

## The Folly of Peering Traffic Ratios?

William B. Norton <[wbn@equinix.com](mailto:wbn@equinix.com)>

### Abstract

*Internet Service Providers peering policies fall roughly into four categories: Open, Selective, Restrictive, and no-peering. Those with an Open Peering Policy peer with anyone in any location(s), while Selective and Restrictive Peers use prerequisites to screen prospective peers. One potential prerequisite is a “Not to Exceed” Peering Traffic Ratio which specifies the maximum allowable asymmetry in traffic exchange to be considered a “peer”. This of course excludes content heavy ISPs and Content Companies from settlement-free peering relationships. But does this make sense? Is there a technical and/or business rationale for using peering ratios in this manner?*

*Several hundred people attended the Peering Birds-of-a-Feather session at NANOG 35 where this issue was debated. During and after the debate, dozens of seasoned peering Coordinators discussed and debated the six strongest arguments for using peering ratios as a peering candidate discriminator. This paper highlights and discusses both sides of this heated debate.*

*(Editor’s note: The NANOG 35 Peering BOF Debate was held in the middle of the Level3-Cogent de-peering event that led to a partition across parts of the Internet. Since then, the Level3-Cogent peering has been re-established with no information released regarding the terms of that peering. However, an observant Peering Coordinator noticed that the Level3 Peering Requirements no longer include explicit peering ratios to be maintained<sup>1</sup>.)*

### Introduction: Peering Prerequisites

Obtaining Peering<sup>2</sup> with companies with an **Open** peering policy can be as easy as a phone call and typing a few lines of router configuration. The technical component of peering is pretty straightforward.

It is slightly more difficult to negotiate peering

---

<sup>1</sup> Ren Provo pointed this out to the author with reference to the new Level3 Peering Policy: <http://www.level3.com/1511.html>

<sup>2</sup> See Appendix A for definitions of Peering, Transit, Peering Ratios, etc.

with a **Selective** Peer - one that has some peering prerequisites<sup>3</sup>. During these initial peering discussions one is likely to hear a variety of peering requirements including traffic minimums, 24/7 NOC and contact information exchange, peering upgrade policies, number of interconnections and geographic distribution requirements, etc<sup>4</sup>. There are rational technical and operational reasons for each of these<sup>5</sup>. Occasionally, peering requirements include a “Not to Exceed” traffic ratio to ensure that aggregate peering traffic ratios will not exceed a particular ratio (quoted in Out:In – not to exceed 2:1 is a common requirement).

The Peering Coordinator Community put on a debate on the rationality of peering ratios as a peering discriminator at NANOG 35 in Los Angeles. During that debate, and during the subsequent informal debates afterwards, the consensus was that this metric was neither technically sound nor business rational<sup>6</sup>.

For the duration of this paper we will present the strongest arguments on both sides of the issue. We present the strongest arguments for using peering traffic ratios as a peering discriminator, and then the counterpoint(s) discrediting the argument.

### Argument #1 – “I don’t want to haul your content all over the world for free.”

There are two hidden assumptions within this claim: First, that there is no revenue realized as a result of hauling this content traffic, and second, that there is an increased load on the backbone by

---

<sup>3</sup> For some sample peering policies with Peering Ratios explicitly mentioned, see MCI (AS701) Peering Policy at <http://global.mci.com/uunet/peering/>, AT&T (AS7018) <http://www.att.com/peering/>, AOL ATDN (AS1668) [http://www.atdn.net/settlement\\_free\\_int.shtml](http://www.atdn.net/settlement_free_int.shtml), Tiscali (AS3257) <http://www.ip.tiscali.net/peering-policy.html> or nac.net (AS8001) <http://eng.nac.net/peering.html>

<sup>4</sup> The gist of these are specified in Appendix B.

<sup>5</sup> Rationale include: decrease time to repair by mutual maintenance and exchange of contact info, improved resiliency and headroom from multiple interconnect points and peering link upgrade policies, etc.

<sup>6</sup> The debate audience vote was about 100 to 3 that the metric was not rational. “Business Rational” is the author’s term to indicate the long-term rationality/irrationality of a practice from an overall business perspective.

# The Folly of Peering Traffic Ratios

“hauling” this traffic over a peering (“for free”) relationship. Let’s consider these broader arguments with a diagram.

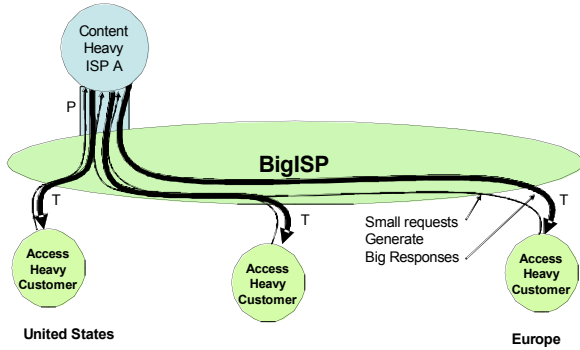


Figure 1 – Content Heavy ISP A Peers with Big ISP.

In this example (see figure 1 above), the Content Heavy ISP A peers with a Big ISP which has a large global network spanning the United States and Europe. In this example, the Content Heavy ISP A is peered with the Big ISP Network in a single location and has a 30:1 outbound to inbound traffic ratio<sup>7</sup>. As is typical with today’s web (http) traffic, we see small requests from the Big ISP access heavy customers generate large responses from the Content Heavy ISP A.

Next (see figure 2 below) let’s consider an Access Heavy ISP A in the exact same situation: Access Heavy ISP A peers with Big ISP, and Access Heavy ISP A Customers send small requests across to Big ISP’s Content customers, which in turn send big responses back.

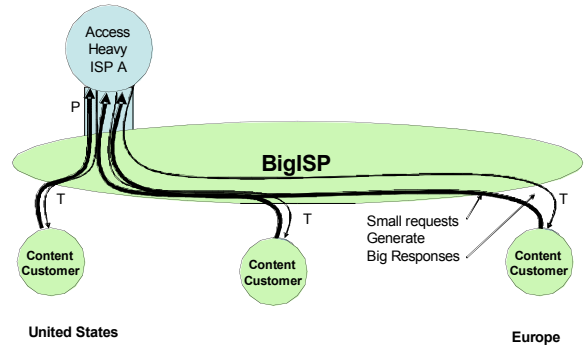


Figure 2 – After Access Heavy ISP A Peers with Big ISP

Notice that the Big ISP Backbone load is **identical** in these two scenarios.

**Counter Argument #1** – Since the load on the backbone is the same in both scenarios, there is no technical rationale for discriminating against one versus the other – the load on the Big ISP is identical!

**Counter-Argument #2:** The Big ISP in these scenarios is hauling the traffic across the globe on behalf of its customers. Massive outbound asymmetry or massive inbound asymmetry is an artifact of the type of customer served, regardless of whether served via peering or transit relationships.

What **is** different between these two scenarios is who is paying and in which direction the traffic is loading the backbone. In the first scenario, the Big ISP **Content Heavy** customers are paying Big ISP to distribute its content to the entire Internet (including to Access Heavy ISP A customers). In the second scenario, the Big ISP **Access Heavy** customers are paying the Big ISP for Transit (access to the global Internet including the Content Heavy ISP A content). In both situations, Big ISP is being paid for the traffic exchange – it is not “hauling the traffic all over the world for free.”

**Argument #2 – “OK, but there is massive asymmetry here. Look at how many bit miles I have to carry your content, while you have only to deliver the traffic across the exchange point.”**

**Counter Argument #1** - This is a valid observation, and perhaps a good argument for requiring a peer to spread the peering load across

<sup>7</sup> This isn’t out of the realm of possibilities. Richard Steenbergen (nLayer) points out "In a purely content to eyeball exchange, a typical outbound packet with payload is 1500 bytes, while the typical TCP ACK response is between 40-64 bytes." This is roughly a 30:1 ratio.

## The Folly of Peering Traffic Ratios

multiple geographically distributed interconnect regions.

This is not however an argument for using Peering traffic ratios to restrict Peering. As demonstrated earlier, Peering traffic ratios are irrelevant; eventually this traffic traverses your network to get to your customers. The asymmetry is merely a side affect of the type of traffic requested by or served on behalf of your customers.

One discussion with a large backbone provider highlighted the folly. This Peering Coordinator was proud that the peering ratio required a peer to “adjust” traffic to maintain an “appropriate” traffic ratio. How did the peer accomplish this shift? They redirected some portion of traffic to a transit provider who in turn forwarded that traffic directly back to the large backbone provider, onto the same router as the peer! There was no change in backbone load, simply a shift of traffic from one interface card to the other. The ISP mandating peering traffic ratios realized no cost savings, enhanced revenue, or performance benefits from this mandate. In fact, traffic now takes a more circuitous path to reach their customers yielding slightly worse performance.

**Counter-Counter Argument** – Multiple interconnect points may make the problem worse. With hot potato routing the large volume of content enters my backbone sooner! On the other end of the spectrum, “cold-potato” routing that some have advocated to distribute the load more evenly across the interconnect points simply does not scale. Cold Potato routing using Multi-Exit Discriminators (MEDs) requires breaking up aggregate routing blocks into more granular chunks that can be routed to a specific interconnect point. This breaks the scaling characteristics of CIDR, and can become an operational nightmare as load-balancing and fail-over scenarios are encountered. So, unless there is a way to effectively spread the traffic across the interconnect points without breaking things, I will have to carry the traffic further; more costs to me.

**Counter-Counter-Counter Argument** – But you will have to pick up that traffic somewhere, either via a transit link that costs you money or a peering link. I can get you that same traffic for free more directly so with lower latency. No matter what, you have to carry that traffic to your customers.

**Counter-Argument #2** – The asymmetry is simply a short-term artifact of the currently popular HTTP web traffic. For the longer term, this will change. For example, there may be more peer-to-peer traffic, or other types of traffic that will swap the

traffic asymmetry. Any policy based on traffic ratios is short-sighted.

### Argument #3 – “I don’t want to peer with Content Heavy ISPs because doing so will screw up my peering traffic ratios with my other peers.”

Let’s understand this argument with another example.




Figure 3 – Before Eyeball Heavy ISP Y peers with Content Heavy ISP A

Pictured here (figure 3) we have a snapshot of a part of the Internet that includes a Content Heavy ISP A, his upstream ISP X, and a peer of ISP X that we call Access Heavy ISP Y. Being Content heavy means that the small requests coming in are answered with large responses. The directed lines and their width highlight the traffic flows between Access Heavy ISP Y customers who request content from Content Heavy ISP A. Assume that the upstream ISP X and ISP Y peer with each other, and that the Traffic Ratio between ISP X and ISP Y is 1:1 in this initial scenario.

Now let’s assume that Content Heavy ISP A peers with Access Heavy ISP Y (see figure 4 below).

## The Folly of Peering Traffic Ratios

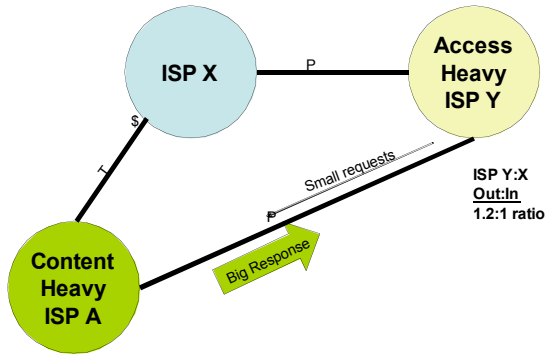


Figure 4 – After Eyeball Heavy ISP Y peers with Content Heavy ISP A, ISP Y ratios with ISP X are *worse* (ISP Y appears to ISP X to be more content-heavy than before). Let's consider the effect.

The traffic between Access Heavy ISP Y and Content Heavy ISP A no longer traverses ISP X but instead traverses the direct interconnection. What is the impact? Migrating those large traffic flows away from the ISP X to the ISP Y interconnect make ISP Y appear to be more content heavy to ISP X than in the previous situation. Why? There is now relatively less content going to ISP Y from ISP X, while the same amount of content is coming from ISP Y. This makes ISP Y's ratios with ISP X slightly worse, maybe 1.2:1 now, after the peering.

**Conclusion 3A: If an Access Heavy ISP peers with a Content Heavy ISP, it can adversely affect its peering traffic ratios with its other peers.**

Now let's consider the opposite case – what happens when a content heavy ISP peers with an Access Heavy ISP? What happens to the Content Heavy ISP's peering traffic ratios with its other peers?

In this next picture (figure 5) we have an Access Heavy ISP A that purchases transit from ISP X, who has a peering relationship with Content Heavy ISP Y.




Figure 5 – Before Content Heavy ISP Y peers with Access Heavy ISP A

Small requests from Access Heavy ISP A's customers generate large responses from Content Heavy ISP Y's content customers. Let's again assume here that ISP X and ISP Y have a 1:1 peering traffic ratio to begin with.

Now ISP A and ISP Y peer with each other, resulting in the following:




Figure 6 – After Content Heavy ISP Y peers with Access Heavy ISP

All traffic between ISP A and ISP Y now traverses the direct interconnection, shifting a large chunk of traffic from the Y:X peering relationship to the Y:A peering relationship. Now ISP Y appears less Content Heavy (from ISP X's perspective), improving his peering traffic ratios with ISP X and others perhaps to .8:1.

**Conclusion 3B: If a Content Heavy ISP peers with an Access Heavy ISP, it can positively affect its peering traffic ratios with its other peers.**

## The Folly of Peering Traffic Ratios

**Counter-Argument:** But herein we reveal the folly.

(False) Conclusion 3: Peering traffic ratios are valid discriminators because they help maintain peering traffic ratios with other (existing) peers.

**While demonstrable, THIS IS A CIRCULAR ARGUMENT for the general question at hand.**

Consider the simplified restatement of this argument: “Peering traffic ratios are a valid metric because they support (other) peering traffic ratios.”

When confronted with this circularity the advocates offer argument #4.

### **Argument #4 – “I want revenue for carrying your packets.”**

This argument is business rational and comes in many forms:

- “You look like a sales prospect – you should buy transit,” and
- “If I peer with you, then you will need what I have less.” Hidden assumption here: direct access to my customers is a major selling feature.<sup>8</sup>
- “Establishing Peering Traffic Ratios allows me to systematically deny peering to my Content Customer”

**Counter-Argument #1** - All peering and transit relationships are negotiable business relationships, but peering traffic ratios are probably not the right metric for determining the terms of the relationship.

**Counter-Argument #2** - As discussed above, the cost of distribution is the same regardless of whether content or access heavy. A more appropriate negotiating discussion surrounds the incremental value provided to the content or access customers by the network in question.

**Counter-Argument #3** - And remember, you will have to carry this traffic by some means (transit or peering) to get it to your customers. So when negotiating, you are just negotiating for a different delivery mechanism to your customers – this may not command a high price.

### **Argument #5 – “My backbone is heavily loaded in one direction – I don’t have the \$ to upgrade the congested portion of the core in that direction without a corresponding increase in revenue.”**

**Counter-Argument #1** - This comes up occasionally, and may appear to be a strong (albeit short term) argument for discriminating against peering with Content heavy ISPs. A company can essentially say we can only peer with Access Heavy networks due to backbone links in one direction being saturated. However, keep in mind what this signals to the Internet community: the congestion on the network is or will soon affect the customers of this ISP. Therefore, this ISP can not meet the needs of its customers and is probably delivering poor performance at peak periods. This may signal that you are neither a good peering nor a good transit provider candidate.

**Counter-Argument #2** - This loading problem can best be solved with better traffic engineering, with the allocation of more resources to the backbone, or by temporarily denying peering until the needed upgrades are done. In any case, denying peering based on the peering traffic ratio rationale is the least attractive of all of these other options.

### **Argument #6 – “I don’t want to peer with anybody else. I don’t have to – I have all the peering that I need. Peering traffic ratios help me systematically keep people out.”**

In this argument (seldom stated out loud) Peering traffic ratios are being used as an arbitrary barrier to peering.

**Counter-Argument** - One could just as easily specify higher traffic volume, more points of interconnect, a larger backbone capacity into more geographic regions including ones that don’t make any sense. This is not a strong defense for the validity of peering traffic ratios as a rational discriminator.

## **The Problems with Peering Traffic Ratios**

Peering Traffic Ratios as a Peering Discriminator present several problems that make them not business rational:

1. Dictating Peering Ratio requirements force non-competitors to build out into many interconnect regions (perhaps into

---

<sup>8</sup> The author notes one comment a few years back from a Peering Coordinator “An email to peering@<ispdomain>.net is a cry for help, and is forwarded automatically to sales@<ispdomain>.net.”

## The Folly of Peering Traffic Ratios

your territories), to acquire (potentially your) local customer traffic from your market in order to meet your ratio requirements; essentially this requirements forces a peer to grow to look like you. So, you are creating a competitor<sup>9</sup>.

2. Dictating Peering Ratio requirements ratios encourage companies to “creatively route” your customers content or access traffic via sub-optimal and often circuitous routing in order to meet these arbitrarily set criteria<sup>10</sup>. Sub optimal performance and customer complaints drive regulatory oversight.
3. Peering directly can increase revenue. Dave Rand (AboveNet) first made this argument to the author. When peering directly, routing generally prefers the direct path. The direct path is under the control of the two parties on either end so packet loss can be monitored and managed. The direct path is also generally a lower latency path. These two combine effects allow the TCP Window to grow faster and therefore more traffic flows between the customers on either side, which in turn, generates more revenue for the peering parties<sup>11</sup>.

### Summary

In the heated discussions surrounding this topic there were six flavors of arguments put forth to support peering traffic ratios as a peering selection criteria.

Argument #1 - “I don’t want to haul your content all over the world for free.” This argument is countered with the observation that, whether peered with a content heavy or an access heavy ISP, the

aggregate load on the ISP backbone is about the same and that the ISP is being paid by its customers for providing access to that traffic. It is not hauling the traffic for free.

Argument #2 – “OK, but there is massive asymmetry here. Look at how many bits miles I have to carry your content, while you have only to deliver your content across the exchange point.” Regardless of whether peered with content or access the load on your network is about the same; the access or content traffic has to traverse your network to get to your customers. The asymmetry is simply a reflection of the currently popular HTTP traffic profile.

Argument #3 – “As an Access Heavy ISP, I don’t want to peer with Content Heavy ISPs, because doing so will screw up my peering traffic ratios with my other peers.” The analysis and diagrammed traffic flows in this paper prove this assertion true, however, the argument is circular: ‘Peering Traffic Ratios are valid criteria because they support Peering Traffic Ratios with others.’

Argument #4 – “I want revenue for carrying your packets.” While this argument is business rational, peering traffic ratios are a poor indicator of the value derived from the interconnection.

Argument #5 – “My backbone is heavily loaded in one direction – I don’t have \$ to upgrade the congested part of my core without a corresponding increase in revenue.” This loading problem can better be solved with better traffic engineering, with more resources for the core or by temporarily postponing all additional peering until the core is upgraded. Introducing a peering traffic ratio requirement is at best a temporary solution and signals a poorly managed network – a bad image for peers and transit customers alike.

Argument #6 – “I don’t want to peer with anyone else. Peering ratios help me systematically keep people out.” This is an understandable (although seldom articulated out loud) argument but is a poor defense for peering traffic ratios per se since it is an arbitrary discriminator; one could just as well have specified the size of the backbone core, the market capitalization or debt structure, or an extremely large number of distributed interconnect points.

Peering Traffic Ratios emerged years ago in the Internet as a peering candidate discriminator, but appear to have their roots in the PSTN world. Ratios reflect the telephony settlement model of buying and selling minutes, where the traffic difference between in and out is used for settling at the end of the month. However, the Peering Traffic Ratio discriminator

---

<sup>9</sup> Richard Steenbergen articulated this argument at the NANOG 35 Peering BOF X debate.

<sup>10</sup> Avi Freedman during the Peering BOF X Q&A following the debate.

<sup>11</sup> A few folks challenged this point. One can peer traffic that leads to more congestion and greater loss, and this has occurred in the past. But, at least this is in the control of the two parties that will suffer as a result of poor performance, as opposed to using a third party to transport the traffic.

## The Folly of Peering Traffic Ratios

model for the Internet interconnection fails because traffic ratios are a poor indicator of relative value derived from the interconnect.

### Acknowledgements

This paper represents the collective set of arguments presented, debated, and discussed at NANOG 35 in Los Angeles. Particularly strong arguments were shared with the author either directly or indirectly from Peter Cohen (Telia Sonera), Richard Steenbergen (nLayer), Patrick Gilmore (Akamai), Vikas Mehta (AOL), Bryan Garrett (BellSouth), Brian Sutterfield (Cox), Kevin Mulvahill (Equinix formerly from @Home), Steve Wilcox (Telecomplete), Louie Lee (Equinix), Brokaw Price (Yahoo!), Joe Provo (RCN), John Curran, Hamish MacEwan (CityLink), Brandon Galbraith, Maureen Carroll (T-Systems), Vince Fuller (Cisco), Joshua Sahala (TWTC), Ren Provo (SBC) .

# The Folly of Peering Traffic Ratios

## Appendix A - Definitions

The Internet is a network of networks, interconnected in peering and transit relationships. Below are the working definitions for this paper.



**Definition: Transit** is a business relationship whereby one entity sells access to the entire Internet to another. Transit is typically sold as a metered service, sold in Megabits-per-second (Mbps) measured monthly at the 95<sup>th</sup> percentile using 5-minute samples every month<sup>12</sup>.

**Definition: Peering** is a business relationship whereby two entities reciprocally exchange access to each others customers.

**Definition: Internet Service Provider** is an entity that provides (usually sells) Transit.

**Definition: Content Heavy** is a characteristic of an entity that has more outgoing traffic than incoming traffic.

**Definition: Access Heavy** is a characteristic of an entity that has more incoming traffic than outgoing traffic.



**Definition: Peering Traffic Ratio** is a measure of traffic asymmetry, measured in the form of a ratio of outbound to inbound traffic. An entity with a 5:1 Peering traffic ratio has five times as much traffic going to particular peer than it receives from the peer.

**Definition: Open Peer** is a peer that will peer with anyone without pre-condition.

**Definition: Selective Peer** is a peer that will peer, but has some prerequisites to be met and/or criteria to be maintained.

**Definition: Restrictive Peer** is an entity that will generally not peer with anyone else. These entities will typically not make publicly available there peering policy, or set the requirements so high as to exclude almost every potential peer. This reflects an underlying desire not to peer with anyone else.

**Definition: no peering policy** is an entity that only purchases transit.

## Appendix B – Common Peering Prerequisites

- Traffic minimums – It isn't worth our engineering and operations resources to turn up a peering session with an ISP with only a few Mbps of traffic. The cost of configuring and maintaining the peering session exceeds the value of receiving your traffic over that peering session. In other words, it is more cost effective for me to hear the traffic elsewhere (i.e. through my upstream transit provider.)
- 24/7 NOC and maintained off-hours contact information – since both parties depend on the proposed peering session operationally, we require that both parties operate a 24/7 Network Operations Center that, among other things, monitors and manages that session. Further, we both agree to exchange each others on-call staff pager and cell phone #s, etc. and keep each other informed of changes.
- Upgrade links at 40% port aggregate utilization – when either side of a public or private peering port reached 40% of the port capacity, we contractually obligate each other to migrate to a dedicated private peering session or to upgrade the public ports such that the port utilization is less than 40% of available capacity.
- Multiple geographically distributed peering interconnect points – we agree to distribute the traffic across the regions in which we operate for load sharing and resiliency reasons.
- Agree to work diligently to repair peering failures – since both parties depend on the proposed peering session we explicitly obligate each other to work

<sup>12</sup> See "Internet Service Providers and Peering" for a more thorough explanation of Peering and Transit.

## The Folly of Peering Traffic Ratios

diligently to repair breakages when they occur.

- Variety of legal clauses – venue for dispute handling and the like.
- “Not to Exceed” Traffic Ratios – to establish (and maintain) the peering we require that aggregate peering traffic ratios will not exceed a particular ratio (2:1 (Out:In) is common).