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[1,3], Vasileios Giotsas[4], Italo Cunha[1,5], Ethan Katz-Bassett[1], Todd Arnold[1]

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

Motivation

Tier-1 and Tier-2 transit providers have traditionally been considered the backbone of the Internet as they guarantee global reachability. In recent years, the Internet has become more densely connected in a process known as Internet Flattening, leading to less reliance on Tier-1 and Tier-2 transit networks. The top cloud providers, which include Google, Microsoft, Amazon, and IBM, have greatly contributed to Internet Flattening by deploying very well-connected, global private wide-area networks (WANs) which allow them to bypass Tier-1 and Tier-2 transit providers.

We wish to quantify how these cloud provider networks are changing the landscape of the Internet by calculating the population footprint around their network points-of-presence (PoPs), and what benefits their private networks have to users.

Cloud Providers WANs vs the Public Internet

The large cloud providers offer two tiers of service to users. We use Google-specific terms to cover all of the cloud providers:

- Premium Tier (PT):
  - Use of private WAN as much as possible. Traffic exits and enters the cloud provider as close to the client as possible.
- Standard Tier (ST):
  - Use of public Internet as much as possible. Traffic exits and enters the cloud provider near the cloud data center.

Measurement Campaign

We conducted a measurement campaign over 2019 and 2020 to analyze different performance metrics. In particular we measured:

- Throughput and ping times, to measure how the performance differs between PT and ST, and thus what improvement the private WAN provides.
- Traceroutes, to gain insight into routes taken by traffic and interconnections between networks.

We measure from Speedchecker and RIPE Atlas vantage points (VPs) around the globe to Google ST and PT virtual machines (VMs) in four locations: Europe-West, US-Central, US-East and Asia Northeast. The locations of the VPs are shown below.

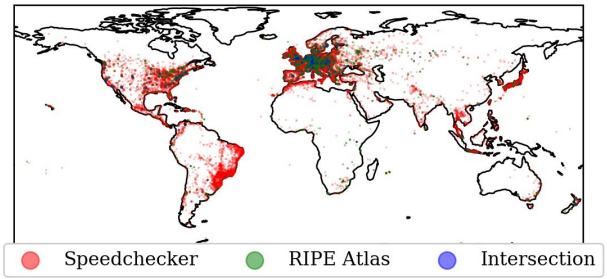

Results

- The individual cloud providers tend to have a greater population footprint at their PoPs than the transit providers.
- In some regions, such as Africa, the transit PoPs as a whole have a greater population footprint than the cloud providers.

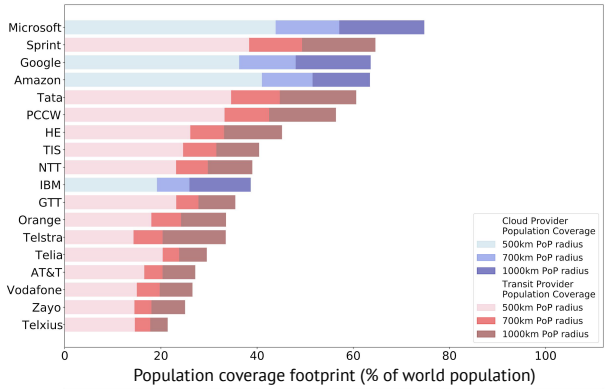

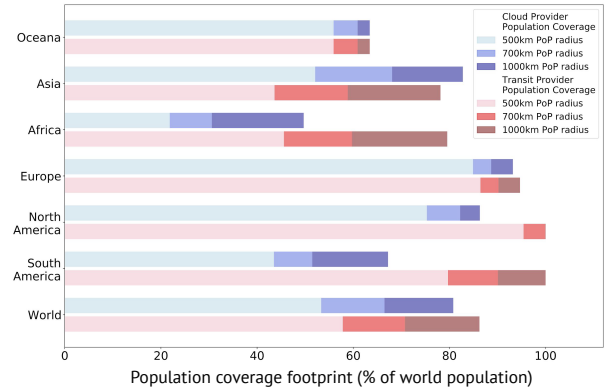

- We define performance improvement as the difference of ST median ping time and PT median ping time.
- We find that the performance difference varies by country and by VM location. By looking at correlation statistics, we can determine what factors can affect performance in certain cases.

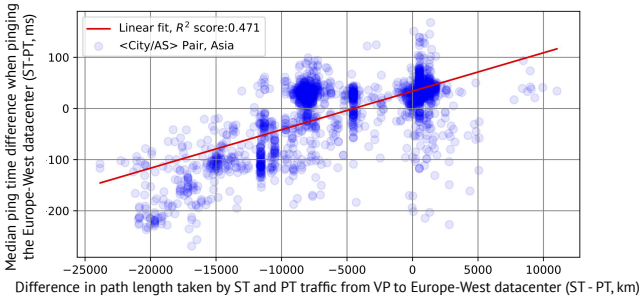

- Individual cases depend on specifics of geographic and network topology. Please see our INFOCOM 2020 paper for case studies.

Conclusions and Future Work

We find that the top cloud providers have established themselves as some of the most accessible networks in the Internet. Their global private WANs have a greater population footprint in densely populated areas than transit providers, and allows them to provide superior PT performance to customers. However, in some cases, the performance with PT can be worse. Measuring these conditional benefits and understanding what impacts them can help customers optimize their cloud platform deployments economically and performance-wise.