The Internet is Flat: Revisited

A Small Transit Provider Case Study

Bob Bender
Steve Surdock
(CTS Telecom)

Amogh Dhamdhere
(CAIDA/UCSD)

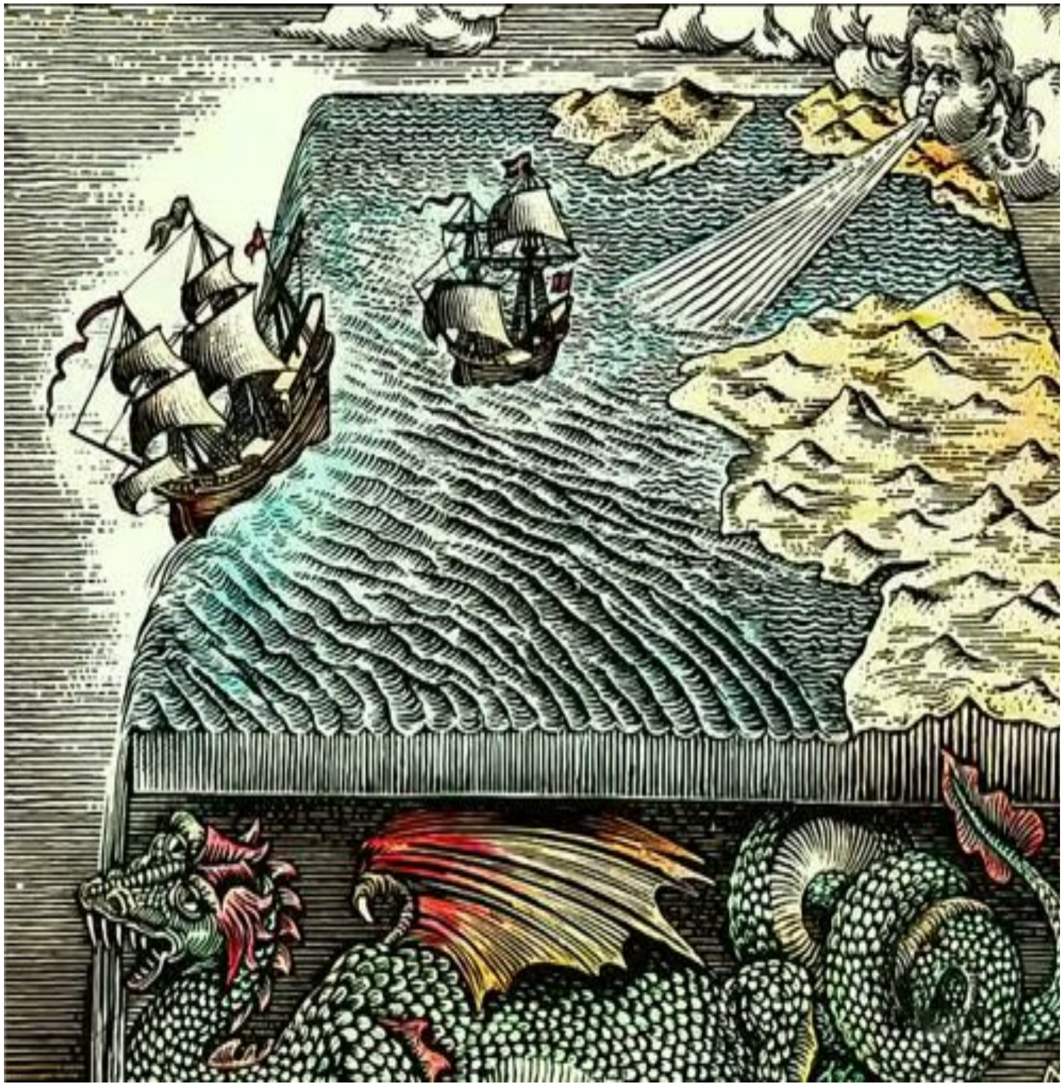

The Internet is Flat: Revisited

A Small Transit Provider Case Study

In 2011 at NANOG52, a small Tier3 ISP AS19653 joined NANOG.

Also in 2011, this small ISP read the paper -
"The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh"

The forecasts in this paper were used to inform business and network planning.

Actual network data was collected from AS19653 from 2010 to present. This small transit provider data is a vignette of the factors that "can transform the Internet ecosystem from a multi-tier hierarchy that relies mostly on transit links to a dense mesh of horizontal interconnections that relies mostly on peering links"

The Internet is Flat:

Modeling the Transition from a Transit Hierarchy to a Peering Mesh

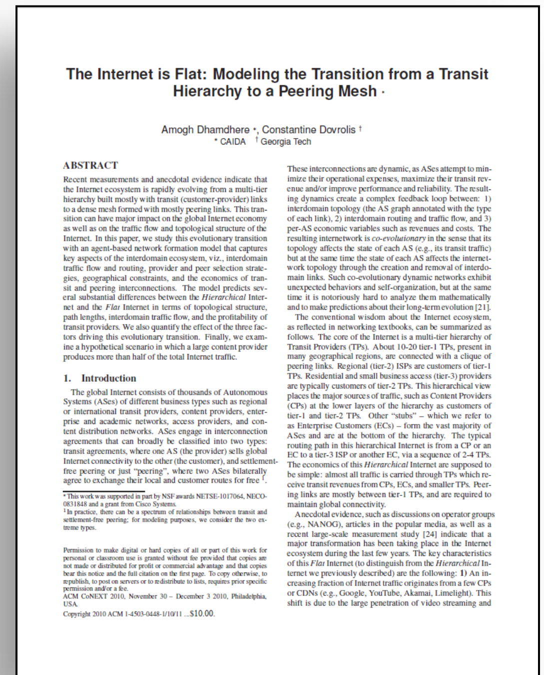

The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh ·

Amogh Dhamdhere *, Constantine Dovrolis †
* CAIDA † Georgia Tech

ABSTRACT

Recent measurements and anecdotal evidence indicate that the Internet ecosystem is rapidly evolving from a multi-tier hierarchy built mostly with transit (customer-provider) links to a dense mesh formed with mostly peering links. This transition can have major impact on the global Internet economy as well as on the traffic flow and topological structure of the Internet. In this paper, we study this evolutionary transition with an agent-based network formation model that captures key aspects of the interdomain ecosystem, viz., interdomain traffic flow and routing, provider and peer selection strategies, geographical constraints, and the economics of transit and peering interconnections. The model predicts several substantial differences between the *Hierarchical* Internet and the *Flat* Internet in terms of topological structure, path lengths, interdomain traffic flow, and the profitability of transit providers. We also quantify the effect of the three factors driving this evolutionary transition. Finally, we examine a hypothetical scenario in which a large content provider produces more than half of the total Internet traffic.

1. Introduction

The global Internet consists of thousands of Autonomous Systems (ASes) of different business types such as regional or international transit providers, content providers, enterprise and academic networks, access providers, and content distribution networks. ASes engage in interconnection agreements that can broadly be classified into two types: transit agreements, where one AS (the provider) sells global Internet connectivity to the other (the customer), and settlement-free peering or just "peering", where two ASes bilaterally agree to exchange their local and customer routes for free [1].

* This work was supported in part by NSF awards NETSE-1017064, NECO-0831848 and a grant from Cisco Systems.
[1] In practice, there can be a spectrum of relationships between transit and settlement-free peering; for modeling purposes, we consider the two extreme types.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ACM CoNEXT 2010, November 30 – December 3 2010, Philadelphia, USA.
Copyright 2010 ACM 1-4503-0448-1/10/11 ...$10.00.

These interconnections are dynamic, as ASes attempt to minimize their operational expenses, maximize their transit revenue and/or improve performance and reliability. The resulting dynamics create a complex feedback loop between: 1) interdomain topology (the AS graph annotated with the type of each link), 2) interdomain routing and traffic flow, and 3) per-AS economic variables such as revenues and costs. The resulting internetwork is *co-evolutionary* in the sense that its topology affects the state of each AS (e.g., its transit traffic) but at the same time the state of each AS affects the internetwork topology through the creation and removal of interdomain links. Such co-evolutionary dynamic networks exhibit unexpected behaviors and self-organization, but at the same time it is notoriously hard to analyze them mathematically and to make predictions about their long-term evolution [21].

The conventional wisdom about the Internet ecosystem, as reflected in networking textbooks, can be summarized as follows. The core of the Internet is a multi-tier hierarchy of Transit Providers (TPs). About 10-20 tier-1 TPs, present in many geographical regions, are connected with a clique of peering links. Regional (tier-2) ISPs are customers of tier-1 TPs. Residential and small business access (tier-3) providers are typically customers of tier-2 TPs. This hierarchical view places the major sources of traffic, such as Content Providers (CPs) at the lower layers of the hierarchy as customers of tier-1 and tier-2 TPs. Other "stubs" – which we refer to as Enterprise Customers (ECs) – form the vast majority of ASes and are at the bottom of the hierarchy. The typical routing path in this hierarchical Internet is from a CP or an EC to a tier-3 ISP or another EC, via a sequence of 2-4 TPs. The economics of this *Hierarchical* Internet are supposed to be simple: almost all traffic is carried through TPs which receive transit revenues from CPs, ECs, and smaller TPs. Peering links are mostly between tier-1 TPs, and are required to maintain global connectivity.

Anecdotal evidence, such as discussions on operator groups (e.g., NANOG), articles in the popular media, as well as a recent large-scale measurement study [24] indicate that a major transformation has been taking place in the Internet ecosystem during the last few years. The key characteristics of this *Flat* Internet (to distinguish from the *Hierarchical* Internet we previously described) are the following: 1) An increasing fraction of Internet traffic originates from a few CPs or CDNs (e.g., Google, YouTube, Akamai, Limelight). This shift is due to the large penetration of video streaming and

Originally Presented at ACM CoNEXT 2010
Nov. 30 – Dec 3, 2010 -- Philadelphia, PA USA

Amogh Dhamdhere (CAIDA)
Constantine Dovrolis (Georgia Tech)

The Internet in 2011: What did the future hold?

The “The Internet is Flat” paper offered an analysis of what we saw happening anecdotally as a small ISP.

As a Tier 3 ISP it became clear that a move to become a Tier 2 ISP would be possible in the new Internet ecosystem.

Most importantly the paper drove home that the importance of Tier1 Transit was diminished and peering with content in the IXP was paramount.

Internet Ecosystem Events

2004 - Google IPO

2007 - Apple iPhone Introduced

2007 - Netflix begins streaming

2008 - Hulu Launched

2011 - Pandora IPO

2012 - Facebook IPO

The ITER Model

Agent-based computational model to answer "what-if" questions about Internet evolution

Inputs

- Network types based on business function
- Pricing/cost parameters
- Interdomain traffic matrix
- Geographical constraints
- Peer/provider selection methods

Output:

Equilibrium internetwork topology, traffic flow, per-network fitness

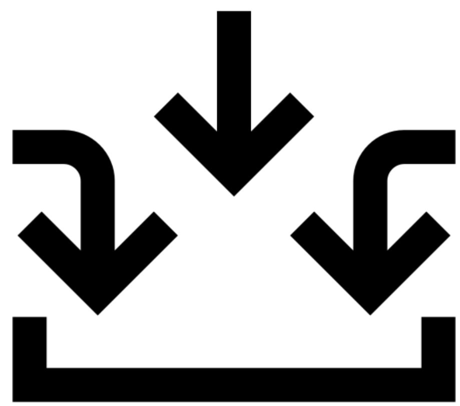

ITER Model Previous Applications are in Demography, Social, Economic and Environmental Sciences.

A. Dhamdhere and C. Dovrolis.
The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh
ACM CoNEXT 2010 – Page 2

The ITER Approach

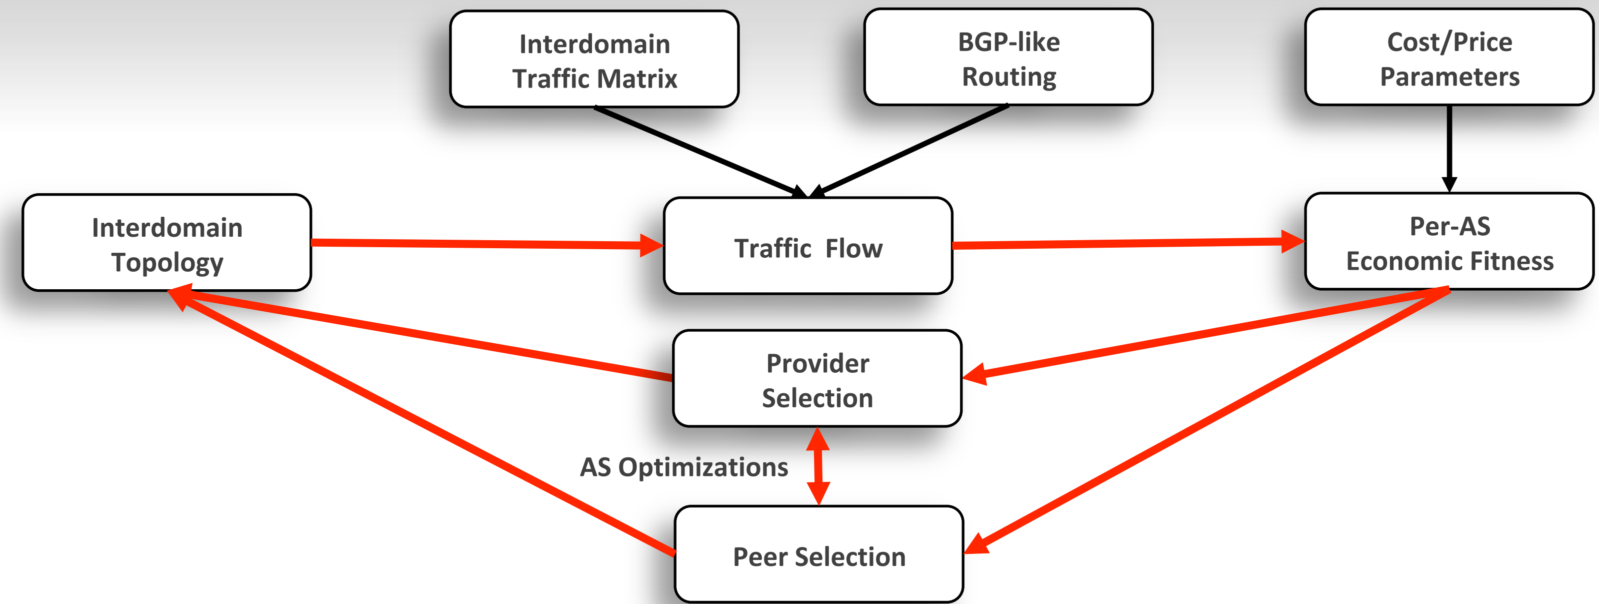

Analytically intractable. Find equilibrium computationally, using agent-based simulations

Equilibrium: no network has the incentive to change its providers/peers

A. Dhamdhere and C. Dovrolis.
The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh
ACM CoNEXT 2010 – Page 3

The "Hierarchical" and "Flat" Internet

The Hierarchical Internet (late 90s – 2007)

- Top content providers generated small fraction of total traffic
- Content providers were typically served from origin
- Peering was restrictive

The Flat Internet (2007 onwards)

- Top content providers generate large fraction of total traffic
- Content providers have expanded geographically
- Peering is more open

Content Consolidation

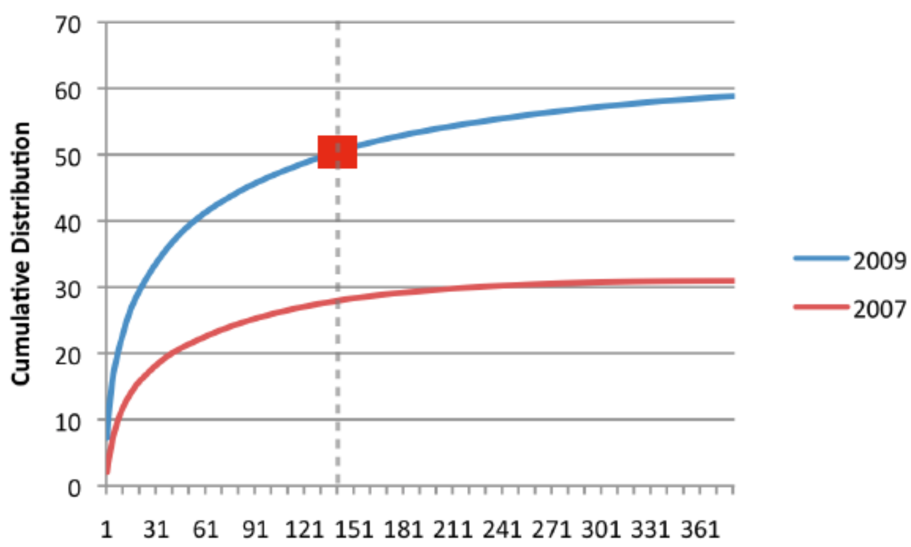

"Internet Interdomain Traffic", Labovitz et al., Sigcomm 2010

Interdomain Routing and Traffic flow

Simulated two "instances" of the ITER model.
First was parameterized to resemble the "Hierarchical Internet".
Second was parameterized to resemble the "Flat Internet".
Then compared various properties of the equilibrium that we get from the two instances of the model.

- More traffic flows over peering links than transit links in the "Flat" Internet
- Traffic follows shorter routing paths due to direct peering in the "Flat" Internet
- This effect is even more pronounced when paths are weighted by traffic volume: paths carrying the most traffic are shorter

A. Dhamdhere and C. Dovrolis.
The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh
ACM CoNEXT 2010 – Page 2

Predictions of Transition Impacts

Content traffic bypasses Tier-1 providers in the "flat" Internet:
Produces conditions for Tier 1 consolidation

It is possible for a Transit Providers to enhance profitability in the "flat" by peering strategically with large Content Providers

Content provider scale promotes peering

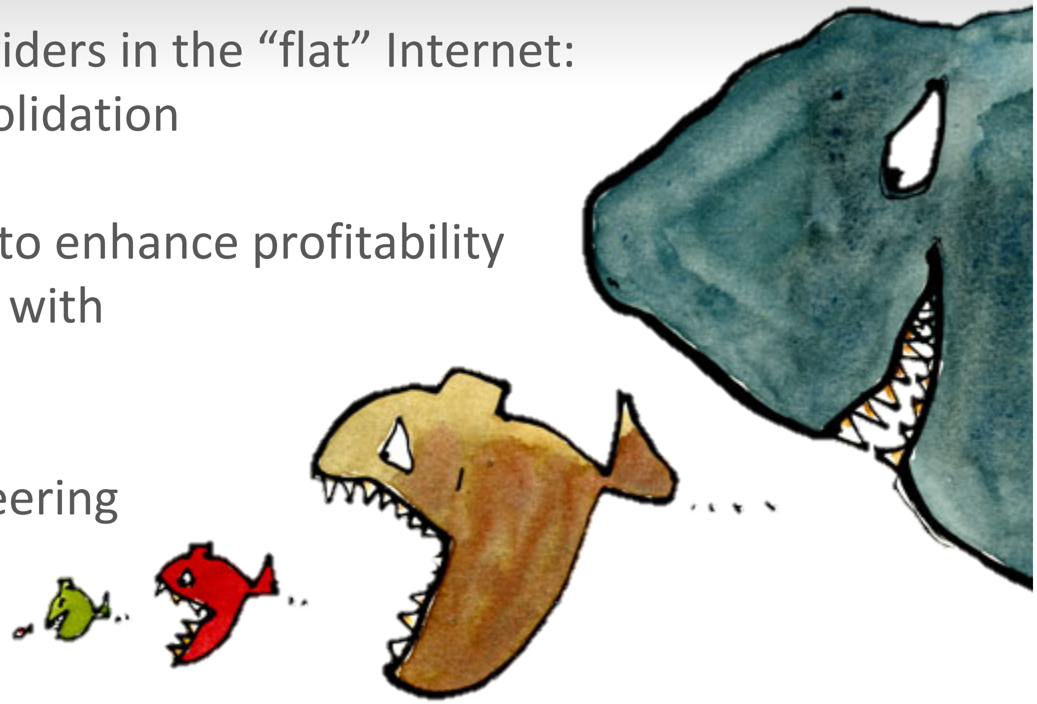

A. Dhamdhere and C. Dovrolis.
The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh
ACM CoNEXT 2010 – Page 12

The Opportunity Presented by Peering Content instead of relying on Tier 1 Transit

In both the Hierarchical and Flat Internet, there is a strong correlation between a Transit Provider’s fitness and the size of its customer base. (need “eyeballs” to peer)

In the Flat Internet, however, strategic peering becomes more important for Small Transit Providers (STP) and LTPs; both can be profitable by peering selectively with the largest content providers.

In the Flat Internet, it is possible for a Transit Provider to transition from unprofitability to profitability by peering strategically, particularly with large Content Providers; such a transition is less likely in the Hierarchical Internet.

A. Dhamdhere and C. Dovrolis.
The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh
ACM CoNEXT 2010 – Page 12

A Small Transit Provider Case Study

AS19653 – Small Transit Provider in Climax, Michigan
Founded in 1911 as Climax Telephone dba CTS Telecom
Independent ILEC-CLEC-ISP.   CLLI = CLMXMIXI

2011 – Joined NANOG

Telephone Company (ILEC-CLEC)
Tier 3 ISP
100% transit (two OC-12s)

2017 – (after 18 NANOGs)

Packet Optical Service Provider
Tier 2 ISP
80% Peering
20% transit
More than 100G in upstream ports

Network Data Source for Graphs

Daily SolarWinds NPM 95th Percentile reports collected since 2010

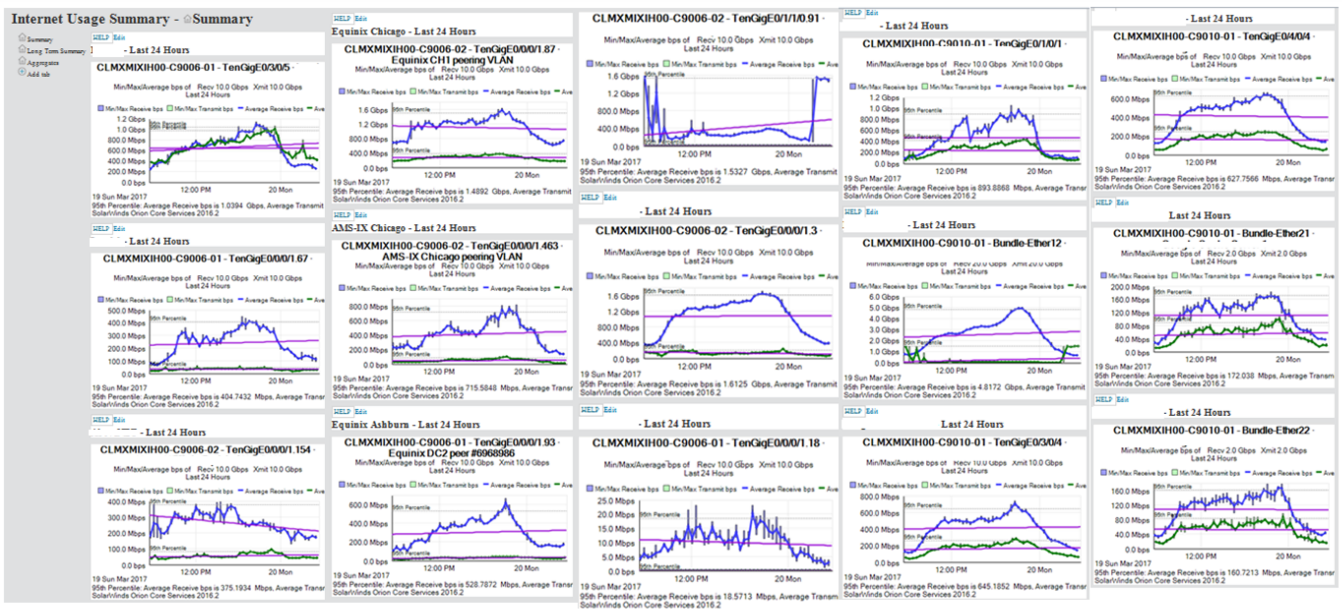

A Small Transit Provider Case Study

SOURCE: CTS Telecom

A Small Transit Provider Case Study

“In the Flat Internet, however, strategic peering becomes more important for Small Transit Providers…”

A. Dhamdhere and C. Dovrolis.
The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh
ACM CoNEXT 2010 – Page 12

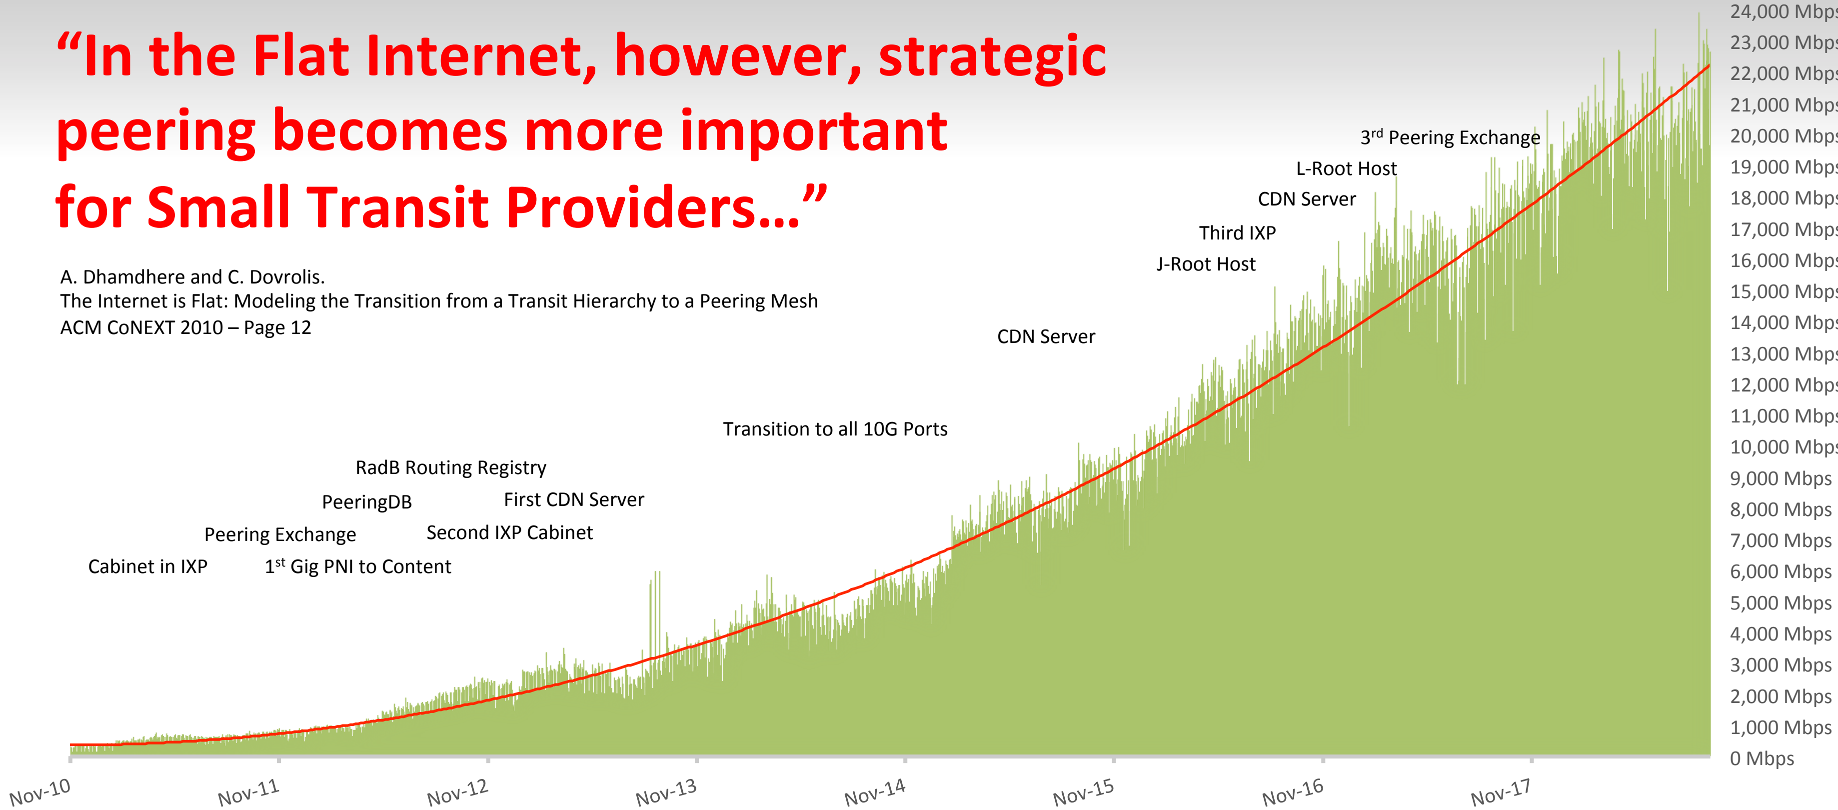

A Small Transit Provider Case Study

SOURCE: CTS Telecom

AS19653 Traffic Mirrors the US IP Traffic Curve

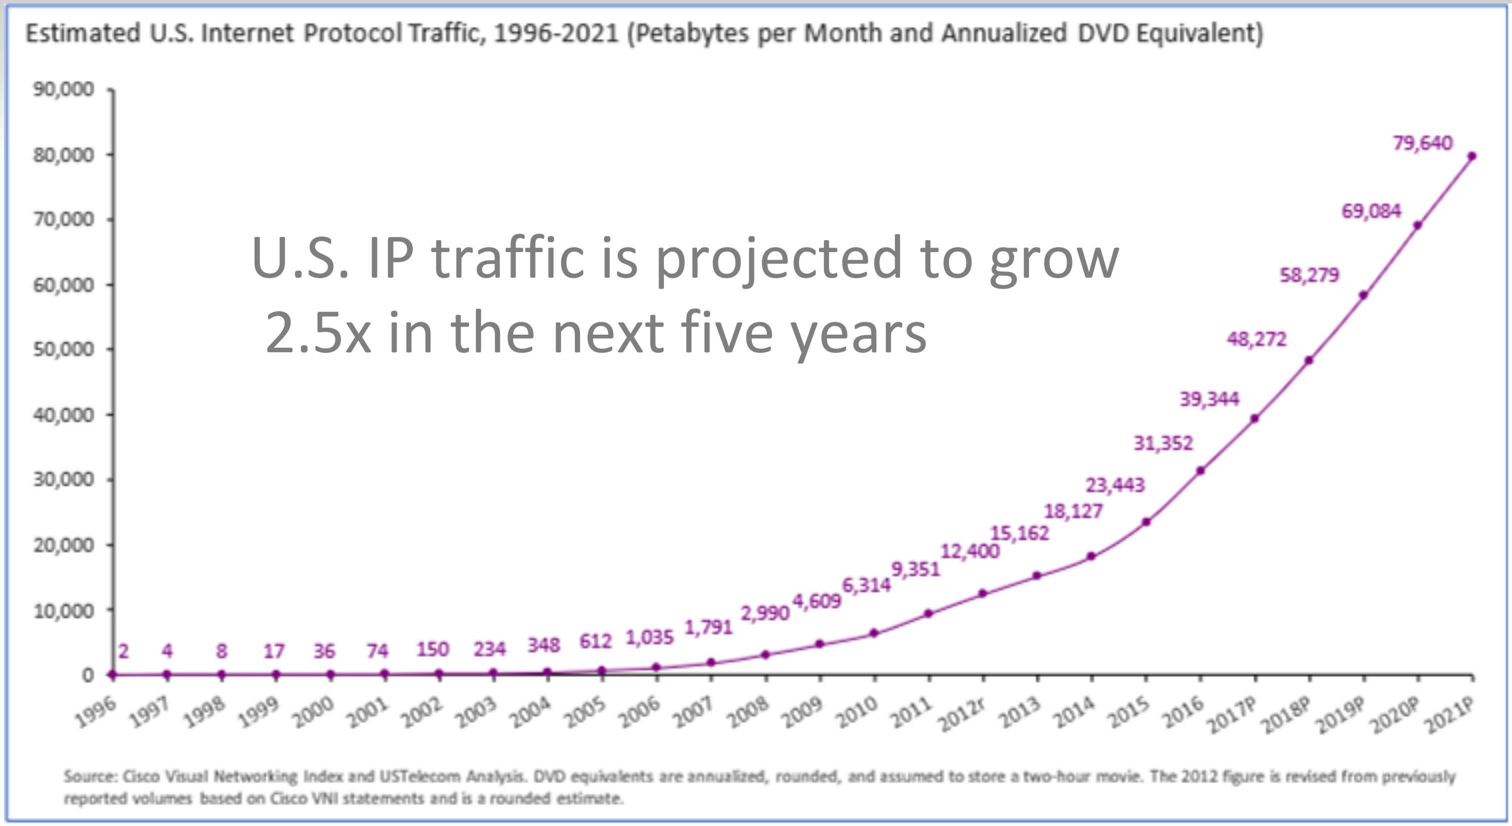

AS19653 Evolution of Transit to Peering

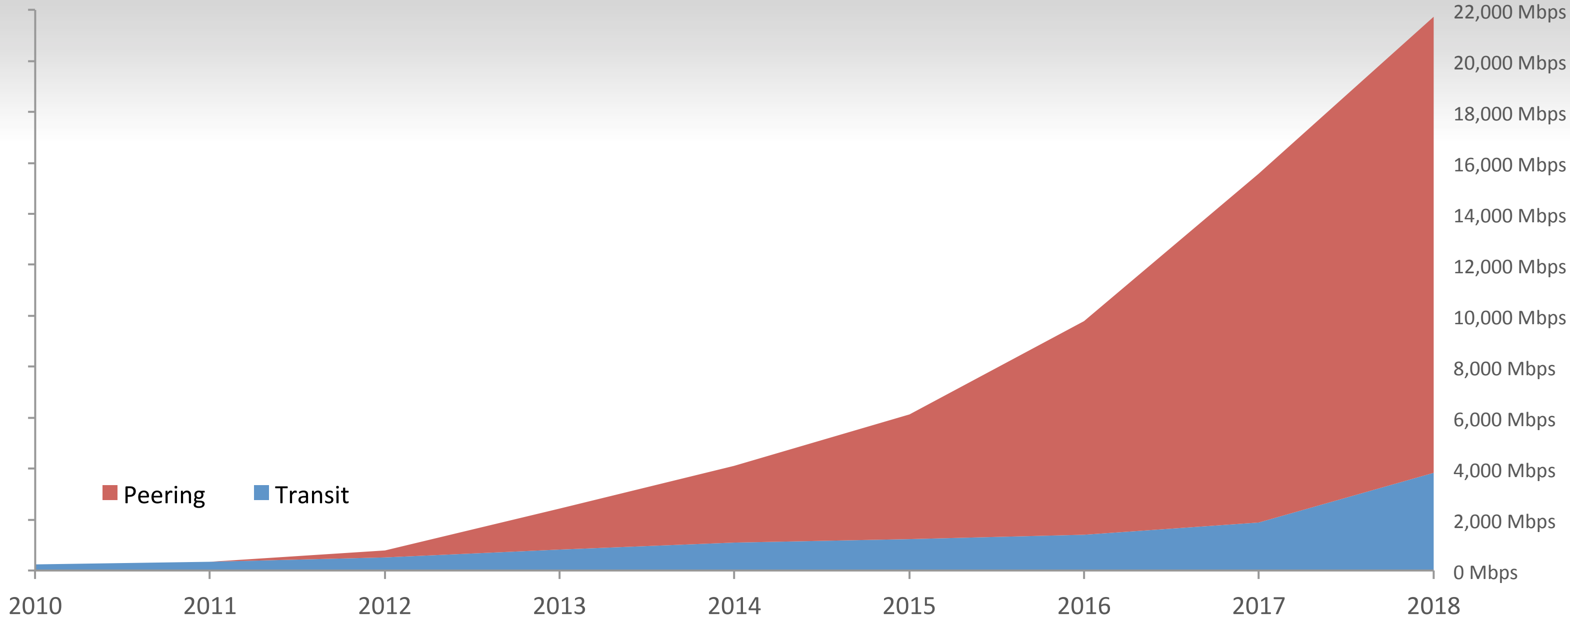

Peering	0%	0%	36%	66%	73%	80%	86%	88%	82%
Transit	100%	100%	64%	34%	27%	20%	14%	12%	18%

A Small Transit Provider Case Study

SOURCE: CTS Telecom

AS19653 From Transit to IXP Peering to CDN

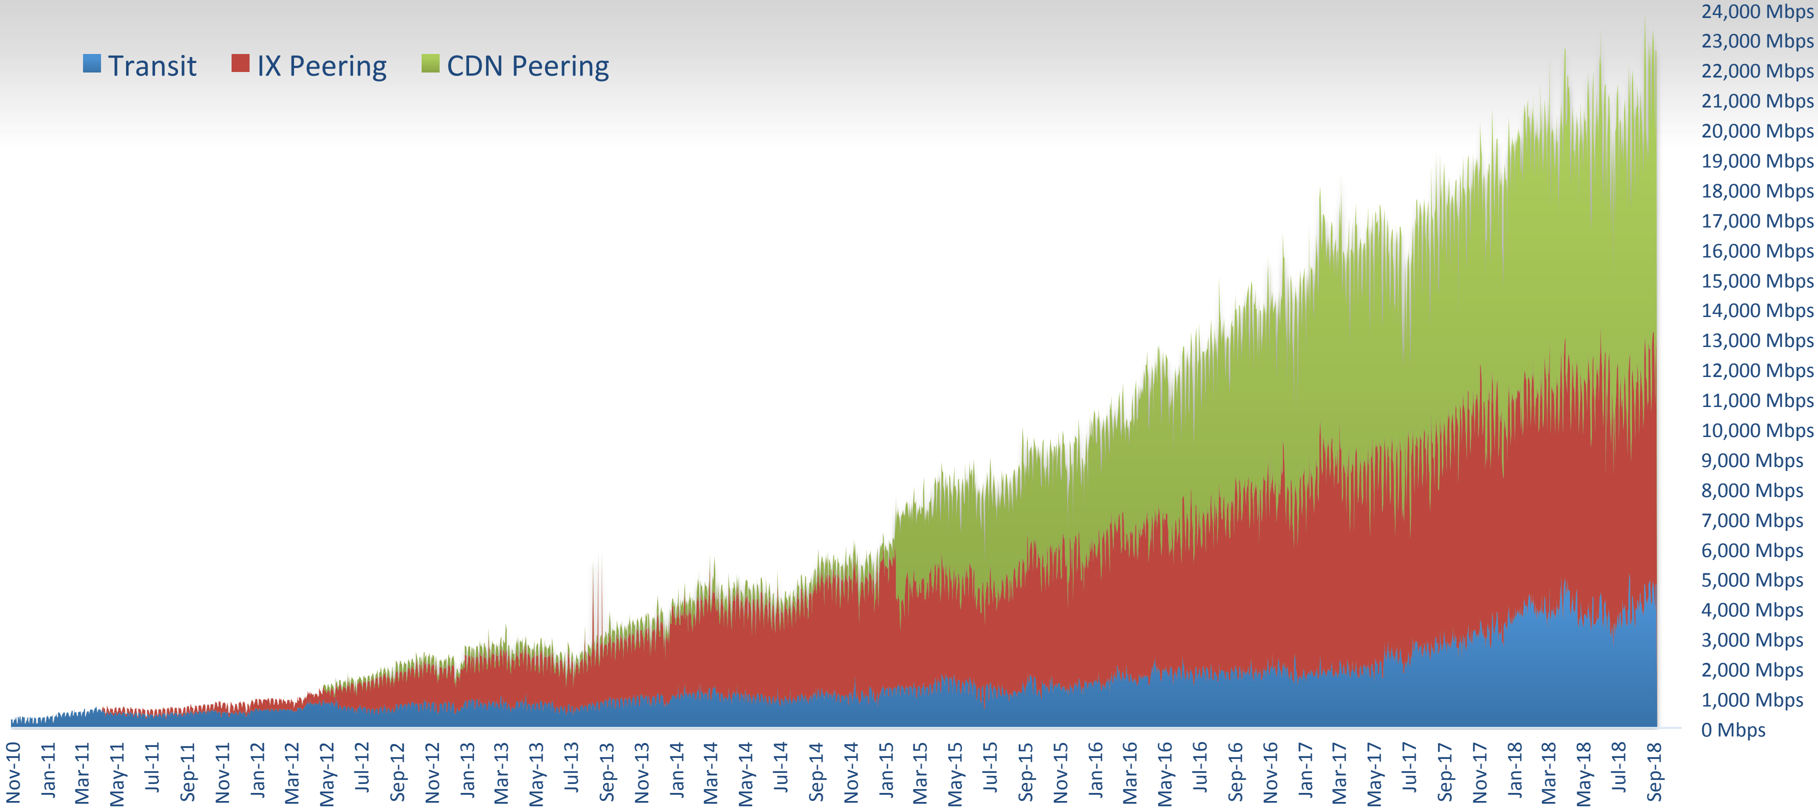

A Small Transit Provider Case Study

SOURCE: CTS Telecom

Percentage of Total Traffic AS19653 Transit/Peering/CDN

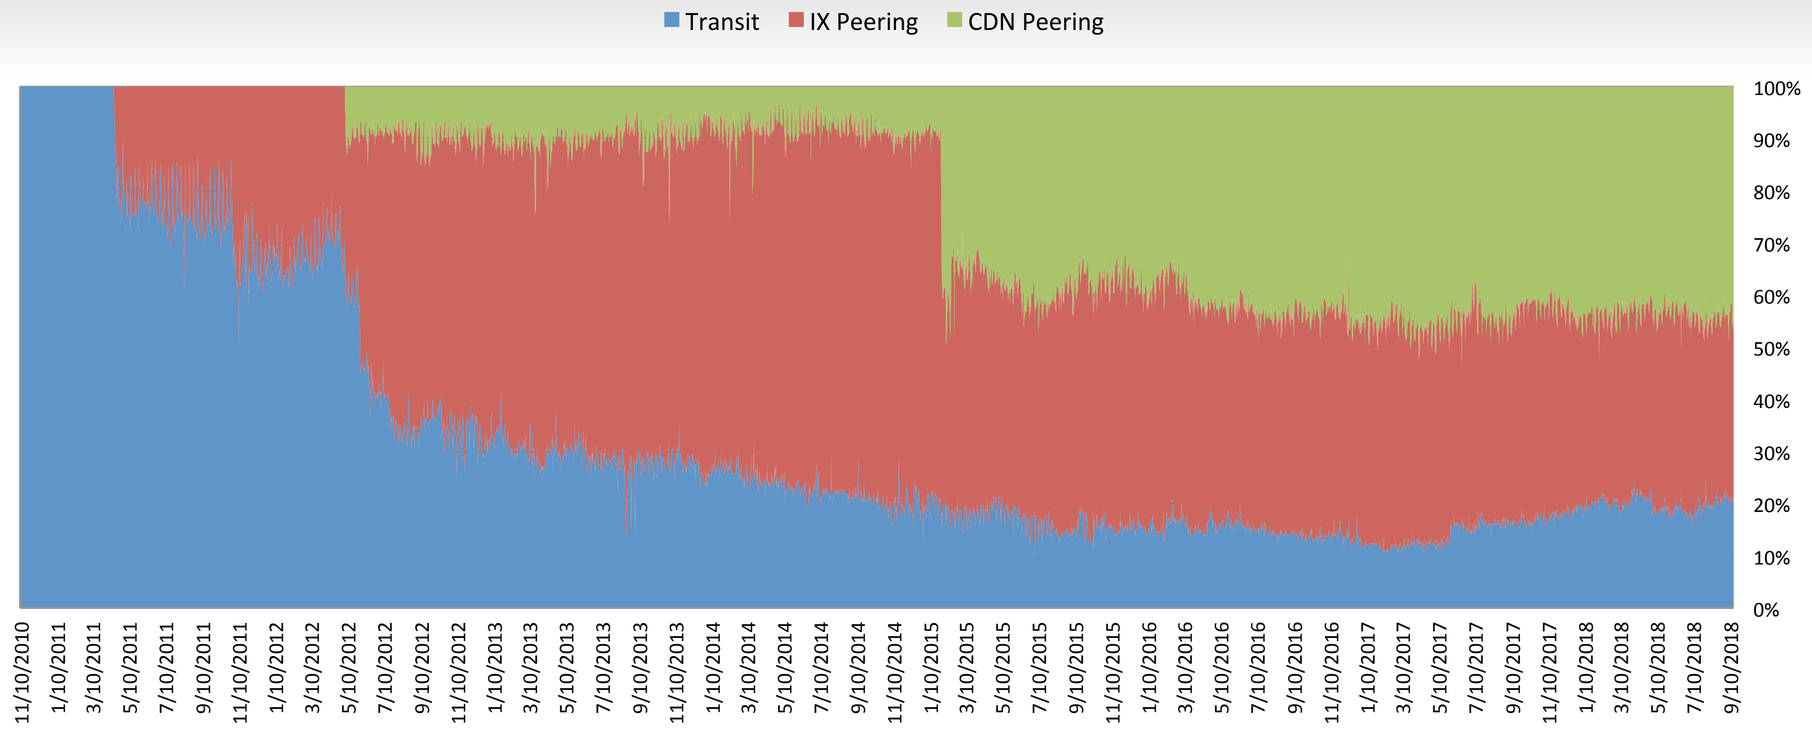

A Small Transit Provider Case Study

SOURCE: CTS Telecom

Network Snapshot AS19653

CAIDA AS rank: 1483
IPs in Customer Cone (v4): 143,104
Internet Exchanges: 3

Prefixes Originated (all): 12
Prefixes Originated (v4): 8
Prefixes Originated (v6): 4

Prefixes Announced (all): 57
Prefixes Announced (v4): 46
Prefixes Announced (v6): 11

BGP Peers Observed (all): 423
BGP Peers Observed (v4): 417
BGP Peers Observed (v6): 261

IPs Originated (v4): 90,624
AS Paths Observed (v4): 91,578
AS Paths Observed (v6): 19,424

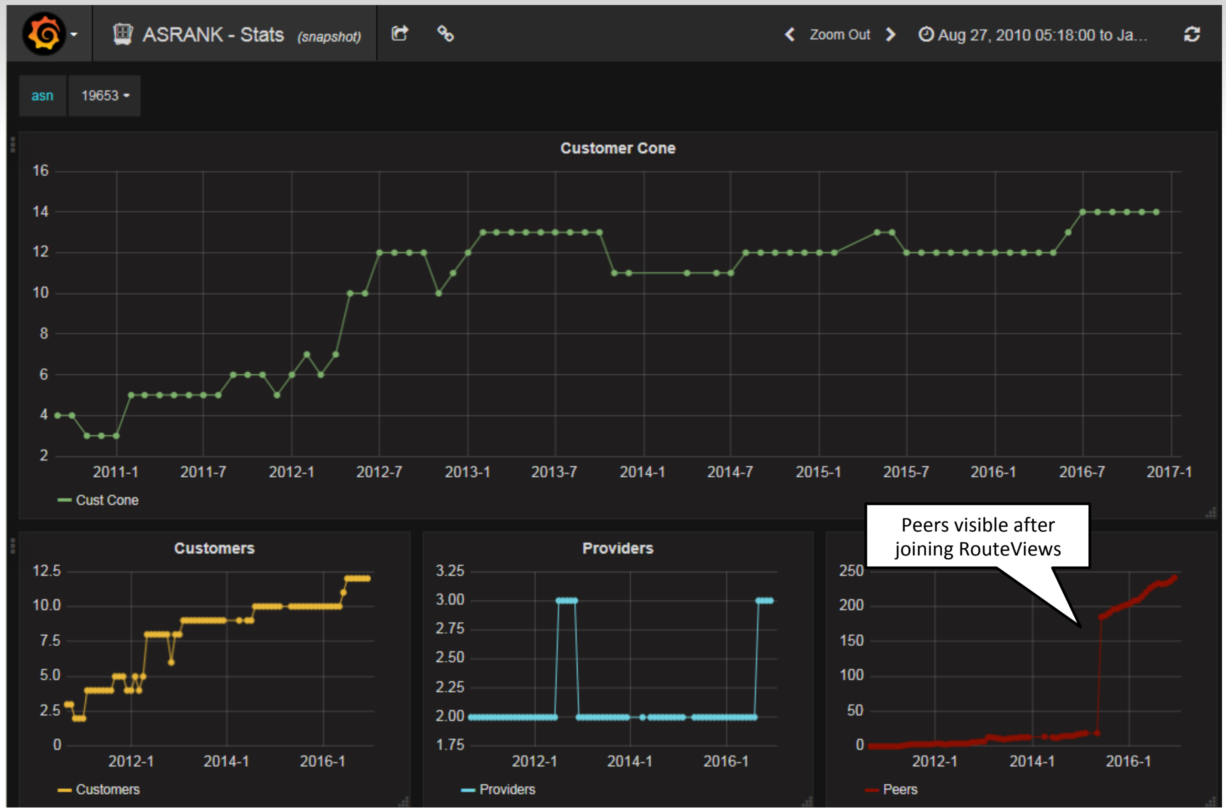

SOURCE: A. Dhamdhere CAIDA

Game Changers

- **Joining NANOG Community**
- **Establishing IXP presence**
- **Joined Peering Exchange**
- **Joined PeeringDB**
- **Read "The Internet is Flat"**
- **Implemented NetFlow analysis**
- **Developing NANOG "savoir faire"**
- **"Dr. Peering" Website** (Thanks to Bill Norton!)
- **Support of Content Providers**
- **Mentoring from the NANOG community**

Challenges and Cautions for Small Providers

- Unless you have a large enough number of "eyeballs" on your network and a high enough traffic level, peering does not make economic sense
- Peering requires a significant amount of technical expertise and commitments of resources.
- Connectivity to Internet Exchange Points is not trivial. Ideally a provider should be at two IXPs and redundant network connections are best. Selective Content Providers require peering at multiple locations.
- The falling price of Transit makes the case for peering for a small provider economically challenging: sometimes buying Transit is easier.
- You must have economical access to fiber transport to reach the IXP.