Global Social Sciences Review (GSSR)
Vol. IV, **No.** III (Summer 2019) | **Pages:** 451 – 461

Internet Infrastructure in Africa: Status and Opportunities

Muhammad Tehsin — Assistant Professor, Department of Defence & Strategic Studies, Quaid-i-Azam University, Islamabad, Pakistan. Email: tehsin@qau.edu.pk

Muhammad Muntazir Mehdi — Works at Microsoft.

Abstract *There is a major information gap in Africa, where access to information is structurally disabled. This study was conducted taking the University of Dodoma in Tanzania as a case study. The causes for poor internet access were identified as: Lack of alignment between last-mile, middle-mile, and long-range network infrastructure; Tyranny of bad on-premises network design; Up to 15 network hops just to leave on-premises network architecture; Lack of best-practice templates and benchmarks for on-premises, middle-mile and national backbone network architecture in emerging and developing markets; and local operators charging exorbitant bandwidth prices. The US Federal Communication Commission's (FCC) definition of broadband is 25 Mbps per host. The World Bank defines broadband as 12 Mbps per hundred consumers. According to the United Nations, broadband is a basic human right and an absolute necessity for productivity and sustainable growth.*

Key Words:
Internet, Infrastructure, Africa, Bandwidth, University.

Introduction

While Microsoft has invested significant resources and capital in building an infrastructure to expand the access to the internet, the customer network remains a black box. Especially in developing markets, the customer network is contributing to more than fifty percent of the latency experienced by the customers in East Africa.

The problem is further exacerbated in the developing markets as cheap (low bandwidth) routing gear means that switches are layered over each other to mitigate the limited port count and port exhaustion while stripping capacity. Further, the investment is staggered over several years, meaning that the infrastructure is carrying over technology debt of low bandwidth gear, at times sitting upstream or downstream of higher bandwidth gear. This results in bottlenecks of congestions, further increasing the latency of the routers.

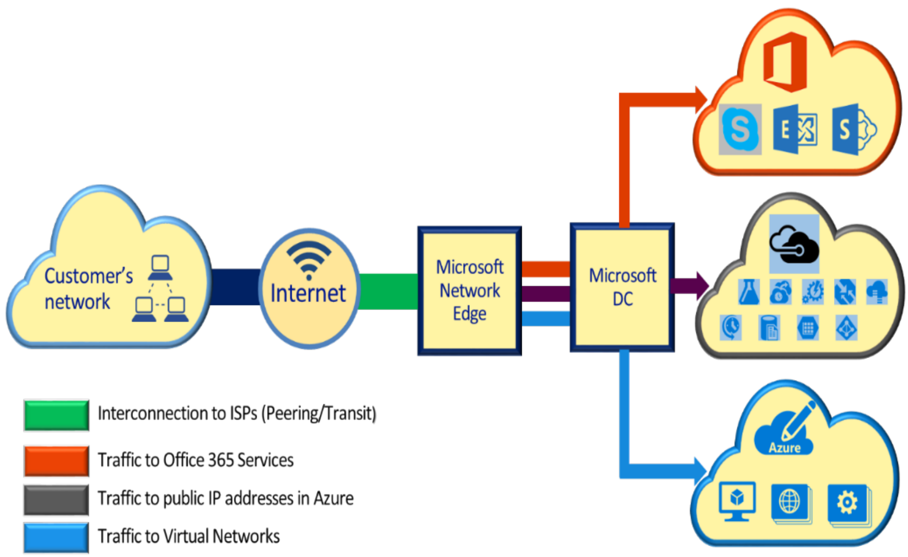

Figure 1: Cloud Assembly Line: From Client to Host

p-ISSN 2520-0348 | e-ISSN 2616-793X | L-ISSN 2616-793X | DOI: 10.31703/gssr.2019(IV-III).56 | URL: http://dx.doi.org/10.31703/gssr.2019(IV-III).56

Latency	Country	France	Kenya	Lesotho	Mozambique	South Africa	South Africa	South Africa	Tanzania	UK	Zambia
Country	**Location**	Marseille	Mombasa	Maseru	Maputo	Cape Town	Durban	Johannesburg	Dar-Es-Salaam	London	Lusaka
Djibouti	Djibouti	62	79	103	78	114	108	98	53	88	118
France	Marseille		141	165	143	167	170	160	115	23	180
Kenya	Momabasa			63	38	71	65	55	83	159	74
Lesotho	Maseru				25	24	18	8	53	168	29
Mozambique	Maputo					33	27	17	25	161	35
South Africa	Cape Town						26	21	61	144	34
South Africa	Durban							10	55	170	29
South Africa	Johannesburg								45	160	19
Tanzania	Dar-Es-Salaam									136	31
UK	London										179

Figure 2: Latency Matrix for Africa

Microsoft's Azure users in Africa enjoy access to applications hosted in Dublin, which gets routed through South Africa. As a case in point, for Tanzania the latency for measurement tests is much higher than average. While this is egregious in itself, the actual customer experience is almost in all cases twice the measured latency as the on-premises network on the customer side adds up to sixteen hops just to reach the access gateway. A TCP session windows scale limits are exceeded at this high level of measured latency, which causes the session to be either reset or terminated.

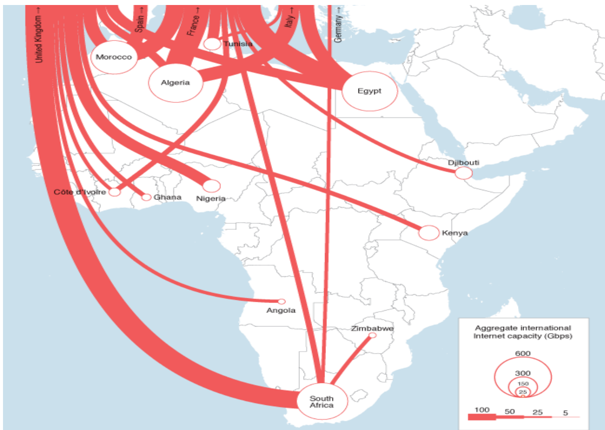

Figure 3: Traffic arteries from Africa to Europe

Most of the traffic to Africa is either coming to or from Europe. South Africa is the largest internet hub in the region. Microsoft has points of presence (POPs) in Cape Town and Johannesburg. This will change in the coming decades as we will see more traffic flows from the Asia-Pacific region to Africa. India and China have been embarking on massive infrastructure projects in Africa, and enterprise business from India, China, Korea, Vietnam, Indonesia, Singapore, etc., is growing. As of now, all traffic from the Asia-Pacific region to East Africa is routed through Suez Canal. In the future broader trans-Arabia and trans-Africa paths can be opened to offer latency and resilience for all Asia Pacific-Africa, North-South Africa, and East-West Africa traffic flows.

Table 1. Internet bandwidth (Mbps)

South Africa						
International Internet Bandwidth (Mbps)						
Country	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**
United Kingdom	24,332	44,572	32,077	127,703	159.615	258.860
France	667	975	13,155	23,045	34,405	61.510
Germany	2.114	7,304	7,769	12,969	30,626	40.937
Kenya	155	155	11,465	12,242	22.203	25 571

Portugal	1,244	1.360	155	9,933	12,613	15.000
Netherlands	1,364	5,922	9,955	11,965	14,000	14.500
India	1,244	1.552	3,110	3,744	13,722	11.800
Namibia	1,010	1.399	889	5,017	6,381	7,123
Singapore	-	-	300	1.455	2,506	2,610
Lesotho	155	322	795	949	1,344	2.345
United Arab Emirates	-	-	310	465	622	622
United States	6.185	2,310	3,120	465	465	465
Nigeria	-	-	-	155	155	310
Ghana	-	-	-	-	155	155
China	4,041	-	311	466	155	155
Belgium	-	155	155	-	-	-
Australia	—	—	1,563	1,563	—	—
National Indicators						
	2011	**2012**	**2013**	**2014**	**2015**	**2016**
Int'l Internet Bandwidth (Mbps)	45,476	71,337	163,704	255.025	385.571	546,415
Broadband Users (thous.)	950	1,150	1,380	1,630	1.735	1,370
Broadband Penetration (%)	7.1%	8.0%	9-5%	11.4%	11.6%	12.2%

Note: All country data as of mid-year

UDOM Case Study

The University of Dodoma is a sprawling campus on the hills of Dodoma, the emerging capital city of Tanzania. The University has 27,000 students and 2,000 staff members. The University consists of several colleges and research institutions, ranging in focus from medicine and basic sciences to telecommunication and information technologies. The University has been able to set up an impressive brick-and-mortar infrastructure. However, the total bandwidth for 29,000 staff/students is 80 Mbps. Based on initial download simulation models and laden by the topology round-trip times, only 12 students can manage to download a research paper simultaneously—and even then, there are considerable buffering and latency delays.

Due to the poor conditions of the network, a need was felt to map the network design. Though the mesh requirement was a relatively simple 14 node network architecture, it was designed in a way that the traffic would keep rebounding and hair pinning through the nodes, requiring up to 14 hops to egress the UDOM campus into external networks.

The network design rendered any amount of bandwidth availability ineffective for the user experience. So, unless there were templates to design smart and efficient on-premises network architecture, the network performance would have remained dismally bad. While the Microsoft Windows Server literature is detailed and prescriptive regarding setting up the network, there is little information on optimal network design best practices, how to optimally peer with external networks, or laying out the network gateways for best network performance and user experience. This means that the most valuable components of the user experience are left to the local system administrator, with little to no guidance about metrics or success criteria.

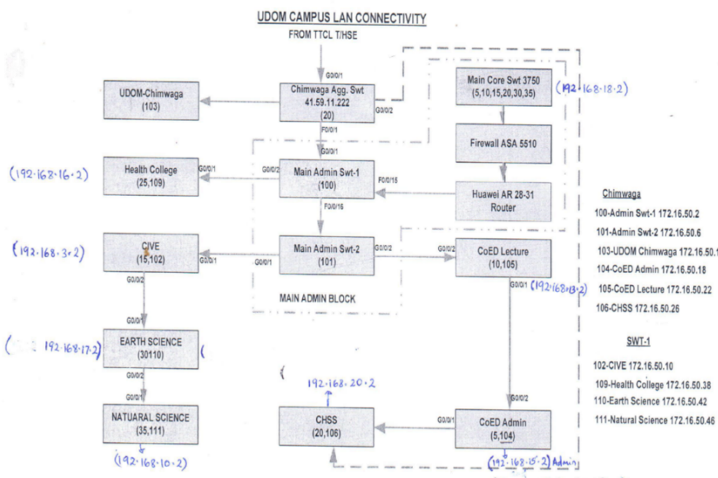

Figure 1: System Architecture for UDOM

Issues with Network Design

In Figure 2, the network requires 9 hops to get from the fourth floor of the building to the adjoining building switch in best-case scenario, wherein the worst-case scenario, it can take up to 16 hops to reach the gateway. This is an issue of professional assistance in understanding network topologies, programmatic discovery of network topology and recommendation on optimal network design.

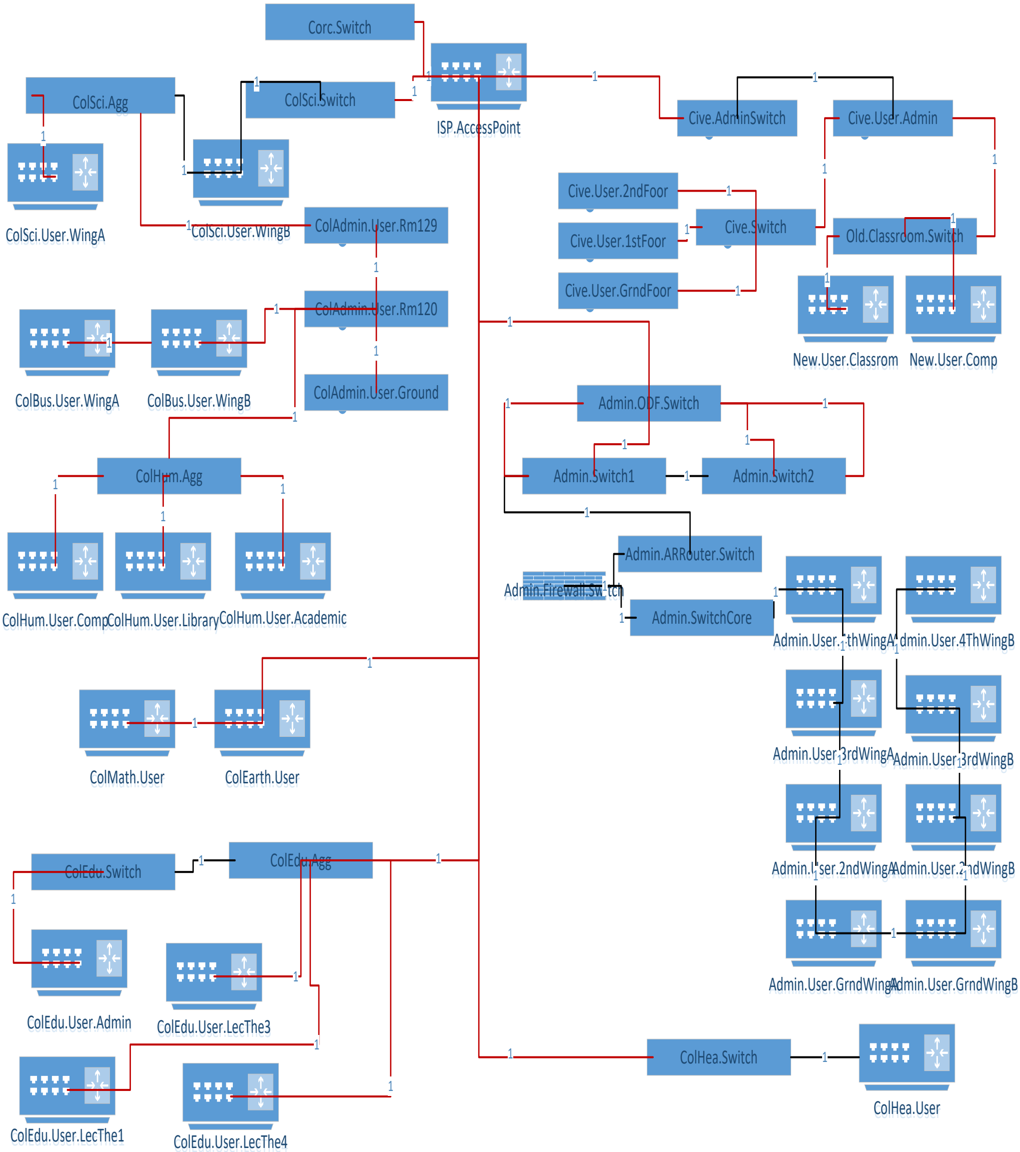

Figure 4: Network Topology for UDOM

Product Development of Network Design & Simulation Tool

Microsoft/Azure tool kit should be provided that can discover the network topology, the number of hops between different endpoints, and generate flow patterns between A-to-Z tuples. It should be built "as-is" on-premises network design, and then provide recommendation on optimal design based on running graph theory algorithms.

The Microsoft tool kit should be repurposed to work for optimization of Microsoft InterDC infrastructure and apply it to customer (on-premises) infrastructures, especially as depicted in the case study of UDOM infrastructure.

The layering of the router switches, port exhaustion and carrying over of technology debt can be resolved using graph theory/simulation analytics.

This approach can be generalized to any topology for small and large enterprises. In order to provide optimal customer experience of Microsoft services, there has to be optimization tools that can offer recommendation on the challenges involved while onboarding to the Microsoft Cloud.

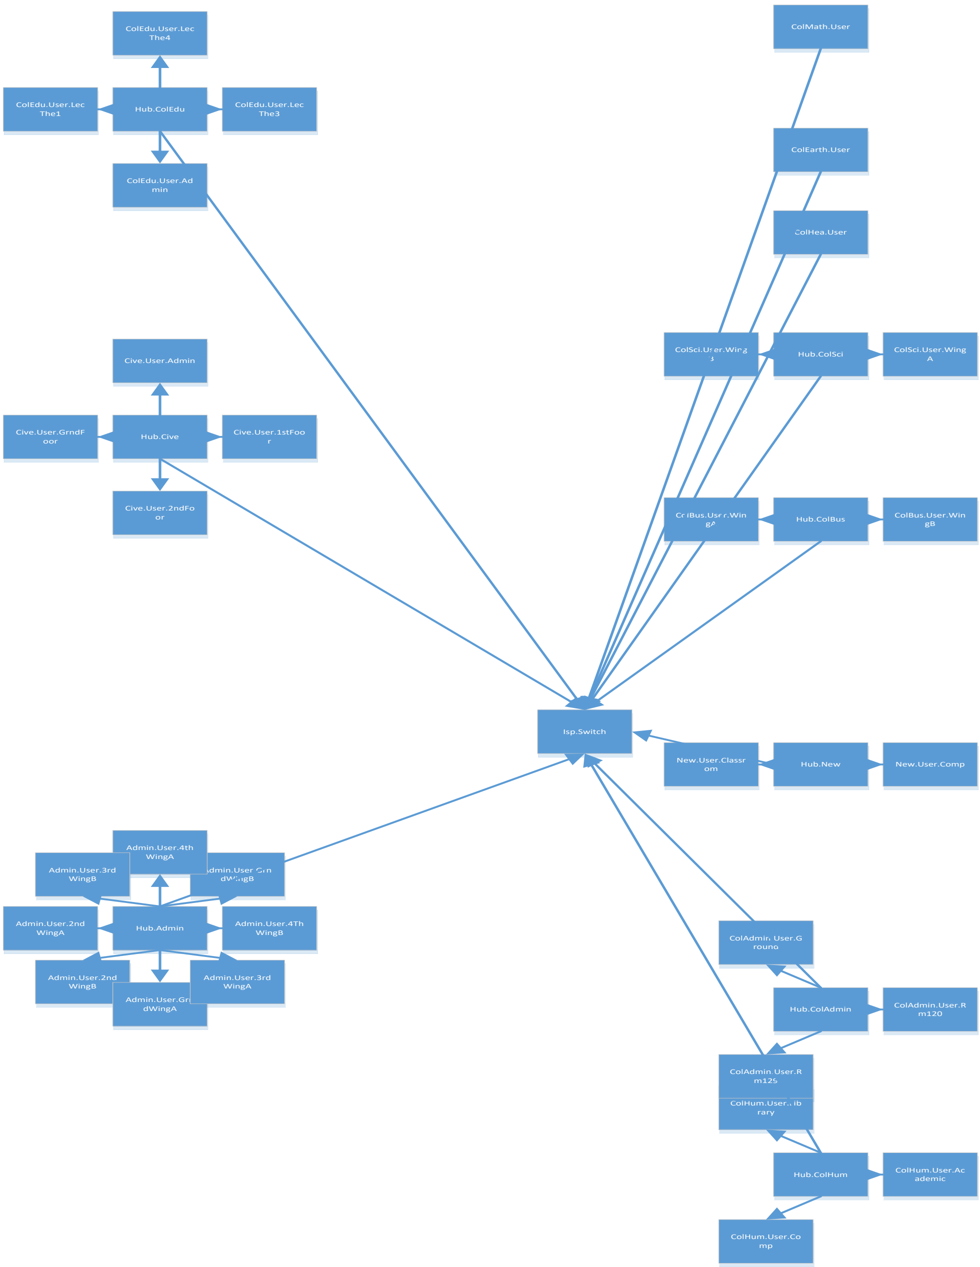

Figure 5: Optimization to Star Schema for UDOM topology

Figure 6 is an algorithmically generated topology that improves information flow permanence by four to five times from all customer endpoints to access gateways.

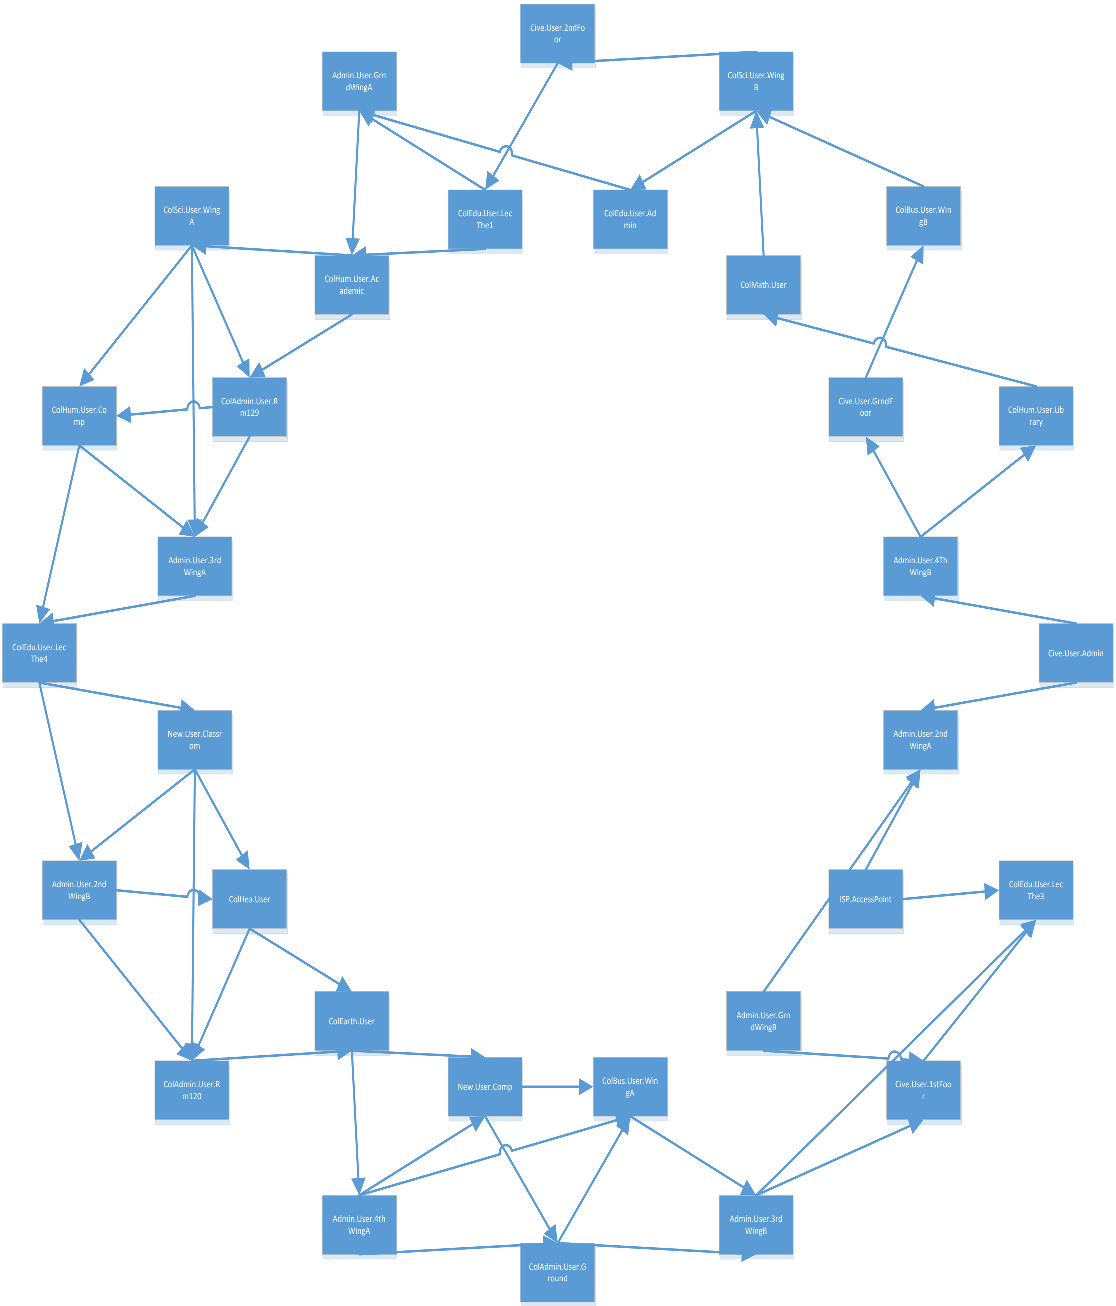

Figure 6: Optimization to Ring Topology using Minimum Spanning Tree with pruning

In Figure 7, the algorithmic output applies Minimum Spanning Tree with pruning, and the programme self-discovers a ring topology which is ten to fifteen times more performant than any comparable topology. It would be several weeks if not months of work for world class architects to lay out such topologies applying iterative approach to network design.

The solution here can improve the bandwidth accessibility for tens of millions of consumers in Africa who are unable to receive broadband connectivity because of congestions and layering of on-premises network design.

The US Federal Communication Commission's (FCC) definition of broadband is 25 Mbps per host. The World Bank defines broadband as 12 Mbps per hundred consumers. According to the United Nations, broadband is a basic human right and an absolute necessity for productivity and sustainable growth.

10 G for 10,000

At the UDOM campus, the University's Information Technology department, local government representatives, and officials from the Ministry of Communication and Foreign Ministry realize that they need to define a new mission for UDOM, which is just as applicable to the rest of Africa: "10 G(bps) for 10,000 (students)." Or simply, "10 G for 10,000." Every educational institution should aspire to a minimum of 10 Gbps of Internet connectivity where the student population is 10,000 persons per campus or facility. For every additional facility or campus, the same relationship would be proportionally applicable.

There are over one thousand universities in Africa that fall into this category. The realization of this vision would provide 10 million researchers and academic users with 10 TeraBytes Per Second (Tbps) of additional network bandwidth for research and collaboration. Once the first phase of 10 G for 10,000 is rolled out, then the downstream impact on enterprise and e-governance can be quantified.

As Dodoma becomes Tanzania's national capital and in accordance with UDOM's status as a flagship research institution, a bandwidth of 100 Gbps for the campus would be appropriate. This would ensure that its future projected growth (a student population of 40,000) would have sufficient bandwidth connectivity. Moreover, it would also ensure that the University was able to provide high-speed internet to other government institutions in Dodoma, furthering e-governance. All e-government initiatives are dependent upon successful and reliable broadband Internet connectivity.

The University of Dodoma is poised to become a flagship institution in Africa. It will likely be viewed by the rest of the continent as a benchmark—a fully integrated student community with a global educational and research communication infrastructure.

However, infrastructure is a prerequisite for all future engagements and collaborations pertaining to technology and digital innovation. The modern infrastructure consists of three ubiquitous grids: a utility grid, a transportation grid, and a communication grid. The Tanzanian government has performed outstandingly well in constructing utility and transportation grids. The logistical arteries connect through air/rail/road travel. Even remote villages have electricity. Yet the country has struggled to develop a communication grid.

While cellular coverage in Tanzania does allow for basic communication, it is woefully inadequate to support education/research institutions, commercial enterprises, and e-government initiatives. This imposes a major digital constraint on the potential of the University of Dodoma. Yet this constraint is artificial. It can be easily eliminated by end-to-end network design and the removal of unnecessary hops for topology optimization.

Examples of Connectivity's Impact on Productivity

i. All teachers, including medical doctors, spend weeks, or even months, preparing and grading exams. In the US, most exams are offered through knowledge management systems (Sakai being an open-source variant).
ii. Most students do not have access to computer resources, so most research is submitted without actual experiments or simulations.
iii. Sixty percent of grocery (fresh produce) prices are due to logistics costs, in comparison to two percent in the US, as the transportation network is not hooked up to the communication grid. This causes under-utilized trucking routes/incomplete truckload runs, which yield empty returns.
iv. Prices for virtually every piece of merchandise/commodity in the market is susceptible to volatile pricing, as demand and pricing information are not readily available. This causes the breakdown of the rules of supply versus demand in the local markets.
v. The government does not have statistics on trade and the migrations of people, so decisions regarding infrastructure capacity planning are carried out almost randomly.

Even state-of-the-art computing labs like the Microsoft Innovation Center (MIC) run on obsolete software. As software is increasingly updated online, a prerequisite for fully operationalizing MIC is a minimum of 100 Mbps (3

Mbps/computer). According to cloud-computing requirements, a benchmark for minimum viable connectivity is 5 Mbps for each host.

MIC should organize and host conferences, summits, and technology events on a weekly and monthly basis. These events should begin by leveraging local capacity and by nurturing UDOM resources. Initiate MIC internship programs and MIC fellowships offered to the students of UDOM and the broader community. Showcase breakthroughs in design and software development through MIC engagements.

Communication Infrastructure Classification

Long-Distance Infrastructure

i. Two subsea (SeemeWe4, and WACA2) and one terrestrial cable (Kenya to Tanzania) to connect Tanzania to the rest of the world.
ii. National ICT Broadband Backbone (NICTBB) connects Dar Es Salaam to Dodoma.

Middle-mile Infrastructure

i. The network interconnectivity and peering to international internet and other local providers is through TTLC.

Last-mile Infrastructure

i. Connectivity of the departments is through a fiber optic link.

The Tanzanian government has invested in long-distance, middle-mile, and last-mile; but now the primary need is to align and stitch together the three classes of infrastructure. This requires redesigning the network and investing in network research within the University of Dodoma.

Framing Principles of Microsoft Innovation Center (MIC) - UDOM partnership

Table 2. Framing Principles for Microsoft Innovation Center Partnership

Principle	Method	Resource
Re-engineering of Network	Network Simulation & Optimization	Microsoft world-class architects and network designers
Broadband connectivity	➢ Achieve 10G for 10000 ➢ Remove TTCL from the equation ➢ Turn UDOM into an e-government internet service provider (ISP)	Team of UDOM researchers, leveraging ICTs and CIVE research arms
Create an organic model for MIC growth (after 10 G connectivity)	Start with concrete innovation initiatives with UDOM students	MIC as a Consulting center MIC as a Data Hub for East Africa Fellowship of Windows Insiders Startup combinator and innovation incubator
Research Collaboration (after 10 G connectivity)	Solution-based and problem-oriented prototyping	Microsoft Research Teams with UDOM faculty

30 day-60 day-90-day milestones and six months-1 year-5 years plan

30 day

Within 30 days, have a dedicated 60 Mbps VLAN connection to MIC from 5 pm to midnight. Based on a peak to trough ratio estimations, this would improve the average utilization for the University links.

Sixty days

Currently, it takes approximately 30 hops for traffic to exit the University campus, with 10-12 hops are spent in leaving the campus. This is because of a multi-layered network design, where switches are buried in unnecessary hierarchies of routers. This calls for the need to prototype a redesigned network topology, complete with simulations and optimization.

90 day

Budgetary approval and assessment realigning the middle-mile and last-mile.

Six months

Full redesign of the intra-campus network with new endpoints and gateways to external networks. Operationalizing of the UDOM internet exchange point (IXP).

1 year

Have 10 Gbps connectivity at UDOM and realize the vision of becoming an ISP.

5 years

Have UDOM connected to a Global educational and research network (NREN).

Digital Divide and Context

While developing markets are spending billions of dollars in marque academic campuses in Africa, LATAM, and Asia, the primary challenge is to create a world-class electricity and transportation grid with little resources. Funding needs to be allocated to a fiber-optic grid to enable high-speed broadband connectivity. Even where a national backbone of fiber-optic connectivity exists, the vagaries of network design affect the end-users in ways that make connectivity painfully slow. This has a significant impact on the ability of users to access global content. It is instructive that how lack of connectivity can impact our consciousness as global citizens. For example, the number of academics that do not know of TED Talk is an indicator.

Product Feature Scenario Scope

There exists price volatility, information asymmetries, and concentration of risk in almost all segments of African commercial activity. From logistics, healthcare, education and tourism, the lack of market dynamics leads to inefficiencies and cost-heavy services.

The very idea of markets is not a western economic construct. Rather the efficient flow of demand signal for new services and ideas is tethered to risk diffusion, concentration factors, local market forces (incumbent effects) and regulations.

Without proper mechanisms to facilitate information flows; people have difficulty handling complicated, abstract, low feedback problems. While technology and western norms may not be applicable across the world – however the following universal principles should not be overlooked:

- Information mechanisms
- Risk (diffused or concentrated)
- Market dynamics

Economies of scale makes monopolies natural – the question is how to turn the forces of monopoly into more dense ecologies.

Everybody benefits from more density. 10 G of bandwidth is being sold at 187,000 US dollars in Africa while in the US it is sold at a hundred dollars. This is a direct effect of market density. Toyama, K. (2010) equates technology to a magnifier of human intent and capacity. This precludes technology as the driving force either contributing or fixing the inequality in the society.

- Digital is infrastructure
- Internet is utility
- 10 G for 10,000
- Internet2 for Africa
- Markets are natural
- Free Markets are super-natural

Conclusion

Azure Networking feature is an opportunity for Africa or for that matter in emerging markets worldwide to help them digitally transform and move to Microsoft Cloud faster. Piloted at University of Dodoma in Tanzania, East Africa, around the capability of Microsoft Cloud and how that can help them transform by leveraging Azure Cloud services.

1. The Microsoft has an opportunity to pilot a customer network optimization tool in Africa. The African Continent has systemic infrastructure challenges. Most customers experience poor performance when accessing Azure Services. However, the cause of the performance challenge is the on-premises topology and last-mile access. In most cases, service experience can be improved by multiple times by just fixing the on-premises, last-mile topology.
2. Azure Networking is uniquely positioned to build a simulation and optimization tool, where customers can visualize their network topology. Identify service performance based on the network topology on-premises, and the last mile network reaching the premises.
3. This work recommends optimal network topology that would improve the network topology by reduction in hops, layers, and loops in network topology leading to gateway with up to 200 times improvement in Azure service performance.
4. None of Microsoft cloud competitors have a viable feature in the space. It is expected that other companies would venture in this market.
5. Another addition of the Microsoft tool is a 1-year to 5-year total cost of ownership and cost of service models that can help in financial modeling of onboarding to the Microsoft Cloud.
6. The Microsoft tool can either be offered as a value-added service with additional revenue stream or it can be a free service embedded into Windows 10, Windows Server, and Azure Portal. This is a feature and capability that could differentiate Microsoft from its competitors, and more importantly help them remove this blocker.
7. The Microsoft tool can help with customers managing their network infrastructure, technology debt, and network technology roadmap with milestones such as when to upgrade, step function of upgrades, minimum network capability requirement based on load/ utilization variables.
8. The value of the Microsoft Cloud is that how users can leapfrog to the digital world given that they do not have massive legacy systems as much as the developed markets.
9. The Microsoft has built a product which either runs within Windows Server, Windows 10 eco-system or is part of Azure portal, which can do the following things:

a. Discover local network topology.
b. Identify bottlenecks and design constraints.
c. Recommend an optimal topology based on standard star-schema and Minimum Spanning Tree with pruning etc. approaches.
d. Microsoft should consider shipping its deprecated network gear or LinkedIn white labelled gear to some of the large research and academic institutions to Africa to jump start the local network infrastructure.

References

Gomez, R. (2012). Users' Perceptions of the Impact of Public Access Computing in Colombia: Libraries, Telecenters and Cybercafés. *Information Technologies & International Development*, 8(3), 19-33

Van Dijk, J. (2005). *The Deepening Divide*. Sage Publishers. (Intro & Ch. 2). Deepening Divide intro,ch2.pdf

Toyama, K. (2010). Can Technology End Poverty? *Boston Review, 36*(5).

Hazeley, D. (2015). Bridging the Digital Divide: The African Condition. http://www.broadbandcommission.org/Documents/reports/bb-annualreport2015.pdf

http://data.worldbank.org/indicator/IT.NET.BBND.P2