

Analysis of peering strategy adoption by transit providers in the Internet

Aemen Lodhi
School of Computer Science
Georgia Institute of Technology
aemen.lodhi@gatech.edu

Amogh Dhamdhare
CAIDA
University of California San Diego
amogh@caida.org

Constantine Dovrolis
School of Computer Science
Georgia Institute of Technology
constantine@gatech.edu

The Internet is a complex and dynamic interconnection of about 40,000 Autonomous Systems (ASes). A link in this context represents a business agreement between the corresponding two ASes that they can exchange traffic under various policy and capacity constraints. Each AS X follows a *peering strategy* (or “peering policy”) that is used to determine whether X will accept to peer with another AS Y . Even though they vary in their details, most peering policies can be grouped in three distinct classes: *Restrictive* (X peers only if necessary to avoid Internet partitioning; typically used by Tier-1 transit providers), *Selective* (X peers only with ASes that have comparable traffic volume), and *Open* (X is willing to peer with everyone, except its customers). The conventional wisdom is that transit providers use *Restrictive* or *Selective* peering, so that they can engage other ASes as their customers, and thus increase their transit revenues. In the last few years, however, there is evidence that an increasing fraction of the interdomain traffic flows through peering links between content providers (or CDNs) and access providers that serve end users. A major fraction of the Internet’s traffic thus bypasses intermediate transit providers, and flows directly from producers to consumers. So, a question that motivated this study is: *How are transit providers reacting to the previous peering trend?*

To answer the previous question, we analyze data from PeeringDB, an online database in which peering coordinators provide information about their AS. Even though the sample of ASes in PeeringDB is small, and not representative of “stub” ASes (such as enterprise networks, universities, or non-profit organizations), it includes a significant number of transit, access and content providers and it is probably representative of the transit providers that are larger in size. The analysis of that dataset shows, surprisingly, that most transit and access providers (70-80%) use *Open* peering. This is counter-intuitive, especially for transit providers, because if they peer openly, how can they attract new or keep their existing customers? This observation raises several questions that we attempt to answer in this work: *Why are so many transit providers using Open peering? What is the process that gives them the incentive to adopt Open peering? What does the gravitation towards Open peering imply for the economics of transit providers, their interdomain traffic flow, and their business role in the broader Internet ecosystem? And finally,*

who are the transit providers that lose or gain most, in terms of profitability, with the transition towards Open peering?

Clearly, we cannot answer the previous questions empirically. Internet providers are secretive about their strategies, economic objectives and operational data. Instead, we rely on modeling, and in particular on GENESIS, an agent-based computational model of the Internet. GENESIS captures all those aspects that can have a significant effect on Internet peering decisions: interdomain routing and traffic flow, highly skewed distributions for the per-AS produced and consumed traffic volume, realistic transit pricing, distinct peering locations, strategic peering decisions. This level of complexity would make any analytical model intractable. Agent-based computational modeling allows us to capture a large part of “Internet reality”, and at the same time answer “what-if” questions that would not be possible in an empirical study.

We summarize our main findings as follows:

- 1) Simulations with GENESIS show that overall, 79% of transit providers are attracted to *Open* peering. This gravitation towards *Open* peering occurs due to the combination of economically-motivated strategy selection, myopic decisions, and a lack of coordination among transit providers. Providers with small traffic volumes and small geographical expanse show the highest affinity for *Open* peering, while providers with large traffic volumes and geographical expanse show the least.
- 2) Gravitation towards *Open* peering results in a loss of fitness for almost 70% of providers.
- 3) The effect on the fitness of transit providers due to *Open* peering is not uniform. Small providers (with small traffic volumes and few customers) stand to gain from *Open* peering. On the other hand, large providers (with large traffic volumes and many customers) lose.
- 4) We examined the economic impact of two *Open* peering variants which are often discussed in operator forums. These variants, based on the simple rule-of-thumb “do not peer with customers of a peer”, introduce an implicit coordination mechanism between transit providers. We find that when all transit providers use one of these variants, the aggregate fitness of transit providers approaches that under *Selective* peering.