Bandwidth Management and Monitoring for Community Networks

Barry Irwin, Ingrid Siebörger, Daniel Wells
Department of Computer Science
Rhodes University, Grahamstown, South Africa
Tel: 0466038291 Fax: 0466361915
email: {b.irwin,i.sieborger}@ru.ac.za, wellsdd@telkom.co.za

Abstract— **This paper describes a custom-built system to replace existing routing solutions within an identified community network. The community network in question shares a VSAT Internet connection to provide Internet access to a number of schools and their surrounding communities. This connection provides a limited resource which needs to be managed in order to ensure equitable use by members of the community. The community network originally lacked any form of bandwidth management or monitoring which often resulted in unfair use and abuse. The solution implemented is based on a client-server architecture. The Community Access Points (CAPs) are the client components which are located at each school; providing the computers and servers with access to the rest of the community network and the Internet. These nodes also perform a number of monitoring tasks for the computers at the schools. The server component is the Access Concentrator (AC) and connects the CAPs together using encrypted and authenticated PPPoE tunnels. The AC performs several additional monitoring functions, both on the individual links and on the upstream Internet connection. The AC provides a means of effectively and centrally managing and allocating Internet bandwidth between the schools. The system that was developed has a number of features, including Quality of Service adjustments limiting network usage and fairly billing each school for their Internet use. The system provides an effective means for sharing bandwidth between users in a community network.**[1]

I. Introduction

Community networks, in South Africa, are often serviced by limited capacity network back-haul connections to the Internet. This component is a relatively significant ongoing expense incurred in their operation. This is particularly so in comparison to the basic cost of living. Even in cases where this back-haul is financed or donated by other parties, it is still a finite resource. The bandwidth needs to be managed, monitored and allocated equitably among the users within the community.

In this paper, the focus of the monitoring and management is on bandwidth usage and throughput, but also monitoring network quality. Network quality can be understood as the degree of satisfaction that the user experiences when making use of the network and its available services. Monitoring reveals which services are used most often and what bandwidth and throughput those services require to make the users' experience more valuable. Monitoring produces statistics which informs the administrator how to best apportion the available bandwidth and throughput between services and at which times certain services should receive preference over others.

[1]This work was undertaken in the Distributed Multimedia CoE at Rhodes University, with financial support from Telkom SA, Comverse, Stortech, Tellabs, Amatole Telecommunications Services, Bright Ideas 39 and THRIP.

In this paper we present a custom-developed system, which can fit into existing community networks and has been constructed to effectively manage and monitor bandwidth usage. The system integrates a number of Open Source Software (OSS) applications to provide monitoring and management of bandwidth usage within the community network. The system discussed replaces the network routers throughout the community network with custom-developed routers which have added functionality. The additional functionality includes a caching proxy server to decrease Internet usage for frequently accessed websites; user- and host-based Internet usage quotas which can be employed to limit users' and hosts' Internet use to acceptable levels; and monitoring critical servers or services within the network, allowing the system to pin-point any faults within the community network. A web-based front end to the system was developed, which provides a single portal for administrators to manage their network.

In addition, we describe an example of a community network which was in need of bandwidth management and details the system that was developed. We then describe the results of the system implementation and discusses their implications. The paper is divided into five sections, beginning with section II which discusses related work, and section III introduces the system. Section IV describes the results of the bandwidth management deployment within the community network. Future work building on the results of this deployment is briefly addressed in section V. The paper ends with section VI which presents conclusions that can be drawn from the research.

II. Related Work

In South Africa, as in most of Africa, using or having an Internet connection remains a luxury and is not as widely used as in developed counties [1], [2]. Privately conducted surveys in 2005 [1] and 2007 [2], compared South Africa's telecommunication prices with those of developed countries and an international peer group. The peer countries selected were considered to have advanced telecommunication infrastructure and were similar to South Africa in terms of geographical dispersion of population, income dispersion and market structure. This research revealed some relevant Internet price differences between South Africa and the other 14 countries. South African prices for international leased lines had reduced by 44 % between 2005 and 2007; however, the price remained 404.7 % higher than the average price in the survey. A national leased line in South Africa had reduced in average price from 102 % higher than the other

14 countries in 2005 to 25.5 % higher in 2007. Business broadband (1 Mbps ADSL) was the third most expensive service in its class of the 15 surveyed countries, and it was 127.2 % more expensive than the average price – this had improved from 148 % more expensive in 2005. Retail broadband (512 Kbps ADSL) was the most expensive service in its class when compared to its peer group and was 130.5 % more expensive than the average price, this had also improved from 139 % more expensive than the average in 2005 [1], [2].

Despite high prices, the demand for broadband Internet at the office and at home is increasing. Between 2003 and 2008, ADSL became the primary form of Internet connectivity in Small to Medium Enterprises (SMEs), with dial-up, ISDN and leased lines decreasing (satellite usage remains small, and wireless technologies are slowly increasing with more options becoming available) [3].

Although the supply of bandwidth is improving year by year, there still exists a great demand for bandwidth management due to the expense incurred in obtaining a connection. Internet access prices are likely to drop in the future in South Africa, but solutions are still required to manage the lack of bandwidth in the interim. As a result of a new undersea cable to South Africa which became operational in June 2009 [4], Internet prices are expected to continue to decrease [5] and new opportunities are expected [6]; however, telecommunication infrastructure in and to rural areas is still being developed.

A. Managing Internet Access

When an Internet connection has a limited amount of bandwidth and throughput, controls can be put in place to maximise and prioritise the use of the resource for everyone. The plethora of information available makes it conceivable that controls may be needed to prevent or limit access to certain inappropriate content, although indiscriminate use of filters may result in blocking legitimate sites [7]. When bandwidth is limited, any content that does not match the ideals or the purposes of the Internet connection can be deemed inappropriate. For example, in a school where technology is to be used to facilitate education, access to pornography or gratuitous violence would be considered inappropriate content.

Many websites exist that are bandwidth-intensive, for example the popular video sharing website YouTube and social networking websites Facebook and MySpace [8], [9]. These websites are generally considered recreational websites, but they can have educational merit in specific circumstances. As these types of websites are bandwidth intensive, they may need to be limited or blocked unless a user has a valid reason for accessing them. Without controls on bandwidth-intensive websites, one runs the risk of a "tragedy of the commons" scenario where a few users ruin the experience for the rest by wasting the shared resource [10].

B. Research Network - The Siyakhula Living Lab

The Siyakhula Living Lab (SLL) research network is situated in the Dwesa/Cwebe area, in the Mbashe Municipality in South Africa's Eastern Cape province [11]. This community network lacked a system to manage and monitor bandwidth usage. Figure 1 shows the logical layout of the community network and is referred to in this section. The Community Wide Area Network (C-WAN) incorporates five schools: Mpume, Ngwane, Mthokwane, Nondobo and Nqabara. Each school has its own Community Local Area Network (C-LAN), which connects its own servers and hosts. Each school is connected to a central point (at Ngwane) which houses the WiMAX (IEEE 802.16) base station technology linking the C-WAN together [12]. The C-WAN has a single Internet connection via Telkom VSAT (a satellite-based Internet connection) which is shared between the five schools currently connected.

Each of the schools has a C-LAN router, which is a customised installation of FreeBSD running on a PC. The FreeBSD routers are low end Intel Pentium III PCs, and each router has two network cards:an internal interface and an external interface. The internal interface is connected to the local LAN at the school and the external interface binds an IP address on the open WiMAX IP network. At four of the schools, their local network is connected to their WiMAX customer premises equipment (CPE), while at Ngwane it is connected to the WiMAX base station. The router at each school routes packets intended for other C-LANs or the Internet to the access concentrator at Mpume via a PPPoE tunnel. The access concentrator then routes the packets to the intended C-LAN over the WiMAX network, or to the Internet via the VSAT link.

Schools are connected using a shared medium, thus authentication and encryption of traffic between them is necessary. A PPPoE service runs on the Mpume access concentrator to establish secure and encrypted tunnels over the WiMAX network connecting the school routers and the access concentrator.

The network is not without its flaws. The community network has a single point of failure at the Mpume school: if the access concentrator is down, the other schools lose connectivity with each other and with the Internet. Likewise, the VSAT connection also represents a single point of failure for Internet access, as no alternative route to the Internet is available. The VSAT connection has limited bandwidth available for the community and latency is very high, and it is for this reason that the community network requires measures to be put in place to manage and monitor the use of bandwidth throughout the community.

A bandwidth management system was developed and later deployed in this network. The next section describes the system that was constructed to replace the existing routers within the Dwesa/Cwebe community network to provide the network with effective bandwidth management and monitoring.

III. System Design and Implementation

The system integrates a number of useful network applications and tools into a single and easily-deployable software package. The integration of the applications and tools is designed specifically to aid in managing and monitoring bandwidth within a community network; however, care was taken to allow for further configuration, extension and customisation for other networks.

The design is that of a client-server architecture and as such composed of two main parts. Figure 2 describes a generic community network based on the project's testbed network (described in section II-B). The Community Wide Area Network (C-WAN) has a single Community Network

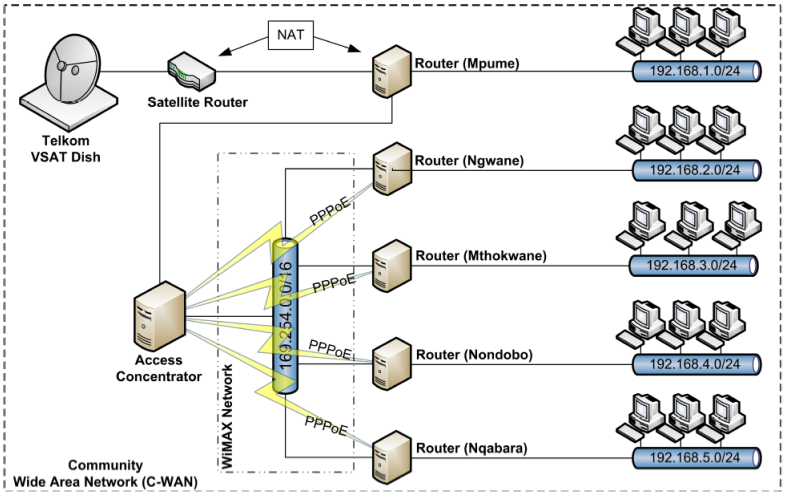

Fig. 1: Dwesa/Cwebe Research Network

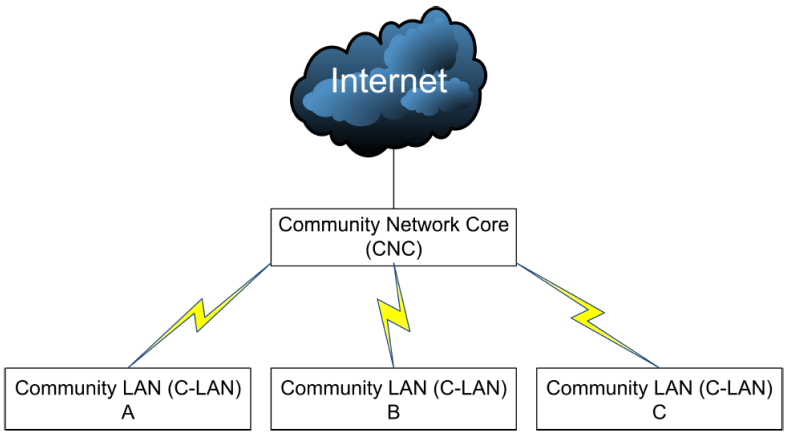

Fig. 2: Generic Community Network

Core (CNC) with multiple Community Local Area Networks (C-LAN) connecting through the CNC (over an IP-based connection) to other C-LANs or the Internet.

The CNC consists of the Access Concentrator (AC), the server in this architecture, which provides a means to access the shared Internet connection. At each C-LAN, the Community Access Point (CAP), the client, provides a gateway through which hosts and servers (within that C-LAN) access the rest of the C-WAN or the Internet. Both devices, the AC and the CAP, were developed on a FreeBSD operating system [13]. FreeBSD is renowned for its reliability, extensibility and adaptability to many situations [14], and thus was chosen for use in this system. In addition, both the AC and CAP implementations have numerous third party OSS applications and tools installed on them to aid in managing and monitoring bandwidth usage within the C-WAN. The CAP and the AC are discussed in further detail in the subsequent sections.

A. Community Access Point (CAP)

The CAP is the default router for the C-LAN to which it is connected. Servers and hosts at the C-LAN communicate with other networks on the C-WAN or the Internet by routing their packets through the CAP. In order to provide authenticated and encrypted communication with the AC, a PPPoE tunnel is established over the C-WAN to the AC.

The CAP collects monitoring data for the C-LAN using an SNMP agent [15] and a NetFlow [16] sensor. The SNMP data is collected from the CAP by the AC and NetFlow data is transmitted to the AC for aggregation, analysis and graphing. The SNMP agent monitors system health and variables such as system up time, secondary storage utilisation, RAM and CPU usage, and numbers of bytes transmitted and received via the network interfaces. To provide finer grain network utilisation data, namely which hosts are using the network and for what purposes, a NetFlow sensor was configured to listen on the internal interface on the CAP and transmit collected data to the AC. NetFlow provides a session-level view of network traffic and records information about every TCP/IP transaction that occurs over the network [17].

A web-based front end is provided for initial and subsequent configurations. Administrators can also access the device via Secure SHell (SSH) for a full-featured shell environment; however, this method should only be used for troubleshooting purposes. The web front end also provides summaries of the monitoring data that has been collected on the CAP. This data is also passed onto the AC for further analysis and management tasks.

B. Access Concentrator (AC)

The AC fulfils two main functions: to route traffic securely throughout the C-WAN and to the Internet, and perform a number of bandwidth monitoring and management tasks. A web based front end is provided for configuration and viewing of monitoring information. The AC is the default gateway for the CAPs so all traffic originating from the C-LANs passes through the AC. Secure and authenticated connections are facilitated between the AC and the CAPs by establishing PPPoE tunnels to them over the C-WAN. The community-wide C-WAN uses an RFC 1918 [18] IP address range (192.168.0.0/16) for internal IP addresses, and these need to be translated to real-world IPs using NAT at the AC for routing across the Internet. Traffic to and from the Internet is restricted by the PF firewall. Direct connections to the Internet from hosts within the C-WAN are not allowed, as the connection must be handled by the caching proxy server. An SNMP agent is installed and configured on the AC similarly to the CAP. The SNMP agent provides critical information about the health of the AC as well as network traffic data. Aggregated flow data is collected from the CAPs by a NetFlow collector.

A core application for the AC is the Squid caching proxy server [19]. Not only does it cache regularly-accessed web pages and therefore reduce Internet-bound traffic but it is configured to provide additional functionality. As mentioned previously, hosts within the C-WAN and the CAPs can not directly access the Internet, as all their connections have to pass through the proxy server. This policy is enforced with appropriate firewall rules. Users must provide authentication credentials to access the Internet, and all users Internet use is logged. Squid has been configured to cache even dynamic content, while delay pools have been configured to control the speed of Internet access, and blacklists have been used to prevent Internet access to specific content. Blacklists are customisable to block many forms of content, including specific web pages, entire domains (ie: ".Facebook.com"; or ".co.za"), IP ranges, web page sizes or web server types. These lists can be modified manually via the web front end or updated using pre-made lists that are freely available on the Internet [20].

By processing the Squid proxy log files, each user's usage of the Internet can be calculated. When a user's Internet usage reaches a particular threshold, they are placed in a delay pool, limiting their throughput until their Internet usage has decreased to below the threshold. The AC has four types of pools that a user can be in: 'no delay', 'slight delay', 'significant delay' and 'no access'. For example, when a user is in the 'no delay' pool, their Internet throughput is not limited, but when they have exceeded the threshold set for the 'no delay' pool, they move into the 'slight delay' pool; if their Internet use continues to rise they will move into a higher delay pool. The pools are calculated over a particular period of time using a sliding window. The window size can be set as per day, per week, per two weeks or per month. A user is allocated a maximum amount of bandwidth they can utilise within that period. If a user uses all their pre-allocated bandwidth, their Internet access is denied by moving them into the 'no access' pool. A user in a high delay pool will have to wait for the sliding window to move to lower their bandwidth use over the period. The thresholds and period at which these delays come into effect can be set via the web front end.

Data received from the NetFlow sensors running on the CAPs is aggregated to provide information regarding network traffic generated by the hosts on the network. The aggregated NetFlow data is used to compile a list of hosts and their generated traffic, including Internet use and this data is used in conjunction with Squid delay pools to limit Internet use originating from a particular host. Thresholds for the hosts delay pools can be configured via the web front end.

LightSquid summarises the users' Internet use from the Squid log files. LightSquid also highlights which frequently accessed websites are consuming the most bandwidth, and can then be placed into a delay pool to restrict access or be placed in a blacklist to prevent access. Blacklists are also configurable via the web front end. Cacti is used to graphically display data received from the SNMP agents on the CAPs and the AC, as well as aggregated NetFlow data. Graphs generated for the CAPs display information such as CPU usage, RAM usage, network interface traffic and secondary storage usage. By using the data received from NetFlow sensors, Cacti generates graphs displaying which protocols (including HTTP, FTP and SMTP) produce inbound and outbound network traffic. Cacti also generates graphs from information in the Squid logs.

A web front end hosted on the AC ties these applications together. This interface provides a number of tools to assist the administrator in managing and monitoring the community network. The front end provides information about the CAPs situated within the C-WAN, including their status, system up time and various system health variables. Using this information, an administrator can identify problems on the network. The front end provides a means for administrators to set thresholds and periods for user- and host-based quotas. It also allows modifications and additions to lists of websites which are denied access or limited in throughput.

IV. Results

This section reports on the results of the management system after six months of use, having been implemented in October 2009. During this initial operational period, the system was able to identify a number of issues relating to the network. Changes were effected in response to these observations in order to provide an improved and more equitable service to the schools and their associated communities connected via the C-WAN within the Siyakhula Living Lab test site. The effects of the system can be summarised into four categories: bandwidth management, awareness, fault finding and robustness.

A. Bandwidth management

During the the first month of operation, the community exceeded their bandwidth quota on the VSAT up-link, resulting in a loss of Internet connectivity. The newly deployed management system was instrumental in finding the root cause of this problem. Detailed logs relating to the traffic generated by each site were analysed, along with the detailed website logs maintained by the proxy server. The automated top-sites report generated by the AC (an abbreviated sample of which is shown in Figure 3) allowed the identification of the major consumers of network bandwidth. Per-site attribution

```
#Wed Apr 28 12:00:36 SAST 2010 - Days,28
# site,traffic,requests
safe-browsing-cache.google.com,96745573,1541
www.elearning-africa.com,62822191,7703
www.yimg.com,61680877,7495
images.supersport.co.za,18043896,957
www.kaizerchiefs.com,14460055,1803
www.yahoo.com,14122382,2237
mail.google.com:443,11909507,168
www.newzimbabwe.com,10889494,557
static.ak.fbcdn.net,10786679,1777
webmail.ufh.ac.za:443,7884592,764
www.facebook.com,7537348,1833
www.google.co.za,6633984,1257
www.fnb.co.za:443,6582971,341
profile.ak.fbcdn.net,6111669,1912
newsrss.bbc.co.uk,4840918,174
www.sabc1.co.za,4489730,518
www.ufh.ac.za,4158025,315
creative.ak.fbcdn.net,4000805,283
www.dwesa.com,3954516,611
www.esu.edu,3692866,1
www.kildareplace.ie,3499314,1
```

Fig. 3: Sample of "top websites" report, values reflect the number of bytes transferred from each site

of traffic could also be performed from these logs. It should be noted that the traffic value recorded in this report is the total traffic served via the caching proxy, and that due to cache hits, the actual volume of traffic traversing the Internet up-link may be less. In the example, this is particularly true where traffic to safe-browsing-cache.Google.com tends to have relatively high hit rates. The sites shown in the example can mostly be categorised as news, mail and e-learning resources. The use of blacklists, as previously discussed, helps in limiting the cases of bandwidth abuse by denying access to sites deemed to be undesirable.

While HTTP and HTTPS traffic constitutes the bulk of traffic on the network, other protocols are allowed and make up a small proportion of the total traffic. The traffic for each CAP's C-LAN network is summarised by the AC on the C-WAN network, with per site details made available on the appropriate CAP. The traffic data collected is utilised in the internal traffic quota system, which limits the monthly Internet usage of each C-LAN. While ideally this should be a completely dynamic system, the quota levels put in place have proved to be sufficient, and have prevented any further outages due to upstream quota exhaustion, while having had a minimal impact on daily activities at the schools.

B. Awareness

Evidence of successful bandwidth management implemented by the system is shown in Figure 4, illustrating usage before and after basic quota management. The traffic shown in Figure 4(b) is somewhat reduced due to a series of power outages in the area early in the month.

The network status overview display provided by the CAPs have aided in informing the technical liaisons within the Dwesa community, in the SLL, as to the probable root-cause of connectivity issues which has lead to a decrease in support calls, and increased the ability to perform remote diagnostics on the system, particularly in cases where the VSAT link has been offline. Reports are also generated and can be viewed on a website within the community which allows open access to view the current status of the network traffic, as well as both current and historical performance metrics.

C. Fault finding

The deployment of the CAP and AC system has, for the first time since the inception of the SLL in 2006, provided homogeneous network infrastructure. This uniformity and the automated collection of monitoring data and metrics related to the network has enhanced the ability to be able to provide remote support and diagnostics to the network – a particularly relevant issue given the geographic remoteness of the area from Rhodes University and the University of Fort Hare.

D. Robustness

The robustness of the network has also been improved, not only due to the introduction of newer homogeneous hardware, but also as a result of the customised uniform software installation. The fact that systems are configured in a similar manner has eased the support load.

V. Future Work

For the deployment of any network-based services in the greater community network, a reliable robust network is required, which, based on the observed data has been achieved. Following the deployment of the developed solution and analysis of the data collected, a number of features have been identified as being desirable further improvements in future iterations of the systems. The current system is due to be extended during the course of 2010 to include additional locations. The primary aims of the extensions to be introduced is to provide for the following key elements:

- improved network redundancy of the Internet up-link,
- improved resiliency in inter-site communications, and
- integration of an alternate means of remote access for support and monitoring.

VI. Conclusions

It is felt that the deployed system has achieved its intended goals of improving the level of service within the target network. While several benefits have been forthcoming, the key issues from the community's point of view are as follows.

- Visibility and awareness of network performance and status has been markedly improved. Information is easily accessible and provides a degree of transparency within the various community sites.
- A functional and configurable management and control system is in place, with much of the functionality being centralised at the AC which takes the role of the network 'core'. This provides for a single point of configuration and management.

Valuable lessons have been learnt in the deployment and will be integrated into future iterations of the network management platform.

References

[1] South African Foundation, "Telecommunication prices in South Africa: An international peer group comparison." Occasional Paper No 1, Apr. 2005.

[2] Business Leadership South Africa, "South African telecommunications prices: An updated international price comparison, with regulatory recommendations." Occasional Paper No 3, Nov. 2007.

[3] SME Survey 2008, "The Findings," Oct. 2008. Last Accessed: Nov 2008 http://www.bizassist.co.za/images/SME_%20Survey_Findings.ppt.

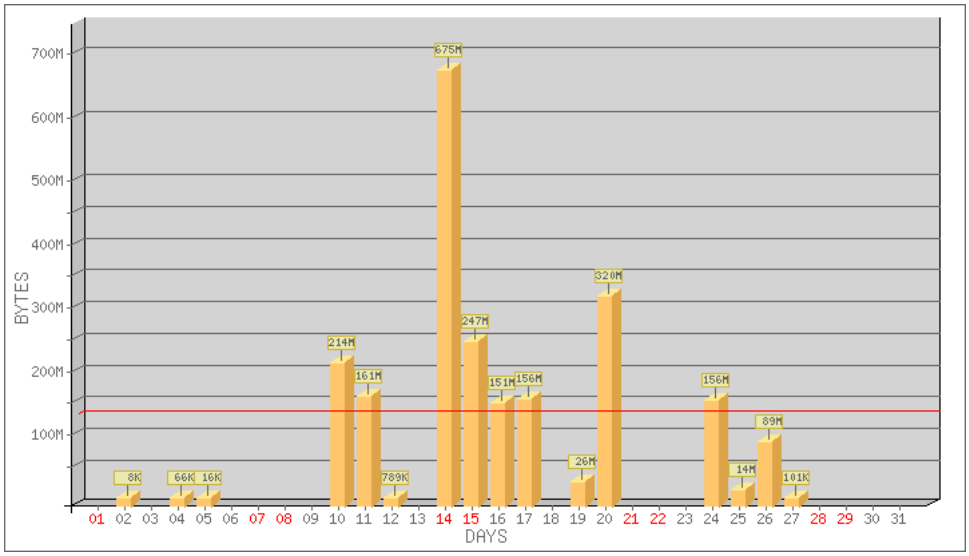

(a) November 2009

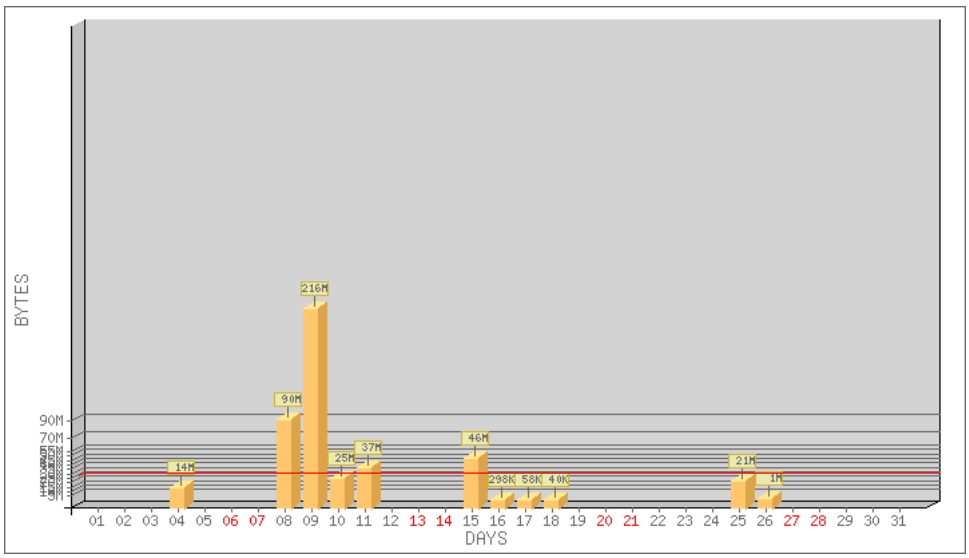

(b) April 2010

Fig. 4: Traffic utilisation graphs

[4] SEACOM, "Connecting Africa to the World," 2009. Last Accessed: Jan 2009 http://www.seacom.mu/.

[5] A. Kayle, "Seacom to lower African bandwidth costs," 2009. Last Accessed: May 2005 http://www.itweb.co.za/sections/computing/2009/0903250905.asp ?S=All%20Africa%20News&A=AFN.

[6] S. Africa, "Seacom, Interoute team up," 2009. Last Accessed: May 2005 http://www.itweb.co.za/index.php?option=com_content&view=article &id=22313:seacom-interoute-team-up &catid=260:telecoms.

[7] R. S. Rosenberg, "Controlling Access to the Internet: The Role of Filtering," *Ethics and Information Technology*, vol. 3, no. 1, pp. 35–54, 2001.

[8] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "YouTube Traffic Characterization: A View From the Edge," in *Internet Measurement Conference (IMC) 2007*, 2007.

[9] C. D. Marsan, "MySpace Threatening Net Bandwidth," 2007. Last Accessed: Apr 2008 http://www.pcadvisor.co.uk/news/index.cfm?newsid=9839.

[10] I. Brandt, A. Terzoli, and C. Hodgkinson-Williams, "Wi-Fi as a last mile access technology and The Tragedy of the Commons," in *Innovative Algorithms and Techniques in Automation, Industrial Electronics and Telecommunications*, pp. 175–180, Springer, 2007.

[11] Siyakhula Living Lab, "We are growing together," 2009. Last Accessed: May 2009.

[12] M. T. Mandioma, G. Rao, A. Terzoli, and H. Muyingi, "Deployment of WiMAX for Telecommunication and Internet Access in Dwesa-Cwebe Rural Areas," 2007.

[13] FreeBSD, "FreeBSD 6.3-RELEASE." Last Accessed: Apr 2009 http://www.freebsd.org/releases/6.3R/announce.html.

[14] M. W. Lucas, *Absolute FreeBSD*. No Starch Press, 2007.

[15] Net-SNMP, "Net-SNMP," 2007. Last Accessed: Feb 2009 http://www.net-snmp.org/.

[16] Softflowd, "Fast Software NetFlow Probe," 2006. Last Accessed: May 2009 http://www.mindrot.org/projects/softflowd/.

[17] M. W. Lucas, "Monitoring Network Traffic with Netflow," 2005. Last Accessed: Feb 2009 http://www.onlamp.com/pub/a/bsd/2005/08/18/Big_Scary_Daemons.html.

[18] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," Tech. Rep. 1918, Feb. 1996. http://www.ietf.org/rfc/rfc1918.txt.

[19] Squid, "Squid: Optimising Web Delivery," 2007. Last Accessed: Feb 2008 http://www.squid-cache.org/.

[20] Shalla Secure Services, "Shalla's Blacklists," 2008. Last Accessed: Mar 2008 http://www.shallalist.de/.

Barry Irwin is a Senior Lecturer in the Department of Computer Science at Rhodes University. He has an active interest in building secure and scalable networks with open source software.