

Internet Inter-Domain Traffic

Craig Labovitz, Scott Iekel-Johnson,
Danny McPherson
Arbor Networks
Ann Arbor, MI
{labovit, scottij, danny}@arbor.net

Jon Oberheide, Farnam Jahanian
University of Michigan
Ann Arbor, MI
{jonojono, farnam}@umich.edu

ABSTRACT

In this paper, we examine changes in Internet inter-domain traffic demands and interconnection policies. We analyze more than 200 Exabytes of commercial Internet traffic over a two year period through the instrumentation of 110 large and geographically diverse cable operators, international transit backbones, regional networks and content providers. Our analysis shows significant changes in inter-AS traffic patterns and an evolution of provider peering strategies. Specifically, we find the majority of inter-domain traffic by volume now flows directly between large content providers, data center / CDNs and consumer networks. We also show significant changes in Internet application usage, including a global decline of P2P and a significant rise in video traffic. We conclude with estimates of the current size of the Internet by inter-domain traffic volume and rate of annualized inter-domain traffic growth.

Categories and Subject Descriptors: C.2 [Computer Communication Networks]: Miscellaneous

General Terms: Measurement.

1. INTRODUCTION

Saying the Internet has changed dramatically over the last five years is cliché – the Internet is **always** changing dramatically: fifteen years ago, new applications (*e.g.*, the web) drove widespread consumer interest and Internet adoption. Ten years ago, new backbone and subscriber access technologies (*e.g.*, DSL/Cable broadband) significantly expanded end-user connections speeds. And more recently, applications like social networking and video (*e.g.*, Facebook and YouTube) again reshaped consumer Internet usage.

But beyond the continued evolution of Internet protocols and technologies, we argue the last five years saw the start of an equally significant shift in Internet inter-domain traffic demands and peering policies. For most of the past fifteen years of the commercial Internet, ten to twelve large transit providers comprised the Internet “core” interconnecting thousands of tier-2, regional providers, consumer networks

and content / hosting companies. Textbook diagrams of the Internet and research publications based on active probing and BGP routing table analysis generally produce logical Internet maps similar to Figure 1a [1]. This diagram shows a strict hierarchy of global transit providers at the core interconnecting smaller tier-2 and regional / tier-3 providers.

Over the past several years industry economic forces, including the continued decline of the price of IP wholesale transit and the growth of advertisement-supported content, significantly altered the interconnection strategies of many providers [2]. In the emerging new Internet economy, content providers build their own global backbones, cable Internet service providers offer wholesale national transit, and transit ISPs offer CDN and cloud / content hosting services [3, 4, 5, 6]. For example, we found that over the last two years Google migrated the majority of its video and search traffic (which we later show constitutes more than 5% of all inter-domain traffic) away from transit providers to its own fiber backbone infrastructure and direct interconnects with consumer networks.

The substantial changes in provider inter-connection strategies have significant ongoing implications for backbone engineering, design of Internet-scale applications, and research. However, most providers treat their Internet traffic statistics with great commercial secrecy as these values reveal insights into market penetration and competitive strategies. As a result, the significant shift in Internet inter-domain traffic patterns has gone largely undocumented in the commercial and research literature.

Most Internet traffic research has typically focused on secondary indicators of Internet traffic such as BGP route advertisements [7, 8, 9], DNS probing [10], broad industry surveys [11], private CDN statistics [12], or traffic measured on an individual provider or enterprise network [13].

A few more closely related efforts have studied global Internet traffic using publicly available exchange point statistics [14] or a small set of residential networks [15, 16, 17, 18]. Still other work used industry surveys and targeted discussions with providers [19, 20, 21]. Finally, traceroute analysis in [22] also identified a topological trend towards a more densely interconnected Internet especially with respect to large content providers.

In this paper, we provide one of the first large scale longitudinal studies of Internet inter-domain traffic using direct instrumentation of peering routers across multiple providers. We address significant experimental data collection and commercial privacy challenges to instrument 3,095 peering routers across 18 global carriers, 38 regional / tier-

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.

SIGCOMM’10, August 30–September 3, 2010, New Delhi, India.
Copyright 2010 ACM 978-1-4503-0201-2/10/08 ...\$10.00.

2, and 42 consumer and content providers in the Americas, Asia, and Europe. At its peak, the study monitored more than 12 terabits per second of offered load and a total of more than 200 exabytes of Internet traffic over the two-year life of the study (July 2007 to July 2009). Based on independent estimates of total Internet traffic volume in [14, 23], we believe the probes directly monitor more than 25% of all Internet inter-domain traffic.

Our major findings include:

- **Evolution of the Internet “Core”:** Over the last two years, the majority of Internet inter-domain traffic growth has occurred outside the traditional ten to twelve global transit carriers. Today, most Internet inter-domain traffic by volume flows directly between large content providers, hosting / CDNs and consumer networks.
- **Consolidation of Content:** Most content by inter-domain traffic volume has migrated to a relatively small number of large hosting, cloud and content providers. Out of the approximately thirty-thousand ASNs in the default-free BGP routing tables [24], 30 ASNs contribute a disproportionate average of 30% of all Internet inter-domain traffic in July 2009.
- **Estimation of Google’s Traffic Contribution:** At an average of more than 5% of all inter-domain traffic in July 2009, Google represents both the largest and fastest growing contributor of inter-domain traffic. Google’s share of all inter-domain traffic grew by more than 4% between July 2007 and July 2009.
- **Consolidation of Application Transport:** The majority of inter-domain traffic has migrated to a relatively small number of protocols and TCP / UDP ports, including video over HTTP and Adobe Flash. Other mechanisms for video and application distribution like P2P have declined significantly in the last two years.
- **Estimation of Internet Size:** Using data from independent known inter-domain provider traffic volumes, we estimate both the volume and annualized growth rate of all inter-domain traffic. As of July 2009, we estimate inter-domain traffic peaks exceed 39 Tbps and grew an annualized average of 44.5% between July 2007 and 2009.

The rest of this report is organized as follows: §2 provides an overview of our data collection infrastructure and analysis methodology. §3 discusses significant changes in Internet topology and commercial interconnection relationships between providers. §4 analyzes changes in Internet protocols and applications. Finally, we conclude with validation of our data and estimates of both the volume of all inter-domain traffic and annualized rate of growth.

2. METHODOLOGY

Our analysis in this paper is based on traffic statistics exported by operational routers from a large and, we argue later, representative sample of Internet providers. Specifically, we leverage a widely deployed commercial security and traffic monitoring platform to instrument the BGP peering

edge routers of 110 participating Internet providers. Based on private commercial sales data, we believe the majority of the probe deployments enjoy complete coverage of the provider’s BGP peering edge. However, we lack specific visibility into the network probe coverage of any individual anonymous study participant.

The instrumented routers export both traffic flow samples (*e.g.*, NetFlow, cFlowd, IPFIX, or sFlow) and participate in routing protocol exchange (*i.e.*, iBGP) with one or more probe devices. A smaller number of providers have deployed inline or “port span” versions of the appliances to monitor traffic payloads and enact security policies. Per our anonymity agreement with participating providers, we did not collect more specific details on deployment configuration (*e.g.*, flow sample rates, router model number, etc.).

While sampled flow introduces potential data artifacts particularly around short-lived flows [25], we believe the accuracy of flow is sufficient for the granularity of our inter-domain traffic analysis. Further, we argue flow provides the only scalable and cost-effective monitoring approach given the scale of our study.

Each probe independently calculates traffic statistics based on user configured information and BGP learned topology. Calculated statistics include breakdowns of traffic per BGP autonomous system (AS), ASPath, network and transport layer protocols, ports, nexthops, and countries. A more detailed description of the probe capabilities is available in commercial datasheets and white papers at [26].

The probe configuration includes user supplied classification of the probe’s primary geographic coverage area (*i.e.*, North America, Europe, etc.) as well as market segment (*i.e.*, tier-1, tier-2, content, consumer or educational). We use the provider supplied self-categorizations in our aggregate data analysis discussed in later Sections.

We worked extensively with the provider community to address commercial privacy concerns. For example, every participating probe strips all provider identifying information from the calculated statistics before forwarding an encrypted and authenticated snapshot of the data to central servers. We also agreed to not publish any per provider traffic rates nor customer data derived from ASPath traffic analysis.¹

We pursued several approaches to mitigate sources of possible error in the data. We began by excluding three ISPs (out of 113) from the dataset that exhibited signs of obvious misconfiguration via manual inspection (*i.e.*, wild daily fluctuations, unrealistic traffic statistics, internally inconsistent data, etc.).

Unfortunately, our measurement infrastructure suffered from the real-world operational exigencies of providers. Throughout the course of the study, providers expanded deployments with new probes, decommissioned older appliances and otherwise modified the configuration of their probes and backbone infrastructure. As a result, the absolute traffic volumes reported by probes exhibited occasional discontinuities. For example, one probe consistently

¹While we discuss several Internet providers by name in this paper, we base all provider-specific analysis on anonymized ASN and ASPath datasets aggregated across all study participants. Any overlap or correlation with providers who may (or may not) be sharing data or have research or commercial affiliations with the institutions or authors of this paper is unintended and coincidental.

Segment	Percentage
Regional / Tier2	34
Global Transit / Tier1	16
Unclassified	16
Consumer (Cable and DSL)	11
Content / Hosting	11
Research/ Educational	9
CDN	3

(a) Market Segment

Region	Percentage
North America	48
Europe	18
Unclassified	15
Asia	9
South America	8
Middle East	1
Africa	1

(b) Geographic Region

Table 1: Distribution of anonymous Internet provider participants in our study by market segment and geographic region.

reported hundreds of gigabits of traffic until dropping to zero abruptly in early 2009 as the provider migrated traffic to other routers and newer probe appliances.

The probe data exhibited less variance with respect to traffic ratios (*i.e.*, the ratio of ASN, port, protocol, etc. to all inter-domain traffic in each deployment). Specifically ratios such as TCP port 80 or Google ASN origin traffic remained relatively consistent even as the number of monitored routers, probe appliances and absolute volume of reported traffic fluctuated in a deployment. Given the relative consistency of ratios and our inability to distinguish changes in absolute traffic volumes from artifacts due to provider measurement infrastructure changes, most of the analysis in this paper focuses on traffic percentages (*i.e.* share of traffic) rather than absolute traffic values. The focus on ratios also simplifies our aggregate analysis across a large set of heterogeneous providers.

Throughout every 24 hour period, the probes independently calculated the average traffic volume every five minutes for all members of all datasets (*i.e.*, traffic contributed by every nexthop, AS Path, ASN, etc.) as well as the average volume of total inter-domain network traffic. The probes then calculated a 24 hour average for each of these items using the five minute averages. Finally, the probes used the daily traffic volume per item and network total to calculate a daily percentage for each item.

The first chart in Table 1 provides a market segment breakdown of anonymous provider participants by percentage of all deployments in our study. The second table shows a breakdown of percentage of deployments by geographic region. Regional and tier-2 providers comprise the largest component at 34% of anonymous statistics followed by unclassified and tier-1 at 16% each.

We observe that the relative high cost of the commercial probes used in our study may introduce a selection bias towards larger providers. We further note that both analyst data and our study participant set reflect a continued weighting towards North America and Europe both in traffic volume and number of providers [27, 11, 6, 28].

While our study included a large and diverse set of Internet providers, evaluation of sample bias is a challenge given the anonymity of the study participants and the lack

of “ground-truth” quantitative market data (*i.e.*, most available data on provider Internet traffic volumes is based on qualitative surveys [27, 11]).

We evaluated several mechanisms for weighting the traffic ratio samples from the 110 deployments to reduce selection bias. However, the anonymity of the study participants and the narrow scope of our data collection provided a limited number of weighting options. Ultimately, we found a weighted average based on the number of routers in each deployment provided the best results during data validation in §5 and represents a compromise between the relative size of an ISP while not obscuring data from smaller networks.

Specifically, for each day d we calculate the weighted average percent share of Internet traffic $P_d(A)$ for a specific traffic attribute A , where A is an ASN, TCP port, country of origin, etc. The weights are calculated based on the total number of routers reporting traffic on that day at each of the N study participants reporting data for that day. Thus, on day d for participant i with router count $R_{d,i}$ we calculate the weight:

$$W_{d,i} = \frac{R_{d,i}}{\sum_{x=1}^N R_{d,x}}$$

We then calculate day d ’s weighted average percent share $P_d(A)$ based on each provider’s measured average traffic volume for A on day d , $M_{d,i}(A)$, and total average inter-domain traffic for day d , $T_{d,i}$. This gives a weighted average percent share of traffic for A as

$$P_d(A) = \sum_{x=1}^N W_{d,x} * \frac{M_{d,x}(A)}{T_{d,x}} * 100$$

We excluded any provider more than 1.5 standard deviations from the true mean in order to focus on values that were less likely to have measurement errors due to transient provider issues (misconfiguration, network problems, or probe failures). With the exception of Comcast’s peering ratios discussed in §3, we used the sum of traffic both in and out of the provider networks for $M_{d,i}(A)$ and $T_{d,i}$.

In some cases, our analysis may underestimate categories of inter-domain traffic. Specifically, the probes lack visibility into traffic exchanged over direct peering adjacencies between enterprise business partners or between smaller tier-2 and tier-3 Internet edge providers. Similarly, the study may underestimate inter-domain traffic associated with large content providers such as Google who are increasingly pursuing edge peering strategies. We also emphasize that our study is limited to inter-domain traffic and excludes all internal provider traffic, such as intra-domain cache traffic, VPNs, IPTV and VoIP services.

Finally, we validated our findings with private discussions with more than twenty large content providers, transit ISPs and regional networks. These discussions provided “ground-truth” and additional color to better understand the market forces underlying our observed inter-domain traffic trends. We note that our derived data matched provider expectations both in relative ordering and magnitude of ASN traffic volumes. In addition, twelve providers supplied independent inter-domain traffic measurements for validation of our analysis. We use these twelve known provider traffic values in §5 to add confidence to our calculated inter-domain ASN traffic distributions as well as to estimate the overall volume of global inter-domain traffic.

