The Great (Public vs. Private) Peering Debate: Peering at 10Gbps

William B. Norton <wbn@equinix.com>

Abstract

This research paper focuses on an Economic Analysis of 10G Peering from the ISP Perspective. It is demonstrates that even with new and pretty expensive 10G hardware, peering makes sense financially after a few gigabits-per-second of peering traffic. The Public vs. Private peering debate part is required because the Peering Coordinator Community points out that the next best alternative to 10G public peering is peeling off cross connects or circuits for private peering, so the paper compares these two options. This white paper is based on interviews with several dozen Peering Coordinators at Peering Forums and on-line during the first quarter of 2005.

Introduction: Public Peering Model

The Internet is a network of networks, interconnected in peering and transit relationships.

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 95th percentile using 5-minute samples every month[1].

Definition: Peering is a business relationship whereby two entities reciprocally exchange access to each others customers. But, what is *Public Peering*?

Definition: Public Peering is peering across a shared (public) peering fabric, a media with more than two participating parties.

Since Ethernet is the most common public peering fabric in the world, we will assume Ethernet is used for public peering from this point on in the paper. Public peers connect their router Ethernet interface card(s) to the public peering fabric and (potentially) peer with many of the other attached public peers.

The costs of Public Peering

Peers exchange access to each others customers for free, but what are the costs associated with Public Peering? If we assume that the peering is to be done by collocating equipment at an Internet Exchange Point (IX), public peering requires

1) a **router** with an
    a. interface card suitable to connect to the public peering fabric,
    b. interface card suitable to send and receive traffic back to the rest of the network,
2) a **port** on the public peering fabric,
3) **colocation** space at an Internet Exchange Point[2].

Note that we are ignoring the cost of the circuit(s) into the Internet Exchange (IX) in this analysis for two reasons. First, for an ISP or Large Scale Network Savvy Content Provider, these costs would be the same for both public and private peering at the scales that we are discussing. Secondly, companies looking to peering towards 10G are probably also purchasing transit at the IX, selling transit at the IX, or otherwise establishing a presence at the IX for reasons beyond peering. If true, then the transport costs of peering will at least partially be covered by these other activities. And here again, any incremental costs will be incurred whether publicly or privately peering.

10G Public Peering Model Hardware

For both cases (10G public and 1G private peering) we will use the Juniper M320 router, with the pricing

[1] See “Internet Service Providers and Peering” for a more thorough explanation of Peering and Transit.

[2] European IXes use a different model; there is an explicit separation between the collocation space operators and the typically association operated IXes. We are assuming that ISPs are not remote peering – that a collocation presence is required.

provided below[3]. We will allocate costs incrementally; as we increase the number of peering sessions we will apply the costs of the corresponding interface cards as required, and we will amortize the cost of all peering equipment over 36[4] months. Note that we also assume the router vendor discounts the price of routers by 35%, which is typical in the industry.

Juniper M320			
First 10G Peering Capacity			
2*10G LAN PIC	$282,000		
Flxible PIC concentrator	$80,000		
Chassis (20Gbps per slot, 8 slots)	$145,000		
Total Initial	$507,000		
less 35% discount	**$329,550.00**	**$9,154.17**	**per month**
For subsequent 10Gs of Peering Capacity			
2*10G LAN PIC	$282,000		
Flxible PIC concentrator	$80,000		
Total Subsequent	$362,000		
less 35% discount	**$235,300.00**	**$6,536.11**	**per month**

Figure 1 - 10G Public Peering Router Configuration and Pricing Estimates

Graphically, the 10G peering router configuration is shown below.

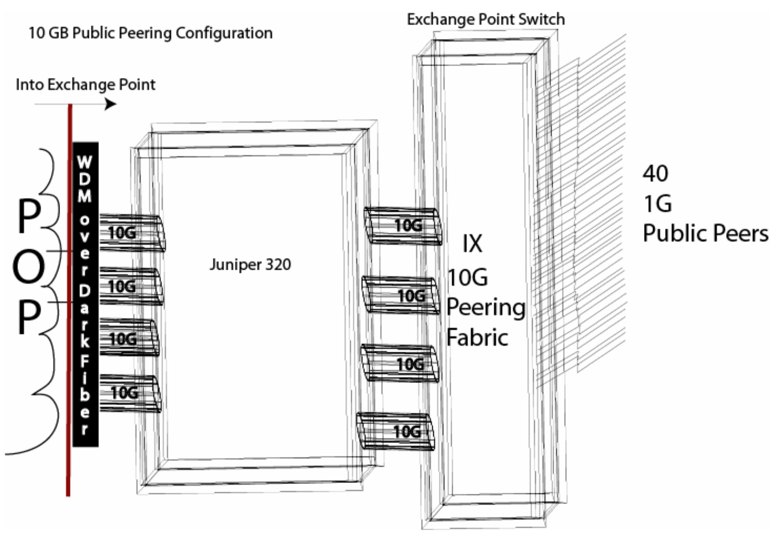

Figure 2 - Public Peering Model Hardware for Publicly Peering with 40+ single gigabit peers

Ingress Peering Capacity. We assume that the ISP anticipates requiring 10G of peering capacity into the Exchange Point and therefore we will start out with 10G of ingress capacity into the IX. From a practical perspective, this will most likely be done by leasing dark fiber and running the 10GE to their metro Point of Presence (POP). We will assume that the ISP will allocate additional strands and/or install Wave Division Multiplexing (WDM) and other equipment needed to obtain additional 10G ingress capacity. Since this capacity is required regardless of the public vs. private peering method, these costs are a wash and not taken into account in this paper.

Aggregation Benefits and Utilization. For modeling purposes we will assume that an ISP has an aggregation ratio of 1.5:1. One peer's peaks may match up with another peer's valleys when the traffic is aggregated out a single large pipe, providing "aggregation efficiencies." This 1.5:1 aggregation ratio means that when an ISP aggregates peering traffic to/from ten 1-Gigabit-Ethernet public peer's traffic, they can statistically multiplex up to 14 1-

[3] Estimates provided by Richard Steenbergen (nLayer)

[4] Vish (NetFlix), Nathan Hickson (eBay), Richard Steenbergen (nLayer)

Gigabit-Ethernet peers' worth of peering capacity and still not need to upgrade the ingress or public peering port capacity. When the 15th peer is added though, the 10G peering port and the ingress 10G port will need to be upgraded and the corresponding costs are added to the model.

Peering Port at IX. We are assuming for modeling purposes that the 10G peering port at the IX costs **$10,000** per month.

Colocation Costs. We are assuming for modeling purposes that racks fees of **$1,000** per month are charged at the IX. Some in Europe suggested that this cost was too high, others suggested that given the power requirements for the equipment described, this cost should be doubled.

These Public Peering costs allocated across peering sessions is shown in the spreadsheet below. Note that the colored lines are highlighting the points where additional hardware is added to the peering configuration.

# Peering Sessions	Effective Peering Bandwidth (in Mbps)	+ Requires Ingress Port	10G Peering Monthly Port(s) fees	1 Rack Monthly Colo Fee	Equipment Monthly Costs	Best Public Peering $/Mbps
1	700	10GE	$10,000	$1,000	$9,154	$28.79
2	1400	10GE	$10,000	$1,000	$9,154	$14.40
3	2100	10GE	$10,000	$1,000	$9,154	$9.60
4	2800	10GE	$10,000	$1,000	$9,154	$7.20
5	3500	10GE	$10,000	$1,000	$9,154	$5.76
6	4200	10GE	$10,000	$1,000	$9,154	$4.80
7	4900	10GE	$10,000	$1,000	$9,154	$4.11
8	5600	10GE	$10,000	$1,000	$9,154	$3.60
9	6300	10GE	$10,000	$1,000	$9,154	$3.20
10	7000	10GE	$10,000	$1,000	$9,154	$2.88
11	7700	10GE	$10,000	$1,000	$9,154	$2.62
12	8400	10GE	$10,000	$1,000	$9,154	$2.40
13	9100	10GE	$10,000	$1,000	$9,154	$2.21
14	9800	10GE	$10,000	$1,000	$9,154	$2.06
15	10500	2*10GE	$20,000	$1,000	$15,690	$3.49
16	11200	2*10GE	$20,000	$1,000	$15,690	$3.28
17	11900	2*10GE	$20,000	$1,000	$15,690	$3.08
18	12600	2*10GE	$20,000	$1,000	$15,690	$2.91
19	13300	2*10GE	$20,000	$1,000	$15,690	$2.76
20	14000	2*10GE	$20,000	$1,000	$15,690	$2.62
21	14700	2*10GE	$20,000	$1,000	$15,690	$2.50
22	15400	2*10GE	$20,000	$1,000	$15,690	$2.38
23	16100	2*10GE	$20,000	$1,000	$15,690	$2.28
24	16800	2*10GE	$20,000	$1,000	$15,690	$2.18
25	17500	2*10GE	$20,000	$1,000	$15,690	$2.10
26	18200	2*10GE	$20,000	$1,000	$15,690	$2.02
27	18900	2*10GE	$20,000	$1,000	$15,690	$1.94
28	19600	2*10GE	$20,000	$1,000	$15,690	$1.87
29	20300	3*10GE	$30,000	$1,000	$22,226	$2.62
30	21000	3*10GE	$30,000	$1,000	$22,226	$2.53
31	21700	3*10GE	$30,000	$1,000	$22,226	$2.45
32	22400	3*10GE	$30,000	$1,000	$22,226	$2.38
33	23100	3*10GE	$30,000	$1,000	$22,226	$2.30
34	23800	3*10GE	$30,000	$1,000	$22,226	$2.24
35	24500	3*10GE	$30,000	$1,000	$22,226	$2.17
36	25200	3*10GE	$30,000	$1,000	$22,226	$2.11
37	25900	3*10GE	$30,000	$1,000	$22,226	$2.06
38	26600	3*10GE	$30,000	$1,000	$22,226	$2.00
39	27300	3*10GE	$30,000	$1,000	$22,226	$1.95
40	28000	3*10GE	$30,000	$1,000	$22,226	$1.90

Figure 3 - Public Peering Costs with Cost Steps highlighted

The last column "Best Public Peering" is a measure of the best cost of peering a 10G public peer can hope for, assuming that the peering session runs to the "Effective Peering Bandwidth", the maximum peak peering before the peering capacity should be upgraded for that peer. This is captured in the graph below which also highlights the incremental hardware additions required for peering at 10G for each additional peer.

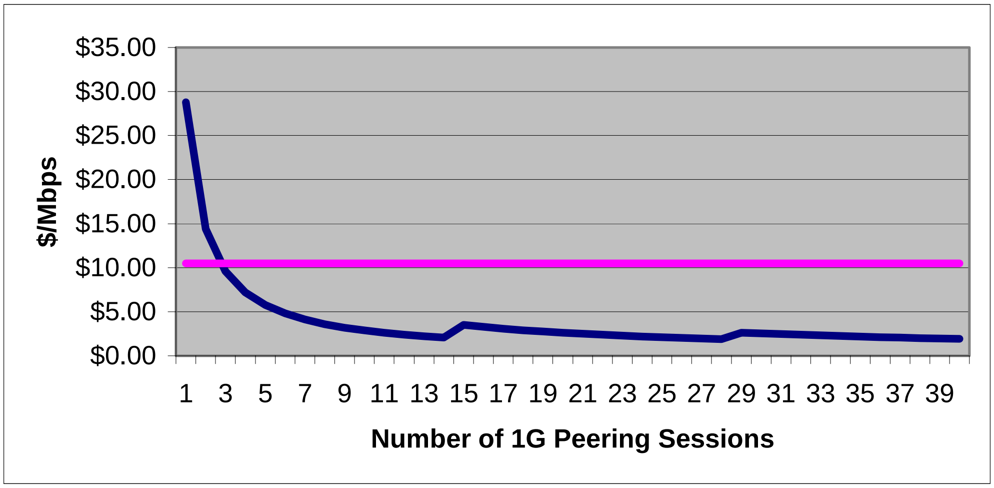

Figure 4 – The 10G Public Peering Model is cost effective after 3 Peering Session

Introduction: Private Peering Model

What is Private Peering?

Definition: Private Peering is peering whereby exactly two parties connect to the peering media.

Private Peering Model Hardware

Here is a sample configuration (see spreadsheet below) for privately peering up to 40 GE peers. For private peering we will add in the Foundry SuperX switch to the 320 Peering Router, and allocate each of two 24-port GE blades to the SuperX as needed. We will also add additional 10G ingress and 10G aggregate peering bandwidth into the equation as needed for the corresponding additional peering sessions.

Private Peering Configuration							
Foundry SuperX				**Juniper M320**			
For first 24 ports				First 10G of Private Peering			
Foundry 8 slot chassis	$14,995			2*10G LAN PIC	$282,000		
Fast Iron 24-port GE SFP mini-GBIC	$6,995			Flxible PIC concentrator	$80,000		
FastIron 2 port 10GE XFP	$5,995			Chassis (20Gbps per slot, 8 slots)	$145,000		
Total Initial	$27,985			Total Initial	$507,000		
less 35% discount	**$18,190.25**	**$505.28**	**per month**	less 35% discount	**$329,550.00**	**$9,154.17**	**per month**
For second 24 ports, add				For subsequent 10Gs of Private Peering			
Fast Iron 24-port GE SFP mini-GBIC	$6,995			2*10G LAN PIC	$282,000		
FastIron 2 port 10GE XFP	$5,995			Flxible PIC concentrator	$80,000		
Total Final Config	$12,990			Total Subsequent	$362,000		
less 35% discount	**$8,443.50**	**$234.54**	**per month**	**less 35% discount**	**$235,300.00**	**$6,536.11**	**per month**

The graphic below shows the components of the private peering model.

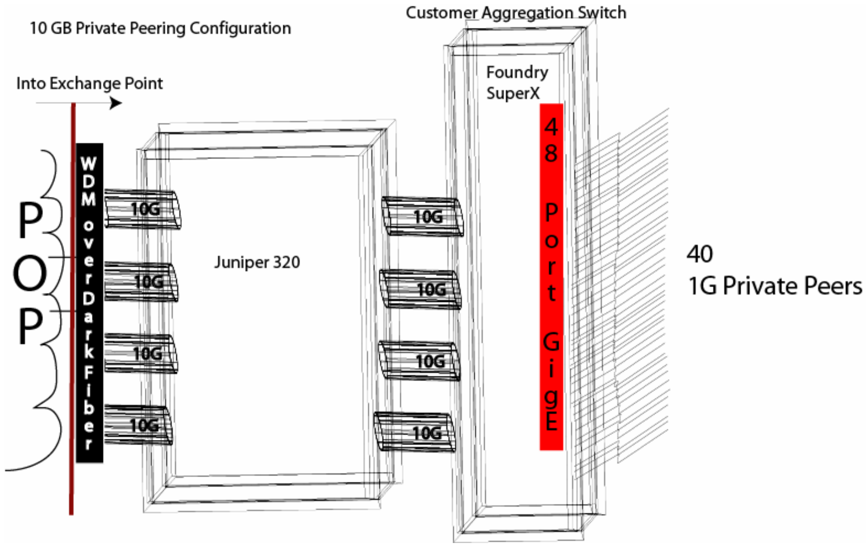

Figure 5 - Private Peering Model Hardware for Privately Peering with 40+ single gigabit peers

Public and Private Graphics look similar from afar! While the public and private pictures may seem similar, they are not. With public peering, the Exchange Point Operator runs the public peering switch on behalf of the IX participants. With Private Peering, the only people connected to the switch are those that peer with the ISP in question, and the ISP in question operates the switch. As you can see in the diagram above, there are no active network elements involved in peering except for those operated by the peers.

These Private Peering costs allocated across peering sessions is shown in the spreadsheet below. Note that the colored lines are highlighting the points where additional hardware is added to the peering configuration.

# gigE cross connects	Effective Peering Bandwidth (in Mbps)	+ Requires Ingress Port	+Cross Connect Fee	1 Rack Colo Fee	Equipment Monthly Costs	Best Private Peering $/Mbps
1	700	10GE	$250	$1,000	$9,659	$15.58
2	1400	10GE	$500	$1,000	$9,659	$7.97
3	2100	10GE	$750	$1,000	$9,659	$5.43
4	2800	10GE	$1,000	$1,000	$9,659	$4.16
5	3500	10GE	$1,250	$1,000	$9,659	$3.40
6	4200	10GE	$1,500	$1,000	$9,659	$2.90
7	4900	10GE	$1,750	$1,000	$9,659	$2.53
8	5600	10GE	$2,000	$1,000	$9,659	$2.26
9	6300	10GE	$2,250	$1,000	$9,659	$2.05
10	7000	10GE	$2,500	$1,000	$9,659	$1.88
11	7700	10GE	$2,750	$1,000	$9,659	$1.74
12	8400	10GE	$3,000	$1,000	$9,659	$1.63
13	9100	10GE	$3,250	$1,000	$9,659	$1.53
14	9800	10GE	$3,500	$1,000	$9,659	$1.44
15	10500	2*10GE	$3,750	$1,000	$16,196	$1.99
16	11200	2*10GE	$4,000	$1,000	$16,196	$1.89
17	11900	2*10GE	$4,250	$1,000	$16,196	$1.80
18	12600	2*10GE	$4,500	$1,000	$16,196	$1.72
19	13300	2*10GE	$4,750	$1,000	$16,196	$1.65
20	14000	2*10GE	$5,000	$1,000	$16,196	$1.59
21	14700	2*10GE	$5,250	$1,000	$16,196	$1.53
22	15400	2*10GE	$5,500	$1,000	$16,196	$1.47
23	16100	2*10GE	$5,750	$1,000	$16,196	$1.43
24	16800	2*10GE	$6,000	$1,000	$16,196	$1.38
25	17500	2*10GE	$6,250	$1,000	$16,430	$1.35
26	18200	2*10GE	$6,500	$1,000	$16,430	$1.31
27	18900	2*10GE	$6,750	$1,000	$16,430	$1.28
28	19600	2*10GE	$7,000	$1,000	$16,430	$1.25
29	20300	3*10GE	$7,250	$1,000	$22,966	$1.54
30	21000	3*10GE	$7,500	$1,000	$22,966	$1.50
31	21700	3*10GE	$7,750	$1,000	$22,966	$1.46
32	22400	3*10GE	$8,000	$1,000	$22,966	$1.43
33	23100	3*10GE	$8,250	$1,000	$22,966	$1.39
34	23800	3*10GE	$8,500	$1,000	$22,966	$1.36
35	24500	3*10GE	$8,750	$1,000	$22,966	$1.34
36	25200	3*10GE	$9,000	$1,000	$22,966	$1.31
37	25900	3*10GE	$9,250	$1,000	$22,966	$1.28
38	26600	3*10GE	$9,500	$1,000	$22,966	$1.26
39	27300	3*10GE	$9,750	$1,000	$22,966	$1.24
40	28000	3*10GE	$10,000	$1,000	$22,966	$1.21

Figure 6 - Private Peering Costs with Cost Breaks Highlighted

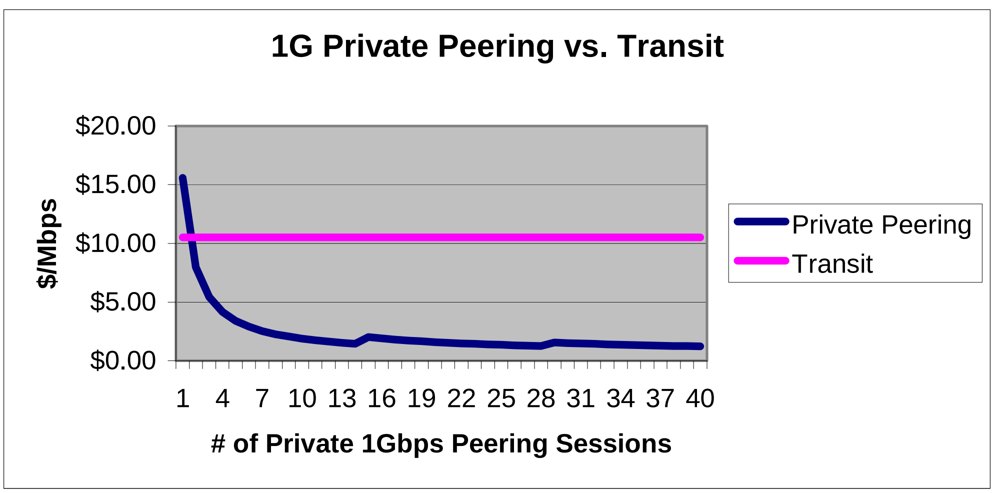

So which is better, Public Peering or Private Peering?

The Great Peering Debate: Public Peering vs. Private Peering

In the Peering Coordinator Community this issue has generated much debate. Let's explore the strongest arguments on each side of the debate, presented below.

The Top 4 Reasons **Public Peering** is better than Private Peering

Here are the strongest arguments presented to the author from public peering advocates.

1. Aggregation Benefits

a. A network can easily **aggregate a large number of relatively small peering sessions** across a single fixed-cost peering port, with zero incremental cost per peer. (Private peering requires additional cross connects and potentially an additional interface card, so there are real costs associated with each incremental peering session.) Small peering sessions often exhibit a high degree of variability in their traffic levels, making them perfect for aggregation. Since not all peers peak at the same time, multiple peers can be multiplexed onto the shared peering fabric, with one peer's peak traffic filling in the valleys of another peer's traffic. This helps make peering very cost effective: "I can't afford to dedicate a whole gigE card to private peering with this guy, but public peering is a no-brainer."

b. Public peering ports usually have very **large gradations of bandwidth**: 100Mbps Ethernet upgrades to 1Gbps Ethernet, which upgrades to 10Gbps Ethernet. With such large gradations, it is easier for smaller peers to maintain several times more capacity via public peering than they are currently using, which reduces the likelihood of congestion due to shifting traffic patterns, bursty traffic, or uncontrolled Denial of Service attacks. "Some peers aren't as responsive to upgrading their peering infrastructure, nor are they of similar mind with respect for the desire for peering bandwidth headroom[5]." The large gradations of public peering bandwidth help reconcile these two issues.

2. Ease of administration

a. Public peering is the **easiest and fastest** way to both turn up and turn down a peering sessions, since no physical work is required. Peering is soft configured by the two parties on the router and the peering session is up.

b. It is common for a network to set up a **trial peering** session to determine the amount of traffic that would be exchanged should a session be turned up. If there is public peering capacity available, there is no incremental cost or extra administrative work required to turn up a trial peer, and the information gathered may prevent choosing an incorrect private peering port size if the traffic is moved to a private peer later.

c. Many Peering Coordinators must work within a budget, and do not have decision making authority for purchases within their company. Once the public peering switch port is ordered, there is **no additional cost** and therefore no additional hurdle to peering for the Peering Coordinator.

d. Public Peering provides **financial predictability**. The hardware requirements and monthly recurring costs of peering are the same every month[6]. This makes planning and budgeting much easier.

e. 10G Public Peering **scales large peering sessions** (those greater than 1Gbps) seamlessly, while private peering beyond gigE capacities requires private peering at 10G (very expensive), or connecting multiple gigEs together, which can be tricky[7].

---

[5] James Rice (LoNap), formerly engineering for BBC Internet.

[6] Modulo the incremental costs for hardware at the cost steps described earlier and the pricing increases at the end of contract terms with the IX Operator. It was pointed out during conversations at the RIPE 50 meeting that the LINX charges a metered rate beyond the flat peering fee, and that some IXes have a fractional gigabit Ethernet rate that causes pricing steps.

[7] Patrick Gilmore (Akamai) in conversation 4/7/2005 regarding load balancing across multiple cross connects.

<wbn@equinix.com>

3. Public Peering is used as Selection Criteria by Customers

a. Corporate and Enterprise customers continue to ask to see the list of the ISP's public peering points[8].

4. Public Peering May Be the only Cost Effective way to Peer across multiple Colos

a. Across Europe, where public peering across multiple collocation centers is the norm, private peering is often a much more expensive solution. Purchasing private peering circuits within a metro is potentially very expensive, while the same traffic can traverse a shared peering fabric for much less.

The Top 5 Reasons Private Peering is better than Public Peering

Here are the strongest argument private peering advocates shared with the author.

1. Private Peering Sessions are Easier to Monitor

a. **SNMP Counters** can be easily collected on each peering port to monitor the utilization of the Peering Session resources. No time intensive Netflow or expensive network analysis software[9] is required to sort through shared peering fabric data to determine per-peering-session traffic volume.

b. **Greater Visibility: No Blind Oversubscription Problem.** With public peering, the remote peer could be congesting his port with the other peering sessions and you have no visibility into their public peering port utilization. Packets could be dropped due to port oversubscription resulting in poor peering performance. Since Private Peering involves only the two parties, when the port reaches an agreed upon utilization (say 60% utilization for example), both parties can see that it is time to upgrade the peering session.

2. Private Peering is Very Cost Effective

a. If an expected peering port and cross connect costs were $400 per month and the parties expected to send 40Mbps to each other, the EPPR would be $400/40Mbps=$10/Mbps, a very attractive price in today's transit market.

b. For those who exchange traffic with a **few large peers, the 80%/20% rule applies**; the majority of peering benefits can be derived by peering with the 20% of potential peers that deliver 80% of your traffic. This suggests fewer larger peers is preferable over picking up lots of small peers across a public peering fabric.

3. Private Peering is more reliable and easier to debug.

a. Private Peering involves **fewer network components** that could break.[10] It should be noted that this argument weakens when the "private" peering are provisioned across VLANs, though optical interconnects, telco provisioned SONET services, or other active electronics.

b. An architecture of private peering **removes the variability** of support processes across IXes[11]. Across Europe, each IX is different, and a NOC Operator may need to understand the processes, the levels of support and debugging capabilities of the switch support staff on call at the IX, and may even need to craft NOC scripts to navigate through the IX operations tasks. A private peering architecture provides consistency that helps the NOC debug and fix things more rapidly.

c. The greater fear is that layer 2 fabrics could be connected through other layer two fabrics perhaps without the knowledge or consent of the peer, resulting in a very difficult debugging and diagnostics situation if a peering

---

[8] Frank Orlowski (T-Systems) in conversation 4/8/2005.

[9] Some of these tools cost $50,000 to license!

[10] Remco Donker (MCI) and Nina Hjorth Bargiser (Tele Danmark) point out that some people view large capacity public peering as too risky; losing a single large public peering port would cause massive disruption to the infrastructure, and would result in much larger convergence times than if a single private peering port went out.

[11] Falk Bornstaedt and Frank Orlowski (T-Systems).

failure occurs.

4. Private Peering Sessions are More Secure

a. A private peering network that is directly connected only with those with whom there is an **explicit peering arrangement** is more secure than a network that connects to a public peering fabric that includes participants with whom there is no relationship with the company. There is some history here; early exchange points were places where “traffic stealing” was accomplished by pointing default at an unsuspecting and poorly secured public peer. Other problems included peers tunneling traffic across the ocean across a peer’s network. These things are explicitly disallowed in most peering and IX terms and conditions and can be further secured through filtering, but are still seen as potential hazards minimized by privately peering.

b. An architecture that solely privately peers is less likely to be compromised. Since fiber has no active components that can be administered, there is nothing that can be broken into. With a switch or other active electronics in between peers, there is the possibility that **traffic can be captured** at the peering point without their detection. It is relatively easy to mirror a public peering port as compared with tapping into private peering fiber cross connects without the detection of the peers involved. A few ISPs pointed to technology that can passively tap into fiber interconnects, which if true, would decrease the strength of this argument.

5. Private Peering Inclination Signals a More Attractive Peer.

a. The “**Big Players**” privately peer with each other and some even loath Public Peering Fabrics for historical reasons. Adopting this attitude puts one in the company of the largest Tier 1 ISPs in the world. “For certain very large networks, public peering makes no sense at all. For certain very small networks, public peering may make perfect sense[12].” Or put more harshly, “if you think that public peering is a good idea, you're just not large enough yet[13]."

Hybrid Approach (Public + Private Peering)

A combination of public and private peering is the most common, where ISPs peer publicly and “peel off” peering sessions to private peering as the volume of traffic to and from those peers increases. The “40% rule” is sometimes used whereby both parties contractually obligate each other to migrate their public peering session to private peering sessions when either party reaches 40% aggregate usage of their public peering port. This provides a safeguard for those concerned about the “Blind Oversubscription Problem” described above[14].

One ISP primarily uses private peering but does maintain public peering for reserve and emergency interconnect capacity. The ability to scale public peering quickly and seamlessly was seen as a key attribute here. History has shown that traffic generally grows, incrementally or sporadically from emergencies and spot events.

At a recent NANOG Peering BOF Peering Debate, Vijay Gill (AOL) pointed out that there may be a “life cycle” of peering inclinations (see Appendix B), where peers migrate from public towards private peering as the scale increases. There is much debate regarding the virtues of de-peering, and of pulling away from public peering.

The Business Case for 10G Public vs. 1G Private Peering

Let’s compare the costs of these two configurations: Public Peering using 10G Ethernet against Private Peering using many 1G Ethernet Private Peering Sessions. In the spreadsheet and graphic below you will see the costs of publicly peering lots of aggregate traffic over large 10 Gigabit-per-second public peering port(s) compared against attaching a switch and privately peering lots of private peering sessions directly to peer routers.

---

[12] Richard Steenbergen conversation 3/23/2005

[13] Anonymous

[14] First heard of this clause from Ren Provo (SBC).

Best Private Peering $/Mbps	Best Public Peering $/Mbps
$15.58	$28.79
$7.97	$14.40
$5.43	$9.60
$4.16	$7.20
$3.40	$5.76
$2.90	$4.80
$2.53	$4.11
$2.26	$3.60
$2.05	$3.20
$1.88	$2.88
$1.74	$2.62
$1.63	$2.40
$1.53	$2.21
$1.44	$2.06
$1.99	$3.49
$1.89	$3.28
$1.80	$3.08
$1.72	$2.91
$1.65	$2.76
$1.59	$2.62
$1.53	$2.50
$1.47	$2.38
$1.43	$2.28
$1.38	$2.18
$1.35	$2.10
$1.31	$2.02
$1.28	$1.94
$1.25	$1.87
$1.54	$2.62
$1.50	$2.53
$1.46	$2.45
$1.43	$2.38
$1.39	$2.30
$1.36	$2.24
$1.34	$2.17
$1.31	$2.11
$1.28	$2.06
$1.26	$2.00
$1.24	$1.95
$1.21	$1.90

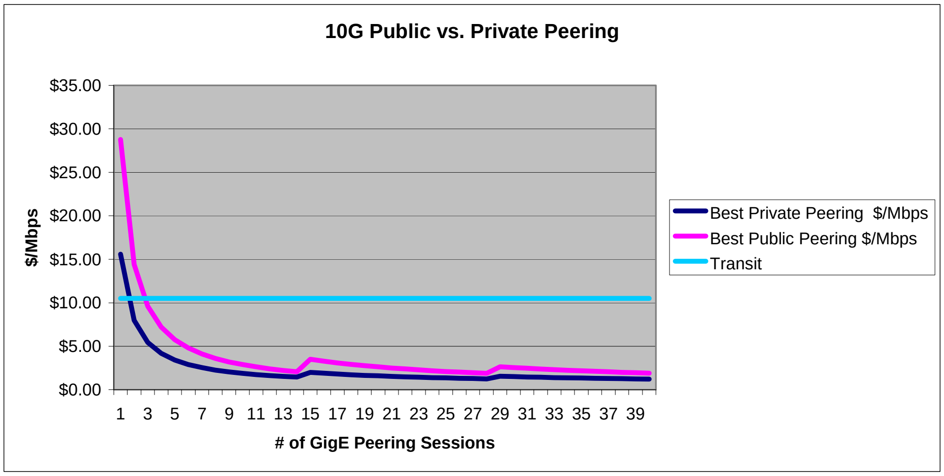

Public Peering vs. Private Peering Analysis.

a. The first thing we notice is the **similarity of the graphs**. Both Private and Public peering routers used in this example are large-capacity, high-performance, and expensive, especially if one only peers with one other party at sub-1Gbps levels. At one peer each, both assuming a fully utilized gigE peering session (at 700Mbps), the cost of traffic exchange is between $12/Mbps and $26/Mbps. These are not bad prices compared against the transit market today.

b. Next, we notice that both configurations require a 10G ingress card and 10G peering card, and we also notice that Private Peering requires a **second piece of hardware**, a comparably inexpensive switch. The steps of the cost functions are highlighted in the graph, but with an extra piece of hardware, why is Private Peering consistently less expensive than Public Peering? The answer is, the IX charges only $250 per month for the cross connect as opposed to $10,000 per month for each incremental 10G peering port. Note that in Europe, the price of a 10G peering port is less than half this price.

c. In the case of public peering, we upgrade the ingress and peering capacity when we reach our 15th peering session, and are set to continue peering until our 29th peering session, explaining the to bumps in the cost graphs.

d. For private peering we upgrade when we allocate our 15th peering session, again assuming the aggregated traffic can fit over the 10G ingress. After that point we allocate a second 10G ingress interface and are all set until we run out of private peering links. For our 25th peer, we deploy a second 24 port Ethernet card. Each of these steps is highlighted in the graphs and spreadsheets.

Business Case for 10G Peering. In both Public and Private Peering scenarios we see a very compelling case for peering at 10G. The cost of peering quickly (for even a few peering session) dips below $10 per Mbps, well below the market price for purchasing transit.

Two 10G Transition Dynamics

Discussions with Patrick Gilmore (Akamai) and Richard Steenbergen (nLayer) highlighted two transition dynamics that have occurred during the past transitions of peering capacity (10M→100M IX Peering Port Size, and then 100M→1G Peering Port Size).

1) When an ISP transitions to the next size up, he will try and encourage his larger peers to transition as well so they both realize the benefits of economies of scale as shown in the graphs above. However, at this point, the peer may respond with sticker shock over the cost of the 10G hardware and, if in the same building, counter propose private peering using dedicated gigE peering ports.
    i. This erodes the network externality benefits for the community as less traffic results in less benefits for the peering community.
    ii. This is only possible if the ISPs are in the same building and cross connects are reasonable inexpensive.
2) When an ISP starts to fully utilize private peering capacity there is a transition step that involves
    i. load sharing across a pair of cross connects which is difficult to manage,
    ii. having both sides upgrade routers or at least interface cards or ports, which can be expensive, or
    iii. migrate the private peering back to a larger public peering port, which can be very expensive.

In these cases, the Peering Coordinators face a transition challenge, and try to manage expectations internally and externally surrounding the timing and cost expected for the peering transitions.

Assumptions

Transport Costs into the IX are ignored. We have, in some white papers in the past, included the cost of the local loop into the IX in the analysis. The assumption then was that the company investigating peering was already installed somewhere else with a transit provider in place. By including local loop in the analysis we compared the current situation against extending a presence at an IX only for peering. Conversations with Peering Coordinators revealed a flaw in previous research papers:

1. Those entering an IX to establish a peering presence, also have access to (potentially many) transit providers and can aggregate peering and transit traffic over a single large capacity transport circuit. In this case, the other transit location is not needed.
2. In the case where both the transit and the peering locations are needed, a local loop is required. For the majority of those participating in 10G peering, the opportunity to sell transit is enough to warrant installation into the IX. In this case, the cost of local loop is included in the cost of selling transit, so should not be a cost allocated solely to peer.

Having said this, it is worth pointing out that there are precious few options for transport into an IX for 10G peering. Most reasonable and widely deployed are 10G peering off an existing multi-OC-192 backbone deployment and 10G directly run over fiber potentially needing long reach optics. When factoring the cost of the transport we see almost no difference between peering privately and publicly – the price differences are negligible compared to the monthly recurring costs of transport[15].

We assume the Peering Equipment as described is capable of handling the peering load and is configurable to the needs of the peer.

We assume that peers publicly peering will continue to publicly peer at 10G levels. One ISP pointed out that each of these steps in peering capacity (10M to 100M peering, 100M to 1G peering, and now 1G to 10G peering) did not always lead to both parties upgrading their public peering capacity. When one party requested the other party to upgrade their public peering capacity with them, the upgrade request would require a capital expense that could not be made at that time. In this case, the other Public Peer would make a counter-offer to convert the public peering session to a private peering session.

We assume that the ISP is not already at the IX.

We assume that the ISP is not remote peering (requires colocation).

We are not defining a redundant router, redundant peering port or redundant peering interface card solution. In conversations with the Peering coordinators we saw great disparity in their views regarding this. Some of the larger ISPs required great levels of redundancy so any component could fail with minimal impact on operations. Others found it more cost effective to handle redundancy by provisioning larger backbone capacity to handle the fail over to another IX, hopefully located close by. For one large European ISP, the thought of putting that much traffic on a single network component was too scary to consider. "If I privately peer and I lose a 1G card, the rerouted traffic is at worst a 1000Mbps. I lose a 10G public peering card and I have a disaster!"

We do not consider backbone costs. Several ISPs pointed out that an ISP that would peer towards 10G capacity would probably need to peer in multiple locations. This would require a broader architecture with a more complex model that would vary based on degree of redundancy but would more realistic.

We are ignoring the revenue producing potential of the 10G infrastructure. Josh Snowhorn (NOTA) pointed out that many ISPs are utilizing the same infrastructure to make money and for them the costs of peering are insignificant incremental costs.

Some have pointed out that the costs of peering should include the operations costs associated with managing peering, including Peering Coordinator salary, NOC time, documentation time, etc. Others argue that the costs are negligible, and that the staff are there anyway. Others still claim that one form of peering or the other saves significant engineering resources. We have not factored these costs into this analysis although we did highlight a few key differences between public and private peering in terms of peering administration.

Further Study

Some feedback from the Peering Coordinator Community included some alternative configurations that should be considered. Eric Troyer suggested for example that switch ports are so inexpensive that it is common for ISPs to used a switch in between their router and the public peering fabric and trunk multiple gigEs together rather than step from a single gigE to a 10GigE connection to the switch. Some suggested that a Wave Division Multiplexer and switch architecture with layer three enabled would provide a less expensive (albeit less feature rich) solution that should be considered. There are many permutations that could be considered and it would be useful to take an ISP survey to identify and model the most common architectures and solutions.

Acknowledgements

The author would like to gratefully acknowledge the Peering Coordinator Community for sharing their insights, and in particular: Ren Provo (SBC), Richard Steenbergen (nLayer), James Rice (LoNAP), Todd Underwood (Renesys), Stephen Wilcox (TeleComplete), Vanessa Evans (LINX), Niels Bakker (AMS-IX), Chris Malayter (TDS Telecom), Patrick Gilmore (Akamai), Frank Orloski (T-Systems), Vijay Gill (AOL), Vish Yelsangikar (NetFlix),

[15] In the past analysis, we found the price of leasing dark fiber for 20 year IRU was approximately $23,000 per month in 1999 from downtown Washington, DC to Ashburn Virginia.

Nathan Hickson (eBay), Steve Feldman (CNet), Lane Patterson (Equinix), Joy Fender, Falk Bornstaedt (T-Systems), Remco Donker (MCI), Danny McPhearson (Arbor Networks), Josh Snowhorn (NOTA), Vince Fuller (Cisco), Philip Smith (Cisco), Andrew Odlyzko (UMN).

Appendix A – 10 G Configuration Estimates

Richard Steenbergen: “Here are some basic minimal example configurations for you to play with. Note that as with ANY complex system, there are many different ways you can configure these boxes using older or newer cards, in order to achieve slightly different pricing. The examples below are main-stream, middle of the road, basic examples of what a network MIGHT purchase to achieve some basic capacity levels (obviously with a lot of room left over for growth). All numbers are list price.”

Carrier class routers		
Cisco GSR 12416 (10Gbps per slot, 16 slots, 1x10GE, 10x1GE)		
GSR16/M320-AC	Cisco 12416 w/1PRP, 2CSC, 3SFC, 2Alarm, 3AC	$130,000
1X10GE-LR-SC	Cisco 12000 1-Port 10GE Card, 1310nm, 10km	$125,000
EPA-GE/FE-BBRD	Cisco 12000 Modular GE w/1GE and 3 EPA Slots	$ 40,000
EPA-3GE-SX/LH-LC	Cisco 12000 3-Port GE PA EPA-GE/FE-BBRD	$ 50,000
EPA-3GE-SX/LH-LC	Cisco 12000 3-Port GE PA EPA-GE/FE-BBRD	$ 50,000
EPA-3GE-SX/LH-LC	Cisco 12000 3-Port GE PA EPA-GE/FE-BBRD	$ 50,000
		========
		$445,000

Some configurations for 10G…

Juniper M320 (20Gbps per slot, 8 slots, 1x10GE, 10x1GE)		
MM320BASE-AC	Juniper MM320 w/CB, 4SIB, RE-1600-2048, 4PEM	$145,000
MM320-FPC3	Flexible PIC Concentator, Type 3 , MM320	$ 80,000
PC-1XGE-XENPAK	1-port 10 Gigabit Ethernet LAN PIC	$141,000
PC-10GE-SFP	10-port Gigabit Ethernet PIC for FPC3	$140,000
		========
		$506,000

Enterprise class router		
Cisco 6509 with SUP720 (40Gbps per slot, 9 slots, 4x10GE, 48x1GE Fiber)		
WS-C6509-E	Cat6509 Chassis, 9slot, 15RU	$ 9,500
WS-CAC-3000W	Catalyst 6500 3000W AC power supply	$ 3,000
WS-CAC-3000W	Catalyst 6500 3000W AC power supply	$ 3,000
WS-C6K-9SLOT-FAN2	Catalyst 6509 High Speed Fan Tray	$ 495
WS-SUP720-3B	Catalyst Supervisor 720 Fabric MSFC3 PFC3B	$ 28,000

<wbn@equinix.com>

WS-X6704-10GE	Cat6500 4-port 10 Gigabit Ethernet Module	$ 20,000
WS-X6748-SFP	Catalyst 6500 48-port fabric-enabled	$ 25,000
	========	
	$ 88,995	

Don't have a price-list for the Foundry IMR640 (which would be the closest competitor to the above config), or even the Extreme Aspen, but the following is an example config of a slightly more low-end box which is less scalable and has less features, but which is currently the cheapest high-density 10GE capable box on the market. I don't have the pricing for the L3 version in front of me (different management card), so this is the pricing for the L2 version only.

Foundry SuperX (25Gbps per slot, 8 slots, 4x10GE, 48x1GE Fiber, 12x1GE)		
FI-SX1-4AC	FastIron SuperX 8-slot chassis w/1 AC	$ 14,995
SX-FI12GM-4	FastIron Management w/12x Combo GE	$included
SX-FI424F	FastIron 24-port GE SFP mini-GBIC	$ 6,995
SX-FI424F	FastIron 24-port GE SFP mini-GBIC	$ 6,995
SX-FI42XG	FastIron 2-port 10GE XFP	$ 5,995
SX-FI42XG	FastIron 2-port 10GE XFP	$ 5,995
	========	
	$ 40,975	

Source: Richard A Steenbergen <ras@e-gerbil.net> http://www.e-gerbil.net/ras

Appendix B – The Life Cycle of Peering Inclinations

Final Analysis on the Public vs. Private Peering Debate. The answer may depend more upon the network asking the question than anything else.

At a recent NANOG Peering BOF, Vijay Gill described various stages in the lifecycle of a peering network, characterized by a change in motivations.

1. During the initial stages of peering, the network is most likely concerned primarily with reducing their IP transit costs by gaining access to other networks "for free". At this point they are usually interested in peering with every other network willing to peer, regardless of the size of the network.

2. As the peering network increase its size, scope, and amount of peering, their motivations begin to include the desire to improve performance and increase their control over their routing. At this stage, the peering network is interested in continuing to expand their peering, and begins to focus their energies on meeting the requirements of more selective networks. Less time and attention is paid to turning up new peering with networks that are considered "less important", though they may still be willing to accept new requests from smaller networks.

3. As the peering network continues to grow, they begin to think strategically about how to continuing scaling their peering. At this stage they often become "selective", establishing a set of peering requirements based on traffic levels, locations, backbone capacity, ratios, or any combination of the above, in order to limit their peering activities to only networks they consider to be important. The peering network begins to realize that they can increase their traffic levels to larger networks they desire to peer with by not directly peering with customers of those networks, and may begin to depeer some of their smaller or more difficult peers.

4. By the next stage, the peering network has already established peering with a large portion of the Internet, including the majority of the other networks selective networks. The only networks left to peer with are the "restrictive" networks who are not interested in establishing any new peerings. The peering network continues to increase its selective peering requirements in order to focus only on peering with the remaining few networks, and is very disinterested in establishing new peerings with smaller networks or potential new competetitors. By now the peering network is usually quite large, and begins to leverage external business relationships such as the purchase of transport or fiber in order to secure peering with restrictive peers.

5. Finally, the peering networks reaches the plateu of the "tier 1" network. They are no longer interested in establishing any new peerings with new networks, and enter the category of "restrictive". By now the network is usually categorized as "huge", and exchanges an enormous amount of traffic with only a handful of other large networks.

During this lifecycle, from a small network just beginning its peering, to a huge network that sends 100% of its traffic to peering, the motivations for public vs private peering changes vastly.

[picture showing the peering evolutionary lifecycle]

In stage 1, the peering network usually has no choice but to peer publicly. The choice of where and how to peer is largely dictated by the other party, and the peering networks usually chooses locations based solely on the number of other peering partners available there.

In stage 2, the peering network is usually enjoying the capacity and burstability benefits of a relatively large peering port. They usually prefer to keep peers on these large public ports, and any private peers they establish will be based on the requirements of the other party, and will be of a capacity smaller than their public peers.

In stage 3, the peering network is moving some of their largest peers to private peering, but continues to use public peering to aggregate most of the smaller peers. While they are often not looking to establish new peering with any new network who connects to a public peer, the public peer is still an excellent location to conduct trials.

In stage 4, the peering network is attempting to divorce itself from its public peering. The majority of its traffic by volume is now being sent via large private peers, and public peers are not being upgraded. As peers increase their traffic levels, they are moved onto private peering. Towards the end of stage 4, the peering network may seek to disconnect public peering ports it considers less important.

In stage 5, the peering network sees almost no value in public peering. Some may retain presences solely for political reasons, such as avoiding the confrontation of de-peering another network and instead simply limiting the growth of new peering. Other networks may disconnect themselves from public peering completely at this point.