Poster: Footprint and Performance of Large Cloud Networks

Jia He[1], Weifan Jiang[1], Ege Gürmeriçliler[1], Georgia Essig[1], Arpit Gupta[1,2], Matt Calder[3], Vasileios Giotsas[4], Italo Cunha[5], Ethan Katz-Bassett[1], Todd Arnold[1]

[1]Columbia University, [2]UC Santa Barbara, [3]Microsoft, [4]Lancaster University, [5]Universidade Federal de Minas Gerais

{jh4001, wj2301, eg2985, ge2211}@columbia.edu, arpitgupta@cs.ucsb.edu, mcalder@microsoft.com, v.giotsas@lancaster.ac.uk, cunha@dcc.ufmg.br, {ebk2141, ta2510}@columbia.edu

Abstract

The Tier-1 and Tier-2 transit providers have historically been considered the backbone of the Internet as they guarantee global reachability. In recent years, Internet flattening has reduced the need for transit providers, an effect greatly contributed to by the top cloud providers, such as Google, Amazon, Microsoft, and IBM. Recently, these cloud providers started offering two performance tiers for routing traffic. One tier, referred to as "Premium Tier" (PT), the Google-specific term, uses the cloud provider's private network as much as possible, while "Standard Tier" (ST), uses the public Internet as much as possible. Through analysis of measurements made to gather performance and connectivity data, we find that the cloud provider networks' points-of-presence (PoPs) tend to be deployed closer to population centers than the transit providers' PoPs. We also find that the performance improvement from PT service is dependent on variables such as the ST/PT path length difference. These metrics demonstrate how cloud providers connect within the Internet, and what benefits their private networks provide to users.

1 Introduction and Motivation

In the past decade, the topology of the Internet has become much more densely connected [7, 9, 11] and paths have become shorter. This change in the structure of the Internet is known as "Internet flattening" and has been brought in large part due to the top cloud provider networks: Amazon, Google, Microsoft, and IBM. These companies deploy global private wide-area networks (WANs) and are very well connected to other networks [1, 4]. As a result, traffic to and from these networks often does not use Tier-1 or Tier-2 networks for transit [17], contributing to Internet flattening.

To understand the impact the top cloud providers have on Internet flattening, we quantify the population footprint within certain distances of their points-of-presence (PoPs). We measure these population coverages using individual network-level graph models we have created for the top cloud providers, Tier-1s, and Tier-2s using data from PeeringDB [14], network maps, and rDNS records.

Cloud providers now offer tiers of service to their users [2, 3, 10]. We use the Google-specific terms for these tiers for all cloud providers. For "Premium Tier" (PT) service, user traffic is routed as much as possible within the cloud provider's private global WAN, and for "Standard Tier" (ST), user traffic is routed as much as possible through the public Internet. We also wish to understand the effect the private WANs have on end users' performance.

2 Measurement Campaign

Measurements were taken over the course of 2019-2020 so different performance and connectivity properties could be analyzed. These measurements were taken from Speedchecker [16] and RIPE Atlas vantage points (VPs) [15] located around the globe to ST and PT cloud platform virtual machine (VM) instances deployed in multiple regions. The measurements consist of traceroutes for path evaluation, pings for latency, and HTTP for throughput.

The traceroutes are used to gain insight into routes taken by traffic. Therefore, we annotate each traceroute hop with the ASN using the Cymru IP-to-AS mapping [6] and a geolocation if it is available. The ping time and throughput measurement are used to analyze the performance differences between ST and PT.

We create network graphs using data from PeeringDB [12, 14] and network maps provided by Tier-1 and Tier-2s and the cloud providers, such as Tata [5] and NTT [13]. In the graph, nodes represent network edge PoPs and edges represent links between edge PoPs. The edges are weighted with the geographic distance between nodes, allowing the use of a shortest path algorithm to find optimal paths and path lengths through each network.

3 Population Coverage

As the cloud networks seek to deliver as many services to as many customers as possible, we expect that their network PoPs are placed to be as close to as much of the world's population as possible. By counting the population within different distances of the network's PoPs, we can gain an understanding of how effective the cloud providers are at being accessible to users.

For every cloud and transit provider network with a network graph, we examine the population coverage at 500, 700 and 1000 kilometer radii around the PoP coordinates, using up-to-date population density data [8]. These population counts are then summed together without double counting overlapping radii and then divided by the world population to calculate a population coverage percentage per network. We also perform a similar analysis, except forming a union of all edge PoPs for transit and cloud provider networks on a per continent basis, to investigate if there are regions that are underserved by either the cloud or transit providers.

Figure 1a shows the results of the population coverage analysis for each network, and we find that three out of the four top networks by population coverage are cloud providers. Figure 1b shows that some continents, such as Africa, are relatively underserved by the cloud providers. For these continents, the union of transit provider PoPs provides greater population coverage.

4 Benefits of Private WAN

We also investigate how the greater proportion of world population covered by cloud provider PoPs can benefit from the PT service. PT service will route traffic within cloud provider's private WAN as much as possible, potentially improving performance [1]. For <City,AS> locations in each country, we can find the difference in

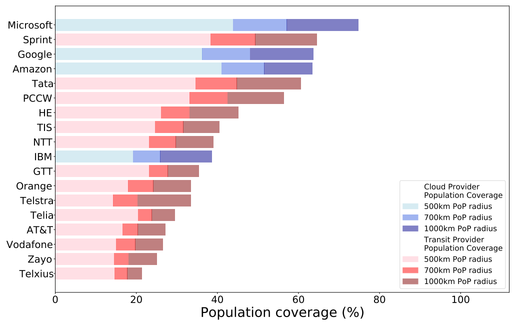

(a)

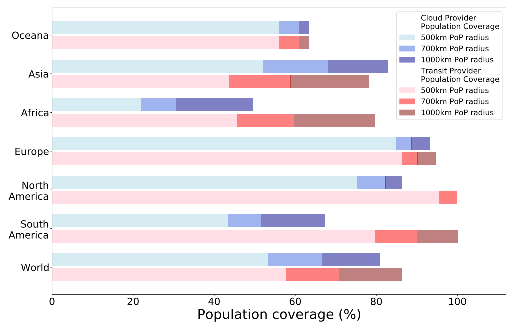

(b)

Figure 1: Percentage of the population within a 500, 700, and 1000 km radius from PoP per per provider (Fig. 1a) and per provider group per continent (Fig. 1b).

median ping time by sending pings to both a ST and PT VM in a given region, to provide insights into performance differences. Figure 2 shows that countries in the Americas, Europe and Oceana see improvement in median ping time for the Western Europe destination VMs when using PT. However, countries in Asia see much worse performance.

Understanding why different regions see different performance benefits can help predict how performance will differ between ST and PT. One potential variable impacting performance is the difference in path length taken by ST and PT traffic. For each <City,AS> location, we can use the geolocated traceroutes along with the network graphs to estimate the ST/PT difference in path length taken by traffic to the destination VM. Figure 3 shows a reasonable correlation between the median ping time and path length differences for <City,AS> locations in Asia and destination VM in Western Europe. This means that the worse performance in this case may be due to the longer PT path length. We find that this effect is caused by Google's private WAN lacking a direct connection between Asia and Europe, meaning traffic must go through the US to reach Europe from Asia. The correlation is not extremely strong due to other factors which may influence the performance difference, such as proximity to a Google PoP. Outliers may also be caused by inaccuracies in the distance measurements.

5 Conclusion

We show that the cloud providers have established themselves as some of the most reachable and well-connected networks in the Internet, a place occupied mostly by transit networks. We also show that PT service offered by cloud providers can provide users

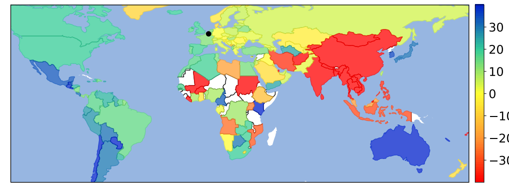

Figure 2: Difference in median ping time (ST-PT), in ms, per country when using ST versus PT to reach the Western Europe VMs. The VM location is shown by the black dot.

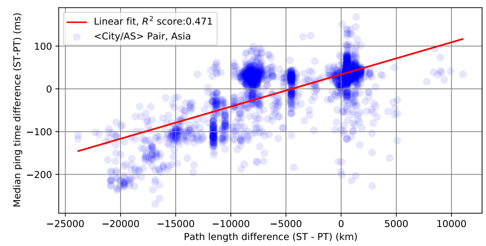

Figure 3: Correlation of median ping time and path length difference for pings originating in <City,AS> locations in Asia with destination VM in Western Europe.

with better performance using their private WANs, but that the benefits depend on multiple variables. Our future work is to form better predictions of performance, as this can be an important consideration for users who wish to deploy cloud services on PT.

References

[1] Todd Arnold, Ege Gurmericliler, Arpit Gupta, Matt Calder, Georgia Essig, Vasileios Giotsas, and Ethan Katz-Bassett. (How Much) Does a Private WAN Improve Cloud Performance? In *IEEE International Conference on Computer Communications (INFOCOM) 2020*, July 2020.

[2] Amazon AWS. AWS global accelerator. https://aws.amazon.com/global-accelerator/.

[3] Microsoft Azure. What is routing preference (preview)? https://docs.microsoft.com/en-us/azure/virtual-network/routing-preference-overview.

[4] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. Are We One Hop Away from a Better Internet? In *Proceedings of the 2015 Internet Measurement Conference*, IMC '15, 2015.

[5] Tata Communications. Network Map. https://www.tatacommunications.com/map/.

[6] Team Cymru. IP-to-ASN Mapping. http://www.team-cymru.com/IP-ASN-mapping.html.

[7] Amogh Dhamdhere and Constantine Dovrolis. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *Proceedings of the 6th International COnference*, Co-NEXT '10, 2010.

[8] Center for International Earth Science Information Network CIESIN Columbia University. Gridded Population of the World, Version 4 (GPWv4): Population Density, Revision 11. Accessed 22 Feb 2020. https://doi.org/10.7927/H49C6VHW, 2018.

[9] Phillipa Gill, Martin Arlitt, Zongpeng Li, and Anirban Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proceedings of the International Conference on Passive and Active Network Measurement*, PAM '08, April 2008.

[10] Prajakta Joshi. Introducing network service tiers: Your cloud network, your way. https://cloud.google.com/blog/products/gcp/introducing-network-service-tiers-your-cloud-network-your-way.

[11] Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. Internet Inter-Domain Traffic. *SIGCOMM Comput. Commun. Rev.*, August 2010.

[12] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, and kc claffy. Using PeeringDB to Understand the Peering Ecosystem. *SIGCOMM Comput. Commun. Rev.*, April 2014.

[13] NTT. Network Map. https://www.gin.ntt.net/about/our-network/.

[14] PeeringDB. https://www.peeringdb.com.

[15] RIPE NIC. RIPE ATLAS. https://atlas.ripe.net.

[16] Speedchecker. ProbeAPI. http://probeapi.speedchecker.com/.

[17] Bahador Yeganeh, Ramakrishnan Durairajan, Reza Rejaie, and Walter Willinger. How Cloud Traffic Goes Hiding: A Study of Amazon?s Peering Fabric. In *Proceedings of the Internet Measurement Conference*, IMC '19, 2019.