the attitude of one of my former employers: "If peering is free, what incentive do my peers have to help solve problems?" This was also sometimes characterized as "why get something for free when you could pay for it?" The thinking here was that with transit, they could call the transit providers and get them to fix things, while peers would have no incentive to help. But it was questionable whether this did

any good, or whether it just left transit providers pleading with their peers for help when a problem came up.

Then there's been the issue of big networks refusing to peer with smaller networks, and the vitriol surrounding that. To some degree, this has just been a matter of networks wanting to make sure they weren't carrying all the long-distance traffic for somebody else, but where that wasn't the case it was a bit more complicated. Big networks often don't want to help their smaller competitors grow, and feel that peering with their smaller competitors will give their competitors an advantage. The less sinister big network versus small network issue involved the bigger network no longer getting any benefit from new peering relationships. If they had already become 100% transit free, new peering simply offloaded traffic from older peering sessions, adding whatever costs were involved in the new sessions without saving them anything on transit.

The lesson from all those fights is that peering rarely happens without some perception of mutual benefit. If a network doesn't think peering is good for them, it probably won't happen, no matter how much the other side stands to benefit.

## **Bill Norton's paper**

At Equinix's peering forum a month ago, Bill Norton presented a paper called *A Business Case for Peering in 2004*. Bill surveyed several network operators to ask how much they were paying for peering and transit, and came up with some numbers on when it makes sense to peer. According to Bill, an ISP not already colocated at an exchange point can break even on peering when they have 204 Mb/s of peerable traffic.

Bill made a number of assumptions to get to that number:

- A gigabit ethernet port at an exchange costs \$5,000.
- Gigabit Ethernet transport to an exchange costs \$4,000.
- A used Cisco 12008 can be bought for \$7,500 per month, and will last three years.
- The total cost of transit at gigabit per second commit rates is \$30 per Mb/s.
- 33% of an ISPs traffic will be peerable at a single exchange point.

Bill also produced calculations for an ISP that is already colocated at an exchange point. In that case, he says peering makes sense when 33.33 Mb/s can be peered. This number relies on assumptions similar to those above, except for:

- A fast ethernet port at an exchange costs \$1,500.
- Hardware required is a fast ethernet interface on a router, of negligible cost.

As a technical correction to the formula Bill used to produce his numbers, it should be pointed out that he left out labor costs involved in peering. Had he included labor, it would have increased the amount of traffic required to justify peering. The more significant caveat required for Bill's numbers are that even in markets where transit costs

are fairly uniform, transport costs can vary considerably depending on where the other end of the circuit is. An ISP in a well-connected telecommunications building will likely pay considerably less, for example, than an ISP in a residential neighborhood. It is therefore important, if you are trying to figure out if it makes sense to peer, to redo these calculations with your own numbers.

## How to do your own calculations

The first thing to figure out is how much traffic could be offloaded to peering at the various nearby exchanges. This requires knowing where your traffic is going, which is most easily measured using Netflow data. Netflow is a protocol that routers use to export information about the packets flowing through them. There are a number of systems for parsing this data.

Arbor and Adlex make what are probably the two nicest traffic analysis systems, both of which can give you up to the minute information on where your traffic is going, what the likely result would be if you turned up peering with a given network, and so forth. Unfortunately, last time I checked, they both had price tags in the range of \$100,000.

Flow-tools is a free package that is very good at giving you the destination addresses and destination ASes of your traffic, among other bits of useful information. This has a problem for doing peering analysis in that what you really need is a number not just for how much traffic you would be sending to a specific AS, but to that AS and everything behind it. This would be solvable using Flow-tools with an appropriate front-end, although last time I checked I couldn't find anybody who had published one.

There was some software that PCH developed several years ago in conjunction with Agilent Technologies that did answer the question of how much traffic was going to everything behind a specific AS. Unfortunately, it was written in some pretty inefficient Perl code, and used cflowd for its flow collection. In testing, it has a tendency to fall over and die when trying to push more than 200-300 Mb/s through it.

There are also several other tools out there for this sort of analysis. If anybody knows of something cheap that meets the requirements I'm suggesting, please let me know.

Once you've got your flow data, you need to compare that to the networks available for peering at the exchange points you're considering. Most exchange operators publish a list of their participants, and for some exchanges you can also go to the PCH looking-glass (<http://lg.pch.net>) to see what routes the various exchange participants are announcing. Since networks have a variety of peering requirements (and some of them just don't answer requests), it's also important to know which of the networks present at an exchange are likely to peer with you. This is a question for which definite answers may be gotten from the networks' peering coordinators, or which those who have been involved in the peering community for a while may be able to help you guess at.

Once you have your traffic numbers, you need to know what peering and transit cost. You probably already know what you're paying for transit. The sales people from your local exchange point should be able to tell you what they'll charge you for a port there (and if the first number they give you doesn't work out, they may be able to come back with better numbers). Other costs that need to be included in both numbers are:

- Transport costs: Are you able to plug your network directly into your transit connections or peering ports, or are circuits required to connect them to your network? How much are you paying for those circuits?
- Equipment costs: Can you plug transit and peering into ports you already have, or will it require new routers or interfaces?
- Labor costs? How much will it cost you to set up new routers for peering, to negotiate peering sessions, and to maintain those sessions and deal with the peers once they're up?

It's important to keep in mind that for transit, you're most likely dealing with incremental costs. Taking some of the traffic of a transit connection probably won't reduce the need for an interface to plug the connection into, unless you can offload enough transit to get rid of the transit circuit. So, your cost calculations may look like this:

Transit per Mb/s = Cost of transit circuit per Mb/s

Or, if it does allow you to reduce your size of transit connections:

Transit per Mb/s = cost of transit per Mb/s + ((cost of circuit to connect to transit provider + (cost of interface card required to plug in that circuit / months of amortization) + cost of labor to deal with new transit circuit) / number of Mb/s)

Peering per Mb/s = (Cost of peering port + cost of circuit to get to exchange point + (cost of equipment / months of amortization) + cost of colocation at exchange point) / Mb/s.

Using this formula, we can go through some example calculations:

### **Medium sized ISP:**

- 500 Mb/s total traffic.
- 20% offloadable to peering.
- Using Bill Norton's numbers:
- Router costs \$7500, or \$312.50 per month over three years.
- Peering port and rack costs \$5,000.
- Transport to exchange costs \$4,000.
- Additionally, assuming \$1,000 per month for labor.

In off-net building – assume transit costs \$60 per Mb/s:

- Peering causes \$4,312 monthly loss.

- Required port cost for break even would be \$688.

In on-net building with transit; off-net with exchange – assume transit cost of \$45 per Mb/s:

- Peering causes \$5,812 monthly loss.
- For break even on a free peering port, transport costs would have to be less than \$3,188.

In Exchange point building – assume transit costs \$45 per Mb/s; no transport costs; router needs \$1,500 GigE card:

- Peering causes \$3,062 monthly loss
- Required port cost for break even would be \$1,938.

**Larger ISP (1 Gb/s total traffic):**

- 1 Gb/s total traffic.
- 20% offloadable to peering.
- Using Bill Norton's numbers:
- Router costs \$7500, or \$312.50 per month over three years.
- Peering port and rack costs \$5,000.
- Transport to exchange costs \$4,000.
- Additionally, assuming \$1,000 per month for labor.

In off-net building – assume transit costs \$45 per Mb/s:

- Peering causes \$1,313 monthly loss.
- Required port cost for break even would be \$3,687.

In on-net building with transit; off-net with exchange – assume transit cost of \$30 per Mb/s:

- Peering causes \$4,313 monthly loss.
- Required peering port cost for break even would be \$687.

In Exchange point building – assume transit costs \$30 per Mb/s; no transport costs; router needs \$1,500 GigE card:

- Peering causes \$62.50 monthly loss
- Required port cost for break even would be \$4,937.50.
- A GigE port cost of \$4,000 would cause peering to make a lot of sense in this scenario.

**Even larger ISP (2 Gb/s of traffic):**

- 2 Gb/s total traffic.
- 20% offloadable to peering.
- Using Bill Norton's numbers:
- Router costs \$7500, or \$312.50 per month over three years.
- Peering port and rack costs \$5,000.
- Transport to exchange costs \$4,000.
- Assuming \$1,000 per month for labor.

In off-net building – assume transit costs \$45 per Mb/s:

- Peering causes \$7,688 monthly gain.
- Port could cost up to 12,687 while allowing break even.

In on-net building with transit; off-net with exchange – assume transit cost of \$30 per Mb/s:

- Peering causes \$1,687 monthly gain.
- Required peering port cost for break even would be \$6,687.

In Exchange point building – assume transit costs \$30 per Mb/s; no transport costs; router needs \$1,500 GigE card:

- Peering causes \$5, 937 monthly gain.
- Required port cost for break even would be 10,937.

It's important to keep in mind that all of these are examples with made up numbers. Any results you get using your own numbers may be considerably different.

Whatever the actual costs are, it's obvious that there are some economies of scale at work here. Even in the best-case peering cost scenarios, there will always be some small networks that will do better to pay some fraction of another network's peering costs (buying transit) rather than peering on their own. In the absence of a complete transit monopoly, those small networks will always benefit from some amount of peering somewhere upstream from them.

## ***Geographic factors***

Transit costs vary considerably in different parts of the world, making peering look considerably more attractive in various other places than it does here. Traffic in the so-called "Internet core" region is relatively cheap. The numbers Bill Norton found in his survey showed even those committing to one Mb/s getting transit for \$125 per Mb/s. In contrast, an ISP in Northwest Montana, with no nearby big backbones and not much telecommunications connectivity coming into their town, says they pay a bit more than \$1,000 per Mb/s. ISPs in Kathmandu, Nepal, are buying their international transit for \$5,000 per Mb/s, over high-latency satellite connections. The high costs and poor

performance are keeping traffic volumes pretty low in those areas, and thus seem to be a significant barrier to Internet development.

While these prices are falling slightly, they seem unlikely to drop anywhere near the levels we see in developed areas of the US any time soon. The more easily solvable problem is that in many places where international transit is very expensive, local traffic goes over the international lines as well. So, not only is an ISP paying \$5,000 per Mb/s in some cases to send local traffic, and another ISP is paying a similar amount to receive it, but if the traffic is going over two satellite links there is also at least a full second of latency involved. It is in these areas where local peering can still be very attractive, even for small ISPs.

Peering becomes financially easier for small ISPs to justify in areas with expensive transit for a number of reasons. Peering can cost more in these areas since the transit it's competing with costs more, but it's much better if it costs significantly less. When peering is significantly cheaper than transit, pipes get bigger, congestion goes away, and traffic volumes increase. End users notice this as a performance gain, and billing departments notice this as more traffic that customers can be billed for.

To use the Kathmandu example again, we get the following prices. Since the Nepal Internet Exchange is run by its member ISPs, switch ports are currently free. They're talking about increasing the switch port price to \$800 per year to cover some spare equipment. 2 Mb/s circuits to get to the exchange cost around \$13 per month. These circuits (dry copper) require about \$1,000 worth of equipment, but when spread out over three years that's \$27 per month. The total cost of connecting to the exchange ends up being \$107 per month for the first two Mb/s, and \$40 per month for each additional Mb/s. So, in the worst case scenario, the cost for local peering ends up being \$53.50 per month per Mb/s, as compared to \$5,000 per month per Mb/s for transit.

This can be expected to have several outcomes. Peering in these areas should make connectivity for end users significantly cheaper, if a substantial portion of their traffic isn't leaving the area. Indeed, one of the broadband providers in Nepal allows its customers to send local traffic for free, while charging for international traffic. ISPs should be able to sell transit at considerably lower prices than those charged by foreign satellite operators, passing savings on even to networks that don't peer directly. Locally hosted content becomes attractive, both because the end users can get to it cheaply and because performance is improved significantly, improving business for local hosting companies. And, assuming there's local access to the DNS, the region becomes much better protected against external Internet outages.

## **How does this affect big American cities?**

As we can see from extreme examples, moving data across the Internet is significantly cheaper when it doesn't have to be carried across long distances. An interesting question is what this means for situations where the distance differences are a bit less. Does this,

for example, mean that transit should cost less in Chicago, where there's lots of local peering, than it does in Detroit, where almost all traffic needs to be hauled to Chicago, 300 miles away?

Here's my hypothesis: Transit is cheap when transit providers can avoid carrying data over expensive circuits. As long as circuit pricing is distance sensitive, networks should be able to sell transit more cheaply in areas where they peer locally. If this isn't the case now, it should become the case as traffic volumes grow, and the cost of carrying the traffic starts to exceed the cost of dealing with local peering.

However, this theory has several problems. The complexity of managing additional peering points may offset the cost savings of local peering. Where peering ports are expensive, aggregating the traffic into a smaller number of peering ports may offset the cost of carrying the traffic over some level of distance. Big carriers may have lots of excess fiber, and may consider its cost to be zero. And, this may all be moot, because networks may not monitor their costs to this level.

I've been doing lots of asking people at ISPs in various areas how much they're paying for transit, in an attempt to substantiate this. Unfortunately, that tends to be data that networks treat as proprietary, and that's covered under a lot of NDAs, so I didn't end up with enough data to be conclusive.

## ***Acknowledgements***

My PCH colleagues Gaurab Raj Upadhaya, Tom Vest, and Bill Woodcock contributed greatly to this paper. Gaurab supplied much of the information about Nepal, while Tom served as an actual economist I could bounce ideas off of. Bill has contributed much of the background for my thinking on this. Bill Norton supplied much of the material for the middle section of the paper. Frank Fifield at Kootenai Valley Internet Service supplied my pricing information for Northwest Montana. Chris Quesada at Switch and Data gave me a forum to present this in, and a deadline for finishing it.

## ***Contact Information***

*Steve Gibbard*

[scg@pch.net](mailto:scg@pch.net)

<http://www.pch.net>

<http://www.stevegibbard.com>

+1 510 528-1263