Performance barriers to cloud services in Africa's public sector: A latency perspective

Josiah Chavula[1], Amreesh Phokeer[2] and Enrico Calandro[3]

[1] University of Cape Town, South Africa
[2] AFRINIC, Mauritius
[3] Research ICT Africa, South Africa

Abstract. Cloud computing allows individuals and organizations to lease storage and computation resources remotely and as needed. For such remote access to computational resources to work efficiently, there is need for reliable and low-delay delivery of traffic. By carrying out month-long Internet measurement campaign, this paper investigates location of cloud-based web hosting and latencies in the public sector of five African countries. Results of the study show that a large percentage of public sector websites are hosted in cloud-based infrastructure that is physically located in America and Europe. Analysis of latencies shows significant differences between local and remotely hosted websites, and where delays are significantly lower for countries that host CDN nodes. The results also indicate high delays for local websites that are accessed circuitously.

Keywords: Latency · Internet Measurements · Cloud Services.

1 Introduction

Cloud computing broadly refers to applications delivered as services over the Internet as well as the hardware and systems software that provide such services [1]. Cloud computing makes possible the provision of large-scale services using remote and shared servers to store and process data. By facilitating ubiquitous and on-demand remote access to externally managed computing resources, the cloud computing helps to eliminate the need for individuals and organisations to own and maintain their own computing infrastructure. This enables organisations that have minimal capital, to deploy their services more rapidly and successfully [2]. By the same token, cloud computing provides an opportunity for developing countries that have limited financial resources, to utilize remote IT infrastructure, such as data centres and applications, to rapidly scale public service provision through electronic means [11]. Despite the perceived opportunities for the public sector, the rate of cloud computing adoption in the public sector has significantly lagged behind that of the private sector. A survey in 2012 [12] found that only 12% of the government departments had spent more than 10% of the total IT resources on cloud services.

Prior studies have suggested that limited availability, accessibility and affordability of the underlying Internet technology is a key hindrance to the adoption of cloud computing developing countries [13, 8]. Apart from the limited access to the Internet infrastructure, concerns related to security and privacy of cloud based information processing have been cited as a critical factors slowing adoption of the cloud. For the South African, survey results by Brenda Scholtz *et al.* [13] suggested that ‘system performance’ and ‘privacy of data’ were the biggest technical concerns that could hamper implementation of cloud computing in the public organisations. Similarly, Gillwald *et al.* (2014) [8] identified limited and costly broadband as major barriers to adoption of cloud computing services in South Africa, Tunisia, Nigeria, Ghana, and Kenya.

The aim of this study was to quantify Internet latency as a barrier to the adoption of cloud computing services by the public sector. This will highlight the quality of service for accessing content hosted in various networks and geographical locations, including cloud-based infrastructure. The rationale is that the success of cloud-based services for the public sector will depend on the ability to achieve high performance connectivity to the cloud. This study therefore provides insight into the readiness of the public sector in the selected countries to interact with cloud-based services. After characterising the network performance the study will also explore the causality behind any observed suboptimal performance (e.g. high delays), and the possible cloud-computing readiness steps that need to be undertaken. In particular, the study seeks to characterise the level of when accessing public sectors online public resources. Performance data was collected via Internet measurement campaigns targeting public sector websites sampled from five countries: Nigeria, Ghana, Kenya, Zambia and South Africa.

The three key questions addressed in this paper are as follows:

1. In which networks and countries are the public sector websites currently being hosted. This question is meant to provide some insight on the current hosting practices in the public sector, particularly regarding the extent to which websites are hosted locally within a country, or remotely. This will also shed light on the locality of networks that are dominating the public sector hosting market in Africa.
2. What are the characteristics of country specific Internet latencies to web servers of the public sector? This question is meant to provide country specific delay characteristics and comparison for public sector websites hosted in various networks and countries.
3. What is the extent of circuitous routing when accessing locally hosted websites, and what impact does this circuitous routing have on the Internet delays?

2 Related Work

Past studies have highlighted general performance challenges in Africas Internet, including high latencies and generally poor quality of experience. A recent study by Formoso *et al.* (2018) showed that many parts of Africa still experience excessively high Internet delays, often exceeding 300ms. A study by Fanou *et al.* (2016) [7] also showed that Internet performance on the continent is largely characterised by slow download speeds and high delays. Other studies have highlighted traffic engineering problems in Africas Internet topology, showing a lack of direct interconnection amongst Africas ISPs, resulting in suboptimal performance for intra-country and cross-border communication, a well as high cost of Internet access [4, 3, 5, 9]. Apart from the lack of direct network-level and country-level interconnections, other studies have looked at inefficient DNS configurations, as well as lack of local content caching servers across the African continent [6, 10, 15].

3 Measurement Methodology

This section describes the process for conducting measurements to gauge readiness of the public sector to utilize cloud-based resources. The first step involves identifying the appropriate set of performance metrics that can reveal the readiness of the public sector to utilize cloud-based services. This is followed by the process of selecting the appropriate tools and platform for conducting the measurements. After selecting the measurement platform, it is necessary to select the locations from which to observe performance (i.e. measurement vantage points), as well as appropriate measurement targets.

3.1 Performance Metrics

The publics high concern for system performance as reported in the study of Brenda Scholtz *et al.* [13] highlights the perceptions that Internet infrastructure in many African countries is not developed and robust enough to provide reliable and high performance access to computing resources and information stored in cloud. One aspect of this relates to high end-to-end delays (latencies) that impact responsiveness and quality of experience (QoE) for online interactions. Latency is an important metric for cloud computing as it gives insight into responsiveness of interactions between cloud servers and Internet clients.

3.2 Measurement Platform and Vantage Points

A number of Internet measurement platforms have deployed thousands of probes in access and backbone networks, as well as behind residential gateways, globally. Researchers are able to make use of these platforms to conduct Internet measurement campaigns, using the specialized network devices (probes) as vantage points from which to launch tests towards specified targets. Recent measurement platforms use dedicated hardware-based probes, and these probes are used to run continuous measurements with minimal end-user participation. When selecting a distributed Internet measurement platform for a large scale campaign, it is important to consider the number and distribution of vantage points in networks that are to be measured [14]. Shavitt *et al.* (2011) showed that extensive topology sampling from a broad and distributed vantage points is required for obtaining an unbiased and accurate topology characteristics. A more widely deployed hardware-based measurement platform is the *Ripe Atlas* [4], which consists of thousands of probes that perform active measurements. As of December 2017, RIPE Atlas had around 230 active probes in 36 African countries. On this basis, this study made use of Ripe Atlas for topology characterization, measuring latency to public sector websites for five African countries. During each measurement episode, 10 Ripe Atlas probes were randomly selected as vantage points for each of the countries.

3.3 Selecting Measurement Targets

The next step in the study was to identify and select prominent websites from the public sector in each of the target countries. This was obtained through the AlexaTop[5], a website that ranks websites based on a combined measure of page views and unique site users. For each country, a filter was applied for Government websites to list the prominent public public sector websites, such as those for government departments and parastatal organisations. For example, 'category/Top/Regional/Africa/Kenya/Government' would list the most popular government related websites in Kenya. In some cases, expert local knowledge was sought to determine prominent public sector in each of the five countries. In total, 86 websites were identified as measurement targets across; 10 in Kenya, 9 in South Africa, 10 in Zambia, and 48 in Nigeria.

3.4 Launching Measurements

Traceroute measurements were then launched from each of the selected Atlas probes, towards each the country's selected public sector websites. The actual measurement was performed towards the IP address of web server for each website. To obtain IP addresses, a DNS lookup was performed for each website domain. In order to mitigate the effects of location based load balancing, where requests for a domain are directed to different web servers based on location of Internet clients, the DNS lookups/resolution were performed from the vantage points (i.e. from each Atlas probe).

For each measurement episode, four Traceroutes were launched successively from each Atlas probe to all of the countrys selected websites, and this was repeated 4 random times a day for one month in December 2017. While not all Traceroute measurements could reach the final destination, a Traceroute measurement was considered successful it is was able to reach the hosting network, i.e. if the last hop in the Traceroute was inside the hosting companys network. In the end, a total of 13790 Traceroute measurements were successfully completed; 2570 in Ghana, 3073 in Kenya, 2790 in Nigerian, 3341 in South Africa, and 2016 Zambia.

[4] https://atlas.ripe.net/

[5] https://www.alexa.com/topsites

4 Results

This section describes results of a measurement study on public sector websites carried out in the five selected countries. The results describe the remote locations and networks where the websites are hosted, the nature of routes for accessing the public sector websites in the respective countries, as well as the performance (delay) observed for each situation. The initial step in the data analysis involved attaching network and geolocation information to each target IP address (obtained from the websites' DNS lookups performed from the vantage point), as well as every router hop in the traceroute data. The RIPE Routing Information Service (RIS) and the MaxMind GeoLite2-City database [6] are used to obtain the Autonomous System Number (ASN) and the geographical location (country) of each IP address. While it is well known that geolocation databases do contain inaccuracies, the analysis in this study is restricted to country-level geolocation which has relatively much higher accuracy.

4.1 Geolocation of Web Hosting for the Public sector

One of the key goals of this study was to examine the ASN and geographical distribution of web hosting servers used by the public sector in Africa and to evaluate the performance implication of such a distribution. This should also reveal the strengths and weaknesses of the web hosting environment in the selected African countries. The first step in the analysis was therefore to compute the geographical and network distribution of public sector websites per country.

The first observation from the analysis was that, on overage, 66 percent of the sampled public sector websites were hosted outside their respective countries, i.e. remote hosting. It was also noted that the level of remote hosting varies widely among the countries, ranging between 4 and 82 percentage. Figure 1 shows the percentages of remote and locally hosted public sector websites in each of the five countries. Of the five, South Africa had the lowest percentage of remote hosting at 4 percent, whereas Nigeria and Ghana had the highest remote hosting of 61 and 82 percent, respectively.

From Figure 2, it can be seen that most of the remote website hosting is situated in USA, UK, Germany, Canada and Ireland. Figure 2a shows that Nigeria, for example, had 75% of the websites hosted in US based companies, mostly through New Dream Network (14.8%), RackSpace (9%), GoDaddy (8%), and Unified Layer (7%). For Ghana, the remotely hosted websites are mostly in the US (26.5%), Germany (25%) and Canada 10%). Kenya and Zambia have a higher percentage of public sector websites locally hosted. Interestingly for the both, local hosting of public sector websites appears to be supported by their respective National Research and Education Networks (NRENs); in Kenya, 37.9% of the websites were hosted in KENET (Figure 2c), whereas in Zambia, 25.3% of the websites were inside ZAMREN (Figure 2d).

[6] https://www.maxmind.com/en/geoip2-city

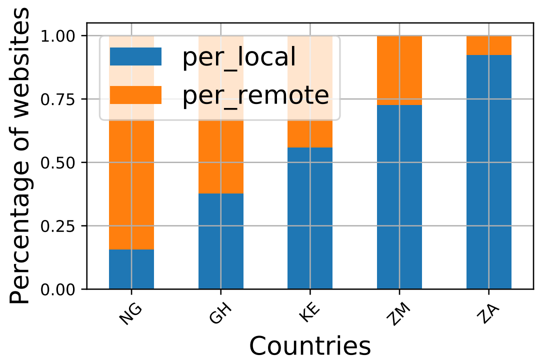

Fig. 1: Percentage of local and remote hosting of public sector websites per country.

South Africa had only one of the sampled websites hosted outside the country, in Ireland (Figure 2e). The remote website was hosted in the Amazon cloud infrastructure, using their data centre in Ireland. The result for South Africa was as expected, particularly considering its significantly more developed web-hosting infrastructure compared to other African countries. It also needs to be noted that the much robust Internet infrastructure in South Africa has also attracted many more foreign and cloud-based hosting companies, the result of which is that, while these companies appear to have presence in the country, some of the content hosted by such networks is physically located in remote data centres. For example, while Amazon Web Services (AWS) has become popular in the South African market, the company does not operate a data centre within the country, meaning content hosted within the AWS infrastructure is remotely hosted.

4.2 Latency to public sector websites

It is important to highlight that choice of hosting provider, and where it is located, can have significant consequences on the level of delay experienced by Internet clients. The websites hosted in more remote places generally experience higher delays. While many of the dominant hosting providers make use of cloud-based infrastructure, and therefore claim to have global presence, the absence of data centres and CDN (Content Delivery Networks) nodes in most of African countries means that hosting in such networks results in high overall delays. The table below shows the range of delays for websites hosted in different networks. Many of the local hosting networks can be seen to have delays that are less than 100ms, whereas the higher delays are mostly for websites hosted by the large international operators. The high delays are particularly prevalent in situations where the remote hosting is not supported by CDN infrastructure that would

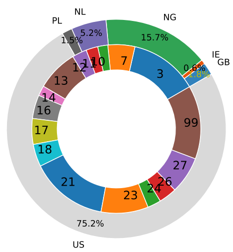

3 - NG/Galaxy Backbone PLC (13.1 %)
7 - NL/Microsoft Corporation (5.2 %)
13 - US/GoDaddy.com (8.0 %)
21 - US/New Dream Network (14.8 %)
23 - US/Rackspace Ltd. (9.0 %)
27 - US/Unified Layer (7.1 %)

(a) Nigeria

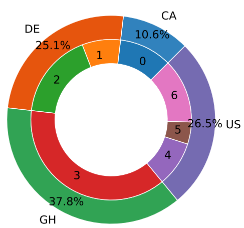

0 - CA/OVH SAS (10.6 %)
1 - DE/ 1&1 Internet SE (7.7 %)
2 - DE/Hetzner Online GmbH (17.3 %)
3 - GH/GGoC1-AS (37.8 %)
4 - US/Codero (9.1 %)
6 - US/Rockynet.com (13.0 %)

(b) Ghana

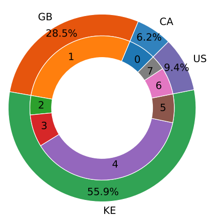

0 - CA/UptimeArchive (6.2 %)
1 - GB/KENYAWEB (28.5 %)
2 - KE/ACCESSKENYA (4.3 %)
3 - KE/JAMBONET (7.2 %)
4 - KE/KENET (37.9 %)
5 - KE/SIMBANET-AS (6.4 %)
6 - US/Cloudflare Inc (6.1 %)

(c) Kenya

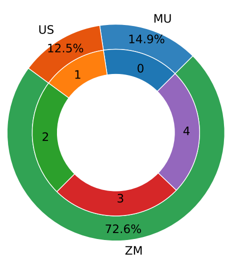

0 - MU/Liquid Telecom Ltd (14.9 %)
1 - US/CyrusOne LLC (12.5 %)
2 - ZM/ZAMNET-AS (22.7 %)
3 - ZM/ZAMREN (25.3 %)
4 - ZM/ZAMTEL (24.7 %)

(d) Zambia

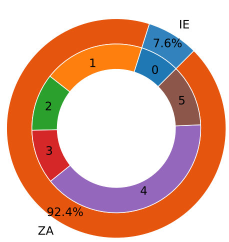

0 - IE/Amazon.com (7.6 %)
1 - ZA/IS (19.3 %)
2 - ZA/Neotel Pty Ltd (11.0 %)
3 - ZA/SITA-AS (10.4 %)
4 - ZA/Telkom-Internet (39.9 %)
5 - ZA/Vodacom-VB (11.8 %)

(e) South Africa

Fig. 2: Pie-charts showing hosting countries and networks most used by each vantage country

otherwise push the content closer to where the intended audience is, i.e. closer to the countries owning the websites.

Figure 3 shows that hosting companies with average delays of over 200ms are mostly based in the US, including *Microsoft Corporation, LunarPages, RackSpace, GoDaddy,* and *New Dream Network*. On the other hand, *CloudFlare*, which is also US-based, runs a number of data centers in Africa, including in South Africa and Kenya, and this is reflected in the lower median delay of 110ms when their content is accessed by Internet clients in Africa.

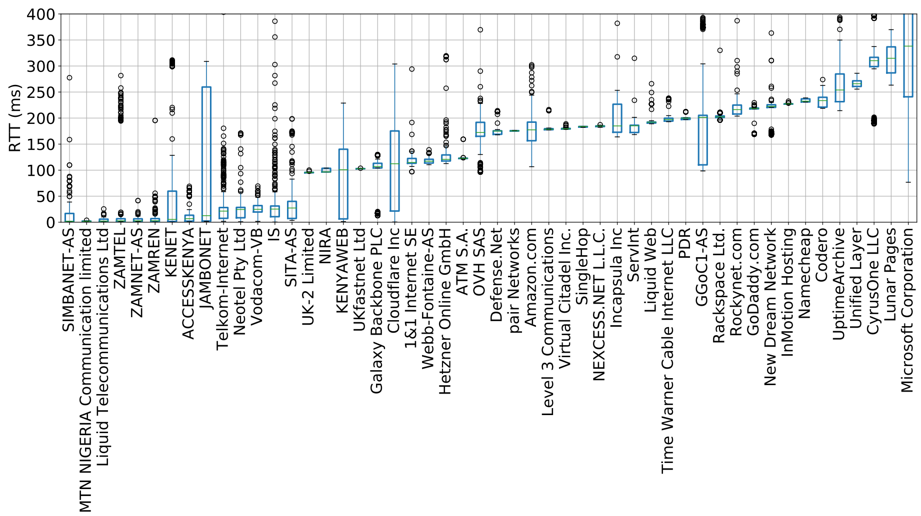

Fig. 3: Distribution of round-trip-times for websites hosted in different networks.

Figure 4a below presents a summary of RTTs to the sampled public sector websites as measured from each of the five countries. As is observed in Figure 4a, Ghana and Nigeria, the two countries with the highest remote hosting percentages in the sample, also had the highest median RTTs of 199ms and 177ms respectively. In contrast, Kenya, which had an almost 50-50 split between local and remote hosting, had a lower mean RTT of 50ms. The best lowest mean RTT of 3ms was observed in Zambia, which also had much lower remote hosting of 25%.

4.3 Impact of Hosting Locations

It is worth analysing the extent to which different host hosting countries result in different levels of delay for the selected African countries. Figure 4b shows the delays to websites hosted in different countries, and illustrates how websites hosted remotely, such as in USA, Canada and Netherlands experience significantly higher delays. As would be expected, the lowest delays were within each of the vantage countries. The only exception was Ghana, whose special case circuitous routing discussed later. In general, the further away a websites hosting country is, the higher the delays. It can be observed that the highest delays are for websites hosted in Canada and the USA (median RTT of 230ms and 220ms respectively), the two countries are geographically the furthest of the hosting countries from any of the five vantage countries in this study. The cumulative distribution in the left plot of Figure 5 shows that about 50% of the delay samples in Ghana and Nigeria were above 200ms. In comparison, only 20% in Kenya, 19% in Zambia, and less than 1% of samples in South Africa were above 200ms. About 10% of the samples in all countries, except South Africa were above 300ms. In terms of hosting countries, the right CDF in Figure 5 shows that the higher delays are more prevalent for websites hosted in USA, Canada, and Netherlands. About 50% of the delay samples to these countries are above 250ms.

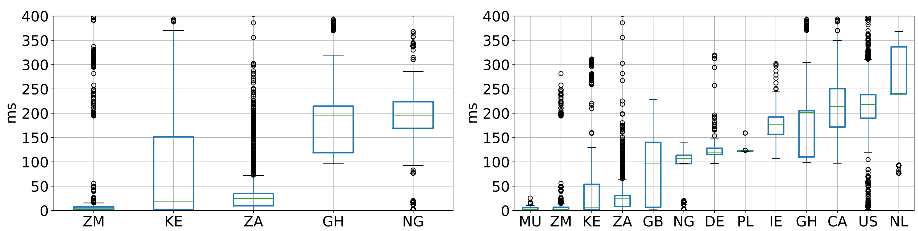

(a) (b)

Fig. 4: Figure 4a showing distribution of delays to per vantage country; and Figure 4b showing distribution of delays per hosting country

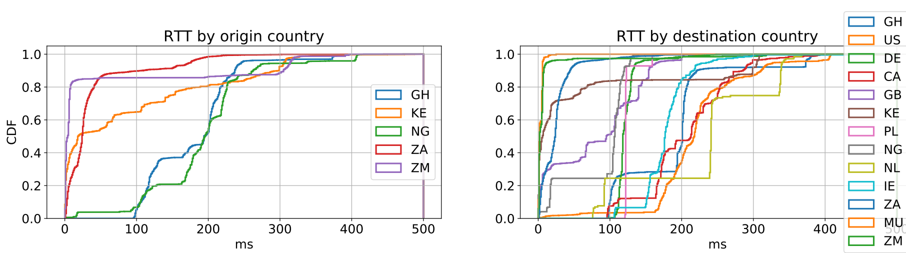

Fig. 5: Cumulative distribution of latency values for each vantage country (left), as well as for each hosting country (right)

Different countries experience different delays to websites hosted in the remote locations. Figure 6 presents delays from the vantage countries to websites hosting countries. The differences can be attributed to the differences in geographical distances, as well as varying logical topologies. Among the countries hosting in the USA, it can be seen that Zambia experiences the highest median delay of 314ms. This should be expected given Zambias geographical distance from the USA, compared to Ghana, Nigerian and Kenya, which experience median delays to the USA of 221ms, 200ms, and 174ms respectively. In the case of South Africa, the only remote hosted website was in the Republic of Ireland, where *AWS* has a data centre. In terms of performance, the median delay between South Africa and the Republic of Ireland is observed to be 180ms.

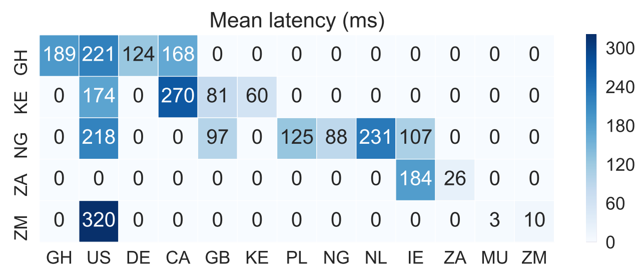

Fig. 6: Matrix of mean latencies from the vantage countries (x-axis) to website hosting countries (y-axis)

4.4 Impact of Circuitous Routing

It is generally the case that local hosting provides lower delay compared to remote hosting, and the dataset from Zambia and Kenya exhibit this expectation, with local delays of 2ms and 9ms respectively. Ghana and Nigeria, on the other hand, go against this trend. In the case of Ghana, there was a median delay for locally hosted websites of 201ms, while the countrys websites hosted in Great Britain had a lower median delay 118ms. Similarly, while the local median delay for Nigeria was observed to be 107ms, a slightly lower median delay of 97ms was observed for websites hosted in Great Britain.

To evaluate this phenomenon, another aspect of this study was to look at the extent and impact of circuitous routing when accessing public sector websites. In this context, circuitous routing is when a website that is locally hosted is accessed by local Internet clients through paths that traverse other countries. Overall, 23% of the websites were accessed circuitously, but this is more prevalent some countries than others. For example, the Ghana dataset had the highest percentage of circuitous routes at about 33%, while Kenya was at 16% and Nigeria at 11%. Figure 7 shows a distribution of RTTs for the three categories of routes. The general distribution does show that locally hosted websites that are accessed through circuitous routes, ie routes that leave and come back to the vantage country, experience higher delays than websites that are remotely hosted.

Circuitous access of locally hosted websites, and the result high latencies are symptomatic of lack local peering of networks within a country. With locally hosted websites appearing to perform worse than remotely hosted websites, the lack of peering has the potential to not only discourage local hosting, but also inhibit success local content initiatives.

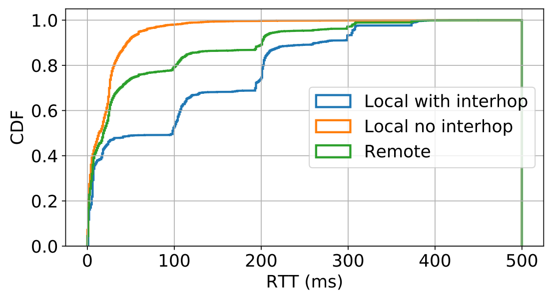

Fig. 7: Cumulative distribution of latencies, indicating differences between local and remote hosted websites, and also for local websites that are accessed via circuitous routes.

5 Conclusion

This study has shown that in some African countries, access to public sector websites is largely characterised by high Internet delays. In addition, a large proportion of public sector websites are currently being hosted remotely, i.e. in other countries. It was also observed that of the five countries surveyed, Nigeria and Ghana had the highest percentage of remote hosting and experienced the highest latencies. A large portion of the remote websites were hosted in USA, Canada, Germany, UK, and South Africa. While the hosting networks generally have global operations, they do not necessarily have physical infrastructure in most of the African countries. This means that while offering the convenience of cloud-based hosting, lack of physical infrastructure in Africa entails that Africa's web content gets to be stored in remote locations. Remote storage of web content has a negative implication on the sovereignty of African countries in that they loose control of their data. In addition, the burden of fetching content from remote locations falls on local network operators, the cost of which get passed on the users. This ultimately has negative implications on the local economies, and also in terms of poor quality of experience due to high latencies as reported in this study.

To reduce these latencies, and to help improve performance of cloud-services in Africa, there is need for leading cloud infrastructure providers to deploy infrastructure in Africa. The advantage of local deployments is demonstrated in the results from Kenya, which although had a relatively high percentage of remote hosting (42%), appeared to have a much lower median web delay of 50ms. In comparison, Nigeria and Ghana had mean latencies of almost 200ms. However, for African countries to fully take advantage of cloud infrastructure that is domiciled on the continent, there is need for better peering and interconnectivity at national and continental level.

Bibliography

[1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, et al. A view of cloud computing. *Communications of the ACM*, 53(4):50–58, 2010.

[2] S. J. Berman, L. Kesterson-Townes, A. Marshall, and R. Srivathsa. How cloud computing enables process and business model innovation. *Strategy & Leadership*, 40(4):27–35, 2012.

[3] J. Chavula, N. Feamster, A. Bagula, and H. Suleman. Quantifying the effects of circuitous routes on the latency of intra-africa internet traffic: A study of research and education networks. In *International Conference on e-Infrastructure and e-Services for Developing Countries*, pages 64–73. Springer, 2014.

[4] J. Chavula, A. Phokeer, A. Formoso, and N. Feamster. Insight into africa's country-level latencies. In *AFRICON, 2017 IEEE*, pages 938–944. IEEE, 2017.

[5] R. Fanou, P. Francois, and E. Aben. On the diversity of interdomain routing in africa. In *International Conference on Passive and Active Network Measurement*, pages 41–54. Springer, 2015.

[6] R. Fanou, G. Tyson, P. Francois, and A. Sathiaseelan. Pushing the frontier: Exploring the african web ecosystem. In *Proceedings of the 25th International Conference on World Wide Web*, pages 435–445. International World Wide Web Conferences Steering Committee, 2016.

[7] A. Formoso, J. Chavula, A. Phokeer, A. Sathiaseelan, and G. Tyson. Dissecting the african internet: An intra-continental study.

[8] A. Gillwald and M. Moyo. The cloud over africa. *Research ICT Africa*, 2014.

[9] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett. Peering at the internets frontier: A first look at isp interconnectivity in africa. In *International Conference on Passive and Active Network Measurement*, pages 204–213. Springer, 2014.

[10] M. Kende and B. Quast. Promoting content in africa. *ISOC Report*, 2016.

[11] N. Kshetri. Cloud computing in developing economies. *Computer*, 43(10):47–55, 2010.

[12] B. L. Sahu and R. Tiwari. A comprehensive study on cloud computing. *International journal of Advanced Research in Computer science and Software engineering*, 2(9):33–37, 2012.

[13] B. Scholtz, J. Govender, and J. M. Gomez. Technical and environmental factors affecting cloud computing adoption in the south african public sector. In *CONF-IRM*, page 16, 2016.

[14] Y. Shavitt and U. Weinsberg. Quantifying the importance of vantage point distribution in internet topology mapping (extended version). *IEEE Journal on Selected Areas in Communications*, 29(9):1837–1847, 2011.

[15] Y. Zaki, J. Chen, T. Pötsch, T. Ahmad, and L. Subramanian. Dissecting web latency in ghana. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 241–248. ACM, 2014.