Internet measurements on public cloud platforms

Ricky Mok

CAIDA/SDSC

August 9, 2022

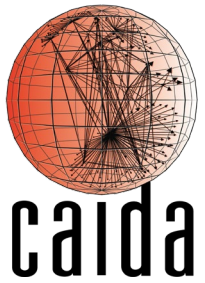

The Ubiquitous Cloud

- Video conferencing
- Distance learning
- Entertainment

➔Important to understand the network performance between the cloud and ISPs

Cloud is not “invincible”

AWS suffering EC2 and EBS performance issues in Northern Virginia

Storage coordination issue affecting EC2 and EBS instances, issues still ongoing for some

September 27, 2021 By: Dan Swinhoe Comment

https://www.datacenterdynamics.com/en/news/aws-suffering-ec2-and-ebs-performance-issues-in-northern-virginia/

Google’s London data center outage during heatwave caused by “simultaneous failure of multiple, redundant cooling systems”

Search giant investigating what went wrong, will be auditing all cooling systems globally

August 02, 2022 By: Dan Swinhoe Comment

https://www.datacenterdynamics.com/en/news/googles-london-data-center-outage-during-heatwave-caused-by-simultaneous-failure-of-multiple-redundant-cooling-systems/

Major Microsoft Teams and Azure outage in South Africa

Jan Vermeulen 29 July 2022

https://mybroadband.co.za/news/internet/454688-major-microsoft-teams-and-azure-outage-in-south-africa.html

Research questions

- How do the cloud connect to the rest of the Internet? [1]
  - Infer the number and location of the interconnection between cloud regions and networks
- How can we capture network performance variations between the cloud and different ISPs? [2]
  - Estimate the performance users perceived in cloud-based applications
- How can we perform bandwidth measurement efficiently?
  - Lower measurement overhead (traffic and cost)

[1] Alexander Marder, kc claffy, and Alex C. Snoeren, “Inferring Cloud Interconnections: Validation, Geolocation, and Routing Behavior”, *In PAM 2021*.
[2] Ricky K. P. Mok, Hongyu Zou, Rui Yang, Tom Koch, Ethan Katz-Bassett, kc claffy, "Measuring the network performance of Google Cloud Platform", *In ACM IMC, 2021*.
Image source: https://www.thecloudpeople.com/blog/move-everything-to-the-cloud-or

Our approach

- Use the cloud to measure the cloud!
- Deploy virtual machines in multiple cloud regions to run “speed tests” from the cloud to speed test servers in ISPs
  - Download and Upload throughput
  - Latency
  - Network paths (Traceroute, a lot of it!)

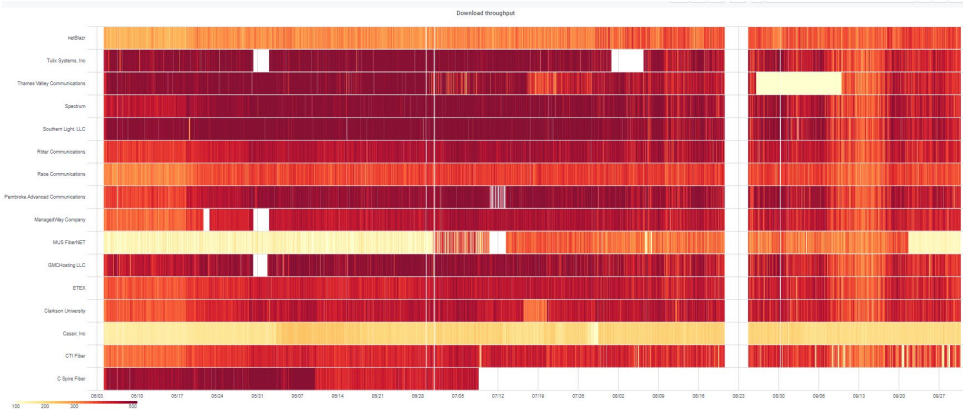

Heatmap of download throughput measured from a GCP VM to speed test servers in different networks.

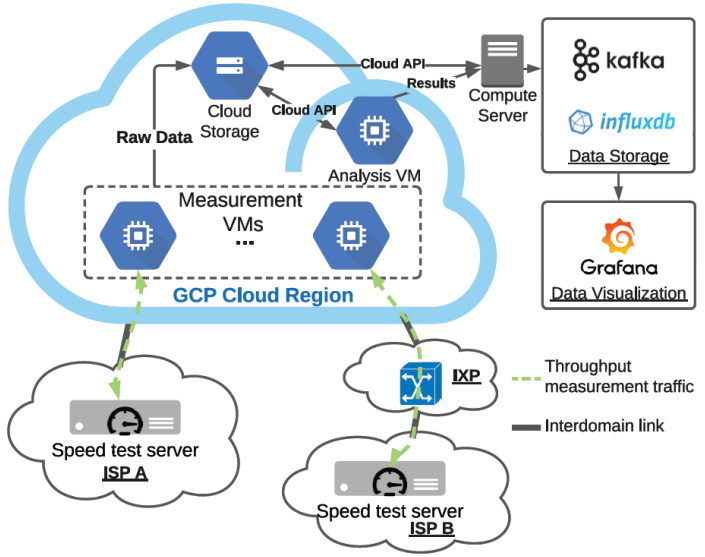

The architecture of **CL**oud-based **A**pplications **S**peed **P**latform (CLASP) [2].

Challenges and our cloud-based solutions

- Set up multiple VMs across cloud regions in different cloud
  - Develop scripts using cloud platform's APIs to automate the deployment
- Large amount of raw measurement data
  - Leverage cloud storage buckets (e.g., AWS S3) to store multi-TetaBytes of data
- Cost control
  - VMs and Egress bandwidth could be costly
  - Uses Nutanix Beam to monitor spending

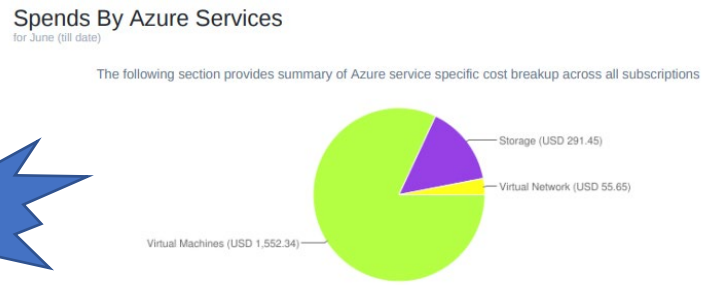

Quality of Experience measurements

- Study the impact of network performance degradation on the quality of experience (QoE) of cloud applications
  - Conduct measurements from many end-users *in the wild*
  - Design and implement a gamified web platform to crowdsource the experiment
  - Trigger targeted measurements from the cloud
- *CNS Core: Small: A Unified Approach to Internet Performance Measurement* https://www.caida.org/funding/cns-quinceng/

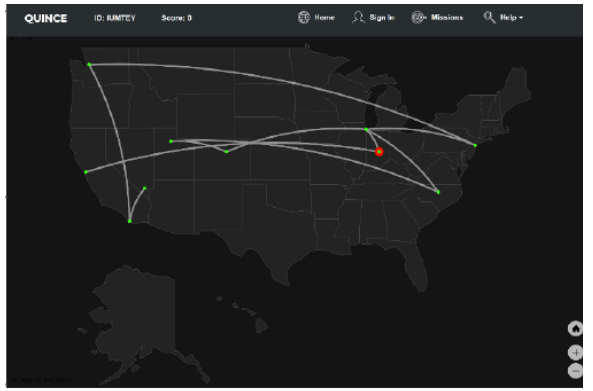

Gamified user interface of QUINCE

Detecting cloud-Internet bottlenecks

- Improve the efficiency of bandwidth measurements in the cloud
  - High-speed network paths
  - Virtualized environment
- Identify the bottlenecks between the cloud and end-users
  - Router
  - Physical location (city)
  - Number of affected users

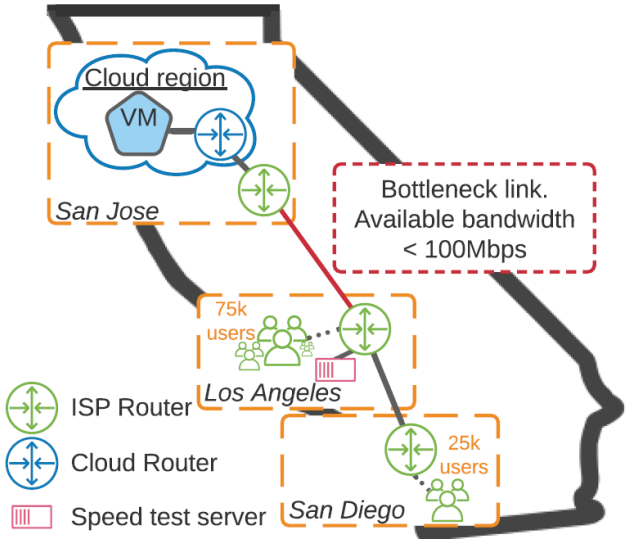

- *CNS Core: Medium: Detection and Analysis of Infrastructure Bottlenecks in a Cloud-Centric Internet*

Thank you

We thank the technical support from the UCSD ITS team.

We acknowledge the support from NSF (CNS-2028506, CNS-1835253, OAC-1724853), DARPA (Cooperative Agreement HR00112020014), Comcast Innovation Fund, and Google Cloud credit grant.

cskpmok@caida.org