Express or Local Lanes – On Assessing QoE over Private vs. Public Peering Links

Walter Willinger, NIKSUN, Inc., wwillinger@niksun.com
Anja Feldmann, Philipp Richter, TU Berlin, {anja,prichter}@inet.tu-berlin.de
Georgios Smaragdakis, MIT/TU Berlin, gsmaragd@csail.mit.edu
Fabian Bustamante, Northwestern University, fabian@cs.northwestern.edu

ABSTRACT

As more large content/cloud and service providers are making colocation facilities that house an Internet eXchange Point (IXP) their location-of-choice for interconnecting with one another in an effort to shrink the physical as well as network distances between where content resides and where it is consumed, the providers of these facilities and/or the operators of the co-located IXPs are reacting by constantly innovating and expanding their interconnection service offerings. These developments directly impact the long-standing public vs private peering debate, and performing a rigorous measurement study that provides a solid understanding of the actual rather than the perceived cost-performance trade-offs between the different interconnection service offerings that are available to networks in today's Internet in one and the same colocation facility would go a long way towards putting this debate on scientifically solid foundations. In particular, we argue that an important first step towards achieving this goal is to establish a proven set of measurement methods and techniques to infer both the existence and type of the different interconnections that the networks in a colocation facility with an IXP presence have established with one another. In fact, any realistic assessment of application-level performance and user-perceived quality-of-experience of the traffic that is routed through such a facility will rely critically on our ability to infer the existence and usage of the established interconnections in that facility.

1. PROPOSED RESEARCH

Unless the content requested by an ISP's customer is served from within that ISP (e.g., ISP-owned or third party-provided cache servers), the resulting traffic will be affected by the conditions of the network interconnections or peering links that this traffic experiences en route from where the content is available to where it is consumed, impacting application-level performance and customer-perceived quality of experience. These network interconnections that facilitate the seamless transfer of data between diverse players across today's Internet are typically established inside commercial colocation facilities or data centers that may or may not house an Internet eXchange Point (IXP). The presence of (parts of the distributed infrastructure of) an IXP in a given colocation facility enhances the interconnection service offerings at that facility by providing public peering options in addition to private peering options.

Until recently, the distinction between public and private peering was straightforward. Whereas private peering relies on establishing a *dedicated* physical connection (i.e., private interconnect or cross-connect, typically purchased directly from the colocation facility provider or a colocated IXP) between two networks that exchange traffic, public peering is based on a *shared* virtual connection and is implemented through leased ports on a public switching fabric (layer-2) infrastructure at an IXP [11]. While the decision making process (also known as "the art of peering" [16]) that different network operators use to choose between these two peering options differs from operator to operator, it typically includes network-specific aspects (e.g., type of business, expected traffic volume, profile of a potential peering partner) as well as factors such as performance, cost, reliability, and security. Nevertheless, there has been a long-held perception – especially in the US where the interconnection marketplace has historically been dominated by a small number of companies that have been eager to promote the lucrative private peering option – that public peering is the "poor man's" approach to interconnecting networks (i.e., low cost and subpar performance). This view of public peering as a "second-class" citizen is perhaps best captured by the comment "if you think that public peering is a good idea, you are just not large enough" attributed in [15] to an operator of a US Tier-1 ISP.

However, despite enormous recent advances in collecting a myriad of different Internet-wide measurements, there exist no rigorous measurement studies that we are aware of that either confirm or refute this common "rich vs poor man's" perception of private vs public peering; that is, the salient assumption that while establishing public peering links costs less than purchasing private peering connections, the performance over the former is in general worse than the performance over the latter, suggestive of the highway/motorway analogy of local vs express lanes. More importantly, during the last few years, the Internet interconnection marketplace in the US and elsewhere has experienced drastic changes, and some of them are significant enough that they necessitate revisiting key elements of the long-standing public vs private peering debate. Some of them also pose new problems and novel challenges for establishing a scientifically sound empirical-based foundation for such a debate and argue strongly for a concentrated community-wide effort to provide the empirical evidence required for a principled approach to assessing and comparing the properties of the different Internet peering options available in today's Internet.

To illustrate some of the problems and challenges that such an effort will face, we first note that in recent years, the successful IXPs have grown enormously in size and geographic reach. Today, the largest exchanges in Europe have more than 500-600 members and handle traffic rates that peak at multi-Tbps. To sustain annual growth rates in membership and traffic volume of 10-20%, these IXPs have deployed the latest switching technologies capable of supporting a range of port speeds, from 1 Gbps and smaller to the latest 100 Gbps offerings (e.g., see [4]). Also, in addition to increasing their own capacity, these IXPs have also expanded beyond their original home base (i.e., city) and have been busy deploying their proven technology in other countries and continents. For example, DE-CIX in Frankfurt, Germany, not only owns and operates two stand-alone regional IXPs in Hamburg and Munich, but has also aggressively expanded beyond the country's border, managing since late 2012 UAE-IX, the first carrier-neutral IXP in the Middle East; opening in late 2013 DE-CIX New York, launching in early 2015 new IXPs in Palermo (Italy), Marseille (France), and Istanbul (Turkey), and announcing in mid-2015 the opening of DE-CIX Dallas in Texas [1]. These and similar expansion efforts by other IXPs have gone hand-in-hand with a concentrated community effort to promote public peering in markets where that interconnection remains underutilized [10, 12].

During the same time, IXPs have also been busy extending their customer base. In fact, recent announcements how some of the major players use IXPs are a telling sign that the popular view of public peering as the "poor man's" choice for establishing Internet interconnectivity is due for a major revision. For example, Netflix's list of peering locations around the world includes most of the major IXPs [9]. Similarly, Akamai, a global CDN that delivers between 15-30 percent of the world's web traffic across its global platform relies critically on extensive Internet peering as part of its distribution strategy. In its effort to help enterprises to distribute their media, cloud and web content as quickly as possible around the world so that end users have a uniformly fast and low-latency Internet experience and in anticipation of constantly growing traffic demands by their customers, Akamai recently announced that it has upgraded its capacity at DE-CIX in Frankfurt to 12x100 GE connections, capable of delivering 1.2 Terabits per second which is the largest service provider bandwidth at any Internet exchange worldwide to date (August 2015) [3]. More broadly, some two years after its opening in late 2013, DE-CIX New York announced that all major CDNs are already connected to the exchange, in addition to large content providers such as Google, Netflix, and Twitter [5].

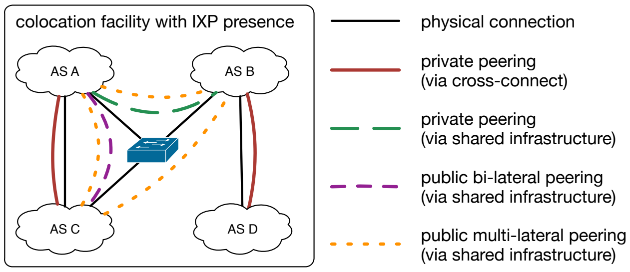

Figure 1: Interconnection options in colocation centers that host IXPs.

The impressive growth rates and world-wide expansion efforts of some of the largest European IXPs have been partly driven by innovative new service offerings provided by those IXPs (e.g., 100 GE ports, remote peering, IXP resellers, free use of route server). On the one hand, these new offerings are mainly intended to make these IXPs even more attractive to yet a wider spectrum of networks. On the other hand, some of these new services require a careful re-framing of our understanding of Internet peering as they impact the meaning, deployment, operations, monitoring and third-party measurement of currently available interconnection services. For example, in an effort to drastically reduce their members' overhead and management complexity associated with establishing and maintaining public peering links (i.e., actively managing BGP sessions) with possibly hundreds of different peers, many IXPs have stratified public peering into "multi-lateral" peering and "bi-lateral peering." While these IXPs support multi-lateral peering through a new service that consists of offering the free use of the IXP's route server to their customers [17], bi-lateral peering reflects the original concept of public peering but acquires a more specific meaning – public peering that does not involve the IXP's route server. Note that this stratification of public peering creates new challenges for discovering or identifying bi-lateral peerings via control plane measurements by third-parties (e.g., researchers with no access to IXP-internal data) [17].

Similarly, the deployment of advanced switching technologies by some of the most successful IXPs has also led to a refinement of the traditional concept of private peering, with direct implications for inferring the performance of the traffic that utilizes this service. In fact, until recently, private peerings have been typically established by either purchasing physical cables (i.e., cross connects) directly from the colocation or data center provider that hosts the peering partners or, if the peering parties are members at the same IXP, buying the cross-connect from the IXP (which in turn may "outsource" the task of physically connecting the parties to the colocation or data center provider at hand). In either case, it is generally understood that private peerings use dedicated physical connections and allow no sharing of the physical link with other parties. In the case where these links are established at an IXP, this means that they are managed and operated completely separately from the IXP's public switching fabric, thus complicating the management and operations of the IXP. However, in a number of IXPs, this picture of strict physical isolation of private peering traffic from other IXP-handled traffic has started to blur with the service offering of "Private Interconnection (PI)" [2]. This new service enables the direct exchange of traffic between two member's networks by utilizing a dedicated VLAN on the IXP's public switching infrastructure, instantiating in practice the virtual equivalent of a physical cross-connect – a "virtual private peering."

In view of these and other recent developments that have impacted the Internet interconnection marketplace in fundamental ways, it is important to first establish a proven set of measurement methods and techniques to infer both the existence and type of the different interconnections that the networks in a colocation facility with an IXP presence have established with one another. In fact, the joint presence of large content and eyeball providers at more and more IXPs makes these locations ideal rendezvous places for shrinking the physical as well as network distances between where content resides and where it is consumed (also referred to as "flattening" of the AS topology [14]). In the process, the interconnections established and used in these locations become increasingly more central and critical for assessing application-level performance as well as overall (i.e., end-to-end) user-perceived QoE in today's Internet. Thus, performing a rigorous measurement study that provides a solid understanding of the actual rather than the perceived cost-performance trade-offs between available interconnection service offerings in one and the same colocation facility would go a long way towards putting the ongoing public vs private peering debate on scientifically solid foundations.

An illustration of the different interconnection services that are available to today's customers of a colocation facility with an IXP presence is shown in Figure 1. Note however that this picture is significantly simplified by omitting interconnections that the different colocation facility customers may have with their upstream or transit provider(s). Aspects of the complexity that is added when accounting for this upstream connectivity are discussed in detail in [13] and will necessarily have to be part of any future measurement studies that intend to inform the public vs private peering debate in which IXP-provided interconnections are often viewed as "shortcuts" when compared to conventional upstream connectivity. The need for such empirical studies has also been echoed in recent FCC documents [7, 8] where IXPs are considered as candidate locations for assessing QoS as part of the FCC's Universal Service Fund/Intercarrier Compensation (USF/ICC) transformation efforts [6]. However, there are a number of important steps that need to be performed (in succession) before the various properties of such co-located interconnection service offerings can be fully understood and their impact on QoE-related issues can be assessed. These steps will require a concentrated community effort and a close collaboration between Internet researchers with expertise in different aspects of network and service performance measurements and include:

- Step 1 (Connectivity): Use a combination of active and passive measurements to discover the different interconnections that the customers of a colocation facility with an IXP presence have established between them.
- Step 2 (Usage): Perform third-party measurements to identify how the different interconnections inferred in Step 1 are used (i.e., what kind of traffic traverses which interconnection).
- Step 3 (Performance): Design localized measurement experiments to compare the performance (e.g., delay, available bandwidth) of the different interconnections inferred in Step 1 that a pair of customers of such a facility have purchased and actively use (including upstream connections).
- Step 4 (Quality-of-Experience): Assess the contribution of an inferred and used interconnection to user-perceived QoE by combing the interconnection's inferred performance (Step 3) with its inferred usage (Step 2); e.g., is the observed performance good enough to support high-quality video streaming.

2. REFERENCES

[1] About DE-CIX - Where all networks meet. https://www.de-cix.net/about/.
[2] AMS-IX Services – Private Interconnect. https://ams-ix.net/services-pricing/private-interconnect.
[3] DE-CIX - News and Events. https://www.de-cix.net/news-events.
[4] DE-CIX Apollon. https://apollon.de-cix.net/about-apollon/.
[5] DE-CIX New York city. https://nyc.de-cix.net/.
[6] Federal Communications Commission, 13-161. https://apps.fcc.gov/edocs_public/attachmatch/FCC-11-161A1.pdf.
[7] Federal Communications Commission, DA 13-1016. https://apps.fcc.gov/edocs_public/attachmatch/DA-13-1016A1.pdf.
[8] Federal Communications Commission, DA 13-2115. https://apps.fcc.gov/edocs_public/attachmatch/DA-13-2115A1.pdf.
[9] Netflix peering locations. https://openconnect.netflix.com/peeringLocations/.
[10] Open IX – Building Community and Consensus to Foster Data Center and Interconnection Standards. http://www.open-ix.org/.
[11] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *SIGCOMM*, 2012.
[12] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. Quo vadis Open-IX? Trying to boost public peering in the US. *ACM CCR*, 5(1), January 2015.
[13] P. Faratin, D. Clark, S. Bauer, W. Lehr, P. Gilmore, and A. Berger. The Growing Complexity of Internet Interconnection. *Communications and Strategies*, 2008.
[14] C. Labovitz, S. Lekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-Domain Traffic. In *SIGCOMM*, 2010.
[15] W. B. Norton. Public vs Private Peering : The Great Debate. http://drpeering.net/white-papers/Public-vs-Private-Peering-The-Great-Debate.html.
[16] W. B. Norton. The Art of Peering: The Peering Playbook, 2010.
[17] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *IMC*, 2014.