An Evaluation of Peering and Traffic Engineering in the Pan-African Research and Education Network: A case for Software Defined Internet Exchange Points

Josiah CHAVULA[1] and Clifford C.L SIBANDA[2]

[1]*Mzuzu University, Private Bag 201, Mzuzu, Malawi*

Mobile: +27847083172 Email: josiahchavula@gmail.com

[2]*University of Science & Technology, Box AC939, Ascot, Bulawayo, Zimbabwe*

Email: clifford.sibanda@nust.ac.zw

Abstract

Despite an increase in the number of Internet eXchange Points (IXPs) in Africa, as well as proliferation of submarine and terrestrial fibre optic cable systems, a large proportion of Internet traffic exchanged among Africa's Internet users is exchanged through higher tier transit providers at IXPs in other continents. This work attempts to quantify this problem and its impact on intra-Africa Internet performance, especially with regard to end-to-end latency. The recent developments in technology such as the Software Defined Network offer Africa and the developing world a chance to implement optimal traffic engineering solutions for NRENs at lower deployment and maintenance costs. This paper makes a contribution in two ways; first through active topology measurements, it provides an assessment of traffic routing and its impact on latency. Secondly, through simulation, it shows performance benefits of implementing an African Internet exchange, and discusses possible software defined mechanisms that could be used to manage traffic in such a setting.

Keywords

NREN, e-Infrastructure, Software Defined Networking, Latency, Traffic Engineering, Internet Exchange Points

1. Background

National Research and Education Networks (NRENs) are vital for research collaboration and e-resource sharing, as many education and collaboration oriented applications, such as video conferencing, remote lecturing as well as simulations for scientific study have quality of service (QoS) requirement that may not easily be met with the commodity Internet (Ahmad & Guha 2012). Many universities are making efforts to provide E-learning platforms, availing the teaching and learning material beyond the classroom time, and beyond campus boundaries. Other universities are beginning to run in completely virtual environments, allowing students in different geographical locations (cross-border) to remotely attend and participate in live lectures. In Europe for example, the concept has been extended to allow students to create their own curriculum across different universities even in different countries. A virtual university form allows the students to pick modules based on known university strengths in particular fields or subjects. Also, universities nowadays aim to provide virtual libraries and online digital repositories, and for collaborating universities, they aim to make such libraries available to students that are located in other distant institutions. NRENs have therefore been conceived to provide specialized internetworks, dedicated to linking research and education institutions for efficient exchange of data. The main goal therefore is to meet the quality of service requirements of educational and research applications through provision of dedicated network backbones to interconnect education and research institutions, as well as through collaborative bandwidth management and peering agreements (Andronico et al 2011).

However, with limited funding, most pan-African NRENs are formed with very little or no equipment and infrastructure of their own to route or switch Internet traffic. As such, many NREN exist just as entities, with a block of IP addresses that are identified to belong to a single Autonomous System (AS). Currently in the Southern African region, the largest and mostly perceived to be operational NRENs are the South African Tertiary Education Network (TENET) operated by South African National Research Network (SaNRen) and Kenya's Education Network (KENET). However, these NRENs are very important as there are already impeccable benefits even for emerging NRENs, such as through concerted lobbying and bandwidth sharing. For example, newer NRENs, such as the Zambian NREN (ZamRen), while hardly with the infrastructure to fully interconnect its research and educational institutions, already benefit from the excess international bandwidth available for the South Africa's TENET.

Research collaboration and e-resource sharing in sub-Sahara African continues to be hampered by the limited interconnectivity which is not only expensive, but fails to meet the quality of service required for important collaboration applications among the NRENs. Many universities still obtain their Internet connectivity from commercial Internet service providers (ISPs) that have no peering relationships amongst themselves (Barry et al 2010; Steiner et al 2005). As a result, the level of direct interconnectivity and peering among NRENs remains low, and the traffic that is exchanged among NRENs in sub-Sahara Africa continues to be exchanged at long distance Internet eXchange Points (IXP) in Europe and North America, resulting in high data transmission costs (Gilmore et al 2007) and sub-optimal performance (e.g high latency). Recent work on the African internet topology (Gupta et al 2014) has shown that about 66% of traffic between South African Internet users and Google cache servers located in Africa is routed outside the continent. The same work also characterized the IXP peering situation in Africa and showed that most African ISPs do not peer among themselves at national or regional IXPs, but rather prefer to peer at larger European IXPs such as London and Amsterdam. The problem is exacerbated by the lack of traffic engineering strategies that could optimize the utilization of the limited communication links and bandwidth.

According to data available in PeeringDB (PeeringDB, 2014. [Online; accessed 13-May-2014]), an open database of peering relationships, most of the IXPs in Africa have no more than 2 participants, suggesting very low levels of peering. The exception is South Africa, which with 5 IXPs and an average of 30 networks per IXP, has the highest level of national peering. In comparison, a study by Ager et al. (Ager et al 2012) has shown that a single European IXP ecosystem has over 400 networks, with over 50 thousand actively used peering links, exceeding the total estimated number of all non-IXP peering links in the entire Internet. The fact that virtually all the major IXPs are located in Europe and America is particularly disadvantageous for southern Africa, as it is geographically the furthest from the Internet Exchanges. The fibre optic cable that runs from southern Africa tip in Cape Town to London, is close to 15000 km, implying that traffic that originates in southern Africa destined for southern Africa, exchanged through London, covers a distance of about 30,000km. For acknowledgement-based TCP/IP traffic, this translates to a round-trip of at least 60,000km, and based on speed of light in fibre optic cables, a minimum RTT of about 300ms.

Latency is an important metric for measuring the performance of Internet, as it affects the performance and responsiveness of Internet applications especially real-time and interactive applications (Landa et al 2013). Latency is measured as round trip time (RTT) for a data packet to move from source to the destination, and for the acknowledgement packet to be received by the sender. With high latency, it is difficult for research communities to make use of Internet-based collaborative tools such as video conferencing or remote sharing of virtualized computer resources such as computer processors. This paper attempts to quantify the performance cost in terms of latency, for sub-Sahara NRENs to exchange traffic through IXPs that located in Europe. Furthermore, this paper discusses possibilities for improvement of the traffic engineering environment through the use of Software Defined Networking (SDN).

2. Network measurements

Active measurement techniques attempt to exploit network management to solicit responses from a set of network destinations, and then use such responses to infer topological characteristics such as route paths, RTTs and packet loss. Common active measurement techniques make use of Traceroute, a tool for discovering IP paths between a host and some destination. The tool works by sending IP packets with increasing time-to-live (TTL) values, in such a way that packets continually expire on their way and cause routers to respond with ICMP time-exceeded messages.

To assess performance of traffic exchange among the African Research and Education Networks, a two (2) week experiment was conducted to probe IP paths to 35 universities across 12 countries. The target universities are within the UbuntuNet Alliance area. Using Scamper (Luckie, 2010) network tool, Internet probes were performed every day to each address for 14 days, using ICMP-based Ping and Traceroute network probes from two main vantage points, in South Africa (SA) and Malawi. Further probes were conducted from vantage points in West Africa (Senegal), and East Africa (Rwanda).

2.1 Geographic location of IP paths

Using IP geolocation database, Geoip2 (Online: MaxMind, “Geoip”2010), the IP hops obtained from the traceroute were mapped to their corresponding geographic locations. Traceroute results show that most of the inter-NREN traffic in UbuntuNet Alliance is routed at the Amsterdam and London exchange points, with other prominent hops being in Lisbon Portugal, Marseille France. SanRen traffic is exchanged through a local Internet exchange point, the Johannesburg Internet Exchange (JINX). Furthermore, there also appears to be direct logic links between JINX and SA's neighbouring countries such as Mozambique, Zambia, and Zimbabwe.

On average, 75% of the traces from African vantage points to African NRENs traversed inter-continental links through PoPs in Europe, such as Amsterdam, London, Lisbon and Marseille.

However, depending on the geographical location of the vantage points, different levels of inter-continental traffic is observed. For example, the vantage points along the north west coast of Africa used inter-continental links for as much as 95% of the traces, whereas vantage points in Central and Southern Africa had a relatively lower usage of inter-continental links. The South African vantage point had only 60% of the traffic using inter-continental links.

The lower usage of inter-continental links by the South African vantage point can be attributed to the direct logic links observed between South Africa and some of its neighbouring countries such as Mozambique, Zambia, and Zimbabwe, as well as links to East Africa via the EASSY submarine fibre optic cable.

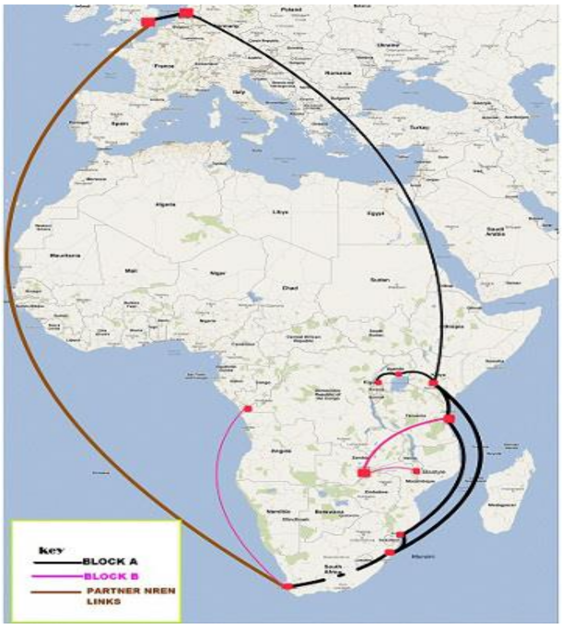

Figure 1: Adapted from Africa Connect Launch Document

2.2 Round-trip Times (RTT)

From the South African vantage point, ping results show an average RTT of 250msec to other sub-Sahara NRENs. Within the SanRen, a lower average RTT of 20msec is obtained, and this lower RTT can be attributed to the fibre SanRen interconnection through TENET, and the routing of traffic through the Johannesburg IXP (JINX). Also achieving low RTTs from South Africa are universities in the neighbouring countries where there are direct fibre links, such as Mozambique (~45msec), Zambia (~55msec) and Namibia (~80msec).

Interestingly, some universities within the same country achieve remarkably different RTTs depending on how their traffic is routed. A good example is Kenya, as depicted in Figure 2, where one university has its traffic from Johannesburg circuitously routed through Amsterdam, back to South Africa (Cape Town), before being forwarded to Kenya, and achieves an RTT of about 400ms. In comparison, another 2university in the same country of Kenya has a direct logical link from Johannesburg to Kenya and achieves an RTT of only 80ms. The highest average RTT from the South African vantage point to the target universities is 954 ms to a university in Malawi that is connected to the Internet through a satellite link to Norway. Results from Malawi vantage point show a similar trend, with a higher average RTT of about 380msec to other sub-Saharan NRENs.

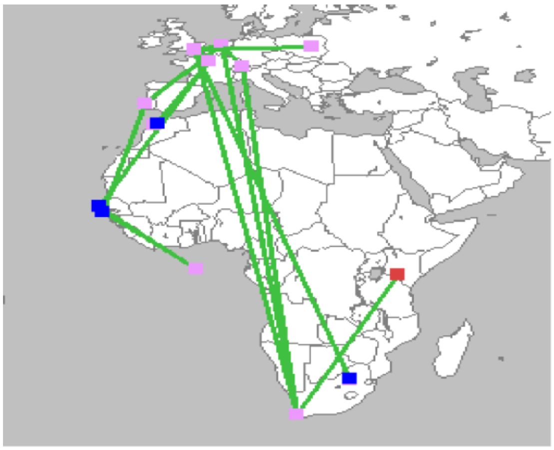

Figure 2: Traceroute from Johannesburg to Kenya showing traffic routed through Amsterdam and Cape Town

2.3 Simulation

This section presents a simulation that depicts the pan-African NREN connectivity with two scenarios, first as is the case at the moment, with major interconnection in Europe, and secondly as it would be with an Africa based Internet exchange. The links in the simulation are configured with administrative distances commensurate with actual geographical distances between the sites. In the first scenario, traffic from African NREN transits through Europe even if it is destined for networks that are within Africa. This scenario is shown in Figure 3. Results from this scenario indicate more transit hops between a source and destination, with a higher average RTT. The extracted results of traceroute and ping are shown in Figure 5.

The second scenario, depicted in Figure 4, has an African Internet Exchange (AIX) entity central to the major Regional Education Networks (RENs), ASREN, UbuntuNet and WACREN. In this scenario, regional RENs cater for their member NRENs through their routers connected to the AIX. These regional RENs could be located in Johannesburg, Lusaka, Nairobi and or Lagos. The AIX is envisaged to be an entity resident centrally to the RENs and solely functions to switch traffic between the connected entities. The scenario can be compared to the launch map available from Africa Connect project which is shown in Figure 1. One can clearly map out connection points in East, West and Southern Africa to the European Internet Exchanges.

The scenario where African traffic does not leave the African network but instead is exchanged through the AIX, suggests major improvement in the traffic flow within Africa. End-to-end latency is reduced by half on average.

 ISSN 2223-7062 Proceedings and report of the 7th UbuntuNet Alliance annual conference, 2014, pp 59-70

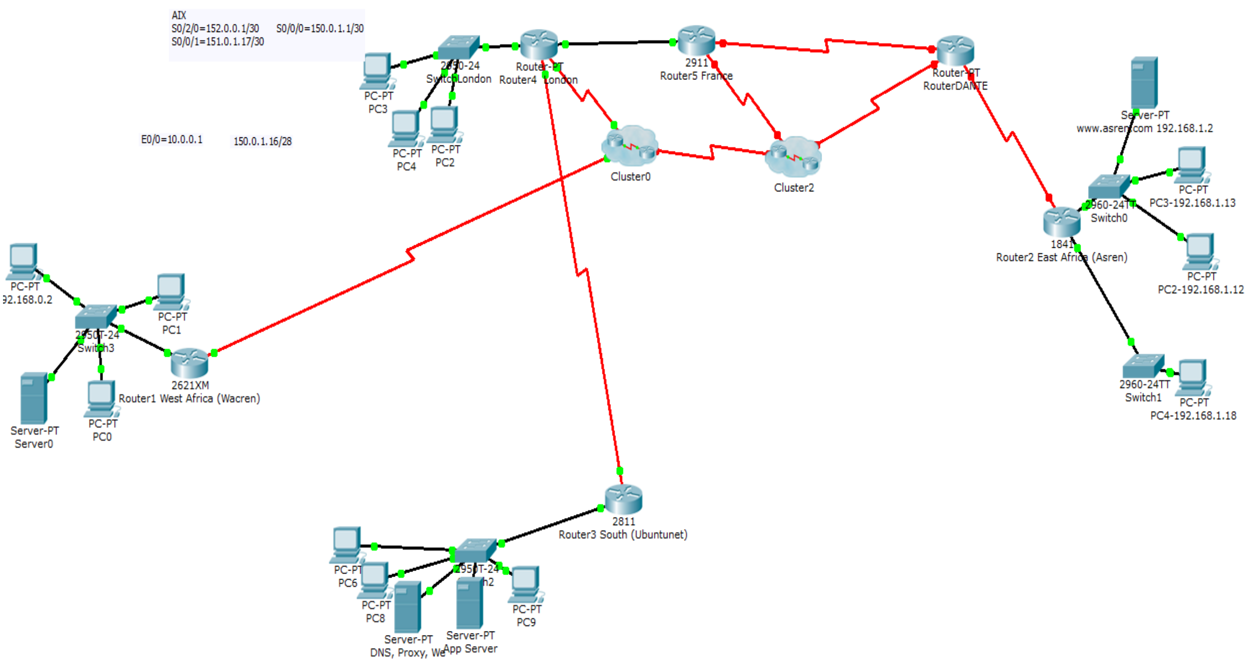

Figure 3 : Current scenario simulation topology

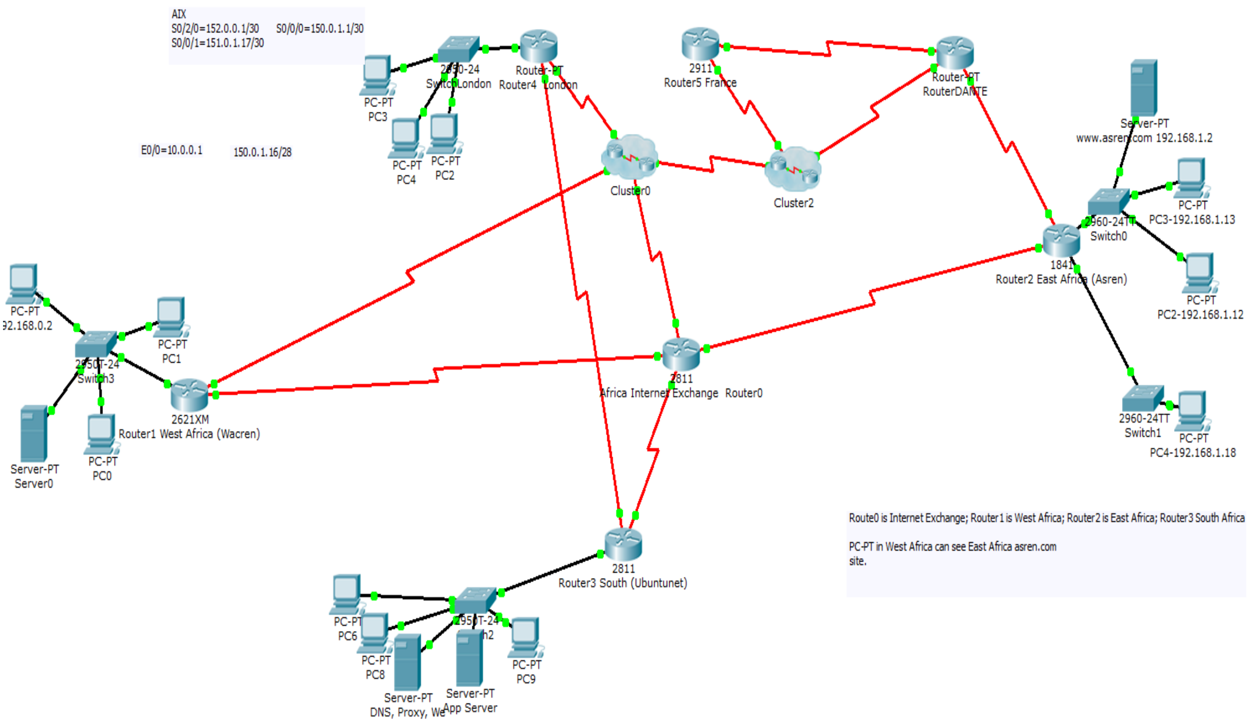

Figure 4: Proposed/Comparative scenario with AIX and peered traffic between African RENs

The number of hops and average RTT for the traceroute packets are obtained using a traceroute tool. In the scenario with a European interconnection, simulation results shows the number of hops and latencies are about double those of the proposed scenario. Using ping tool we get similar results with RTT min/avg/max values being approximately double those of the AIX scenario.

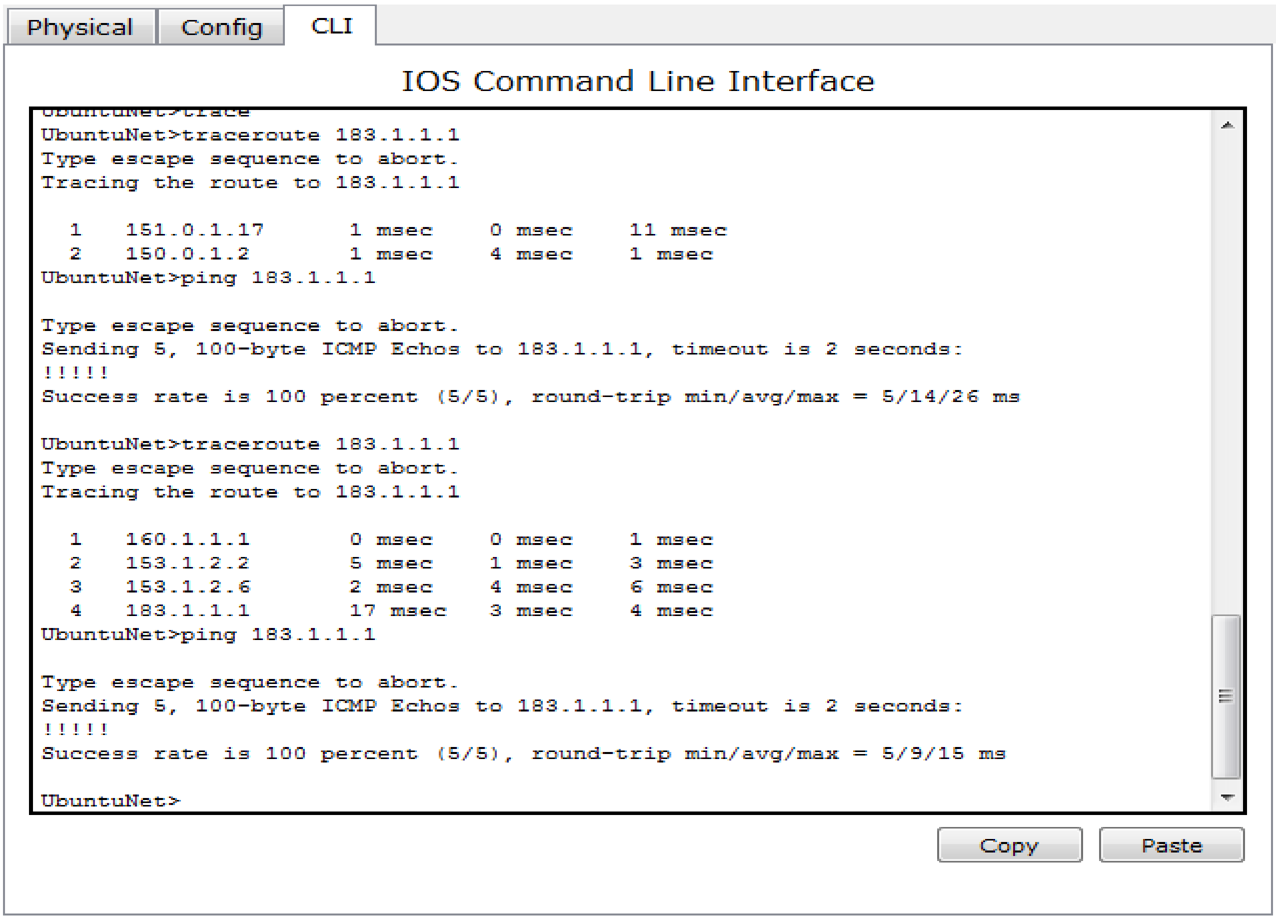

Figure 5: Router UbuntuNet Traceroute and Ping results for current scenario topology and a proposed scenario

3. Discussion

Dealing with the latency problem requires designing cross-border interconnectivity and traffic engineering mechanisms that can identify and use optimal low latency paths. This requires gaining understanding of the cross-border topology as well as having routing and traffic engineering systems that are able learn and make use of the inter-domain topology information in selecting end-to-end data paths. With many edge networks having multiple connection points to the Internet through multi-homing and also through peering relationships, there exist in the Internet topology multiple end-to-end paths. This gives a good opportunity for solving the problem of circuitous routes, through effective traffic engineering techniques. In the context of pan-African NRENs, this could entail optimising usage of inter-continental links as well as cross-border terrestrial links, taking into consideration factors as QoS requirements of network applications, provisioning on the links, as well as cost of data transmission.

3.1 Traffic Engineering

A possible way of achieving optimal traffic exchange is through collaborative routing and traffic engineering among NRENs. A pair of NREN may dynamically [re]route the traffic between them with the aim of minimising latency for certain types of traffic such as video conferencing, while minimising transit cost for non-delay sensitive traffic. This can be achieved with mechanisms for dynamic end-to-end path reconfiguration, and performing collaborative and dynamic load balancing using network metrics. For example, an NREN

 ISSN 2223-7062 Proceedings and report of the 7th UbuntuNet Alliance annual conference, 2014, pp 59-70

may want to dynamically [re]distribute its traffic with the aim of minimising latency for certain types of traffic, while minimising transit cost for the rest of the traffic. This can be achieved if NRENs have mechanisms for dynamic path re-configuration, which can be achieved much more easily through software-defined networking (SDN).

3.2 SDN for Dynamic Network Flow Control

Software Defined Networking (SDN) is an emerging paradigm for network design which decouples the network's control plane from the forwarding devices into a software based network controller (Hyojoon & Feamster 2013). This makes possible the programmatic management and flexible control of networks, coupled with a global view of the network. The network controller uses a network decoupling communication protocol OpenFlow to access to the network devices' forwarding plane and to perform network configuration functions, including setting up packet forwarding rules. The data plane implementation, which is responsible for forwarding the data packets remains within the hardware,

A number of SDN traffic engineering (TE) implementations have been deployed (Akyildiz et al, 2014), and if properly explored, could play an important role for improving pan-African NRENs traffic management. For example, Bailey et al. (Bailey et al 2013) have proposed a software defined Internet exchange that is based on decoupling the IXP's BGP control plane into software based network controllers that have direct access to the forwarding tables in SDN cable switches. This allows different domains to apply custom route selections mechanisms. This is aimed at allowing the domains to select multiple best routes to a destination, as opposed to BGP's convention of selecting only one best path to a destination.

At the global scale, a good example of SDN deployment is Google's B4 (Jain, 2013), a global private WAN that connects Google’s data centres and edge deployments for cacheable content across the globe. The architecture uses OpenFlow to centrally control WAN switches and to split application data flows among multiple paths, taking into consideration capacity and application priority/demands. B4 implements a centralized TE solution with three key characteristics: ability to balance competing demands at the network edge during resource constraint; using multipath forwarding/tunneling to leverage available network capacity in accordance with application priorities; and dynamically reallocating bandwidth in the face of link/switch failures or shifting application demands.

Bell Labs have used an approach similar to B4, by leveraging the centralized controller to implement dynamic routing for SDN with the aim of gaining improvements in network utilization, and to reduce packet losses and delays even in cases where there is only a partial deployment of SDN capability in a network(Agarwal, S et al, 2013). Another example is Microsoft Corporation's implementation of an SDN based WAN (SWAN) (Hong, et al 2013) , where a central controller determines when and how much traffic each network service is able to send, and frequently reconfigures the network's data plane to match current traffic demand. SWAN utilizes policy rules to allow inter-data centre WANs to carry significantly more traffic for higher-priority services, while maintaining fairness among similar services. Making use of SDN controller's global network, SWAN is able to optimize the network sharing polices, thereby being able to carry more traffic and support flexible network sharing. It is reported that SWAN is able to carry about 98% of the maximum allowed network traffic, whereas in contrast, traditional MPLS-enabled WANs are only able to carry about 60% of the maximum allowed network traffic.

SDN, through the OpenFlow protocol, could provide important improvement for NRENs network management, including the ability to centrally and automatically manage heterogeneous devices from different vendors using a standardized Application Programming Interface (API). In terms of traffic engineering, SDN provides the ability for fine grained control of flow-based packet forwarding at various levels of granularity, including session, user, device, and application. SDN could also eliminate the need for manual reconfiguration of devices when there are changes in policy or network structure. With the separation of the control plane from the data plane, the SDN paradigm allows for network virtualization, making possible the creation of multiple separate logical networks over the same physical architecture. This could allow NRENs to create and manage their own custom traffic engineering strategies across the same existing IXP infrastructure.

4. Conclusion

This paper has discussed performance challenges, in terms of latency, for Africa's NREN traffic to be exchanged in Europe. Internet probes from different locations in Africa has shown that on average, over 75% Africa's NREN traffic originating and destined for Africa is routed outside the continent. Using simulation, it is shown that latency for Africa's inter-NREN traffic could be reduced by 50% by implementing a central exchange entity and introducing peering for the African NRENs. Lastly, the paper has discussed the potential for improving the pan-African NRENs traffic exchange through traffic engineering and the use of software defined networking.

References

Agarwal, S., Kodialam, M., & Lakshman, T. V. (2013). 'Traffic engineering in software defined networks' *INFOCOM, 2013 Proceedings IEEE.* pp. 2211-2219.

Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., & Willinger, W. (2012) 'Anatomy of a large European IXP' *ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication, Proceedings.* pp. 163–174.

Ahmad, M. & Guha R. (2011) 'Internet exchange points and internet routing' *Network Protocols (ICNP), 19th IEEE International Conference.* pp. 292–294.

Ahmad, M. & Guha, R. (2012) 'A tale of nine internet exchange points: Studying path latencies through major regional ixps' *Local Computer Networks (LCN), IEEE 37th Conference.* pp. 618–625.

Akyildiz, I. F., Lee, A., Wang, P., Luo, M., & Chou, W. (2014). A roadmap for traffic engineering in SDN-OpenFlow networks. *Computer Networks.* 1-30.

Andronico, G., Ardizzone, V., Barbera, R., Becker, B., Bruno, R., Calanducci, A., Carvalho, D., Ciuffo, L., Fargetta, M., & Giorgio, E. (2011) 'e-infrastructures for e- science: a global view' *Journal of Grid Computing.* vol. 9, no. 2, pp. 155–184.

Bailey, J., Clark, R., Feamster, N., Levin, D., Rexford, J., & Shenker, S. (2013) 'SDX: A Software Defined Internet Exchange' *Open Network Summit (Research Track)*, 2013.

Barry, B., Barton, C., Chukwuma, V., Cottrell, L., Kalim, U., Petitdidier, M., & Rabiu, B. (2010) 'egy-africa: better internet connectivity to reduce the digital divide' *IST-Africa.* pp. 1–15.

Donnet, B. & Friedman, T. (2007) 'Internet topology discovery: a survey' *Communications Surveys Tutorials, IEEE.* vol. 9, pp. 56–69, Fourth 2007.

Gupta, A., Calder, M., Feamster, N., Chetty, M., Calandro, E., & E. Katz-Bassett, (2014) 'Peering at the Internet's frontier: A first look at ISP interconnectivity in Africa' *Passive and Active Measurement Conference, Proceedings.*

Gilmore, J., Huysamen, N., & Krzesinski, A., (2007) 'Mapping the African Internet' *Proceedings Southern African Telecommunication Networks and Applications Conference (SATNAC).*

Hayes, J. 'Connected continent' *Engineering & Technology.* vol. 8, no. 5, pp. 42– 43.

Hong, C. Y., Kandula, S., Mahajan, R., Zhang, M., Gill, V., Nanduri, M., & Wattenhofer, R. (2013) 'Achieving high utilization with software-driven WAN.' *ACM SIGCOMM 2013 conference on SIGCOMM, Proceedings.* pp. 15-26.

Jain, S., Kumar, A., Mandal, S., Ong, J., Poutievski, L., Singh, A., & Vahdat, A. (2013). 'B4: Experience with a globally-deployed software defined WAN' *ACM SIGCOMM 2013 conference, Proceedings.* pp. 3-14.

Kim, H., & Feamster, N. (2013) 'Improving network management with software defined networking' *Communications Magazine, IEEE.* vol.51, no.2, pp.114,119.

Landa, R., Araujo, J., Clegg, R., Mykoniati, E., Griffin, D., & Rio, M. (2013) 'The large-scale geography of internet round trip times' *IFIP Networking Conference, Proceedings.* pp. 1–9, May 2013.

Luckie, M. (2010) 'Scamper: a scalable and extensible packet prober for active measurement of the internet' *10th ACM SIGCOMM Conference on Internet Measurements, Proceedings.* pp. 239–245.

MaxMind. LLC, 'Geoip' (2014) Online: https://www.maxmind.com/en/home - Accessed 1 September 2014.

PeeringDB, "http://www.peeringdb.com". [Online; authenticated access- [guest] accessed 13-May-2014].

Steiner, R., Tirivayi, A. , Tirivayi, N., Jensen, M., Hamilton, P., Buechler, J., Jeffries, A., Emdon, U. H., & Ibrahim, I. A., (2005) 'Promoting African research and education networking' *International Development Research Centre.*

Surborg, B., Carmody, P. (2013) 'Of cables, connections and control: Africa's double dependency in the information age' *Enacting Globalization: Multidisciplinary Perspectives on International Integration.* p. 240.

Biographies

Josiah Chavula is an ICT lecturer at Mzuzu University in Malawi, and a PhD student in Computer Science at the Center for ICT4D, University of Cape Town. His research focuses on Internet performance and traffic engineering in Africa. He is interested in approaches to achieve flexible Internet peering with a goal to improve performance in the pan-African National Research and Education Networks (NRENs). He obtained a Master of Science degree – Networking and Internet Systems – from Lancaster University (UK), and a BSc – Computer Science – from the University of Malawi.

Clifford Sibanda is a PhD student at the University of Cape Town in the Telecommunications (Electrical Engineering). He completed his MSc Engineering (Telecom) at University of Cape Town and a BSc (Hons) in Computer Science in NUST, Bulawayo.