Master of Science Project Proposal

Dan Wells

July 15, 2008

1 Details of the Research Proposal

Student Full Name:	Daniel David Wells
Student Number:	603W0418
Degree:	Master of Science
Department:	Computer Science
Project Start Date:	February 2008
Project Completion Date:	December 2009
Supervisors:	Barry Irwin and Ingrid Siebörger
Provisional Project Title:	Network Monitoring and Accounting for Shared Community Networks
Date:	12 June 2008

2 Problem Component

Community networks are often serviced by limited capacity network backhauls. In relative terms this is an expensive ongoing part of their operation. In cases where this backhaul is financed or donated by other parties it is still a finite resource. It needs to be managed and allocated equitably. The focus of the monitoring and management should be on bandwidth, but also allow for monitoring of network quality. Network quality can be understood as the degree of excellence that the user experiences when making use of the network and its available services. Monitoring will reveal what services are used most often and what those services require in terms of bandwidth to make the user's experience more valuable. Monitoring will produce statistics which will inform the administrator how to best apportion the available bandwidth between services and at what times certain services should get preference over others.

A problem with current network deployments of this type is ensuring that there is an equitable distribution of the scare resource between sites. A simple division of the resource (such as each site only having $1/n$th of the total bandwidth as a maximum rate) may seem appropriate, but could lead to significant under-performance and under-utilization of the resource. This division leads to an inefficient division of bandwidth, especially when multiple parties are not utilising the shared link. Division needs to be fair for each site as each site may have different requirements of the resource. Different site patterns could include a school housing many computers in a laboratory all concurrently using the link or a community centre with a couple of shared computers or even a small centre with only email requirements.

Another factor that needs to be considered is that while most traffic may be billable other traffic may not be, such as a local wiki, administrative work or email. In addition, measures might need to be put in place to limit access to certain bandwidth intensive Internet websites such as Facebook, MySpace and YouTube [5, 16], while others may have to be restricted completely such as websites containing illegal or pornographic material. Regularly updated blacklisted domain names [23] are freely available to prevent

access to a range of different categories of information. A usage policy will have to be drawn up which will include prioritization for specific network traffic.

3 Research Component

Two networks have been identified, each with a need for a system to be deployed to monitor and manage traffic. Firstly, the Telkom Centre of Excellence (housed in the Department of Computer Science at Rhodes University) testbed, the Settler City Wireless (SCW) network, which utilises many different types of technologies to connect devices, including Ethernet, DSL, WiFi and WiMAX [19, 11]. The SCW network stretches over Grahamstown to provide access to the Rhodes University network and the Internet. See section 3.2 for a further technical explanation of the SCW network. Secondly the Dwesa-Cwebe test network, which is located on the Wild Coast in the Eastern Cape of South Africa [14, 12]. This area lacks technical personnel, has a low level of economic activities, low income per capita, poor electricity availability and is geographically mountainous. These conditions have led to the deployment of a WiMAX network connecting five schools to a central base station, which allows sharing of the single broadband connection to the Internet via a Very Small Aperture Terminal (VSAT) connection [15]. Section 3.3 explains the Dwesa-Cwebe network in greater technical detail.

The technology required for the solution to the problem discussed above is already existent, but needs to be integrated into a single appliance. A client-server architecture will be used, with a central server managing information about the clients and the network as a whole, and the clients transmitting monitoring data to the server. The clients will have a dual purpose, to perform as the router for that particular site and to transmit usage data to the server. In addition, the clients will receive updates for their local settings. A form of centralised authentication will be integrated to aid the solution, providing users with a single username and password.

Existing technologies and products that should be investigated include Cisco's NetFlow [4], Simple Network Management Protocol (SNMP) [7], Cacti [28], RADIUS [21] and the Squid caching proxy server [25]. These technologies and how they link together, are discussed in detail below.

NetFlow will provide network traffic usage statistics. NetFlow provides more fine grained data than SNMP but not as detailed and high volume as packet sniffers [24]. It is an open protocol developed by Cisco for collecting IP traffic information. A flow is described as "active as long as observed packets that are meeting the flow specification are observed separated in time by less than a specified timeout value" [3]. A NetFlow enabled router regularly exports aggregated flows to some predefined collector host using UDP. NetFlow is a widely used tool for visualisation, accounting and traffic analysis [20]. Using NetFlow to aggregate traffic information it will be possible to understand how the network is being used and how bandwidth should be separated between sites.

Key devices, such as hosts and routers may be equipped with SNMP agent software so that they may be managed from a management station. "The management agent responds to requests for information from a management station, responds to requests for actions from the management station, and may asynchronously provide the management station with important but unsolicited information" [27]. Resources in the network are referred to as an object, each object is a data variable that represents one aspect of the managed system. The collection of these objects is referred to as the Management Information Base (MIB). Monitoring of the managed system can be done by retrieving the value of MIB objects. A management station can then cause an action to take place at an agent or can change the configuration settings of an agent by modifying the value of specific MIBs [27]. SNMPv1 and SNMPv2 both lack security features, namely authentication and privacy [26]. SNMPv3 will be used in the project to provide authentication, privacy and access control when managing the SNMP agents [26]. SNMP will allow the project to manage all devices from a singular central point.

Cacti is a complete front-end to RRDTool [22], which stores the collected data in a MySQL database. It uses the collected data to create and populate RRDTool graphs. Cacti is capable of handling many data sources. Of specific interest is that it can accept external scripts as input, and then Cacti will gather this data in a cron-job (schedule commands to be executed periodically) and populate a MySQL database. Due to the fact that Cacti can accept scripts allows it to accept any data input, which provides the capability of using any monitoring software to provide data [28]. By combining Cacti with NetFlow and SNMP outputs, a totally customisable monitoring system can be created.

RADIUS is a widely used protocol in network environments and is commonly used for embedded network devices such as routers. Generally embedded systems cannot deal with large numbers of users each with distinct authentication information. RADIUS facilitates centralised user authentication and consistently provides some protection against active sniffing attackers [8]. This project hopes to bring authorisation, authentication and accountability (AAA) to the networks it is deployed within. A centralised authentication server will go a long way in effectively charging for services used by specific users.

To aid in efficient Internet usage, a Squid caching proxy server will be combined with the project. It reduces bandwidth usage and improves response times by reusing frequently-requested web pages. "Squid optimises the data flow between the client and the server to improve performance and caches frequently used content to save bandwidth" [25]. Squid can also be used to restrict or limit certain types of traffic, such as the access to illegal or pornographic content discussed earlier.

Extensive research regarding the FreeBSD [13] Operating System (OS) needs to be undertaken as this will be the platform on which the project will be implemented. FreeBSD provides control over web, mail, file and support services. FreeBSD is an OS which allows dedication of your hardware to complete the tasks you require, and the control to restrict tasks that are not necessary [13]. This OS has been chosen as the routers and access concentrators at the test sites are running FreeBSD. This OS is also ideal for running on embedded PCs which have lower performance hardware that traditional PCs.

3.1 Embedded PC

Various low power embedded PCs need to be considered and preferably locally procured. Although they are low power solutions, a sufficient amount of processing will be required (400 MHz+) for routing, along with a minimum of two Ethernet ports. The embedded PCs should also ideally run off a Solid State Disk (SSD) and preferably have the option of linking to an IDE or SATA hard disk drive. The device will most likely be set up with a read only file system, where in the event of an unexpected power down the operating system can easily recover. Collected network usage data will either be sent to the server to be processed or stored locally for later transmission. No VGA output will be necessary, input/output will occur via SSH or through a serial port. An embedded PC would be better than a standard PC as it takes up very little space, provides no visible input/output method to tamper with and requires reduced power to provide services.

MikroTik provides numerous hardware and software solutions for monitoring and managing networks [17]. MikroTik produce an embedded PC, the 'RouterBOARD' [18], in a variety of different hardware arrangements. These devices run a proprietary OS, 'RouterOS', which is a highly customised OS with many features (such as routing, firewall, bandwidth management, wireless access point, backhaul link, hotspot gateway and VPN server). MikroTik has recently released a new network monitoring application called 'The Dude', which automatically scans all devices on a given subnet, provides a layout map of the network and monitors services of devices and alerts you in case some service has a problem [17]. 'The Dude' also triggers alarms if any Intrusion Detection Systems (IDS) are running on your subnet due to the way that it scans for devices, and the way it scans those devices for services they might be running.

By comparing this projects final solution with an already existing application it will be possible

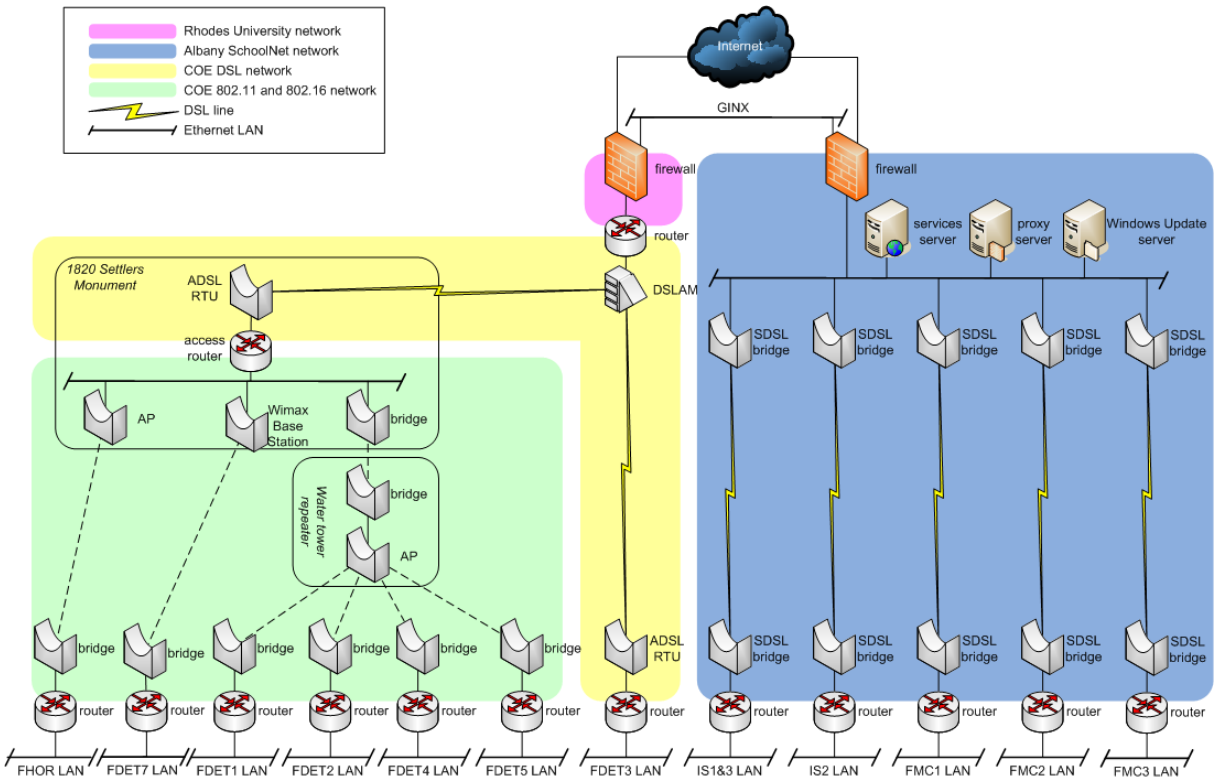

Figure 1: Grahamstown Network

to benchmark the results obtained. In this way it will be possible to thoroughly evaluate the project application developed.

3.2 Settler City Wireless (SCW) Network

The Telkom Centre of Excellence (CoE) in the Department of Computer Science of Rhodes University has been conducting research into cost effective last mile Internet access solutions since its inception in 1998. One of the aims of the CoE was the identification of affordable solutions for previously disadvantaged schools. The research over the years has included Digital Subscriber Lines (DSL) [6, 2], WiFi [9, 1] and WiMAX [11, 19] type connections. DSL was limited in that it cannot be used over a maximum of 5km and many disadvantaged schools are further away from Rhodes University than this [6]. Over time WiFi had benefits over DSL in that it was quick and easy to install, the equipment was cheap and didn't rely on any previous infrastructure. WiFi had limitations in that it was not specifically designed to connect computers at distances over 100m, requires line of sight between devices and is very susceptible to interference [2, 9]. The research logically progressed to deploying WiMAX in the experimental network. WiMAX bypassed most of the problems brought about when using WiFi, providing better transmission rates, greater resistance to interference and connectivity to locations that do not have direct line of sight [10].

The SCW test network has been deployed similarly to that of the Dwesa-Cwebe network (see Section 3.3). Refer to Figure 1 which shows the Grahamstown Schools Network. This project will be focusing on the IEEE 802.16 (WiMAX) section of the network and perhaps the IEEE 802.11 (WiFi); as WiFi is being phased out due to its limitations.

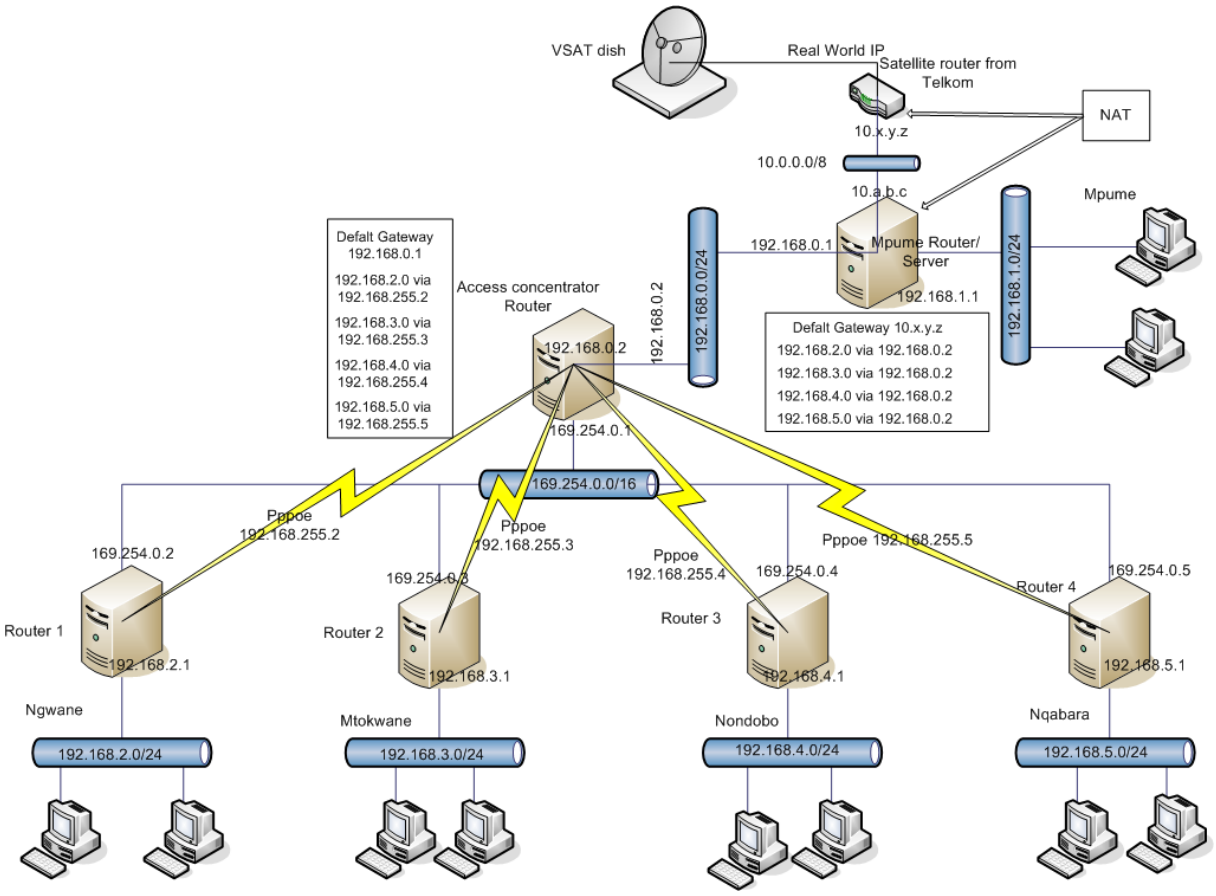

Figure 2: Dwesa Community Network

3.3 Dwesa-Cwebe Network

Figure 2 refers to the current working layout of the Dwesa-Cwebe Community Network. The network is composed of five schools; Mpume, Ngwane, Mtokwane, Nondobo and Nqabara. These schools are connected to each other via a WiMAX test bed network. The base station for the WiMAX network is situated at Ngwane. The Telkom VSAT Internet connection (the local network's backhaul) is located at Mpume, also located at Mpume is the Dwesa Access Concentrator.

Ngwane, Mtokwane, Nondobo and Nqabara each have a FreeBSD 6.1 PC acting as a router. They are all low end Intel Pentium III PCs. Each router has two network cards (an internal interface and an external interface). The internal interface is connected to the local LAN via a switch at the school, the external interface is bound to an IP address on the "raw" WiMAX IP network (169.254.0.0/16). This means at Mpume, Mtokwane, Nondobo and Nqabara their local network is connected to their WiMAX CPE, while at Ngwane it is connected to the WiMAX base station.

Each of the routers are configured that when they receive packets from within their LAN that are meant for computers on another network, including the Internet, they will forward them on to the next known network/router that they are connected to. The access concentrator runs a PPPoE service. Each of the routers use the "raw" network to establish a PPPoE session. Routers authenticate themselves with a username and password to the access concentrator which checks the credentials and establishes the PPPoE session. All traffic from the school's LAN to the rest of the network or the Internet is routed via the PPP tunnel. This allows each of the routers to forward traffic from their local network intended for the rest of the network (or the Internet) securely to the access concentrator. The access concentrator will then route it onto the correct network, either to another school in Dwesa-Cwebe or to the Mpume router

to be sent to the Internet.

Although out of scope of this project, ideally the WiMAX base station, access concentrator router and Internet connection should be located at the same point. By situating these devices at the same location it would be possible to prevent the number of hops required to make a connection between the schools.

4 Project Outcomes

With a view to producing sustainable community networks and ICT infrastructure, quotas and logging of traffic per user and per machine would go a long way in charging for services and managing bandwidth usage. Single sign-on for the network quotas would also be useful so that community members can access connectivity from anywhere that there is networking.

This project hopes to produce a working demonstrable system that provides easy monitoring of community networks, including quality, bandwidth management and authentication. It is planned that the system will generate reports for network evaluation; reports per site, per user and per device. The system will be designed with a management interface to allow authorised administrators to modify the constraints and parameters of the system (specifically those of the site, user and device). The system will be developed using Free and Open Source Software and will be licensed as such.

Initially the system will be produced on a virtual machine as an installable appliance, after initial testing and evaluation of the application the project will be re-deployed onto an embedded PC. The application end product will be distributed as source archives to permit end users to compile and install them on their target systems. Alternatives would be to distribute them as FreeBSD ports and packages, however, this would add complexity as they would need to be committed to the FreeBSD ports tree. Additionally, this would limit their use to FreeBSD rather than offering the possibility of alternate platforms such as OpenBSD and NetBSD.

To deploy the project to the test sites, the embedded machines (running the application) will need to contain the routing logic of the existing routers and access concentrators in order to ultimately replace them.

As an extra, I hope to develop a plugin for Mozilla Firefox to inform the user of their current quota usage (per user/per device). A simple plugin can easily inform the user in the status bar of the browser of their quota usage either on an individual or site basis.

4.1 Project Deliverables

- Overall design of system
- Develop system to design guidelines
- Evaluation of system in laboratory experiments and anomalies fixed
- Deployment of system in SCW Network and then Dwesa-Cwebe network
- Analysis and evaluation of system using collected data from deployments

5 Provisional Contents of Thesis

Chapter 1: Introduction

The introduction to the project with the goals it seeks to accomplish. Should contain background, project objectives, research questions and outline of thesis.

Chapter 2: Related Work

Relevant work in the field, discussing similar projects and their outcomes. As well as all necessary explanations about the various technologies used in the final solution. Discussion on the technologies used in the project will also be in this section (such as SNMP, NetFlow and Cacti and how they can fit together).

Chapter 3: System Design

This chapter will describe the overall system and its high level design. The design provides the plan on which the developed system is based. The separate application designs will be discussed and their user interfaces developed.

Chapter 4: Implementation

Development of the system is discussed with how the relevant technologies are combined; the implementation is planned in the System Design chapter. This chapter will discuss the finer points of implementation with snippets of interesting code to discuss.

Chapter 5: Deployment

(Implementation and Deployment chapters may be merged at a later stage)
Once the system has been developed, it needs to be tested in a controlled environment to ensure correctness and tested against a baseline benchmark. Testing in the lab is an ideal environment, as testing in either Dwesa-Cwebe or SCW could prove disastrous. Once the system has been proved to be accurate in the lab, it needs to be deployed. First to the SCW network and then to the Dwesa-Cwebe network.

Chapter 6: Evaluation and Results

The system has now been deployed, the system has been running smoothly for some time and much data has been collected. Discussion on the data collected will be presented in this chapter.

Chapter 7: Conclusion

Conclusions to the project. Goals from Chapter 1 are revisited and the outcomes of the project are discussed. Research question(s) are also answered and assess if the outcomes are met and therefore provide some academic worth. Possible future extensions to the project are discussed as well as methods for overcoming any shortfalls that were encountered during the project.

References

[1] B. Whittington, G. Halse, and A. Terzoli. Secure, extensible and heterogenic wireless networks: A model for community orientated wireless Internet in South Africa. Computer Science Honours thesis, Rhodes University, Grahamstown, South Africa, 2003.

[2] Brandt, I. Models of Internet connectivity for secondary schools in the Grahamstown Circuit. Master's thesis, Rhodes University, Jan 2006.

[3] Claffy, K.C. Braun, H.W. Polyzos, G.C. . A parameterizable methodology for Internet traffic flow profiling. IEEE Journal on Selected Areas in Communications, Oct 1995.

[4] Claise, B. Cisco Systems NetFlow Services Export Version 9. RFC 3954 (Informational), Oct. 2004.

[5] Gill, P., Arlitt, M., Li, Z., Mahanti, A. YouTube Traffic Characterization: A View From the Edge. Internet Measurement Conference, 2007.

[6] Halse, G. A., and Terzoli, A. Open Source in South African Schools: Two Case Studies. Online: http://eprints.ru.ac.za/100/01/HALSE-Highway-Africa-2002.pdf, Accessed: 12/06/2008, 2002.

[7] Harrington, D., Presuhn, R., and Wijnen, B. An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks. RFC 3411 (Standard), Dec. 2002.

[8] Hill, J. An analysis of the radius authentication protocol. Online: http://www.untruthorg/~josh/security/radius/, Accessed: 11/06/2008, 2001.

[9] I. Brandt, A. Terzoli, and C. Hodgkinson-Williams. Wireless Communication for Previously Disadvantaged Secondary Schools in Grahamstown, South Africa. in SATNAC 2005, Convergence - Can technology deliver?, Sept 2005.

[10] I. Sieborger, A. Terzoli and C. Hodgkinson-Williams. The development of ICT networks for South African schools: Two pilot studies in disadvantaged areas. Learning to Live in the Knowledge Society, the TC3 Conference in WCC 2008, Milan, Sept 2008.

[11] I. Sieborger and A. Terzoli. Field testing the Alvarion BreezeMAX as a last mile access technology. 10th Annual Southern African Telecommunication Networks and Applications Conference (SATNAC) 2007, Sept 2007.

[12] L. Dalvit, M. Thinyane, A. Terzoli, and H. Muyingi. The deployment of an e-commerce platform and related projects in a rural area in South Africa. 3rd Annual International Conference on Computing and ICT Research - SREC07, Kampala, Uganda, 2007.

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

[14] M. Thinyane, H. Slay, A. Terzoli, and P. Clayton. A Preliminary Investigation into the Implementation of ICTs in Marginalised Communities. in SATNAC 2006, Next Generation services - the network @work, Spier Wine Estate Western Cape, South Africa, Sept 2006.

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

[16] Marsan, C. D. MySpace Threatening Net Bandwidth. Online: http://www.pcadvisor.co.uk/news/index.cfm?newsid=9839, Accessed: 07/04/2008, 2007.

[17] MikroTik. MikroTik Routers and Wireless. Online: http://www.mikrotik.com/, Access: 12/06/2008, 2008.

[18] MikroTik RouterBOARD. RouterBOARD. Online: http://www.routerboard.com/, Accessed: 12/06/2008, 2008.

[19] P. Beyleveld. Implementation and testing of WiMAX wireless network technology. Computer Science Honours thesis, Rhodes University, Grahamstown, South Africa, 2006.

[20] Plonka, D. Flowscan: A network traffic flow reporting and visualization tool. LISA '00: Proceedings of the 14th USENIX conference on System administration, 2000.

[21] Rigney, C., Willens, S., Rubens, A., and Simpson, W. Remote Authentication Dial In User Service (RADIUS). RFC 2865 (Draft Standard), June 2000. Updated by RFCs 2868, 3575, 5080.

[22] RRDtool. Rrdtool logging & graphing. Online: http://www.rrdtool.org/, Accessed: 11/06/2008, 2008.

[23] Shalla Secure Services. Shalla's blacklists. Online: http://www.shallalist.de/, Accessed: 10/03/2008, 2008.

[24] Sommer, R., and Feldmann, A. NetFlow: Information loss or win? Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment, 2002.

[25] Squid. squid: Optimising Web Delivery. Online: http://www.squid-cache.org/, Accessed: 10/02/2008, 2007.

[26] Stallings, W. SNMPv3: A Security Enhancement for SNMP. IEEE Communications Surveys & Tutorials, 1998.

[27] Stallings, W. SNMP and SNMPv2: the infrastructure for network management. IEEE Communications Magazine, Mar 1998.

[28] The Cacti Group. Cacti: The Complete RRDTool-Based Graphing Solution. Online: http://cacti.net/, Accessed: 10/02/2008, 2007.