Measuring the impact of COVID-19 on cloud network performance

Ricky K. P. Mok and kc claffy
cskpmok@caida.org, kc@caida.org

CAIDA/UC San Diego

1 Motivation

In the crisis of the COVID-19 pandemic, the Internet is facing unprecedented surges of traffic induced by the use of cloud-based applications, such as online shopping, video conferencing, and video streaming. Measuring network performance between cloud platforms and access networks can shed light on the quality of experience (QoE) that end-users are perceiving in cloud-based applications.

End users often use network throughput measurement (or speed measurement) services to understand their performance of their last-mile links. However, these test results are not representative of cloud-based application performance, because web speed test servers are usually located within access networks. To better understand the performance impact of congestion between cloud platforms and access ISPs, we developed the *CLASP* (*CLoud-based Applications Speed measurement Platform*) to conduct throughput measurements *from* speed test servers *to* clients in the cloud. These measurements are representative of a component of video conferencing: sending video (and audio) data from access networks to the cloud. We use a topology-aware approach to select a small set of test servers, such that our speed test measurements traverse interconnections to different networks.

2 Measurement Approach

To create this platform, we deployed 80 VMs in U.S. regions of Amazon AWS, Google Cloud Platform (GCP), and Microsoft Azure. We used CAIDA's `bdrmap` tool [1] to discover interconnections between cloud regions and different networks. We selected 640, 576, and 412 speed test servers in AWS, GCP, Azure, respectively, to conduct hourly download and upload throughput tests. We started these measurements in mid-May 2020, and targeted capture of changes in performance as the U.S. economy began to reopen.

3 Results

Although we expected the capacity between cloud platforms and large access ISP to be well-provisioned, we observed large diurnal throughput variations in our measurement. Figure 1 shows a diurnal congestion example visualized using CLASP's Grafana interface. Observed download throughput from the Cox test server to our client on GCP dropped significantly, from 400 Mbps to less than 50 Mbps between peak hours of 11am to 5pm PT.

Figure 1: Download throughput from Cox test servers in Las Vegas to GCP West 1 region showed diurnal patterns, revealing evidence of network congestion.

To reveal network performance changes correlated with COVID-19 pandemic social distancing policies, we analyzed the data on a macroscopic scale by comparing the median download throughput for the first week of each month. We observed that throughput performance was relatively stable from June to August. However, we found that performance dropped on the paths from speedtest servers in access ISPs to some cloud regions. Figure 2 and 3 show download throughput to clients in {AWS Ohio region/Azure Central 1} from {five servers in Comcast/nine servers in Spectrum}, respectively.

We can see that the throughput generally decreased across the three/four-month time period. In Figure 3, 7 out of 9 servers showed the median throughput dropped more than 80 Mbps from June to Sept.

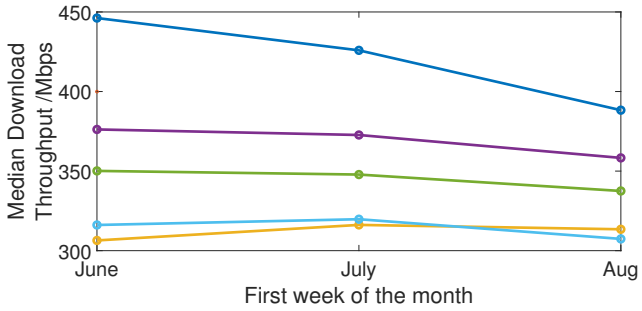

Figure 2: Download throughput measured from test servers in Comcast's network to AWS Ohio region.

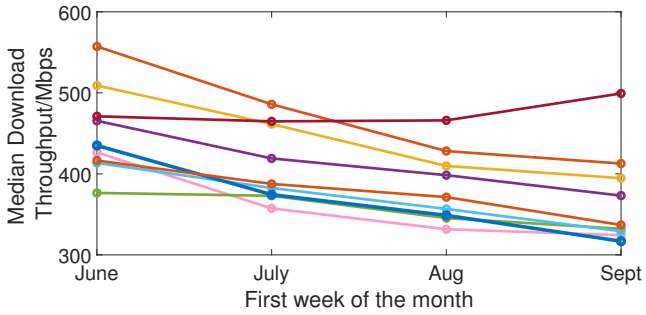

Figure 3: Download throughput measured from test servers in Spectrum's network to Azure Central 1 region, showing a general trend of decreasing throughput over the summer months.

We also observed a downward trend in performance from web speedtest servers located in educational networks. Figure 4 shows the throughput we observed from the three speedtest servers in universities to the GCP East 1 region. Even though most university students were on summer break between June and August, universities employed remote instruction for summer classes, increasing the use of cloud-based applications to support distance learning.

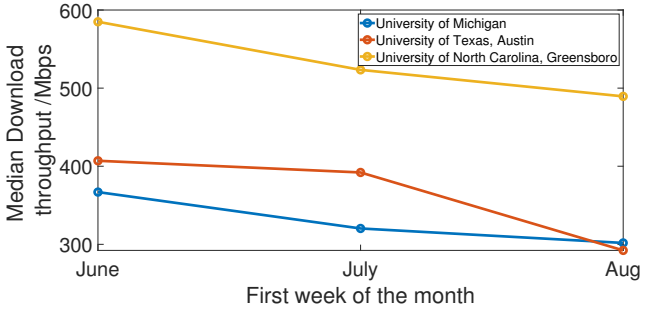

Figure 4: Download throughput measured from three servers in universities to GCP East 1 region.

4 Conclusion

We are pursing this work as part of an NSF RAPID project to gain insight into the Internet's performance in the face of COVID-related demands for Internet services [2]. We are still refining methods as we gather and analyze data. We see evidence of end-to-end throughput degradation across some paths over the last 4 four months, but have not begun efforts to correlate these metrics with possible impairments to QoE. We will continue to gather these measurements at least until the end of 2020.

We will make our measurement data publicly available and will provide fine-grained network performance metrics, such as latency and packet loss rate, to enable better characterization of performance of these critical components of the Internet.

Acknowledgments

Funding for this work was provided by NSF CNS-2028506, NSF OAC-1724853, and DARPA Cooperative Agreement HR00112020014.

References

[1] M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, and k. claffy. bdrmap: Inference of Borders Between IP Networks. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, Nov 2016.

[2] R. Mok and kc claffy. Measuring critical infrastructure for coronavirus-related congestion. `https://www.caida.org/funding/rapid-speedtest/rapid-speedtest_proposal.xml`, 2020.