

The Benefits of Local Content Hosting: A Case Study

May 2017



Internet
Society

Table of Contents

| | |
|-------------------------------------|----|
| Acknowledgments..... | 2 |
| Executive Summary..... | 3 |
| 1. Introduction..... | 5 |
| 2. Results..... | 8 |
| 2.1 Research method and data..... | 8 |
| 2.2 Results..... | 9 |
| 3. Conclusions..... | 17 |
| 4. Lessons learned..... | 19 |
| 5. Recommendations and Roadmap..... | 20 |

Acknowledgments

We would like to thank Sally Wentworth, Karen Rose, Konstantinos Komaitis, Jane Coffin, Dawit Bekele, Kevin Chege, and Michuki Mwangi for their input and feedback. We would also like to thank the Honourable Jean Philbert Nsengimana, Minister of Youth and ICT of Rwanda, for his vision and leadership; Ghislain Nkeramugaba, CEO, Grace Ingabire, ccTLD System Administrator, and René Manzi, Network and System Engineer, of RICTA, for their insight and assistance, and James Cowie, Andrew Sullivan, Matt Shoemann along with Dyn, for their detailed network measurements.

Michael Kende – Senior Fellow; The Internet Society; kende@isoc.org

Bastiaan Quast – Economics Fellow; The Internet Society; bquast@gmail.com



Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License
https://creativecommons.org/licenses/by-nc-sa/3.0/deed.en_US

Executive Summary

The availability of locally relevant content is increasingly seen as a limiting factor in the widespread adoption of the Internet in emerging regions such as sub-Saharan Africa. In every country, there are, of course, already websites that do target local audiences, although not in abundance. However, one aspect of local content often overlooked is that those existing websites are generally hosted in the US or in Europe, because it is often more affordable. This overseas hosting has several direct and indirect adverse effects that should be understood and addressed.

Hosting abroad adds cost and latency to downloads. The loading time for websites increases dramatically, often going from several milliseconds to half a second for a very simple page. This in turn leads to visitors spending less time on websites, visiting fewer pages, as well as returning less often to the site.¹

Additionally, hosting abroad adds costs, as Internet Service Providers are forced to purchase greater capacity on submarine cables to carry the traffic. These costs are passed on to consumers in the form of data charges, which increases the cost of Internet access and usage, which in turn has a direct negative effect on the number of individuals being able and willing to use the Internet.

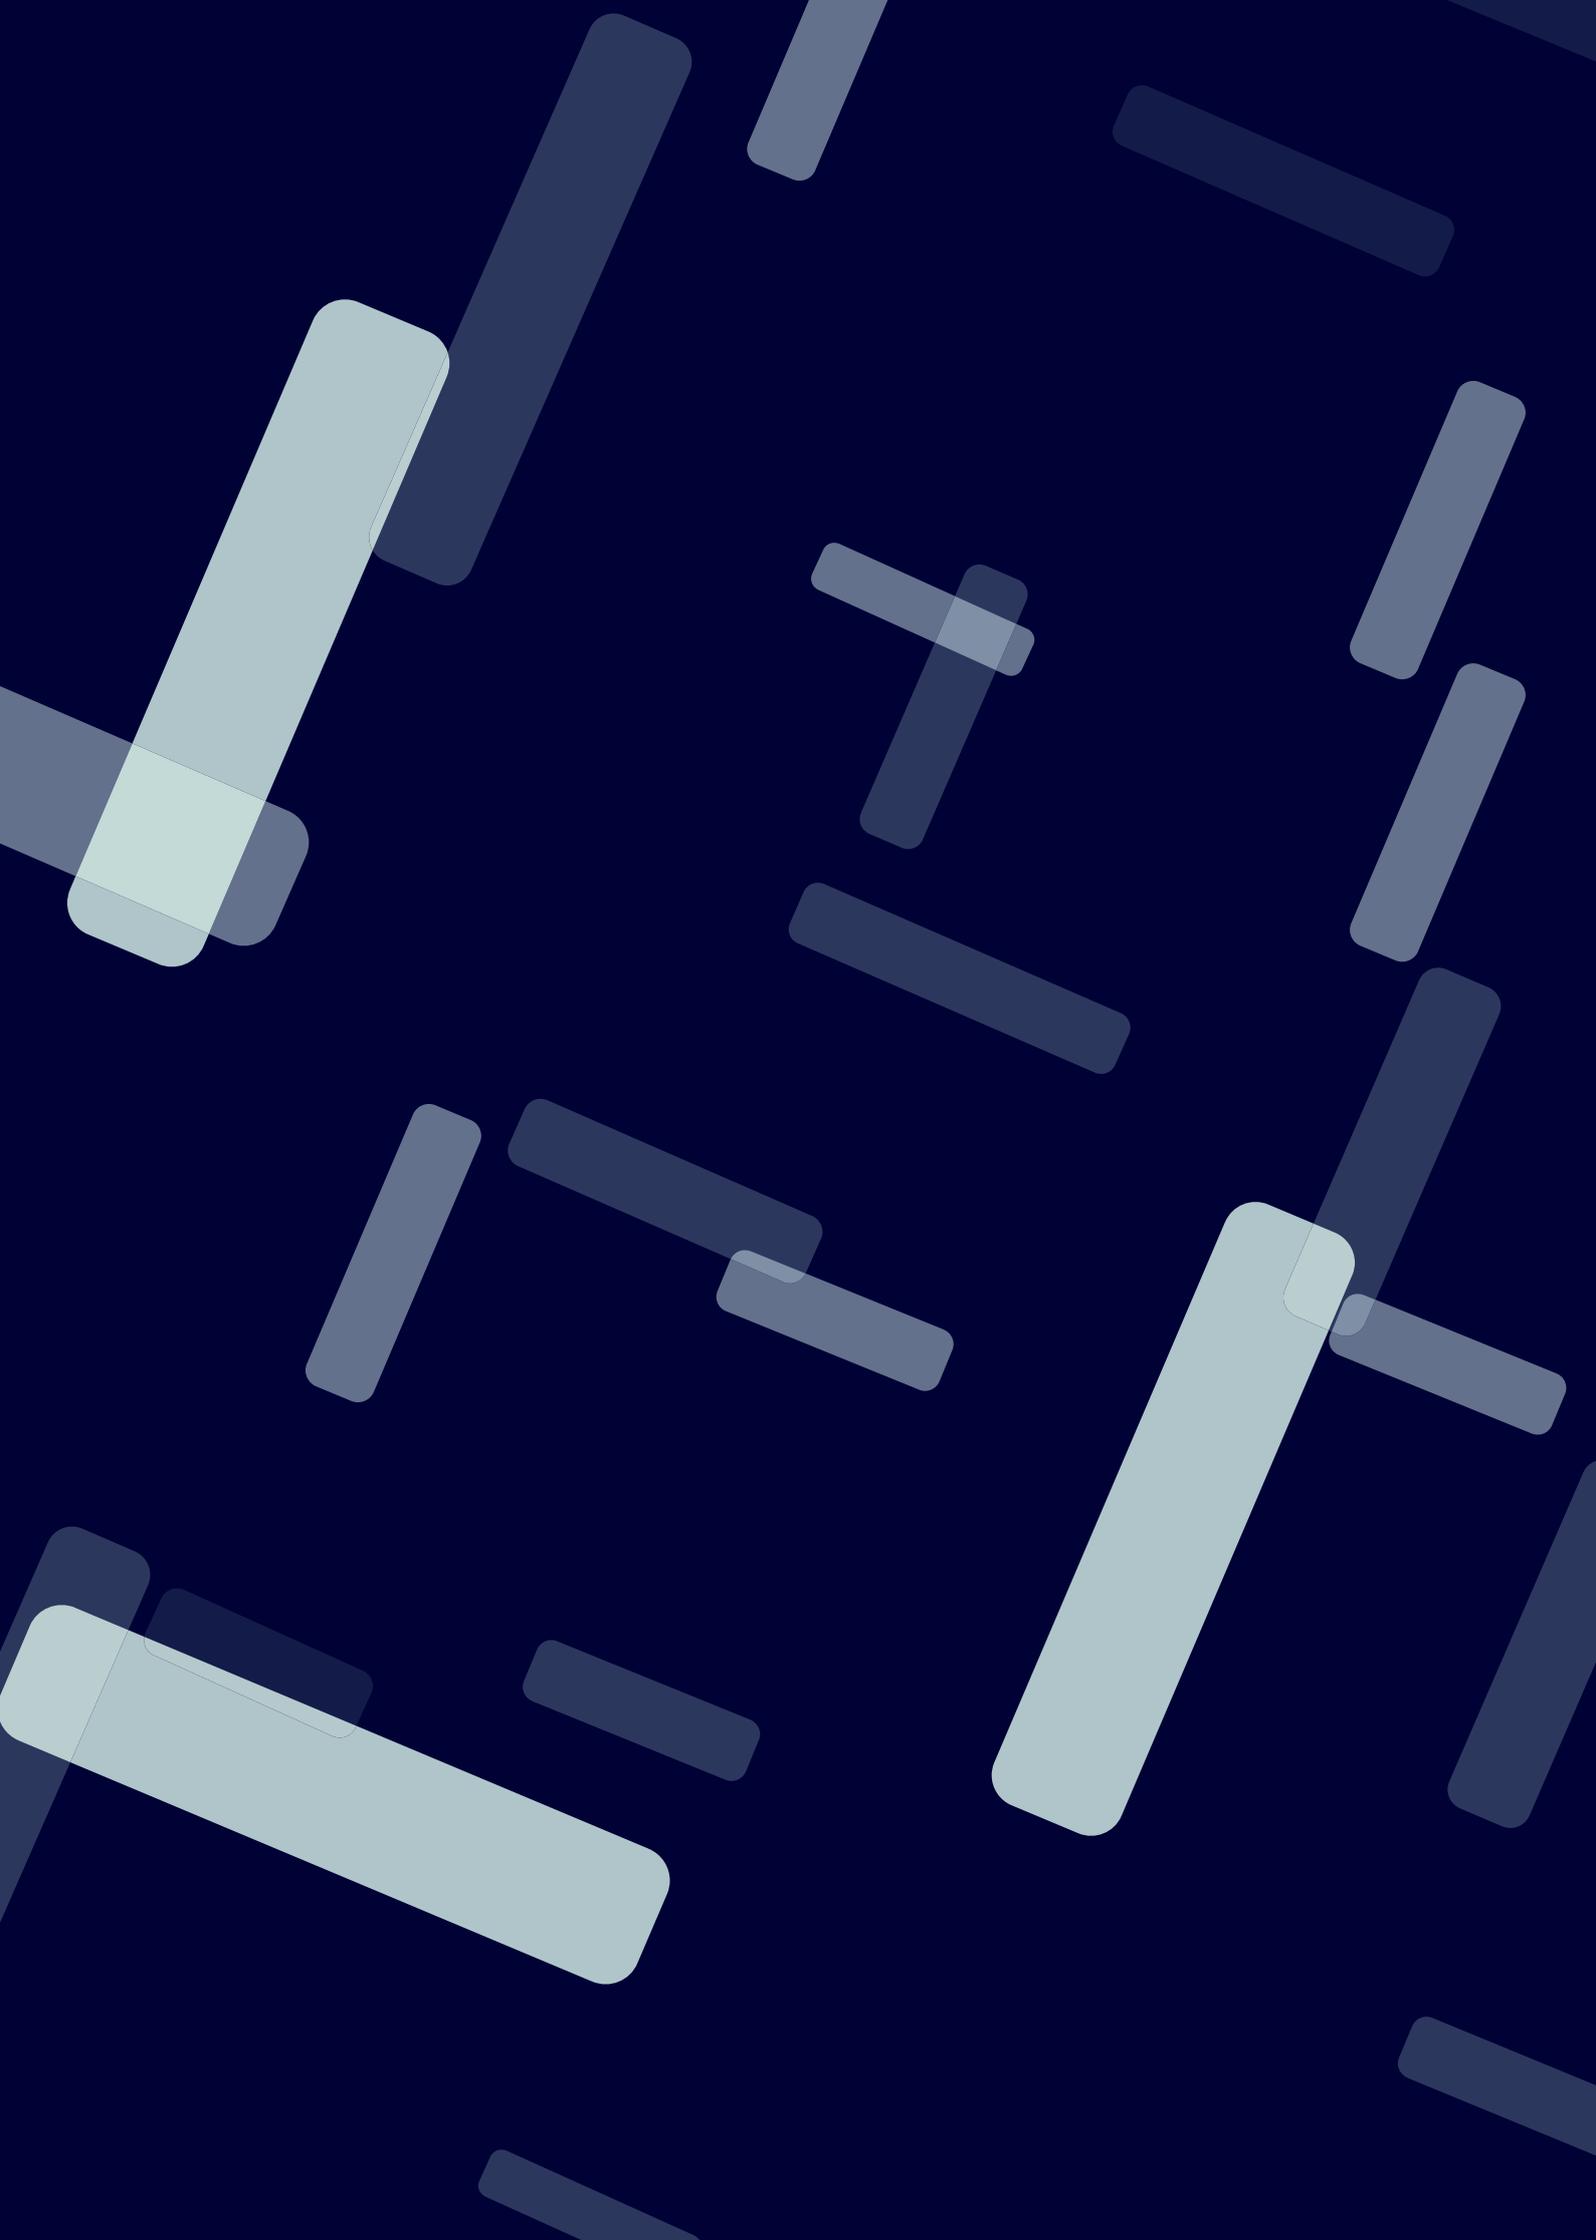
After a previous Internet Society project highlighted these costs of hosting local websites abroad, the Ministry of Youth and ICT in Rwanda decided to conduct a pilot study to repatriate a selection of Rwandan websites back to Rwanda. The pilot project was conducted by the Rwanda ICT Association, with assistance from the Internet Society.

The study confirmed that moving to local hosting leads to several benefits:

- For website owners, a dramatic drop in loading times, which increased the number of visits to the websites that were moved, and fostered greater visitor engagement, both in terms of more page views per session, as well as more return visits.
- For website developers, expanded business opportunities as hosting resellers, and more business for local data centres.
- For government, progress towards the goal of increased local content and corresponding increase in local Internet usage, as a means to develop a robust local Internet economy.

Overall, participants in the study were very positive, which bodes well for Rwanda as the pilot project continues to expand, and highlights the importance of local content hosting for other countries.

¹ The Internet Society does not endorse forced content localisation in any shape or form. Arguments presented here show the benefits of choosing to host content locally. A further discussion of why government attempts to localise content backfire is available in the introduction.



1. Introduction

In most African countries, the clear majority of websites targeting a local audience are hosted abroad, generally in North America or Europe. This has negative impacts on users' engagement with these websites, which we detail here. More broadly, hosting websites abroad stifles the growth of a local content infrastructure, which further stifles the growth of local content and services. As such, local content hosting forms an important stepping stone to the development of a robust local Internet economy.

This report discusses the results of a pilot project conducted by the Rwanda ICT Association (RICTA), with the assistance of the Internet Society, in which websites were moved from international to domestic hosting. We measured the positive impacts on the websites, both in terms of latency measurements, as well as in terms of visitor behaviour.

Local content hosting offers several key advantages, such as greater visitor engagement due to shorter loading times, as well as a possible reduced cost of access for consumers, due to lower Internet Service Provider (ISP) spending on international bandwidth. The paper "Promoting Local Content Hosting to Develop the Internet Ecosystem", released by the Internet Society in 2015², which serves as the foundation for this project, sets out these advantages in full detail (see box).

The effort to localise hosting in Rwanda is part of a larger move by the Government of Rwanda to increase the amount of content served locally. Rwanda already has a Google Global Cache, as well as a node for the Content Delivery Network (CDN) Akamai, both generally serving international content. The Government of Rwanda chose RICTA, the Rwandan registry, as well as the manager of the Rwandan IXP, to manage the pilot phase of the local hosting project under the name "Rwanda Content Hosting"³. The Internet Society provided technical expertise for the analysis of the impact of the project.

The first Internet Society paper on local hosting established the following observations:

Costs

- International hosting is likely to be less expensive than local hosting
- However, international transit is more expensive than local connectivity
- Choosing to host overseas for cost reasons can increase the costs for ISPs to bring the traffic back to Rwanda

Latency (Delay)

- Hosting distance adds latency; international links can be congested and add more latency – slows the user experience
- Latency can diminish the advantage of high-speed bandwidth
- Usage
- Latency and high costs reduces usage, impacting user experience, ISPs, and content providers

The results of the pilot project described in this report bear out the observations that were established in the first report, highlighting the practical benefits of local content hosting.

² Michael Kende and Karen Rose, "Promoting Local Content Hosting to Develop the Internet Ecosystem", 2015, <http://www.internetsociety.org/sites/default/files/Promoting%20Local%20Content%20Hosting%20to%20Develop%20the%20Internet%20Ecosystem.pdf>

³ RICTA has also released a separate report on this study. "Rwanda Web Hosting Project" RICTA Management, <http://www.ricta.org.rw>

The structure of the market in Rwanda is such that, generally, developers provide a full range of services for their clients – they develop the websites, they resell hosting, and they provide the domain names as registrars who sell the ccTLD (here notably .rw and .co.rw). For website hosting they generally turn to a large US-based hosting provider, such as GoDaddy, and for website design, they often use a Content Management System (CMS) such as WordPress.

Various reasons are given by developers for why hosting in the US or in Europe is preferred to hosting in Rwanda. Primarily, the monthly cost of hosting is lower, especially in the United States. Indeed, hosting in the United States is the most common option, with over half of all .rw and .co.rw websites being hosted there⁴. In addition to this, in interviews with developers, historic doubts about reliability and uptime of hosting in Rwanda were mentioned.

The issue of reliability of local hosting is something that is frequently mentioned in countries like Rwanda that rely on overseas hosting. As a precaution, countries that wish to improve or begin local hosting should focus on removing any unfounded fears of reliability. We note that there were no issues with uptime for hosting in Rwanda during the year-long pilot phase. Favourable uptime information along with the advantage of faster loading times can help to persuade website owners to use local hosting platforms.

The study finds that moving websites to Rwanda dramatically lowered the latency and loading times. As a result, websites that moved from abroad to Rwanda improved their visitor engagement in various metrics such as return visits, pages viewed and overall number of visitors. This is likely to be a result of the decreased loading time of the websites. In addition, website developers have more business opportunities by being able to act as a reseller of hosting capacity and having the opportunity for greater skill development. The developers also expressed to us a wish that they had learned about the benefits of local hosting sooner.

We note that, with its tourism orientation, Rwanda has a relatively high number of websites under the ccTLD (.rw) that do not target the domestic market. For instance, there are websites that advertise hotels and safaris to foreign audiences. The majority of these visitors tend to be in the US or in Europe. For tourism-related websites, it is therefore best to remain hosted in the US or Europe, near the target audience.

The main problem that came forward during the pilot phase of the project was the issue of developers using outdated software on their websites, which was exposed when the websites migrated to the local server. This weakness enabled outsiders to send spam messages that appeared to come from the migrated websites, which eventually led to one of the three servers being blocked by companies such as Google, which classified all email originating from this IP address as spam in its Gmail service. This issue, including the prevention measures taken by RICTA to solve the problem, is discussed in detail in a box in section 3.4.

⁴ Dyn measurements are presented in the results section.

It should also be noted that although there are significant benefits for website owners to move their sites back to their home country, these should not be used to justify forced data localisation, be it for economic, security, or other reasons. Attempts by government to use data localisation requirements to exercise greater control over information often backfire, as discussed in the box.

Negative Consequences of Data Localisation Requirements

Although voluntary data localisation can have a variety of positive effects, the involuntary localisation of commercial data has adverse effects, which make such policies ineffective and undesirable.

Attempts to force data localisation by governments, sometimes justified with security considerations, or for economic reasons, are counterproductive for several reasons:

- Draws negative international attention (Streisand Effect).
- Raises costs for international providers, who may choose to base their operations elsewhere.
- Potential negative consequences for economic growth.⁵
- Does not necessarily provide any real security, as concentration of data is more prone to security attacks.

An extensive discussion of the negative consequences of forced localisation can be found in the background paper by the World Economic Forum.⁶

Local content hosting is important because of the direct and measurable impacts on user engagement with local websites, which we document here. More broadly, however, local content hosting helps to develop vibrant data centres and web hosting providers, which in turn help to support the development of more local content and services. The availability of this local content brings new users online, and also contributes to jobs and revenues for local entrepreneurs. As such, governments should see local content hosting as a key ingredient in developing a local Internet economy.

The next section of this report offers a brief overview of the methods used to collect the information referenced in this report along with the presentation of some of the key results, in Section 2. Section 3 presents some conclusions, Section 4 highlights the main learnings, and Section 5 provides recommendations and a roadmap based on these results.

5 <http://ecipe.org/blog/the-costs-of-data-localization/>

6 http://www3.weforum.org/docs/Background_Paper_Forum_workshop%2009.2016.pdf

2. Results

We first detail the data that was gathered, and then the results.

2.1 Research method and data

For this pilot project, servers were set up by RICTA through a sponsorship by BSC Ltd. in the Telecom House, a key hub for telecommunications in Rwanda. Three Rwandan web hosting providers, each with an existing portfolio of websites, volunteered to take part in the pilot. Each was provided with server capacity, in the form of three virtual private servers (VPSs). The server capacity was used to migrate a selection of websites previously hosted in the US (and in some cases Europe) to Rwanda. An additional VPS was created to enable Dyn⁷ to conduct latency measurements from within Rwanda.

There are two main forms of metrics that are used to examine the impact of website migration.

First, latency measurements were gathered, using a Rwanda-based node operated by Dyn. This allowed Dyn to test the access time to all websites using a Rwandan extension, such as .rw, .co.rw, or .org.rw, from a computer based in Rwanda. To do this, Dyn first had to locate where each website was hosted, and thus could tally where websites were based throughout the pilot project.⁸

Second, using web analytics tools Google Analytics⁹ and AWStats¹⁰, it was possible to track the number of user visits to each website, as well as a number of related metrics, such as: number of pages viewed, return visits (an indication of positive perception of the website), and total time spent on the website. These metrics were collected in anticipation of the stronger engagement with these websites, as a result of reduced loading times for users in Rwanda.

Using a combination of the above-mentioned visitor statistics and the latency measurements, this study provides an overview of the impacts of the migration of websites to Rwanda. These results are presented in the next sub-section.

⁷ Dyn is an Internet Performance Management Company measuring online infrastructure.

⁸ Latency from Rwanda is measured using a VPS operated by DYN from the same facility where the websites are hosted (Telecom House) and the connection is established through an IXP also hosted there. Presumably the latency with other facilities in Kigali or Rwanda is slightly larger, though still orders of magnitude smaller than with US-hosted websites.

⁹ Google Analytics was collected by adding a small snippet of code to every page on a website, every time this page is loaded on a user's computer, this sends a message to the Google server, which records visit and related metrics, such as mentioned above.

¹⁰ AWStats provides similar metrics, such as user visits, but uses a different data collection method. It uses the webserver's information logs, in order to compile reports on visits and other metrics.

2.2 Results

This section provides an overview of the results of this study. Section 2.2.1 compares the hosting location for all .rw and .co.rw domain websites, which was measured using Dyn’s infrastructure. Section 2.2.2 presents the results of latency measurements conducted by Dyn from the newly installed node in Kigali. Visitor engagement results, such as page views, bounce rates, and return visits for selected websites are presented in section 2.2.3. Finally, in sections 2.2.4 and 2.2.5 the outcomes of the interviews with developers and the hosting provider (RICTA) are discussed.

2.2.1 Hosting Location

As mentioned in the introduction, the migration of websites from foreign hosting entails moving from a hosting provider located abroad to a hosting provider based in Rwanda. The below tables 1 and 2 present the physical hosting location for websites with a Rwandan ccTLD, broken down into .co.rw and .rw.

| Table 1: .rw Websites Location by Share | | | |
|---|------------------------|------------------------|-------------------------|
| Region/Country | March 2016 (pre-pilot) | September 2016 (pilot) | March 2017 (post-pilot) |
| Africa | 13.8% | 16.5% | 17.0% |
| Europe | 18.9% | 19.9% | 19.7% |
| North America | 61.2% | 59.7% | 57.8% |
| Rwanda | 11.2% | 13.9% | 14.9% |
| United States | 60.8% | 59.4% | 57.3% |
| United Kingdom | 7.8% | 8.6% | 7.5% |
| France | 3.5% | 3.2% | 4.5% |
| Germany | 2.7% | 2.9% | 2.3% |

Table 1 presents the results of the location analysis for websites that use the general Rwanda ccTLD, .rw, which is also the most popular Rwandan ccTLD. As can be seen from the row in bold, the percentage of websites that use the .rw domain hosted in Rwanda before the start of the pilot in March 2016 was 11.2%. This number quickly rose to 13.9% in September. In the last measurement, the percentage of websites using the .rw extension exceeded the pre-pilot share, with a new total share of 14.9%. This is significant growth, albeit from a low base.

Overall, hosting in Africa reflected this upward trend, starting at a value of 13.8%, before the start of the pilot in March 2016. This value also rose quickly, to 16.5% in September 2016. During the last measurement, conducted in March 2017, the share rose further to 17%. This positive trend in local and regional hosting represents more websites than those involved in the pilot, and thus hopefully points to increased awareness of the benefits of local hosting, which this paper hopes to build upon.

| Table 2: .co.rw Website Location by Share | | | |
|---|------------------------|------------------------|-------------------------|
| Region/Country | March 2016 (pre-pilot) | September 2016 (pilot) | March 2017 (post-pilot) |
| Africa | 26.2% | 24.7% | 23.6% |
| Europe | 24.3% | 23.8% | 21.6^ |
| North America | 48.3% | 49.8% | 53% |
| Germany | 1.4% | 2.4% | 2.3% |
| France | 7.8% | 8.6% | 7.2% |
| United Kingdom | 13.1% | 9.5% | 10.0% |
| Kenya | 5.9% | 3.7% | 3.1% |
| Rwanda | 18.7% | 19.8% | 18.7% |
| United States | 46.8% | 48.6% | 50.7% |

As Table 2 shows, the trend in hosting location of Rwandan websites using the commercial second level domain, .co.rw, which is less popular than the .rw ccTLD, is somewhat different from the .rw trend. For the .co.rw websites being hosted in Rwanda, the pre-pilot measurement from March 2016, was 18.7%. In the post-pilot measurement from September 2016, the value for Rwanda had gone up somewhat, to 19.8%. In the final measurement, taken in March 2017, the value for Rwanda had fallen back to its original level of 18.7%.

In Africa as a whole, there was a general downward trend in hosting of .co.rw websites, with a decrease from 26.2% in the March 2016 pre-pilot measurement, to a 24.7% value in the September 2016 measurement. The final value recorded in March 2017 was 23.6%, showing a consistent downward trend. At the same time, the share hosted in the United States increased – showing that there is still work to be done in terms of increasing the number of websites hosted locally.

Overall, these values all suggest a general upward trend for regional or local hosting, with the more popular .rw increasing hosting in both Africa in general and Rwanda in particular, and .co.rw decreasing in Africa while remaining stable in Rwanda. It should be noted that the total number of websites using the Rwandan ccTLD rose during the period in which measurements were taken, from 2300 at the beginning, to 2485 at the end, suggesting that the percentage gains in local hosting represent an even larger gain in total numbers.

2.2.2 Latency Measurements

One of the key benefits for website owners of hosting a website locally is the fact that loading times are much lower. In particular, modern websites have a high number of interactions between the server and the visitors, during which the website cannot yet be displayed. It is therefore very beneficial to have low latency in the communication between the server and the visitor. To measure these latencies, a Dyn node was installed in one of the VPSs on the hosting server in Kigali, to conduct latency measurements from Rwanda to Rwandan websites, hosted in Rwanda and internationally.

The below table shows the latency measurement for websites with a .rw or .co.rw domain, plus a number of Rwandan websites that use other TLDs, such as .com or .org. Websites hosted in the United States or Europe generally have a latency upward of 200 milliseconds from Rwanda, with websites hosted in Rwanda having a very low latency, below five milliseconds, as seen in the table below.

As mentioned above, the Dyn node for latency measurements was hosted on the same server as the VPSs provided to the developers. This server was located in Telecom House, where the Rwanda IXP operated by RICTA is also located. This tends to decrease the latency measurements, and indeed several websites have very low latencies of less than one millisecond, including the websites hosted on the developers' VPSs.

| Latency | March 2016 (pre-pilot) | September 2016 (pilot) | March 2017 (post-pilot) |
|---------|------------------------|------------------------|-------------------------|
| >300 MS | 20 | 9 | 117 |
| >200 MS | 416 | 496 | 431 |
| >100 MS | 122 | 82 | 138 |
| >50 MS | 659 | 615 | 638 |
| <50 MS | 760 | 669 | 917 |
| <10 MS | 66 | 181 | 48 |
| <5 MS | 13 | 14 | 77 |
| <1 MS | 1 | 1 | 0 |

As can be seen from the table, the number of websites with very low latency measurements of less than 5 milliseconds went up from 13 to 77. At the same time, we see the number of websites with latency measurements upward of 200 milliseconds remained about level. This reflects the results from the previous section, that the number of Rwanda-oriented websites hosted in Rwanda increased.

2.2.3 Visitor Engagement

As mentioned in the introduction and methodology sections, in order to gauge the effect of changing hosting location on visitor engagement, we monitored a number of typical websites using Google Analytics and AWStats. The expectation was the faster loading pages foster improved visitor engagement.¹² This expectation is borne out in the following case studies.

¹¹ Note that this does not measure the total number of domains registered, but rather those for which latency measurements were available.

¹² <https://www.ampproject.org/latest/blog/new-industry-benchmarks-for-mobile-page-speed/>

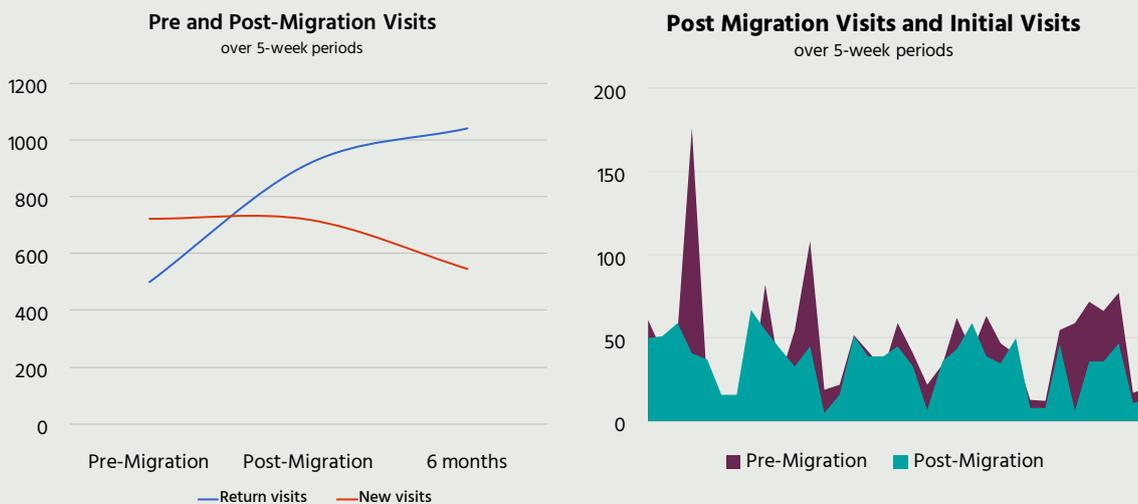
Case Study: Professional Association

Our first case study is a professional association that provides members with updates on events and developments in the industry. The fact that the website targets paying members of an association and provides them with updates relevant to their profession suggests that visitors would be relatively tolerant even to slow loading times, as compared to competitive news or entertainment pages, for instance. This would suggest that even a dramatic increase in speed might not necessarily improve engagement, but that turns out not to be the case.

In the first latency test by Dyn, prior to the migration, this website had a latency time of 269.901 milliseconds (about 1/3 of a second) measured from Rwanda. It is important to note that this is not the actual loading time of a website, it is merely the time it takes to get a response from the server. The actual loading of the website can involve about 5-8 sequential interactions, plus the time for actual data transfer, so that the actual load time experienced by a user could have been significantly longer. After the migration of the website, latency was measured again, at which point it was measured at less than one millisecond (1/1000th of a second).

We compared the impact of the latency on the number of visits, and included both new visits to the website, and return visits. We compared these numbers in the five weeks directly prior to the migration with the five weeks directly after the migration, followed by a final analysis of a five-week period, six months after the migration. This shows the following results:

Figure 1. Professional Association Website (initial 5 weeks vs. final 5 weeks)



The above figures show two key results. First, in the top graph the number of return visits – one of the most important indicators of visitor engagement – saw an immediate and substantial increase after the migration. When returning to this statistic six months after the migration, this number was even higher. While the number of new visitors declined, this could be because there are a limited number of members of this profession, so the number of potential new visitors declined and more became return visitors. This is confirmed in the bottom graph, which overlays the number of daily visits before and after the migration, over a five-week period. The comparison shows that the post-migration visits outnumber the pre-migration visits almost every day over the five-week period, showing a total increase in visitors.

Overall, this is a very encouraging result, with overall visits increasing, mostly as a result of increased return visits. This is all the more encouraging because the target audience is 'captive' in the sense that the website is important for their profession, making them potentially more tolerant of delays in loading the website. The results show that they nevertheless responded very positively to the reduced latency.

Case Study: Cultural Association

Our second case study is a cultural association. The latency of this website – which was hosted in the United States – prior to the migration was recorded at 306.477 milliseconds (about 1/3 of a second). After the migration to Rwanda, the latency was also recorded at less than one millisecond (1/1000th of a second).

In this analysis, we compare three weeklong periods, once prior to the migration, once directly after the migration and once again six months after the migration. The top graph shows the number of visits in each period – prior to the migration, right after, and six months later. The bottom graph shows how these visits divided between new visits and return visits.

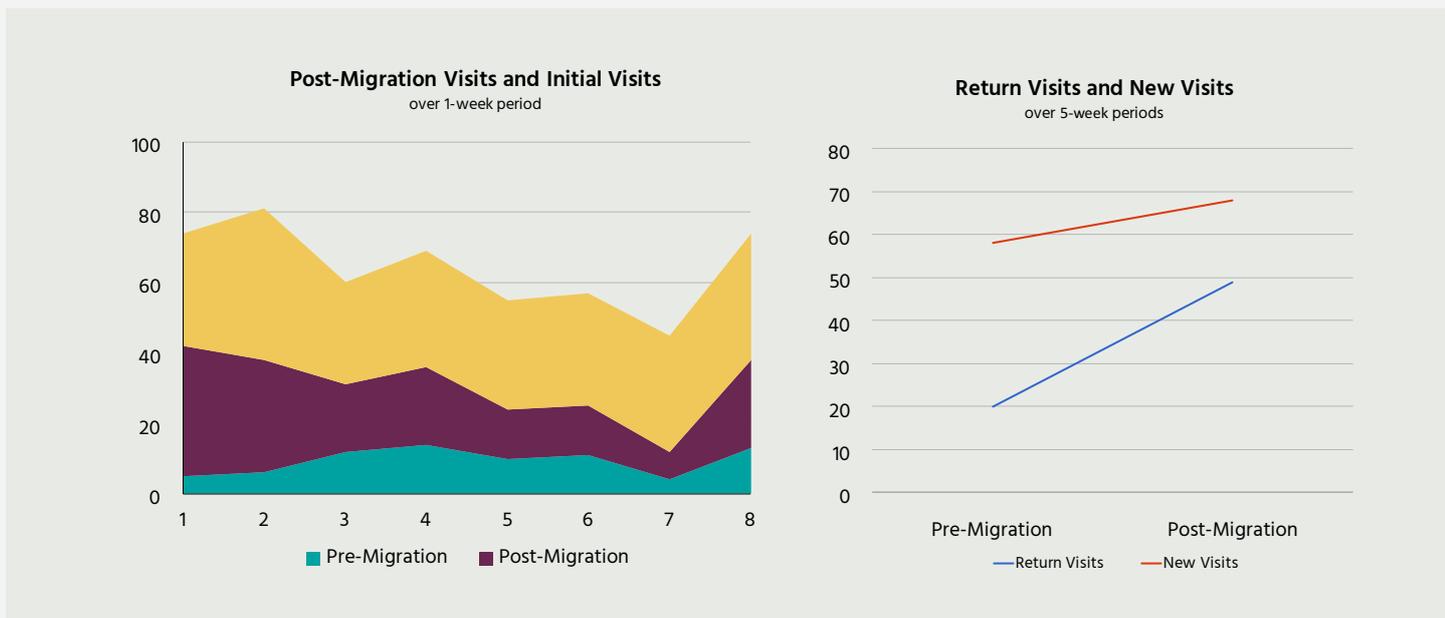


Figure 2. Cultural Association (visits per day over a one-week period)

The key result to observe here is in the overlay on the top graph, showing that after the migration, daily visits were higher every single day of the week. However, even more interestingly, this upward trend continued, with visits still increasing six months after the migration. Additionally, daily visits from this period suggest that the visits are more stable day-to-day. When comparing the new and returning visitors in the bottom graph, both increased, with return visits rising faster. This is another indication that visitor engagement has improved.

2.2.4 Developer Impressions

The overall impression of the developers who migrated their websites was very positive. Crucially, in interviews after the conclusion of the initial phase, the developers reported no issues with downtime, which goes against the prevailing belief among both developers and customers in Rwanda alike, that hosting overseas is more reliable than hosting in Rwanda.

Moreover, two of the three developers reported no issues with the local hosting capacity with which they were furnished. These developers initially migrated a small number of low-volume websites from servers located in the United States, to the Virtual Private Server (VPS) in Rwanda. In order to break the migration down into small steps, many developers chose to first maintain the original copy of the website on the server in the United States.¹³

In addition to the above mentioned impacts, such as reduced loading times and greater visitor engagement, developers reported a number of other benefits relating to the nature of their role as a reseller and server administrator. Prior to managing a VPS in Rwanda, developers would purchase shared webhosting packages from US-based hosts, which both limits developer freedom and becomes costlier as more websites are hosted. One developer reported about the migration:

“It’s made me become a better admin. Usually everything was taken care of by Hostgator. All we needed to do was open up a chat session and voila, the issue is sorted. No skills gained there.”

By being trained in the server management side of webhosting, developers came to understand more of the whole process, which includes a critical understanding of the need to run security updates on websites.

In addition to this, the developers mentioned that controlling the VPS in Rwanda increases the business potential to act as a reseller, by being able to fully utilise the capacity available. In the previous setting, the webhosting packages were limited to one single website per package. As such, unused capacity could not be repurposed to add an additional website. In Rwanda, leftover capacity can be used for other websites, making the hosting more efficient.

An additional benefit that developers mentioned was the ability to meet face-to-face with the server administrators, which was not possible when the websites were hosted overseas. Developers found that this made it easier to address issues that arose. This suggests that local hosting links together new components of the Internet ecosystem in Rwanda, completing the connection from registry, to hosting, to developers.

Unfortunately, one of the developers did not continue his use of the local VPS because Google blacklisted the server’s IP address from sending emails. The reason for this blacklisting was the server was sending out spam messages from third parties. The Internet Society’s Africa Bureau stepped in and helped solve the problem, but that developer did not continue the migration after the problem was solved.

¹³ By updating the domain settings, all web traffic was directed to the new Rwandan server. All other forms of traffic, such as email, was still directed to the original server in the United States. This also made sure that, although low volume websites were used, there was no risk of customers’ critical information being lost. Eventually, after successful testing of the website on the Rwandan server, all traffic including email was eventually directed to the Rwandan server.

The conclusion from this is that security issues should be addressed as part of the initial setup of the VPSs. While this includes addressing the spam issues that were identified and resolved in this pilot study, similar vulnerabilities can expose the websites to other malicious attacks such as modifications of content and breach of any sensitive data. The technicalities of the spam issue are discussed in the box below.

Spam

Security was highlighted as a particular issue, after the migration of outdated Content Management Systems (CMSs) to the Rwandan VPSs. These outdated CMSs had known vulnerabilities, which led to the local VPSs being used by botnets for spam distribution. When this happened to one VPS, it led to the blacklisting of its IP address.

Before the migration, the developers hosted their websites on very large-scale hosting services based in the United States or Europe. These hosting providers have a number of systems in place that mitigate common security issues. In particular, all outgoing email is routed through a spam filter. When the developers migrated their websites to the VPSs in Rwanda, such a filter was not yet in place. Since some of the Content Management Systems were outdated, hackers could use known vulnerabilities in the CMSs in order to compromise them.

Once compromised, malicious scripts were installed, which used the server's email service to send out spam to a large number of email addresses. This led large email providers such as Google, Yahoo, and Microsoft to block the IP address of the VPS, which also led to the blocking of legitimate emails sent by the VPS. This problem was solved by routing all outgoing email through a spam filter, as well as removing and periodically scanning for malicious scripts.

To prevent this problem from occurring, a number of steps to improve the CPanel system should be taken. This process is known as hardening. In addition to this, it is important to train developers about the problems that could occur when systems are not updated.

Lastly, it is advisable to use the automatic updating of CMSs through systems such as the CPanel plugin Installatron or the open-source Bitnami (example email from an automatic update included below). The Bitnami toolbox has the advantage that it is open source and therefore free to use. Bitnami is also used by large organisations such as the United Nations.

Example of automatically updated Drupal CMS by Installatron

This is an automated email from Installatron. To unsubscribe from these emails or to change notification settings, login to your web hosting control panel, navigate to the Installatron tool, and select the installed applications you wish to modify.

The update process for OpenPlatforms.org located at <http://www.openplatforms.org> has succeeded.

New installed version: 8.1.10

End of report.

Still, clients of web hosting providers are reported to have mixed responses to the idea of hosting in Rwanda. Roughly half of the clients had concerns about the reliability of the hosting capacity in Rwanda. The other half were reported to be enthusiastic about the idea and are keen to migrate their websites to Rwanda soon. This indicates that informing clients about the proven reliability of hosting in Rwanda acts as an important demand driver.

One of the developers summed up the migration as successful, but he wished he had known about it earlier.

2.2.5 Host Impressions

The server that the developers used in Rwanda was provided by RICTA through a sponsorship of BSC Ltd.. The general impression from the server administrators at RICTA was that the project went well, provided significant lessons, and is ready to move from the pilot phase to a more wide-scale migration.

Their impression was that this was a large project to undertake for a relatively small organisation, but there was a lot of domestic interest in pursuing this. Although Rwanda is probably ahead of most other countries in the region, in terms of its Internet development strategy, many neighbouring countries are also showing interest in adopting similar strategies. As the regional market grows based on the success of the East Africa Community Integration initiative and other integration efforts in the region, this will provide greater economies of scale for data centres, more competitive hosting offers will become available, and local hosting will be yet more attractive.

The existing local hosting project in Rwanda can fulfil a key role here. One of the barriers reported by RICTA was that, at the time of starting this project, it was something that nobody in the region was familiar with, since no similar projects existed. The general view of the Internet ecosystem was that website hosting was something done abroad. In this sense, the Rwanda local hosting project functions well as a proof of concept. Additionally, there is new capacity developed with the RICTA server administrators, as well as with the developers, who developed expertise in terms of managing their own Virtual Private Servers.

The RICTA administrators did report that the spam issue was a big obstacle in the overall course of the project. The issue reaffirmed preconceived ideas with developers that the capacity for hosting simply was not available in Rwanda, even though the problem was overcome and that no further problems with reliability occurred. RICTA reported one aspect initially overlooked by developers was that the eventual goal would be to migrate entire websites, including email traffic. For future projects, this should be made clear from the outset.

A final comment was that the hardest obstacle was overcoming the scepticism towards hosting in Rwanda, such that any similar project should start with information and education for all stakeholders in the Internet ecosystem, including developers and their clients. This can avoid scepticism during the initial project phase, which could lead to half-hearted attempts at migration, which RICTA said are never likely to be successful. Having a full understanding of the ability of RICTA to manage such a project would help developers to fully commit to the project, which would make migration more successful.

Any similar future projects in the region can use the successful migration in Rwanda as a reference and draw on the newly developed capacity. We hope that this paper helps to make the case for local hosting more compelling.

3. Conclusions

The Rwanda local content hosting project is part of a larger push in Rwanda to serve more content to Internet users from within the country, to build on the benefits of the local Google Global Cache and Akamai node by increasing the amount of local content hosted locally. Generally, there are several reasons for wanting to serve more content, in particular, local websites, from domestic servers.

This report shows that locally hosting websites dramatically reduces the latency of a website when accessing it from Rwanda. This implies shorter loading times for domestically hosted websites, which help websites achieve greater visitor engagement, such as more page views, longer time spent on the website, and more return visits. This was observed for a number of different types of websites. Additionally, it enables more interactive low-latency content on websites, such as videos or video games, which cannot be achieved with the high latency that is observed when hosting websites abroad.

The results show how the discussion for local hosting needs to move beyond the mere cost of hosting locally. In order to overcome the concerns over costs, website developers and hosting providers should convey to their clients the greater quality of the hosting that they can provide from Rwanda, which will allow them to pass on any higher cost of hosting locally. In order to succinctly convey this idea to clients, this can be broken down as follows:

Closer ▶ Faster ▶ Happier ▶ Rewards

Closer: The websites are physically closer to the user.

Faster: The proximity enables faster loading times and more types of content.

Happier: Visitors will be happier with the website experience.

Rewards: This will lead to more return visits and more page views per visit.

The example that the first movers set, such as in the pilot in Rwanda, can be used to encourage more localising of websites, creating a virtuous circle.

In terms of costs, however, the reliability and cost of electricity remains an issue that will continue to keep hosting in Africa expensive. Particularly when power outages are frequent, data centres need to rely extensively on generators for continued uptime, which adds costs. While this challenge also creates an opportunity for Africa to leap frog traditional power sources with green data centres and other innovations, in the short run it represents an increase in cost for locally hosted data.

On the other hand, as more traffic originates from local servers, this allows ISPs and telcos to fulfil demand with less capacity on international cables, thereby reducing costs. Ideally, savings in the purchase of international bandwidth would offset the higher electricity costs of local hosting, thereby lowering the cost of access. In addition to the increased content relevance¹⁴ from the locally hosted websites, the reduced cost of access will translate to greater Internet adoption, enhancing the virtuous circle of a greater number of users and a greater amount of relevant content. As shown in the box, there are other means by which local hosting costs could be decreased.

¹⁴ For more information on content relevance see: Promoting Content in Africa <http://www.internetsociety.org/PromotingContentAfrica>

Software Licence Price Differentiation: Lessons from the Markets for Medicine

Software licenses for applications such as web hosting control panels (such as cPanel, DirectAdmin, Webmin) constitute a substantial part of the cost of a standard Virtual Private Server. A typical VPS configuration costs about USD120-160 a month, with software licence fees – mostly for the control panel – contributing about USD20.

At the moment, the price of these software licences is not differentiated on the basis of destination country by most vendors. This seems to be because the majority of the hosting takes place in developed economies.

The cost of software is a fixed one, incurred almost entirely in development, with the cost of making an additional copy being close to zero. As such, software is similar to medicine, where cost is largely based on development and intellectual property.

In the markets for medicine, such as for HIV, it took outside intervention and market shaping from various UN organisations for manufacturers to realise the potential for country-based (or region-based) price differentiation (third degree price differentiation)¹⁵

While medicine can be resold back to developed countries where it costs more, web hosting related software is inherently well suited for location-based price differentiation, as it is tied to an IP address, which can be geolocated. Therefore, there is great potential for country-specific pricing for server software such as the control panel, as well as other tools such as install scripts.

As most of the software in the server market is run on top of the Linux operating system, a logical focal point for the organising of such an initiative would be the Linux Foundation, which frequently coordinates between software developers. This would enable individual software vendors, such as the cPanel control panel, to provide discounted access to a full software stack required for server management.

Alternatively, open source software packages can also be used to replace commercial solutions. In the case of the web hosting control panel suites, open source alternatives include: Webmin¹⁶, Froxlor¹⁷, or CentOS Webpanel (free but not open source, only for CentOS)¹⁸.

Local content hosting is an important component in improving the local content ecosystem, given its location in the value chain between the national registry and the website developers. In the Internet Society's "Promoting Content in Africa"¹⁹ report released in 2016, we discuss a number of barriers to the emergence of a larger local content ecosystem. Having more local content will help incentivise more people to start using the Internet, which increases the audience and potential revenue for local content creators and hosts. This, in turn, contributes to developing a vibrant local Internet economy.

¹⁵ For more information on content relevance see: Promoting Content in Africa <http://www.internetsociety.org/PromotingContentAfrica>

¹⁶ <http://apps.who.int/medicinedocs/en/d/Js18253en/>

¹⁷ <http://www.webmin.com/>

¹⁸ <http://www.froxlor.org/>

¹⁹ Michael Kende and Bastiaan Quast, 2016: <http://www.internetsociety.org/doc/promoting-content-africa>

4. Lessons learned

The following lessons were learned and are relevant for different stakeholders and issues.

Website providers

- Localisation reduces the latency of websites by orders of magnitude.
- Reduced loading times of websites has a positive impact on website engagement in metrics such as page views, time spent, and return visits.
- Quality of local hosting may be as high as abroad, and should not be dismissed without investigation.
- Websites focused on international users, such as tourists, should tailor their approach to their audience, perhaps using a CDN to host content in multiple locations and serving it from a location closest to the user.

Developers

- Local content hosting enables developers to learn new skills and become better admins.
- Developers can expand their role as resellers of web hosting, and make full use of the capacity available on a VPS, thereby reducing the cost of local hosting.
- Given the doubts about reliability, it is vital to avoid any downtime on the server.
- By installing CMSs using new plugins, websites can easily be migrated from remote servers, using the built-in export-to-file function on the existing website and importing these files on the new local installation of the CMS.

Security

- Developers should be aware that outdated CMSs have known exploits that will inevitably be exploited by botnets and how this can have significant impacts, such as blacklisting of the IP address.
- CMSs should be made easy to install and update, using software such as Installatron or Bitnami, making the process more accessible to website developers and hosting providers, and provide better overall security to avoid attacks such as the outgoing spam that occurred in the pilot project.
- To prevent spam, an outgoing spam filter is needed to make sure websites compromised by botnets do not get their IP address blacklisted.
- Developers should receive training on how to update CMSs with security updates and eventually migrate to new major versions when support for the previous versions expires.

Local Internet Ecosystem

- A greater proportion of websites hosted domestically might eventually reduce the cost of access for consumers, improving Internet adoption.
- Having a local-hosting infrastructure creates high-quality jobs and expands the local Internet ecosystem, providing further opportunities for innovation, including developing green data centres to lower costs.
- Occupying a position between the national domain registry and the website developers, local website hosting closes a missing link in a national Internet ecosystem.
- Migration only works well when it is voluntary.
- Education of all stakeholders can help address uncertainties regarding reliability – this should be done even before the start of the migration.
- Certain components of the server software stack are not open source and therefore require licensing fees. Software suites for web hosting should lower the cost for local hosting by using location-based price differentiation.

5. Recommendations and Roadmap

Based on our work on local content hosting in Rwanda, here is a roadmap of recommendations for any country seeking to increase the number of websites hosted locally.

The first step should be to determine what the reasons are for hosting websites abroad. In Rwanda, as the Internet Society began its work on this project, the Minister of Youth and ICT convened a workshop to help answer this question. This workshop involved all the relevant stakeholders: ISPs, content providers, website developers, web hosting providers, data centres, and others from government and the technical community, who might not otherwise come together to discuss issues around local content and hosting.

The workshop provided an opportunity to discuss the benefits of local content hosting in terms of lower latency and cost of traffic; it also provided various stakeholders the opportunity for an open discussion about the potential impediments of local hosting. We learned that some web hosting providers had not heard about recent data centre price reductions for hosting, while others raised issues about the reliability of the data centres. Such discussions further provide room to learn whether there are questions about content regulations, skills development and other issues.

A workshop such as the one convened in Rwanda is likely to raise a number of actual issues, as well as perceived concerns, and provides an opportunity to address these. For instance, content providers may not be clear on current content regulations, which the government stakeholders can address (we note this was not an issue in Rwanda, but may be elsewhere). Likewise, web hosting providers may base their hosting decisions on outdated prices charged by data centres, which the data centres can address.

Other concerns may be more difficult to address, however. For instance, costs in local data centres may be higher because of higher energy prices, and the need for backup power, as well as lower economies of scale. Staff at the local data centres may not have the level of training common in larger data centres overseas, and it may be difficult to assess the reliability of the data centres.

A pilot project, as was undertaken in Rwanda, may be one way to address particular concerns, while also identifying others, such as the issue that arose with spam described above. Other countries may not be able to run such a pilot, however, but hopefully can still benefit from the lessons of the Rwandan experience detailed here.

Based on these considerations, our recommendations for helping to create an enabling environment for local content hosting are as follows:

All stakeholders

First, it is important to have an awareness of the issues, starting with the benefits of local content hosting described here in this paper.

Second, it is important to have a dialogue between stakeholders, such as the local content workshop described above.

Third, it is important to present data and information about key issues, including content regulations (or lack thereof), and cost and reliability of local data centres.

Government

To help to lower the cost of creating and running data centres, tax breaks on equipment imports and energy costs could play an important role in helping to match the prices of overseas data centres, particularly before the data centres have significant economies of scale.

Understanding content regulations is also important, so that content providers have certainty about the rules that may apply to them. This includes, in particular, issues of intermediary liability – it is important for developers and hosting providers to know whether they are liable for their clients' content.

Training and capacity development is also important for content providers, website developers, hosting providers and data centre operators, to help develop a vibrant local content ecosystem.

Developing a regional market, such as for instance building on the regional integration being fostered by the Northern Corridor initiative in East Africa, can help to develop scale for data centres, to increase investment and lower the cost of hosting.

Businesses

Data centres should ensure that their pricing is as competitive and transparent as possible to help attract customers, and provide information on their uptime reliability and security levels, using third-party certification if possible.

ISPs are the indirect beneficiaries of local content hosting, because it saves them buying expensive international capacity. While they do not make the decisions on where content is hosted, they can play an important role in helping to raise awareness, and could also go so far as to offer to subsidize short trial periods of local hosting (particularly if, as is often the case, they own their own data centre – the savings on international capacity could be used to subsidize local hosting).

Website developers and hosting providers should understand the benefits of local content hosting, and also ensure that the Content Management Systems underlying the websites they develop are kept fully up to date, using upgrades if possible, to protect against spam or other attacks. Automated systems such as Installatron or the open-source Bitnami can help with this.

Finally, content providers should take an active role in hosting decisions because they are the ultimate beneficiaries of local content hosting, in terms of increased page views and customer engagement.



**Internet
Society**

internetsociety.org

Galerie Jean-Malbuisson 15,
CH-1204 Geneva, Switzerland
Tel +41 22 807 1444

1775 Wiehle Avenue, Suite 201
Reston, Virginia 20190, U.S.A.
Tel +1 703 439 2120