Rebuilding the Internet eXchange Point in Uganda

Diarmuid Ó Briain*†, David Denieffe*, Yvonne Kavanagh* and Dorothy Okello†
***GameCORE Research Centre**,
Department of Computing & Networking,
Institute of Technology, Carlow, Ireland
Email: diarmuid.obriain@itcarlow.ie
†**netLabs!UG Research Centre**,
Department of Electrical and Computer Engineering,
College of Engineering, Design, Art and Technology,
Makerere University, Kampala, Uganda

***Abstract*—Uganda is a land locked country in South East Africa. It is separated from the Indian ocean by Kenya to the east and Tanzania to the south. It is separated from South Atlantic by the Democratic Republic of the Congo to the west. Uganda's connection to the Internet is overland to Mombasa, Kenya where the Seacom cable installed in 2009 connected Mombasa to Europe with a capacity of 640 Gb/s and the World Bank funded Eastern Africa Submarine Cable System (EASSy) undersea fibre optic cable connects Eastern Africa to South Africa and Europe with a capacity of 1.4 Tb/s. [1]. Additionally until recently Uganda was dependent upon a single overland link to Kenya via Uganda Electricity Transmission Company Limited (UETCL) power lines whereas today there are a number of options such as the Seacom cable fibre cable completed in 2015. During the 2010 FIFA World Cup Final a terrorist attack on the Kyandondo Rugby Club in Nakawa, Kampala brought all Internet Service Providers (ISP) in Uganda off-line as the rugby ground sits directly underneath the UETCL power lines bring Internet services from Kenya.**

This leaves Uganda at significant disadvantage regionally and to-date the technology industry has tended to converge at Mombassa and Nairobi as well as in Dar es Salaam, Tanzania. Recently however as the modern Internet services require Global Service Providers (GSP) to get closer to the end-user so they can avail of lower latency, higher bandwidth applications, the justification for serving all of South East Africa from data centres in Narobi, Mombassa, Dar es Salaam and even from South Africa is becoming less valid. In order to facilitate the deployment of locally hosted content and reduce the country's reliance on international cables, additional capacity was required at the country's primary local interconnection point, the Uganda Internet Exchange Point (UIXP).

UIXP is currently developing its infrastructure to support a larger membership encompassing growing demand from both ISPs and Application Service Providers (ASP). This paper serves to outline the developments that are being put in place to take UIXP from a local Internet eXchange Point (IXP) supporting local ISPs and e-government services to the next level where it is in a position to support GSP Content Delivery Networks (CDN) and therefore pave the way for the next phase of development of the Internet in Uganda.

I. INTRODUCTION

A. The make-up of the Internet

The Internet is built on a hierarchical model with local, retail focused tier 3, ISPs purchasing upstream transit from higher tier ISPs. Regional tier 2 ISPs provide transit to tier 3 ISPs and business customers while at the top national and international tier 1 ISPs provide transit to tier 2 and tier 3 ISPs. These tier 1 ISPs are connected to each other and are considered the global Internet backbone. Each ISP as a single technical administration entity represents an Autonomous System (AS) and therefore is assigned a unique Autonomous System Number (ASN). Using the inter-AS routing protocol Border Gateway Protocol version 4 (BGP4) [2] the ISPs exchange network reachability information with each other. Such BGP relationships are called BGP peering sessions.

Taking the example of a tier 3 ISP in Uganda solely connected to an international tier 1 ISP. All of its traffic not destined for customers on its own network will transit to the tier 1 ISP and onto the Internet. Should the destination of that traffic be local and the end customer supported by a different tier 2 ISP who has purchased transit from a different tier 1 ISP provider, then this traffic will hairpin to at least the nearest point of intersection of these two tier 1 networks, that could be in Kenya or even back in Europe. This process is referred to as ***tromboning***.

Tromboning has the effect of imposing three constraints on a national Internet ecosystem and the Service Providers (SP) that provide it [3].

- **Latency**: As traffic must circumvent the long path to the network convergence point between the tier 1 ISPs it therefore adds a significant amount of time for the delivery of the traffic due to the distance and the number of node hops that must be traversed. This latency has a significant impact on time sensitive traffic such as Voice over IP (VoIP), multimedia video and online video

gaming. End-users have a reduced experience and often cannot benefit from video based services like YouTube.
- **Cost**: The fact that the traffic must traverse international transit, twice, adds significantly to the cost which reduces the margin the ISP can make and ultimately is passed on to the consuming end-user.
- **Growth of ecosystem**: In such an environment there is little encouragement for an ASP to host their services in the country. Even locally based content ASPs will choose to host their content in Data Centres at intersection points between tier 1 ISPs to maximise their impact within a larger region.

II. The Internet eXchange Point

"An IXP is a network facility that enables the interconnection of more than two independent Autonomous Systems (AS), primarily for the purpose of facilitating the exchange of Internet traffic." [4].

If two ISPs peer with each other at an IXP then traffic between them does not need to traverse via the higher tier providers, saving both ISPs on cost and latency. The creation of an IXP in a region has the following short-term benefits:

- **Reduction in latency**: As domestic traffic is now passed between SPs locally it avoids tromboning which reduces the distance and number of nodes it must pass through, thereby reducing the latency and improving the effectiveness of time sensitive applications like VoIP and multimedia video applications. The impact of improved latency is greater usage as end-users gain confidence in the VoIP and video services, this increases revenue for local ISPs and ultimately can be passed on to the end-user in reduced service fees.
- **Reduction of costs**: The fact that significant traffic now avoids tromboning over international transit reduces the ISP costs over transit networks of higher tier ISPs. Considered alongside the increase in usage as end-users gain confidence in the services can have a very positive impact on revenues.
- **Increased autonomy**: As significant portions of national traffic is now contained locally, there is a corresponding reduction in dependency upon external, uncontrollable factors. Network outages on international transit has a reduced effect on the local Internet experience.

A. The history of Internet eXchange Points

The first IXP in 1992 was created by the Swedish University Computer Network (SUNET). Telia and Tele2 had expressed a need to exchange traffic with each other so SUNET developed a solution allowing all three operators to connect to what became the first neutral Internet eXchange Point in Stockholm called D-GIX [5].

The Swedish model was quickly followed in the United Kingdom (UK) with the formation of the London Internet eXchange (LINX) in 1994. Local traffic in the UK was actually tromboning via the United States (US) and British Telecom (BT), Demon Internet, PIPEX, UKnet and the UK academic network (UKERNA) decided that the creation of an IXP would keep local traffic local and save the expense and time delay of the tromboning [6].

In Ireland the Internet Neutral EXchange (INEX) became operational in May 1997 with four members; EUnet Ireland, HEAnet, Indigo and Telecom Internet with the objective of locally exchanging IP traffic between Irish-based ISPs, today INEX has 117 members [7].

The UIXP was established in 2001 with three initial participants located at the former Uganda Communications Commission (UCC) head-quarters in Communications House in Kampala. Today the eXchange has 24 members and traffic peaked beyond 2 Gb/s for the first time in January 2017. The initial eXchange was a Layer 2 LAN allowing a small number of members to bi-laterally peer with each-other. As membership increased a Route Server (RS) was added and in 2012 to deal with power problems a 3 Kilo Volt Amperes (kVA) inverter and two 100 Amphere/Hour (AH) accumulator batteries were installed. At that time the eXchange also implemented Internet Protocol version 6 (IPv6) peering for the first time at the RSs.

III. Internet eXchange Point technical primer

A. Border Gateway Protocol version 4

BGP4 is the inter-AS routing protocol used between SPs with ASNs. BGP provides a speaking system to exchange network destination-based forwarding reachability information over the Transmission Control Protocol (TCP) port 179. This information includes ASes traversed with other BGP systems. This information has routing loops pruned and provides a point for the enforcement of policies. Each BGP speaker maintains a Routing Information Base (RIB) that stores routing information learnt from other peer BGP speakers, local routing information to which local policies have been applied and the information the BGP speaker will share with other BGP peers via advertisements. Advertisements share the following main attributes as well as a number of other optional attributes:

- **ORIGIN**: defines the origin of the path information. It is never altered as it passes through BGP speakers.
- **AS_PATH**: sequence of AS path segments through which the routing information carried in an UPDATE message has passed.
- **NEXT_HOP**: defines the IP address of the next hop router.

BGP therefore allows each AS to exchange AS routing information between peering entities. In this way the Internet routing information is shared and dynamically updated to allow the routing of traffic across it [2].

B. BGP at the Internet eXchange Point

IXPs provide SPs the facility to interconnect. This is typically provided by a shared layer 2 switched network. Of course members could decide to peer directly with each-other on the peering Local Access Network (LAN) and in some cases that is exactly what happens if either the IXP is very small or the level of traffic between SPs warrant it. What is more typical however and to reduce the number of direct peering relationships for members, the IXP establishes a RS. The RS acts to reflect routes received from each member to the other members. It is not a router function as it does not record routes learnt in the hosts routes table and it does not carry any traffic. It uses BGP to exchange reachability information with IXP members and then passes the NEXT_HOP attribute unmodified to the other IXP members. The RS therefore does not prepend its own AS number to the AS_PATH segment. For all intents and purposes it appears as if IXP members are peering with each-other. The RS is simply facilitating that. [8]. A typical implementation of a RS at IXPs is using the BIRD Internet Routing Daemon. BIRD was developed as a college project at Faculty of Math and Physics, Charles University Prague. Currently it is developed and supported by CZ.NIC Labs. [9]. Figure 1 is a short example extract showing 3 ISPs connected to the RS. In this case routes learnt from each are reflected to the others.

```
bird> show protocol
name   proto  table   state  since       info
ISP_1  BGP    master  up     2016-12-31  Established
ISP_2  BGP    master  up     2016-12-31  Established
ISP_3  BGP    master  up     2016-12-31  Established
```

Figure 1. BIRD BGP connections

IV. IXP DEVELOPMENT

IXP development can typically be described in five developmental phases. Each of these growth phases can be followed by the number and type of new members that join it. The more members the more valuable the eXchange and the membership becomes.

1) **Initiation**: Despite competition and rivalry, national and local ISPs come together either on their own or with some coaxing from Government organisations or active local individuals to establish an IXP. This is despite the fact that initially there is no real gain for the members as the majority of traffic still traverses transit with tier 1 ISPs. However in Uganda where transit costs are extremly high even a small volume of traffic directed away from transit can mean savings.
2) **Government support**: Government departments realise that to join the eXchange increases the efficient delivery of its e-government services. It is also considered that by supporting the eXchange it will push it towards the next phase with subsequent improvements in Internet access for the citizens.
3) **Enterprise returns home**: As the success e-government services becomes obvious, enterprises serving the same citizens move their Internet based services from foreign hosting providers to local ones to gain from the lower latency and thereby improving customer access to their services.
4) **Global Service Providers**: International ASPs consider the market improvement as a result of the increased activity. They decide that it is now worth locating their services in country and move their Content Delivery Network (CDN) caches to the eXchange. The IXP provides a central point of interconnection within the country otherwise they would need to deploy equipment in multiple ISPs to reach one market.
5) **International peers**: As the exchange grows international ISPs from neighbouring countries see value in peering at the eXchange. This allows them to gain similar benefits for their customers in their own country as they see citizens in their neighbouring country receive.

V. UGANDA INTERNET EXCHANGE POINT IN 2017

As can be seen from Figure 2, the first half of 2016 saw the UIXP network averaging approximately 300 Mb/s [11]. It should be pointed out that this is *average* traffic and on any particular day the peak traffic can be much higher than that stated. Over that period there was a gentle increase in traffic levels. However the graph demonstrates a marked change at the end of June 2016. This change co-incided with the commissioning of an Akamai Technologies, CDN cache at the eXchange. Akamai Technologies are a GSP and this cache joined the most pervasive, highly-distributed CDN with more than 216,000 servers in over 120 countries and within more than 1,500 networks around the globe [12]. It became self evident however that a network based on 2 HP ProCurve 3400CL switches would not support the new levels of traffic. Another major limiting factor was the interconnect between the switches. This was a single physical 1 Gb/s trunk. The distribution of members between the switches meant this Ethernet link rapidly approached its limits of bandwidth capacity. Additionally Google are currently in talks with UIXP to install a Google Cloud CDN to accelerate content delivery for websites and applications served out of Google Compute Engine and Google Cloud Storage [13]. It is expected that this CDN is commissioned in the first half of 2017. This clearly demonstrates that UIXP is moving from *Phase 3 - Enterprise returns home* to *Phase 4 - Application Service Providers* and this has triggered the need for a network overhaul. Another area of development at the eXchange will be the addition of an Internet Corporation for Assigned Names and Numbers (ICANN)"L"-root server. This is one of 13 named authoritative name servers that serve the Domain Name System (DNS) root zone. The addition of such a root server to Uganda will reduce the time for DNS lookups significantly.

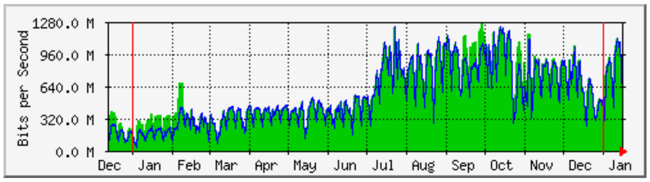

Figure 2. UIXP 2016 Graph (1 Day Average)

VI. UIXP PHYSICAL NETWORK REBUILD

A complete rebuild of the Core and Peering elements of the network and a separation of these functions to separate cabinets was the first step in the eXchange rebuild. Figure 3 outlines the various elements that have been added to the network as part of this major upgrade.

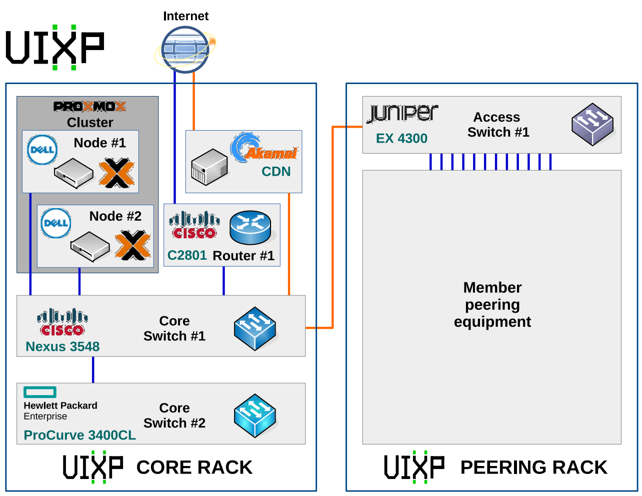

Figure 3. UIXP Physical Network.

- **Core switch**: A Cisco Nexus 3548 was donated to the eXchange by Packet Clearing House (PCH). This switch was installed as the new core switch. It offers 48 fixed Small Form Factor Pluggable plus (SFP+) ports. These give the ability to enable either 10 Gb/s fibre, 10 Gb/s copper, 1 Gb/s fibre or 1 Gb/s copper depending on the SFP module employed in each port. If the Gigabit Line Card (GLC) SFPs are employed then lower speeds are configurable when necessary.
  At the present time the core switch has a HP ProCurve 2524 connected as a secondary core switch to cater for lower speed interfaces within the core network. This reduces the number of SFPs necessary in the Cisco Nexus switch, a significant cost saving to the eXchange. This switch will be replaced by one of the HP ProCurve 3400CL switches once the eXchange members migrate to the new peering cabinets and free up the hardware in the old cabinet.
- **Peering switch**: A Juniper EX4300 was donated by the Uganda Communications Commission (UCC) and the eXchange has employed it as a peering access switch. It offers over 4 times the throughput of the HP ProCurve 3400CL switches, 4 SFP+ ports that support 10 Gb/s interfaces as well as 48 port 1 Gb/s copper interfaces for member peers. This switch has been placed as a Top of Rack (ToR) switch in the Peering cabinet. It facilitates the connection of eXchange members as either 10, 100 Mb/s or 1 Gb/s.
  The Virtual Chassis configuration feature of the EX4300 is attractive given the potential to connect a second such switch in the second peering cabinet at some time in the future. This switch is interconnected to the Cisco Nexus 3548 core switch with a 10 GB/s link configured as a Virtual Local Access Network (VLAN) trunk.
- **Layer 3 gateway**: An old Cisco 3500 router that provided the eXchange with transit has given way for a less power hungry Cisco C2801 router in the new core cabinet. As this router function is to facility the distribution of traffic between the internal UIXP networks and the Internet, the bandwidth requirement is actually quite small and the C2801 is quite adequate for the function. It was considered that this routing function could be replaced by a Virtual Network Function (VNF) on the Proxmox cluster but considering the eXchange already had the Cisco C2801 and the potential that a Proxmox cluster outage would prevent maintenance access from outside the network it was decided to use the physical router.
- **Akamai CDN Cache**: Akamai Technologies installed a CDN Cache in June 2016 at the eXchange which connects to the Core switch via a 10 Gb/s fibre interface [10] however it is currently limited by upstream capacity so the 10 Gb/s capacity cannot be fully utilised yet.
- **Proxmox Hypervisor**: Currently a Dell PowerEdge 750 server and an old HP ProLiant DL320 provide the hardware nodes that support a Proxmox hypervisor. Each server is connected via 1 Gb/s copper interfaces to the core switch. This supports the virtualised functions within the eXchange.

A. *Infrastructure as a Service (IaaS) platform*

To deliver core services it was necessary to build a robust hypervisor based Infrastructure as a Service (IaaS) that supports the orchestration of both Virtual Machines (VM) and Containers (CT) to support the functions required at the eXchange.

The primary selection criteria for the hypervisor platform was the need for it to be a Free/Libre and Open Source Software (FLOSS) platform that supports High Availability (HA) as well as both VMs and CTs. The options explored were OpenStack [14] and Proxmox [15]. Both meet the requirements of HA and IaaS. OpenStack is released under a FLOSS Apache License, while Proxmox is licensed under the GNU is Not Unix (GNU) Affero General Public License (AGPL) version 3, so both are FLOSS. OpenStack however was considered more suitable for a SP wishing to provide cloud services to end-users. This is not a requirement for the eXchange and would add significant unnecessary complexity. While the Proxmox Virtual Environment (VE) is not as fully featured as OpenStack, it is powerful, simpler to deploy and has all the features required by the eXchange.

Proxmox is Debian GNU/Linux based, uses robust Kernel Virtual Machine (KVM) technology and LinuX Containers (LXC). A major plus of Proxmox is the HA Cluster features. When VM or CT instances are configured as HA and the physical host fails, the virtual instance is automatically restarted on the remaining Proxmox VE Cluster node. It was considered that the Proxmox VE HA Cluster is based on proven GNU/Linux HA technologies and it provides the stable and reliable HA service required.

While currently the Proxmox cluster consists of the Dell PowerEdge 750 donoted by Uganda Communications Commission (UCC) and an old Dell ProLiant DL320 Server, however thanks to a donation by the Internet Society (ISOC) of an additional Dell PowerEdge 750 it will be possible to upgrade the Proxmox cluster hardware to have a matching pair of Dell servers. This hypervisor cluster is an essential element of the eXchange and hosts the various VNF and Server instances as either VMs or CTs.

B. Logical networks

Considering a number of items, the need to separate traffic types and security to name but a few it was decided to split the network into logical elements. These elements are displayed in Figure 4.

- **Peering LAN**: This is the backbone LAN the contains the member peering interfaces as well as the RS and the Autonomous System 112 (AS112) Nameserver. The AS112 service, taking its name from its assigned ASN responds to DNS 'reverse-lookup' requests for IP addresses in private address space [16]. It is not possible to access this LAN from the Internet, it is used for members to peer with the RS and AS112 services only. Peering on this LAN is via ASN 37xxx.
- **Private management LAN**: This is not in place yet as the eXchange wants to upgrade the Proxmox server nodes first. It will act as an internal LAN used for intercommunication between the various functions, there will be no access to this network from the Internet. It will be possible to connect from this network to the Internet via a pfSence firewall using Network Address Translation (NAT) for troubleshooting purposes only.
- **DeMilitarised Zone (DMZ)**: This LAN will act as a management access network to permit controlled access to the various networking devices, the VMs and CTs. This LAN is assigned ASN 37yyy and peers with the RS on the peering LAN via the two networks shared router.

C. Virtualised Services

The minimum services an IXP provides is a Layer 2 LAN to permit members to bi-laterally peer, having said that it is

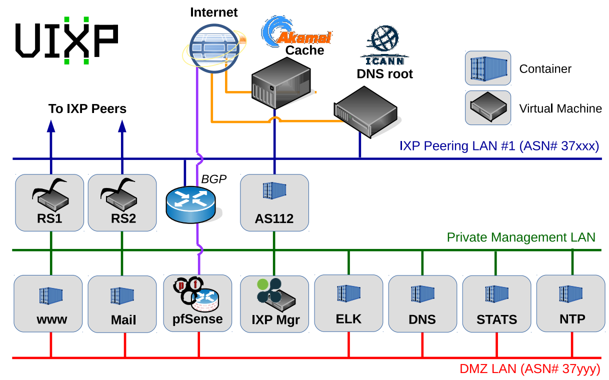

Figure 4. UIXP Logical Network.

Service	Guest	Function
RS1	VM	BIRD service that permits members to BGP peer and reflect routes between them
RS2	VM	Backup RS BIRD service on the second server in cluster
AS112	CT	DNS to handle reverse lookups corresponding to private IP address space and discard them
IXP Mgr	CT	Internet eXchange Point manager version 3
ELK	CT	Elasticsearch, Logstash, and Kibana are three tools to perform log analytics
DNS	CT	Berkeley Internet Name Domain (BIND) DNS server for the uixp.co.ug domain which facilitates Reverse PTR records
MRTG	CT	Multi Router Traffic Grapher (MRTG) is a tool to monitor eXchange and member traffic loads
NTP	CT	Network Time Protocol (NTP) service to provide synchronisation for eXchange elements
WWW	CT	www.uixp.co.ug Apache2 webserver
Mail	CT	UIXP Postfix mailserver
pfSense	VM	Firewall managing access via transit link

Table I
UIXP SERVICES

typical for an IXP with more than a very few members to provide RS and AS112 services. As an eXchange develops there is the expectation that more services are provided either for members or simply for the efficient running of the eXchange itself. With this in mind UIXP have provided services on VMs or CTs and the selection of Proxmox guest type in each case is dependant on the service it supports. Table I is a list of the services provided at UIXP.

D. Current state

Most of the physical network elements have been installed and are in place. The Proxmox cluster is in place and will be further improved by the addition of the second Dell

PowerEdge 750 server. The eXchange awaits the migration of members peers to the new peering cabinets, this will facilitate the activation of the multi-network logical VLANs. The Akamai CDN cluster is in place and active and has made a dramatic, positive impact on the traffic levels through the eXchange.

E. Future work

As the peers migrate to the new infrastructure it will be possible to implement the multi-network logical VLANs and via the addition of the pfSense firewall increase security and control at the eXchange. The soon to be added GSP CDN Cache from Google further demonstrates the confidence these international players have in UIXP. The eXchange will also add the ICANN ”L”-root server in the near future. Testing has begun on version 4 of IXP Manager and it is hoped that following the completion of the testing a project will be put in place to upgrade the existing IXP Manager.

VII. Conclusion

In the wider Information and Communications Technology (ICT) sector the synergies and integration between computer networking and cloud computing is leading towards an all software defined Cloud Integrated Network (CIN). The future will see the cloud fragment with edge cloudlets connected to the core cloud nodes by high-speed links augmenting the existing centralised cloud infrastructure.

The IXP has a key place in this future, to facilitate the delivery of next generation intelligent, virtualised, content delivery functions close to the end-user. Such functions will require even lower latency and higher bandwidth that doesn't exist on today's networks. A distributed IXP with nodes at key regional points across the country supported by high-speed, high-capacity links will facilitate Local Service Providers (LSP) to offer GSPs a means to establish and deliver services to their end customers as if they existed on the GSPs own network.

The Uganda Ministry of ICT (MoICT), in its draft National Broadband Strategy for Uganda 2016 - 2020 (NBS2016-20) [17], is preparing for the Internet of today, but it needs to prepare for the CIN of tomorrow. Growth in the Internet ecosystem should be encouraged by way of an enabling policy environment that eliminates barriers to infrastructure and content deployment. This will help to ensure a competitive domestic and international connectivity landscape which, in turn, will make Uganda more attractive to GSPs, and others.

Industry will transform as the all elastic network, the cloud as well as parallel developments in automation, Augmented Intelligence (AugI), Artificial Intelligence (AI) and the Internet of Things (IoT) drive towards the fourth industrial revolution (Industry 4.0) [18]. Industry 4.0 will be governed by Moore's Law and therefore unlike previous linear societal changes in previous industrial revolutions, this change will be exponential and will impact right across society. [19].

Uganda must prepare for that change and current developments at the IXP as well as a national plan to encourage ISPs to expanded capacity to regional centres would lay a solid foundation for CIN as well as pave the way for technology clusters in the regions. The NBS2016-20 report conclusion has a specific reference to move away from ”silo thinking” to a more comprehensive perspective encompassing different sectors and that a broad-based buy-in by different stakeholders is critical to success. The strategy needs to be extended to specifically consider the impact of CIN in the short term and Industry 4.0 in the medium term.

Acknowledgment

I would like to acknowledge the directors at the UIXP for their support of my work at the eXchange.

References

[1] M. D. J. Williams, R. Mayer, M. Minges, V. Foster, C. Briceo-Garmendia, and W. Bank, Africas ICT Infrastructure: Building on the Mobile Revolution. World Bank Publications, 2011.
[2] S. Hares, Y. Rekhter, and T. Li, 'A Border Gateway Protocol 4 (BGP-4)'. [Online]. Available: https://tools.ietf.org/html/rfc4271. [Accessed: 16-Jan-2017].
[3] 'Background - Netnod', netnod. [Online]. Available: http://www.netnod.se/about/organisation/background. [Accessed: 10-Jan-2017].
[4] 'Definition of an Internet Exchange Point', IX-F Internet eXchange Federation. [Online]. Available: http://www.ix-f.net/ixp-definition.html. [Accessed: 10-Jan-2017].
[5] Andra Utgvan, De Byggde Internet i Sverige. 2009.
[6] 'History of LINX - LINX', London INternet eXchange (LINX). [Online]. Available: https://www.linx.net/about/history-of-linx. [Accessed: 10-Jan-2017].
[7] 'History of INEX - INEX', Internet Neutral EXchange (INEX). [Online]. Available: https://www.inex.ie/about-us/history/. [Accessed: 10-Jan-2017].
[8] N. Bakker, E. Jasinska, R. Raszuk, and N. Hilliard, 'Internet Exchange BGP Route Server'. [Online]. Available: https://tools.ietf.org/html/rfc7947. [Accessed: 16-Jan-2017].
[9] 'The BIRD Internet Routing Daemon Project'. [Online]. Available: http://bird.network.cz/. [Accessed: 17-Jan-2017].
[10] UIXP, 'Strongly welcome Akamai to the UIXP', UIXP Twitter feed. [Online]. Available: https://twitter.com/UIXP/status/754726081278668800. [Accessed: 17-Jan-2017].
[11] 'UIXP Stats - Aggregate Traffic'. [Online]. Available: https://stats.uixp.co.ug/mrtg/aggregate.html. [Accessed: 17-Jan-2017].
[12] 'Content Delivery Network (CDN) and Cloud Computing Services — Akamai'. [Online]. Available: https://www.akamai.com. [Accessed: 17-Jan-2017].
[13] 'Google Cloud CDN - Low Latency Content Delivery — Google Cloud Platform'. [Online]. Available: https://cloud.google.com/cdn/. [Accessed: 17-Jan-2017].
[14] 'Home OpenStack Open Source Cloud Computing Software'. [Online]. Available: https://www.openstack.org/. [Accessed: 17-Jan-2017].
[15] 'Proxmox - Powerful Open Source Server Solutions'. [Online]. Available: https://www.proxmox.com/en/. [Accessed: 17-Jan-2017].
[16] J. Abley and W. Sotomayor, AS112 Nameserver Operations. [Online]. Available: https://tools.ietf.org/html/rfc7534. [Accessed: 17-Jan-2017].
[17] MoICT, 'The National Broadband Strategy for Uganda - 2016 - 2020', Ministry for Information and Communications Technology, Jul. 2015.
[18] M. K. Weldon, The Future X Network: A Bell Labs Perspective. CRC PressINC, 2015.
[19] D. Ó Briain, D. Denieffe, Y. Kavanagh, and D. Okello, 'The move to a software defned future and the implications for Uganda', presented at the National Conference on Communication (NCC), Mbarara, Uganda, 2016.