# The Evolution of Internet Interconnection from Hierarchy to “Mesh”: Implications for Government Regulation

May 5, 2014

# The Structure of the “Early” Internet: Hierarchical

- In order to connect with other ISPs, an ISP *purchased transit* from Tier 1 backbones that, in turn, engaged in *settlement-free peering* with one another.
- We have characterized this arrangement as hierarchical because ISPs occupied a “lower” place in the Internet than did backbones.

# The More Recent Internet: Mesh

- Although a significant amount of Internet traffic still follows the hierarchical structure, this has been supplemented by *secondary peering*, direct connections between ISP that do not use transit, that is, do not require the participation of Tier 1 backbones.
- We have characterized this structure as a “mesh”.
- Some, but not all, secondary peering, is *settlement-free*, as were the transactions among Tier 1 backbones, but there is also *paid peering*.

# How Much Internet Traffic Involves Secondary Peering?

- It is difficult to tell because many peering links are invisible to the monitors that collect the data used to measure their importance.
- Some estimates
  - Available data tend to miss more than 75% - 85% of the links for large content networks (Dhamdhere et al, 2010; Olivera et al, 2007).
  - The largest percentage of content providers do not use transit (Dhamdhere and Dovolis, 2008).
  - Microsoft had ≥ 24 direct connections, Google had ≥ 23, and Yahoo! had ≥ 18 (Gill et al, 2008).
  - Netflix had one private peering connection in 2010 and 27 in 2013 (Lodhi et al 2014).

# How Much Secondary Peering is Paid?

- We have even less information about this than about the amount of secondary peering.
- However, we do know that network operators rely on several criteria in determining whether they will peer and whether they will require payment for doing so. These include the:
  - traffic ratio between the two networks
  - geographic diversity of the other network
  - traffic volume exchanged between the two networks
  - minimum backbone capacity of the other network and the number of points of interconnection

# Examples of Peering Policies

- XO: “The total 95% percentile aggregate traffic ratio shall not exceed 2.0 to 1.”
- AboveNet: “The ratio of the aggregate amount of traffic exchanged between the Requester and the AboveNet with which it wishes to interconnect shall be roughly balanced and shall not exceed bidirectional ratio of 2:1.”
- AT&T: “… a new peer must have:…No more than a 2.00:1 ratio of traffic into AT&T: out of AT&T, on average each month.”
- Level 3 and XO Communications agreed to settlement-free peering if “both networks carry approximately the same bit-miles of data, a model that…[ensures] a balanced cost burden across each network.”

# Peering Disputes Have Been Rare

- Level 3 briefly “depeered” Cogent Communications in 2005 because “Cogent was sending far more traffic to the Level 3 network than Level 3 was sending to Cogent's network….”
- Level 3 and Comcast reached an agreement in 2011 that “changes how Level 3 routes traffic across Comcast’s network, sharply cutting the fees [Level 3] must pay when traffic overwhelms certain connections….”
- Netflix and Level 3 proposed (in 2011) that the FCC mandate settlement-free peering, what they called the “open, regional no-charges, interexchange model”.
- More recently, Netflix complained about Comcast’s decision to charge for peering although Netflix eventually agreed to what are apparently paid peering arrangements with Comcast and Verizon.

# Constraints on the exercise of market power

- A Content Delivery Network or other IP network has alternative paths into an ISP's network and thus they can reroute traffic among these paths in real time.  Importantly, they can always use transit.
- …the complex mesh of interconnections, with diverse pricing models, constrains the range of negotiating positions that can be sustained by [an access network]….the limit on the payment that [an access network] can extract from [a content delivery network] will be related in some way to the customary price for transit, which is a commodity product….“ [Clark, Lehr, and Bauer]
- Indeed, in negotiations with an ISP, a CDN can threaten to exploit transit alternatives that would leave the ISP worse off than if it had entered into a reasonably priced paid peering relationship with the CDN.

# What Would be the Harm from Regulating Peering, Including Mandating Settlement Free Peering

- There would be reduced incentives to minimize total cost
- Rates to end users would rise
- Significant burdens would be placed on regulators

# There Would Be Reduced Incentives to Minimize Total Cost

- Price regulation may lead to higher costs when different costs are subject to the control of different parties.
- Hypothetical example:
  - A new method for delivering traffic that would increase an ISP's costs by $1 million but would reduce a CDN's costs by $2 million will *not* be adopted if regulation limited the payment from the CDN to the ISP to less than $1 million.
- Real example:
  - Netflix reduced the amount of data that it generates in Canada with what it described as a "minimal" effect on picture quality.

# Lower Interconnections Rates Would Increase Rates to End Users

- This is a “two sided market” that is subject to the “seesaw principle”
- According to the seesaw principle, if regulation were to mandate lower revenues from interconnecting CDNs, this would be associated with higher prices charged to subscribers.
- Example:
  - If newspapers could not charge advertisers, prices to subscribers would increase and/or the quantity of news provided would decrease.

# Regulation Would be Difficult

- Regulators would have to determine:
  - which IP networks would be required to peer with which others and on what terms (i.e. paid peering, settlement free peering, or transit).
  - where, and on what technical terms, networks would have to interconnect
  - the obligations of an ISP to maintain capacity in anticipation of changes in the amount and nature of traffic
- These complexities leave two possibilities:
  - Regulators would leave aspects of rules “TBD”, which is likely to create investment deterring uncertainty.
  - Regulators would impose precise rules, which is likely to prevent beneficial experimentation.

# For More Detail

- Stanley M. Besen and Mark A. Israel, The evolution of Internet interconnection from hierarchy to “Mesh”: Implications for government regulation, ***Information Economics and Policy*** (2013).