Optimising Internet Bandwidth in Developing Country Higher Education

Gerhard Venter

The author

This report was written by Gerhard Venter, from AfriConnect Ltd, Cambridge, UK. Africonnect Ltd is an Internet Services Company specialising in providing Internet connectivity and application solutions to Africa and other developing areas of the world.

First published 2003

ISBN 0 902928 21 0

INASP
PO Box 516
Oxford OX1 1WG
UK

Telephone: +44 1865 249909
Fax: +44 1865 251060
E-mail: inasp@inasp.info
World Wide Web: http://www.inasp.info

This report is available free online: see
http://www.inasp.info/pubs/bandwidth

A 4pp summary report (InfoBrief) is also available from INASP – see
http://www.inasp.info/pubs/infobrief

Reprinted upon demand

Contents

1. Executive summary and recommendations 1
2. The need for bandwidth optimization 5
3. Recommendations for librarians 9
4. Network optimization review 12
5. Charging and quotas 19
6. Bandwidth 20
7. Network design 23
8. Usage policies 24
9. Authentication 25
10. Managing the IT department 26
11. Using free software 28
12. Connection options 29
13. Glossary 30

Appendices

Appendix A: Practical technical implementation 35
- 1. Web caching 35
- 2. DNS caching 39
- 3. Content filtering 40
- 4. Monitoring 42
- 5. Bandwidth management 48
- 6. Security 52
- 7. Mail and dealing with spam 56
- 8. Web-based e-mail services 57
- 9. Anti-virus 58
- 10. Charging mechanisms 59
- 11. TCP/IP factors over a satellite connection 60
- 12. Major problem areas 62
- 13. Split DNS and a mirrored server 67
- 14. Bandwidth testing 68
- 15. Authentication 69
- 16. Network layout 69

Appendix B: Connection options 71
- 1. Introduction 71
- 2. Connection via VSAT 71
- 3. Wireless WAN networks 72
- 4. Wireless LAN 73
- 5. Leased line 73
- 6. Peering 73

Appendix C: Usage policies 74
- 1. Malawi College of Medicine 74
- 2. Addis Ababa University 75
- 3. Addis Ababa Library IT-related policy 88
- 4. Makerere University 89

1. Executive summary and recommendations

1.1 Pressures on performance

With wider Internet connectivity, educational institutions in developing countries are beginning to tap the many opportunities offered by today's information societies. These digital connections act as gateways where researchers and librarians can find, download and share world knowledge and learning materials; they can be platforms where local research is published, disseminated and uploaded; and they can facilitate links and collaboration among scientists, promoting discourse and dialogue on shared issues and problems

However, the same connectivity also supports all sorts of applications and behaviours that consume bandwidth. For example, it helps to distribute email spam and viruses, it allows music lovers and sports fans to keep up with their hobbies and favourite artistes from around the world, it can be used to view pornography, it can provide opportunities for computer hackers to practice their skills, and it facilitates cross-border movement by all kinds of 'intelligent' software, agents, worms and spiders harvesting and indexing content or simply tracking usage and offering upgrades. Where these applications hamper the intended uses of the Internet connection, it can be termed abuse.

These, and many other applications, all consume the limited bandwidth of higher education institutions in developing countries. As a result, response times slow down and performance drops, leading to frustrated users and ICT managers. Most researchers, information providers and institutions remain unable to access the high speed, broadband connectivity that is more and more necessary for research, teaching, and learning.

Typical solutions are to upgrade infrastructure, to install faster, larger, and higher performing systems, lines and facilities. However these are costly solutions, and may only solve short-term needs. As demand continues to grow, and as overall available bandwidth shrinks, performance will again decline, necessitating further upgrades

Bandwidth in developing countries is expensive. In a report for the Partnership for Higher Education in Africa, Mike Jensen calculates that Makerere University pays about $22,000/month for 1.5Mbps/768Kbps (in/out), Eduardo Mondlane pays $10,000/month for 1Mbps/384Kbps, while the University of Ghana pays $10,000/month for 1Mbps/512Kbps. These figures indicate that African universities, outside of South Africa, are paying over $55,000/month for 4Mbps inbound and 2Mbps outbound. These figures are about 100 times more expensive than equivalent prices in North America or Europe.

Clearly, one solution to controlling costs and improving access is to press for more affordable access by, for instance: suggesting that governments open up their telecommunications markets; by joining forces with other academic institutions to negotiate better connectivity deals; by encouraging local Internet service providers to set up country Internet exchange points – that route traffic within the country instead of via Europe and North America; and by making use of Open Source systems and software.

1.2 A different approach?

An alternative response is to recognise that 'bandwidth' is a valuable institutional resource or asset that needs to be managed, conserved, and shared as effectively as possible. Instead of simply extending computer and network infrastructure, or finding cheaper providers, this approach puts emphasis on ways to control and manage the many hungry Internet applications, uses, and practices that consume bandwidth.

Such an approach has technical implications regarding network configuration and management. Suitable policies and guidelines are also needed to encourage proper bandwidth saving behaviour. Most critical, it requires that people with the necessary technical expertise and understanding of users' needs are available to the organization.

Making better use of bandwidth helps to ensure that high priority applications get the access and the performance that they need – when it is needed.

1.3 A cross-country study

This report draws on a report commissioned by INASP in response to the concerns of partner organizations in Africa, Asia and Latin America

The report was prepared with input from eight countries (from Africa, Asia and Europe), and has been written for three main audiences – senior management, librarians, and IT managers. It first identifies access problems in university environments, and then explains why access is frequently slow and costly, and how it can be affected by government policies as well as by other regulatory and economic restrictions. Finally, the report identifies how each stakeholder group can influence and improve online access within their institution

For management, the report explains the institutional decisions that need to be made, including the purchase of bandwidth and delivery of Internet services (cable, satellite or wireless).

The question of usage policies is discussed, including enforcement measures and incentives that help to ensure that all users are aware of their responsibilities and how their actions impinge on others. The various possibilities related to charging users for Internet access are also discussed, along with possible revenue generating options to subsidise such services.

The report also looks at the roles that librarians need to play in the provision of Internet access as part of their information services, and what questions they need to ask of ICT and management in order to ensure that their goals are met.

The full report contains eight case studies of Internet optimization. Each gives first-hand experience of bandwidth problems and solutions from various situations (Addis Ababa University, Ethiopia; Malawi College of Medicine; the Multilateral Initiative on Malaria network (MIMCOM); University of Zululand, South Africa; University of Moratuwa, Sri Lanka, University of Dar es Salaam, Tanzania; Makerere University, Uganda; and the University of Bristol, UK).

1.4 Recommendations

The findings of the report are summarised below, presented as recommendations for each of the stakeholder groups: management, information service providers, and IT staff.

1.4.1 Recommendations for senior management

- Make bandwidth management a priority. Ask questions concerning Internet and network usage, security, and user behaviour. Reward staff that install the often unglamorous applications that help to conserve and share this scarce resource.
- Be safe. Insist on strong network and server security, firewalls, and virus protection. Once an undesirable has access to the system, any nightmare is possible.
- Respond to demands. Understand why and how different users make use of the Internet, prioritise among user categories and among applications, and help users to be more effective by providing appropriate training. Expect regular usage reports so that everyone knows what the Internet connection is being used for, and so the right decisions can be made when deciding how much bandwidth to get.
- Encourage positive behaviour. Devise a usage policy to be signed by all users before they use the Internet. This identifies good practices and appropriate ‘netiquette’ and it sets out procedures to be followed. It should be relevant to the average user and it should be enforced. The existence of such a policy (which explains that browsing and other Internet uses can be monitored), can often be a sufficient ‘scare factor’ to quell selfish or illegal behaviour.
- Monitor the IT team. Sometimes the IT experts are themselves the biggest abusers of bandwidth. Competent and dedicated IT personnel are essential, not only for normal ICT functions, but also to keep bandwidth use under control. More broadly, a representative high level governance mechanism or committee can be used to set overall ICT policy goals, directions, and priorities.

- Give everyone an identity. Tracking usage requires that each user of the network have an individual identity or logon that confers rights and responsibilities. It is a relatively small step to also provide each with a unique local email address as well as efficient ways to access it from remote locations. Providing such a facility will help to reduce the use of large bandwidth consumers like hotmail and yahoo.
- Evaluate connection options regularly. Bandwidth becomes cheaper every year. As prices go down, upgrades may enable a university to provide a better Internet service (and to make it available to more users). Moving to another provider or re-negotiating contracts may provide more bandwidth at lower cost. However, the benefits may not outweigh the cost of changing. It is also important to ensure that the contracted and paid for bandwidth is actually being provided.
- Join forces. Explore network integration – create networks between academic institutions within the country that allow for joint bandwidth purchasing, help to keep local traffic local, and lobby for regulatory or other political changes.

1.4.2 Recommendations for librarians

- Obtain an understanding of the technical issues, in order to ask the right questions and to demand an appropriate service from the IT team.
- Expect usage reports and statistics, and encourage appropriate use through training.
- Understand what library users are using the Internet for.
- Get PDF and PostScript reader software installed on library computers, so students can read online journals and other publications.
- Train library users in proper use of the Internet, especially in search techniques. This saves bandwidth by enabling users to find resources more easily. It also enables them to use the Internet productively, and to avoid random browsing.
- Consider setting up library 'portals' or gateway websites where users are quickly directed to relevant and annotated Internet resources. This reduces random searching and, if local caching is used, saves frequently used resources to the local network.
- Where feasible, obtain large electronic resources, full text journals databases, or images for example, on CD-ROM or DVD formats, making them available on the local network. This avoids online downloads that consume bandwidth.
- When downloading is the only option, do it during quiet times (at night for example). Saved files can then be made available locally (where copyright allows).

1.4.3 Recommendations for IT staff

- Install good security and adopt anti-virus practices that prevent situations where malicious actions cause bandwidth to be used up.
- Track all relevant aspects of the network, bandwidth and Internet usage. Regular reports help all concerned to understand usage patterns and trends, to pinpoint problems, and to signal potential bottlenecks.
- Implement content filtering to block undesirable web content such as gaming, pornography and commercial streaming media. Email can also be restricted, to some extent, to exclude certain types of files or files that are very large.
- Use proxy servers and local caching DNS servers to keep local copies of previously retrieved web pages and Internet addresses. This prevents a situation where the same page (or address) is retrieved from the Internet several times per day.
- Use the network to locally manage upgrades and updates. Updates, fixes, and patches for software and systems (as well as commonly used software readers etc.) can be downloaded once to the network and then made available to individual workstations. This avoids situations where users are using international bandwidth to update their computers. Such downloading can be done at night when demand is low; updates to individual workstations can be automated.
- Make sure all user activities can, if necessary, be traced. This is only possible where the

network has an authentication system (users must log in before they can do anything). A usage policy should spell out which data on users is collected and how it can be used.

- Consider implementing a bandwidth manager product that allows you to give bandwidth priority to certain protocols (such as web), and to throttle others (such as Kazaa – an application used to share music files).
- Discourage and control certain types of 'peer to peer' networking. Hungry applications like Kazaa need to be 'rate limited,' using a bandwidth manager, or prevented via network layout changes.
- Offer email addresses and web based email facilities to users, allowing them to access their emails from anywhere if necessary. This reduces the need for staff and students to set up hungry web accounts like hotmail and yahoo.
- Configure the network to avoid open relay hosts (mail servers that will accept connections from anywhere) and open proxies (proxy servers that will accept connections from anywhere). These can easily be abused by outside interests and can lead to the 'blacklisting' of your network. Ensure that these servers only relay mail or accept connections from the University network.
- Train all users to use the Internet safely and efficiently. They should especially not reply to spam, and not launch unknown programs they receive via email. Encourage bandwidth-conserving behaviour. All users should be aware of the impact they can have when competing aggressively for bandwidth, for example by downloading music. They should realise the potential consequences of their actions to the whole network community.
- Consider outsourcing and mirroring options. Depending on which audiences the institution wants to reach, separate strategies for internal and external (international) users can be followed. For example, consider directing international users to an international site or server, guarding local bandwidth for local users and applications.
- Charging for bandwidth may help encourage users to use bandwidth sparingly. Charging should be based on the amount of traffic a user generates on the international link (and not per hour of use, for example).

2. The need for bandwidth optimization

Part of INASP's work enables access to online academic and scientific publications. For these programmes to be effective, the institutions targeted need to have a usable Internet link.

Several universities in this study have an Internet connection of between 512 Kbps and 1 Mbps (as at May 2003).[1] This is about as much as a DSL connection (512 Kbps to 1.544 Mbps), which is typically used to connect a single household in the West to the Internet. Bristol University, by contrast, has a 2.5 Gbps link; which is 5120 times as much as the University of Dar es Salaam has.

Students and researchers in the West tend to take free, fast access to the Internet for granted. While it is not necessary to have very fast access to the Internet for it to be usable, there is a limit below which it becomes frustrating. Usability studies show that an average Web page should load within 10 seconds; if the text starts loading immediately, followed by the graphics, load times of up to 39 seconds can be acceptable.

Unfortunately, proxy servers (described in this document and in the glossary), cause a delay and then the page loads all at once. As will be seen in this report, it is essential to implement a proxy server. Therefore, a typical page-load time of around 10 seconds could be a target for IT departments.

Usability is important for researchers because of the nature of Web searches. A user might have to load many pages and scan through them before finding the right document. If each 'false lead' takes a long time to load, Web searches become a frustrating experience.

2.1 Participants in the study

The Universities of Addis Ababa (Ethiopia), Bristol (United Kingdom), Dar es Salaam (Tanzania), Makerere (Uganda) and Moratuwa (Sri Lanka) were asked to supply information about their networks and optimization efforts. Bristol was included in order to compare the others with a university that has cheap access to high amounts of bandwidth.

Some institutions could not supply all the required information for reasons such as the information requested not being known.

The MIMCOM network (a network for Malaria researchers across several African countries) and the Malawi College of Medicine were also included because they provided interesting examples or were well described and documented.

Information about the University of Zululand is also included because they did interesting work on a charging system, but Zululand was not part of the original case study.

The table below compares the bandwidth situation of the universities and other networks included in the case studies.

[1] These connection speeds are explained in the glossary entry under the entry 'data rate'.

	Users	Computers	Bandwidth	GB/month	Connection	ISP*
Addis Ababa**	4000	Unknown	512 Kbps	106	Leased line	VSAT
Bristol***	22000	16000	2.5 Gbps	4500	Leased line	Cable
Malawi College of Medicine	Unknown	250	128 Kbps	Unknown	Wireless	VSAT
Dar es Salaam	11000	2000	512 Kbps	Unknown	Leased line + own VSAT	VSAT
Makerere	25000	Unknown	1.5 Mbps	144	Leased line	VSAT
MIMCOM	1400	1200	800 Kbps	90	Own VSAT	N/A
Moratuwa	Unknown	Unknown	2 Mbps	Unknown	Leased line	Cable & VSAT
Zululand	6000+	750	1152 Kbps	120	Leased line	Cable

* How the ISP makes its connection to the Internet.

** Addis Ababa has over 20,000 students, but only academic staff, postgraduate students and some undergraduate students in IT-related departments have Internet access.

*** Bristol is connected to the Janet network, which is a network of educational institutions in the UK. Janet is a very large and well-funded network, more a part of the core Internet infrastructure than something that is *connected to* the Internet.

2.2 Why bandwidth is expensive

Bandwidth to developing countries is so expensive that most universities cannot afford more than 1.544 Mbps – equivalent to the average Western household with ADSL connection. The reasons for this situation include the following:

- In many cases, Internet access to the country is available only via satellite connections, which are much more expensive than cable. Reasons why many organizations opt for satellite connections are given below. A map showing the marine cable connections to Africa (and the fact that only a few countries are connected in this way) can be viewed at: <http://network.idrc.ca/ev.php?URL_ID=6568&URL_DO=DO_TOPIC&URL_SECTION=201&reload=1053101936>
- Where marine fibre cables do exist, they may not carry enough traffic to achieve the economies of scale that make transatlantic bandwidth to Europe, for example, so affordable. In some of the countries that are connected via a marine fibre cable, the telecommunications infrastructure for connecting it to most of the country does not exist.
- The wired telecommunications networks in many developing countries reach only a small part of the population, and many areas (even parts of cities) are not covered at all. The development of wired networks cannot follow the same course as it did in industrialized countries owing to small populations or low population densities in some areas, poverty, the rise of mobile and satellite communications.
- Some telephone companies that have telephone lines lack the capacity (owing to low demand) to create leased-line connections. Low demand exists mainly because many companies and institutions bypass the national telecommunications grid by using VSAT.
- Leased lines are sometimes analogue instead of digital. On an analogue line, a modem is used for digital transition (such as connection to the Internet), resulting in a maximum speed of 56 Kbps. Digital lines are capable of much higher speeds.
- Bandwidth is also expensive due to the comparative weakness of the currencies of developing countries that have to pay in US dollars or euros or other major currencies for most or all of their upstream international bandwidth.
- While the cost of the telecommunications link between two countries is generally shared, in the case of African countries (and possibly of many other developing countries) the cost of the international link is paid for entirely by the African country. This amounts to reverse subsidization of developed countries.
  (see <http://www.afrispa.org/HalfwayDocs/HalfwayProposition_Draft4.pdf>).

- Considerable congestion exists at ISPs where many users share a small amount of bandwidth; ISPs simply have too many customers for their capacity.
- Inter-country links do not exist between most developing countries. For example, most communications between African countries must be routed at great expense through Europe or the USA.
- Communications and computing equipment is expensive for African organizations as a result of weak currencies, high transport costs, small budgets and high tariffs. In many African countries computer equipment is classified as a luxury item and taxed accordingly, though this counter-productive policy is likely to change in the medium term.

The regulatory situation in some countries is another problem. Active bureaucratic resistance is often experienced when changes to the telecommunications environment are proposed. For example, it may take a very long time before a licence to operate a satellite system is granted or refused. Refusal may be designed to protect the state-owned telecommunications company, or to protect a foreign company that has bought the original national operator. Governments also consider security issues when deciding whether to grant permission for new communication links. This might concern their own security (by limiting their population's access to outside information and contacts) or in response to the terrorist threat that is affecting ever more countries.

Lack of government investment in telecommunications is widespread because of more pressing priorities such as health and education.

Universities cannot afford a decent link, or in some cases do not see its value.

2.3 Bandwidth problems associated with university campuses

The graphic below shows how Web access through a proxy server at an academic institution can be analysed using inexpensive software (Sawmill in this case). Since the proxy server log keeps a record of every Web site visited, as well as the amount of bandwidth consumed, the top sites by bandwidth can be identified (sites can also be sorted by number of visits, but that is of less interest for bandwidth analysis).

The graphic shows no academic sites in the top ten sites (by bandwidth). The top three sites were not even visited by users: they are connection attempts by worms to URLs on the Internet that were programmed into the worm's code. Apart from that, there is MSN, which is not a particularly useful site, Yahoo's e-mail service, and Gator, a shopping service that installs a 'helper' on unsuspecting users' PCs. This example shows how the bandwidth usage of an academic institution can be monopolized by malicious worms and the commercial interests of large corporations. The academic purpose of the network is not reflected in the bandwidth usage, and it is clear that this network is out of control.

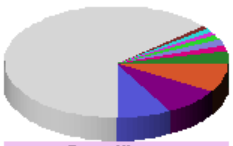

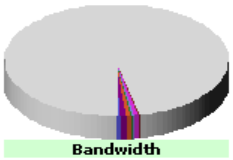

		Page or folder	Page views	Bandwidth	Visitors	Page views bar
1		http://www.opasoft.com/work/scheduler.php?(parameters)	97,331	86.32M	4	
2		http://www.n3t.com.br/work/sscheduler.php?(parameters)	97,272	86.25M	4	
3		http://www.gwmnet.com.br/marcos/gayer.php?(parameters)	96,639	48.34M	4	
4		http://shttp.msg.yahoo.com/notify/ (default page)	43,479	18.06M	16	
5		http://140.239.165.252/http_page.html?(parameters)	21,976	17.14M	1	
6		http://rad.msn.com/ADSAdClient31.dll?(parameters)	19,187	15.08M	73	
7		http://h.msn.com/c.gif?(parameters)	15,745	7.21M	70	
8		http://bannerserver.gator.com/bannerserver/bannerserver.dll?(parameters)	15,044	53.03M	12	
9		http://uk.adserver.yahoo.com/a?(parameters)	13,961	13.33M	38	
10		http://mail.opi.yahoo.com/online?(parameters)	9,460	4.46M	63	
		205587 other items	757,491	9.60G	-	
		Total	**1,187,585**	**9.94G**		

Other problems particularly associated with university campuses include:

- Students typically have more time, are less supervised, and are under less pressure from work targets than, for example, office workers. Therefore, a university network is one of the most challenging environments to manage.
- People use the Internet in many different ways, some of which are inappropriate or do not make the best use of the available bandwidth. For example, while it may not generally be a problem if a student downloads a music file, plays on-line games or experiments with the latest Microsoft service pack, it becomes a problem when the bandwidth consumed by this activity prevents a researcher from downloading or viewing a scientific article
- Even where high amounts of bandwidth are available, control, monitoring and optimization are necessary because users (and especially students) will always find a way to fill the available amount of bandwidth.
- Hacking: students experiment with their computing knowledge, and connect to exposed computer systems both on the campus and elsewhere in the world.
- Peer-to-peer (P2P) networking (using Kazaa and other programs) is very popular among students.
- Universities may need to police the amount of bandwidth they are getting from a shared system because they might be competing for bandwidth with other customers. For example, if a university gets its bandwidth from an ISP, it is likely that the ISP also sells bandwidth to other users, such as local companies. An organization should have a clear understanding of the nature of the shared system – how much minimum bandwidth they are paying for (the Committed Information Rate, CIR) and how much Burst Excess (BE) they can get. When there is a BE, what is the contention ratio? (These issues are explored in more detail in Section 6.1.)
- 'Long fat pipe network' factors affect TCP/IP performance in networks that have relatively large bandwidth but high latency (delays), as can be found in satellite networks. The high latency is due to the long distance that signals must travel from the VSAT dish to the satellite, and in some satellite networks this requires steps to enable TCP/IP to make full use of the available bandwidth. This issue is discussed in more detail in the glossary and Appendix A.
- Latency due to a connection via satellite also makes it important that functions such as DNS (see Section 4) are provided locally rather than across the Internet link. When the DNS server is on the local LAN, most DNS queries will take less than 10 ms (milliseconds). The satellite link adds at least another 550 ms because of the long distance between the earth and the satellite.
- Staff issues. Without proper management, IT staff may themselves become part of the problem (see Sections 10 and 11).
- Control and monitoring is necessary in order to make informed decisions about how much bandwidth is needed. If the graphs show that bandwidth is used mainly for recreational activities, or is consumed by virus activity and Windows updates, then control is more urgently required than additional bandwidth.
- The after-effects of a new installation. Soon after any organization connects its staff to the Internet and e-mail, it can expect to see its IT-support problems spiral. It takes a long time to get these problems under control, even for a rich private company. A university's IT environment can easily get even more chaotic. Management needs to plan for this transition.

3. Recommendations for librarians

Librarians often find themselves in charge of rooms located in the library where students can access the Internet. In addition, they are required to provide library services and information distribution over the Internet. For these reasons, librarians typically have an interest in network optimization and control, even if it is at a non-technical level.

If the library has enough computers, the librarian might be able to make a convincing case that the library should have its own proxy server (and even its own IT technician) for reasons of optimization and control.

Librarians who are in charge of computers connected to the Internet should have access to daily or weekly Web access logs and statistics of the usage at the library. They should be able to tell what the computers in the library are being used for without having to look over the user's shoulder.

Webalizer statistics of a library's proxy log can show the top sites visited, as shown in the screenshot below. In this example, the librarian might want to know why Windows updates and Symantec (Norton Anti-virus) as well as Web-based e-mail services are the top bandwidth destinations at the library. (They might also ask why the IT team is not distributing Windows and Anti-virus updates from a local server.) Also, they might wonder why are there so few academic sites on the list.

Since this program is free, this information is the minimum that a librarian should expect from the IT team.

#	Hits		KBytes		URL
1	7451	22.00%	10321	7.63%	http://windowsupdate.microsoft.com/ident.cab
2	184	0.54%	6309	4.66%	http://www.yahoo.com/
3	171	0.50%	4045	2.99%	http://www.msn.com/
4	152	0.45%	3362	2.48%	http://mail.yahoo.com/
5	3	0.01%	2212	1.63%	http://liveupdate.symanteclive update.com/ennfull2.x86
6	3	0.01%	1530	1.13%	http://defender-downloads.eacceleration.com/defender/UPDATES/eAnthologyApp_Update.exe
7	1	0.00%	1417	1.05%	http://www.reebok.com/x/us/rbk/features/atrRemixed/atrRemixed.swf
8	20	0.06%	1403	1.04%	http://www.sph.umich.edu/epid/GSS/application/gssapp.pdf
9	85	0.25%	1348	1.00%	http://www.google.com/search
10	3	0.01%	1282	0.95%	http://defender-downloads.eacceleration.com/defender/UPDATES/defscan_setup2.exe

In cases of abuse, other products (such as Sawmill) may be used to indicate which users have used which sites. This will only be possible on a network where users have to log on.

Librarians may also see the need for a usage policy or 'Internet usage agreement' if the institution doesn't already have one, or if they want to address certain library-specific issues. Appendix C includes Addis Ababa University Library's 'IT-Related Policy', which co-exists with the University's general 'University Use and Security Policy'.

3.1 Bandwidth

Librarians need to be aware of bandwidth and other technical issues in order to make good use of the Internet. The section about bandwidth in this document shows, for example, how the actual amount of bandwidth to the library can be tested. Armed with these figures, and comparative bandwidth figures from similar libraries, they may be able to make a case for more or guaranteed bandwidth to the library. The section on bandwidth management in this document (and the greater detail in Appendix A) describes methods of prioritization that can be used to guarantee a minimum amount of bandwidth to the library.

3.2 Access to journals

Many journals are available via the Internet for free (for example the *British Medical Journal* <http:// http://bmj.com/>), and a comprehensive list of free resources can be found on the INASP website. In addition, many publishers make their journals available at reduced prices to developing countries. Librarians can obtain further details from INASP. Examples of initiatives to provide content to developing countries include the following:

- The World Health Organization has created the Health InterNetwork (also called Hinari). Articles from many health-related journals are available from <http://www.healthinternetwork.org>. Since these journals are normally very expensive, the content is password-protected, and must be accessed via a secure connection (HTTPS, see glossary). However, researchers from certain countries can get access to this massive knowledge base without charge. The secure site is at <https://hin-sweb.who.int>. In order to access secure sites, the network layout and firewall needs to be set up to allow this.
- INASP manage the PERI project (Programme for the Enhancement of Research Information), which negotiates reduced rate (or sometimes free) access to online journals from a variety of publishers in different subject areas. Access to these resources is through the publisher's own website, and involves either an assigned password, or a recognised URL.

Universities can provide their own front page, providing links to all the journals that the university has access to. It is also possible to create a Web-based interface from which users can request journal articles, which the library could then send to the user via e-mail. An example of such an interface can be seen at <http://www.nlm.nih.gov/mimcom/document_delivery.html>.

Where universities have their own electronic copies of journals, these could be made available via the Internet, particularly for students in continuing education who may live and work away from the campus. These students may benefit greatly from online access rather than having to travel to the library. If no Web-based system exists, the librarian or study leader may be able to send articles via e-mail.

All university computers should have PDF and PostScript reader software installed. This is important to prevent users downloading these (large) programs and using bandwith to do so. Most journals and other articles are distributed via PDF and are often also available as PostScript. Software that can read PostScript includes Gsview PostScript Previewer <http://www.cs.wisc.edu/~ghost/gsview>; Gsview can also read PDF files. Another PDF reader is Acrobat Reader <http://www.adobe.com/acrobat>, but it cannot read PostScript files. Where material is available in multiple formats, users should be encouraged to download the smallest files for reading (PDF is usually smaller than postscript).

If students and researchers are unable to access journals and other online resources because of network or bandwidth problems, the Internet link is failing in its core function. In this case, heads of departments, including the librarian, are entitled to ask questions such as these about the Internet link:

- Are connections to journals failing because of abuse elsewhere on the network?
- How much is spent on bandwidth? Can we afford or get funding for more?
- What steps are taken to maximize the available bandwidth?
- Why was the Internet connection so slow at a particular time?
- Can the IT department solve the problem of library computers automatically wasting bandwidth by accessing Windows updates, anti-virus updates, etc.
- What proportion of the total university bandwidth is used by the library?
- Can the IT department create an off-line downloader, to enable users to schedule a large download for night-time usage?

Other questions to ask include

- How many PCs can we support on our link?

- How do we control users and what should they be allowed to do and not to do?
- Who should set controls up and what will the costs be?
- How can we make the resources we have on CD available more widely? Can we put them on a server's hard disk and make them available on the campus network (e.g. through read-only file sharing or a Web interface).
- Can we get a CD writer to distribute study material on CD to users on remote campuses? (Where copyright restrictions allow.)

4. Network optimization review

4.1 Web caching

A Web proxy server is a server on the local network that keeps copies of recently retrieved or often-used Web pages or parts of pages. For example, if one person on the network has visited a Web page, that page is stored in the cache, and if someone else later visits that page, it will actually be delivered from the local server instead of from the Internet. This results in two major advantages: faster Web access and less bandwidth usage.

A proxy server can be used for additional features such as caching often-used Web sites during periods of low usage and keeping a log of Web sites visited, which enables administrators to understand what the Internet was used for and how much bandwidth was used by each user, and to notice any possible misuse of the system. A site that has a proxy server will also have tight control over which sites can be blocked, and will have the ability to identify users who are abusing the system.

Effect on response times: Using a proxy server vastly improves the response times of pages that have been cached and may also slightly affect un-cached pages. Without caching, pages typically take between 4 and 30 seconds to load. A cached page should take less than a second. Un-cached pages might also load faster because some elements of the page, such as graphic elements and logos, may already be cached, or because more bandwidth is available on the network because of caching.

Effect on bandwidth usage: When a proxy server is introduced, the bandwidth usage at a site will initially rise until the cache is 'populated' and all the optimum settings are achieved, but thereafter the bandwidth usage will go down to previous or lower levels. The ultimate effect depends on human factors. In theory, bandwidth use should go down if everyone keeps using the Internet in the same way as before. In practice, bandwidth use tends to increase slightly, because users may be inclined to make greater use of the Internet as a result of the better response times. The 'pre-fetching' features of proxy servers also cause additional use of bandwidth, but this is designed to get new versions of popular pages during periods of low usage, so that, although overall bandwidth usage may increase, a better spread of this usage is achieved, along with access that is faster most of the time.

When implementing a proxy server, it is important to design the layout of the network in such a way that users cannot simply bypass the proxy server. Methods are outlined in Appendix A.

Is the proxy server a bottleneck? Proxy servers need to be powerful machines, and should have as much memory as possible. Also, the larger the disk space allocated to caching, the better. On a university campus network, for example, there should be more than one proxy server, and with today's cheaper and larger hard drives, powerful proxy servers can be built relatively cheaply with, for example, 50 GB of disk space allocated to the cache and memory sizes of 1 GB or more.

4.2 DNS caching

DNS (the Domain Name System), is used every time an e-mail message is sent, and once or more for every Web page visited. While computers use IP addresses to communicate (e.g. 216.239.51.100), humans prefer to use names (e.g. google.com). So if someone wants to connect to the server at <http://www.google.com>, the domain name service matches ('resolves') that name to an IP address to enable the connection to take place.

If the DNS server is far from the user (for example, if the DNS request has to travel via a satellite link to a DNS server in London, and the response travels all the way back via satellite), the service will be slow, and it will take at least one second longer for every DNS request to be answered.

For any organization that is large enough to have a server, it is useful to configure it as a caching DNS server. This server is not an authoritative source for a domain name; it simply resolves the DNS name on behalf of local clients. But since it keeps a cache of all the names it has resolved before, these names can be resolved much faster than if they had to be resolved every time by a DNS server on the other side of a satellite link. This causes Web browsing to be noticeably faster. Ways to implement DNS caching are outlined in Appendix A.

4.3 Content filtering

All sizeable organizations find that, soon after introducing Internet access, the access logs show that users use it for a variety of unintended purposes, such as downloading large music or movie files, pornography, executable files or large software programs. Where these activities hamper the intended use of the Internet link, it can be regarded as abuse. The screenshot below shows typical usage patterns in a US-based company.

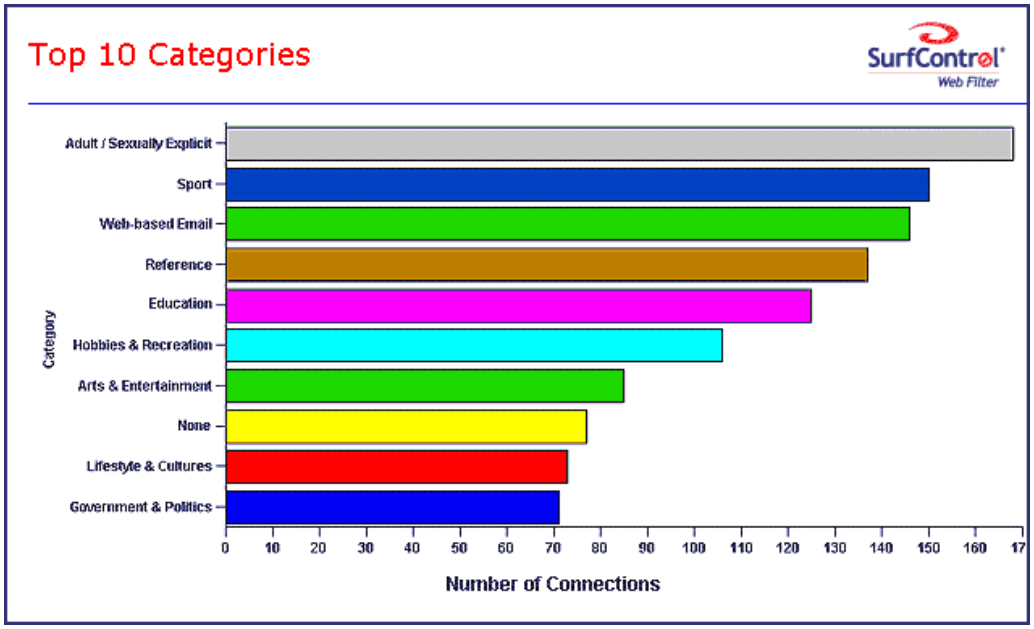

Content filtering can be used to prevent users from accessing certain categories of Web sites or from downloading certain types of files. The programs that can be used for content filtering are typically implemented along with the proxy server, and the measures for ensuring that the proxy server is not bypassed will also ensure that content filtering is not avoided.

When a user tries to access a site that is blocked, a content-filtering application simply displays a message to say that access has been barred, as shown below.

The main aims of content filtering are:

- Saving bandwidth. Where bandwidth is scarce, saving bandwidth by implementing content filtering is entirely justifiable, though using content filtering to achieve this has not turned out to be as effective as had been expected. This is because people (especially students) who are not using the Internet for a specific purpose but are simply exploring ('surfing') tend to find alternative activities that also consume bandwidth. However, one cannot only look at total bandwidth consumption. For example, a student may use up most of the bandwidth for three hours downloading pornography, but this may hardly be noticeable from the monthly statistics. However, during those three hours, the Internet link may be all

but unusable for the rest of the university community. It is wrong to allow users to do this if their activity prevents researchers from, say, downloading journal articles.

- Preventing the installation of unauthorized software on university computers. It can be argued that protecting university computers against unauthorized software use can be achieved in other ways, such as restricting users from installing any software, with only the IT department having the log-in access to install anything. However, protecting a university's computers against harmful software should not have the effect of denying students access to useful software.

In the MIMCOM network, which is a serious research network, there has nonetheless been abuse of the network for entertainment purposes. There has seldom been an academic site among the top 20 Web sites in terms of bandwidth usage. This is not to say that researchers were not accessing journals (they were), but the abuse that was going on, from researchers and staff, was threatening to drown legitimate uses. This is a typical experience of almost all organizations that have installed an Internet connection.

The introduction of content filtering caused a reduction in traffic. However, since the network was growing, the introduction of content filtering caused only a once-off drop in what is a continuously rising pattern of traffic. The reason for this was that when the abuse by the few was stopped, the rest of the staff could make greater use of the Internet link because the connection was faster.

The representation of research or academic sites among the top Web sites visited increased noticeably. It can be concluded that when content filtering is introduced, serious work is less-often frustrated by people who abuse the system.

The screenshot below shows the blacklist (sites blocked by administrators) of the MIMCOM network.

The screenshot below shows the categories of sites that are blocked by N2H2's filtering system. Most content-filtering packages have similar categories that can be selected by administrators. Management, or the ICT implementation committee, rather than IT technicians, should decide which categories to block. There must also be a mechanism for a user to request the unblocking of a certain site, because even the best content-filtering software sometimes blocks a site in error.

Category		Category		Exception	
adult	☑	mbb	☐	allowterms	☐
alcohol	☐	news	☐	education	☑
auction	☐	nudity	☑	filteredsearch	☐
chat	☐	personalinfo	☐	forkids	☑
discrim	☑	personals	☐	history	☑
drugs	☐	porn	☑	medical	☑
ecommerce	☐	recreation	☐	moderated	☐
freemail	☐	schoolcheat	☐	textonly	☑
freepages	☐	search	☐		
gambling	☑	searchterm	☐		
games	☐	sex	☑		
gross	☑	sports	☐		
illegal	☑	stocks	☐		
jobsearch	☐	suider	☑		
jokes	☐	swimsuit	☑		
keyword	☐	tobacco	☐		
language	☐	violence	☑		
lingerie	☑	weapons	☑		
loophole	☐				

A list of content-filtering packages appears in Appendix A, including software for blocking advertisements. Many Web sites carry advertisement banners or pop-ups, and these consume considerable unnecessary bandwidth.

4.4 Monitoring

Optimization of bandwidth usage is only possible in a network where the administrator has full knowledge of the usage patterns and everything else that is going on. Monitoring can be achieved by looking at the Web and e-mail logs, or by using specific software to analyse them. In addition, other tools that monitor the traffic load on network links can be implemented. A list of monitoring tools is included in Appendix A.

4.5 Bandwidth management

There are bandwidth-management tools that can be very useful in controlling traffic in a low bandwidth environment. Prioritization, quality of service (QoS) and 'Traffic shaping' are terms that refer to various techniques in the area of bandwidth management. Some of the many functions that these products can provide are designed to

- ensure equal distribution of bandwidth to various parts of a campus, or provide a means of guaranteeing more bandwidth to a certain prioritized area, such as the library;
- prioritize or deprioritize some types of traffic. For example, most people would agree that it does not matter if an e-mail message is delayed by one or two minutes. Mail can be deprioritized to ensure faster Web access, for example. Bandwidth-management tools are also useful for detecting and limiting peer-to-peer file sharing.

A variety of free and commercial products are available to implement prioritization; some of them are listed in Appendix A.

4.6 Security

Apart from being important in its own right, security is also important because there are people on the Internet who 'hack' into a vulnerable machine and then use it for sharing things like

illegal copies of software, music and movies. These people hide the fact that they have compromised a machine because they want to keep using it for their own purposes. Illegally shared files are often called 'warez', and participants use IRC to tell each other where to download the files from. Kazaa, BearShare, LimeWire, IRC's download facility and FTP are used to download these files. Once a host has been compromised, a university or other organization may find that its bandwidth is completely consumed by these activities, leaving nothing for its intended use.

This tends to happen more often at educational institutions because they seldom have the level of manpower needed to handle security issues that is usually found at commercial companies.

The approach to security should be comprehensive. No network that relies on a single approach (such as protecting the network by a firewall) is safe.

The following security measures should be considered:

- Protect the institution's network with a firewall.
- Harden all hosts (all Web or file servers), even if they are inside the firewall. Hardening a host includes steps such as disabling all unused services and setting strong passwords. A Web server that is visible on the Internet should not be running Windows file sharing services, for example.
- The IT department should subscribe to reputable security mailing lists. These lists will alert them to security flaws that are regularly found in operating systems, Web servers, other applications and even in firewalls. Details of some of these mailing lists appear in Appendix A.
- There is a security-related mailing list for high-level staff. The SANS NewsBites is a weekly high-level executive summary of the most important news articles that have been published on computer security during the previous week (see <http://www.sans.org/newsletters>). This sort of mailing list will enable managers to ask IT staff relevant questions.
- The SANS Institute has a top 20 list of security issues. This list can serve as a good starting point from which to secure a network – not only for the network team, but also for management who may want to use the list as a set of targets to achieve. The list can be found at <http://www.sans.org/top20>.
- Implement all security bug fixes that apply to the institution's equipment.
- The IT department should test the security of the network using penetration tools, and keep an eye on system logs and sudden increases in bandwidth usage.
- There should be sensible password and user policies, including disabling or deleting accounts of users who have left, and ensuring that users don't share passwords.
- Intrusion detection tools.

See Appendix A for more detail on these recommendations.

4.7 Dealing with spam

Unsolicited e-mail sent to many addresses ('spam') is becoming a serious problem, particularly for educational institutions, not only because it wastes bandwidth but also users' time.

User education used to be sufficient to deal with this, but when users begin to receive too many spam e-mails, a filter is needed. Client-side spam filters are available, but it might be best to address the problem for everyone, which means that a server-side product needs to be installed. A list of products that can be used can be found in Appendix A.

User education consists primarily of teaching users never to reply to a spam e-mail. Some spam messages have a line that invites you to reply if you want to be 'removed' from their list. In many cases, this is in fact a way to confirm that your address exists. Users should simply delete spam messages.

Another form of spam avoidance now appearing is to list e-mail addresses on Web sites in a form different from their correct appearance. For example, avoiding the symbol "@" and

giving <jsmith at abc.ac.uk> instead of <jsmith@abc.ac.uk>. This is because some companies that sell lists of e-mail addresses to spammers have software that searches the Web and collects e-mail addresses from Web sites, mailing-list archives, etc.

Since this approach does have its drawbacks, one could also put a Web contact form on a Web page instead of an e-mail address. Another approach is to use a Javascript that hides or encodes an e-mail address (see Appendix A).

4.8 Web-based e-mail services

Web-based e-mail services are very convenient for people who don't have their own computer, who travel, who access the Internet from different places, or who have to use university computers. Unfortunately, they cause many problems (too many to list here – see Appendix A). The main problems are related to a waste of bandwidth and spam, and the fact that every thing the user does, causes international Internet traffic (most commonly to the USA and back). For example, if users of Web-based services want to send e-mail messages to their colleagues, this will cause traffic on the Internet link, whereas if a local e-mail system is used, these messages will remain on the local servers.

Commercial Web-based e-mail services such as Hotmail are also prime targets of spam: on average, 80% of e-mail on Hotmail is spam. Dealing with that requires bandwidth, which is a major reason for avoiding (and even blocking) Commercial Web-based e-mail. Web-based e-mail also uses much more bandwidth than regular e-mail programs because of the advertising, etc., that is included on the Web pages. For these and other reasons (see Appendix A), Web-based e-mail services in a low bandwidth environment should be banned, or only allowed to be used outside working hours, as is done at Makerere University.

But before banning commercial Web-based e-mail, the IT department must make sure that a stable e-mail system is running, and one that has proper protection with a UPS and a back-up system. The mail-server administrator should be very competent. Nothing is more important to users than their e-mail; they need to have confidence in the e-mail system. The reason why some users switched to commercial Web-based e-mail systems at some of the institutions in the case studies was precisely because they had lost e-mail messages (precious scientific data in some cases) through mail-server problems.

There is a growing acceptance at universities that blocking commercial Web-based e-mail requires a usable alternative, and many of them therefore implement their own Web-based e-mail service. An organizational Web-based e-mail system is preferable to commercial offerings, because it does not carry advertisements. However, it is still less efficient than regular e-mail, because Web-based mail is used in real time and mail doesn't just come in the background when bandwidth is available. Using http to carry e-mail traffic may also work counter to any prioritization measures that may be in place. Therefore, it is recommended that Web-based e-mail be used only as a last resort.

Other steps might also be taken to enable users who are away from their primary computer to access their mailboxes remotely. This can be achieved by allowing port-forwarding of the POP or IMAP protocols through the firewall (see Appendix A for more details of this and Web-based e-mail packages).

4.9 Anti-virus software

A virus can spread quickly through an institution's computer network and cause major problems, including consuming most of the available bandwidth in its attempts to find other hosts to connect to, or by sending out many e-mail messages, etc. In addition to educating users about viruses, it is necessary to provide two levels of virus scanning: server-based e-mail virus scanning, and file-based anti-virus software protection on each machine. The anti-virus software market is very competitive, and new, very low-cost competitors are now available that meet industry standards. See Appendix A for a list of anti-virus packages.

The main way in which viruses are spread is though the Microsoft Outlook and Microsoft Outlook Express e-mail programs. This is because these applications have capabilities that are tightly integrated with the Windows operating system, which unfortunately makes them 'insecure by design' and very commonly targeted. This has led some institutions to actively discourage or prohibit the use of Microsoft Outlook, and to require users to use other e-mail client software such as Eudora or Mozilla Mail. Bristol University is one institution that takes

this approach, and they provide no support for Outlook.

In addition, it may be necessary to disable certain features (such as VBScript) on workstations, and to teach users not to launch any executable files that they may receive via e-mail.

4.10 Major problem areas

- Hosting a Web site locally. If the Web site is hosted locally, international users to the site use up the institution's bandwidth. Connections via satellite are often asymmetric, with the uplink smaller than the downlink. International visitors might totally congest the smaller uplink, making it very slow for local users to access the Internet. This also puts international visitors off because the organization's Web page might load too slowly for them.
- Open proxies. There are people on the Internet who find and use open proxies (proxy servers that can be accessed from anywhere on the Internet). They use them for a variety of reasons, such as to avoid paying for international bandwidth or to hide their identity.
- Open relay hosts. An incorrectly configured mail server will be found by unscrupulous people on the Internet and used as a relay host to send bulk e-mail and spam. They do this to prevent getting caught. See Appendix A on how to test for and prevent open relay hosts.
- Peer-to-peer (P2P) networking. Programs such as Kazaa, Morpheus, WinMX and BearShare (successors of Napster) enable users to share files on the Internet. People use them to share things like music files with other users on the Internet, thereby consuming large amounts of bandwidth. The resources shared by other computers are searchable, so the searching and communication with other P2P computers also constantly consumes a large amount of bandwidth, even if no file is actually being downloaded. See Appendix A for details on how to deal with Kazaa.
- People downloading large files such as videos, music or ISO images of software CDs.
- There are programs that are automatically installed if a user is not alert, and then keep on using bandwidth – for example, the so-called Bonzi-Buddy, the Microsoft Network, and some kinds of worms. Some programs are 'spyware', which keep sending information about a user's browsing habits to a company somewhere on the Internet.
- Windows updates. The latest Microsoft Windows operating systems assume that a computer with a LAN connection has a good link to the Internet, and automatically download security patches, bug fixes and feature enhancements from the Microsoft Web site.
- In addition to Windows updates, many other programs and services assume that bandwidth is not a problem, and therefore consume bandwidth for reasons that the user might not predict. For example, the Windows SMB protocol broadcasts information about available shares at regular intervals, and Windows computers hold 'elections' to determine which computer should distribute the list of Windows resources that all users can see when they look in the 'Network Neighbourhood' or 'My Network Places' of the computer. These protocols should be kept off the Internet link using the firewall or router.

4.11 Training

It is possible that training may result in the reduction of bandwidth usage to some extent. For example, if users have an idea about file sizes, they will know what size of file they can send via e-mail (less than 1 MB, preferably), which ones they should upload via FTP (files less than 10 MB in size, for example), and that they should get the IT department to upload larger files overnight. They should also know how to compress files before they are transferred.

Knowing all the techniques for sending advanced queries to a search engine may enable users to find what they need more quickly.

Knowing about viruses that spread via e-mail (such as the 'ILoveYou' virus) will prevent some virus problems.

5. Charging and quotas

Given that university education (including Internet access) is free in some parts of the world, universities in developing countries should aim ultimately to provide Internet access to all students for free. However, this section describes why a case exists for charging for Internet access while bandwidth is still scarce and expensive.

A university may have enough funding to get a modest amount of bandwidth, and will allow as many students and staff as possible to use it. However, as is shown in this study, some people use large amounts of bandwidth, and the activities that take up the vast majority of their high usage are not academic (or even worthwhile) in many cases.

To limit the activities of these users, and to generate additional funding for bandwidth, it may be useful to give all users a bandwidth quota, and then allow them to pay if they need more during that month. In this way, users will not waste bandwidth with unnecessary activities, but will try to conserve their bandwidth in order to stay within their monthly quota.

If a prepaid system is implemented, the inclusion of a subsidized quota is essential because students must have some free access to the Internet.

The danger with this system, however, is that, in a university that really needs more bandwidth, management may feel that a technical solution could be implemented and that it is therefore not necessary for them to try to obtain increased funding for bandwidth. On the other hand, experience at the University of Zululand shows that prepaid access, and a system that produces good statistics, enables the IT department to present a case for more bandwidth to management because it gives them a sense that bandwidth issues are under control and money is well spent.

A few charging mechanisms are presented in Appendix A.

6. Bandwidth

As it is an expensive resource, organizations should ensure that they are purchasing the right amount of bandwidth, and test whether they are getting what they pay for. A university should also aim to use the available bandwidth as best it can, not only by limiting abuse but also by providing access to as many people as possible in order to maximize the use of the resource. For example, if there are periods (such as during weekends) when there is low usage of bandwidth, a university could consider providing free or paid access to the local community. Unused bandwidth is wasted money.

6.1 Bandwidth testing

Bandwidth testing is an issue for both management and accountants, because when institutions buy an expensive commodity like bandwidth they need to determine if they are getting what they are paying for. ISPs might neglect to explain properly the difference between Committed Information Rate (CIR) and Burst Excess (BE). In particular, the link might be sold on the BE capacity of the link, without the organization realizing that it is not getting that amount of bandwidth some (or even most) of the time. This happens because the ISP's other customers are contending for the same bandwidth.

In a shared system, CIR is the guaranteed amount of bandwidth that the institution should get, while BE is everything above that, depending on how much traffic the ISP's other customers are generating. A shared system is generally better value for money, because more bandwidth than if only CIR was obtained would be available some of the time. However, some CIR is necessary in order to guarantee some bandwidth all of the time. It is also necessary to determine whether the institution is getting its fair share of BE, because if there is too much competition for the BE, it is not worth having a shared system.

Typical link speeds* (in order of speed):

Ethernet Local Area Network	100 Mbps
T3, DS3 (North America)	44.736 Mbps
Older type Local Area Network	10 Mbps
T1, DS1 (North America)	1.544 Mbps
DSL (ADSL)	512 Kbps to 1.544 Mbps
Typical African university's Internet connection (in 2003)	512 Kbps to 1 Mbps
ISDN	64 to128 Kbps
Dial-up modems	Up to 56 Kbps

There are tests that can be used to give users or librarians an idea of how much bandwidth they are getting – for example, the Web-based test at <http://www.bandwidthspeedtest.com>, or MicroTik's free Windows-based bandwidth testing utility (see <http://www.microtik.com/download.html>). If there is congestion on a link (i.e. if other people are using it) the user cannot expect to achieve the full bandwidth that the institution has access to. Therefore, accurate tests can really be done only overnight. These tests are described in Appendix A.

6.2 How much bandwidth is needed

Institutions will need to increase their bandwidth from time to time. Typical reasons include:

- Student numbers tend to grow, and universities increase the number of computers they own.
- The volume of resources on the Internet keeps growing, and tends to become ever more bandwidth-hungry.

* See Glossary for an explanation of the difference between data rate and link speed.

- New services on the Internet, such as streaming media, may present new opportunities for education, though it is also possible that streaming media will prove to be useful for entertainment only.

The price of bandwidth is falling, even in developing countries where access is only possible via satellite, so increasing bandwidth is more achievable. However, few institutions undertake studies on how much bandwidth is needed. Normally, an institution simply gets as much as it can afford, or increases its bandwidth because a new cheaper rate enables them to do so, or because Internet access simply becomes too slow. There is nothing wrong with this approach, but clear thinking on the subject enables authorities to decide whether they need more bandwidth or whether they need more control over usage (or both). What follows are some pointers from Hughes Network Systems, the MIMCOM network, and Moratuwa University.

A graph of bandwidth required (and bandwidth required per user) against number of users can be expected to look something like this:

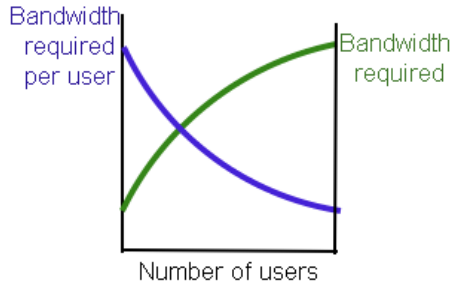

Two reasons why bandwidth per user reduces (but overall bandwidth demand increases) are as follows:

- The higher the available bandwidth, the less is needed per user because requests can be satisfied faster, leaving the 'big pipe to the Internet' open for the next user. Also, the waste of bandwidth due to retransmissions is reduced.
- Not all connected users use bandwidth all the time, because in between requesting Web pages they are also reading them.

For these reasons, an ISP such as Hughes Network Systems make the following assumptions for their DirecPC satellite Internet service:

- ***175 subscribers per 128 Kbps***, which gives 0.73 Kbps per user. This is the same as the MIMCOM network, which currently has between 1000 and 1200 users. The Hughes network can make use of such a low bandwidth per user because it has so many users (see the graphic above).
- ***10% of subscribers will be logged on, and 5% of logged-on users will be active – actually using bandwidth***. This is not a good assumption for the MIMCOM network, where many more than 5% of users are active during peak times, and some sites make large data transfers.

An amount of 0.73 Kbps per user is not much, and in the MIMCOM network it was felt that this should be increased to about 1 Kbps per user; but the MIMCOM network has only about 1400 users. In a larger network, 0.8 Kbps may well be enough if sufficient control and optimization is in place. Of course, the amount of usage also depends on the kind of activity undertaken.

In Moratuwa university, it was found that 1 Kbps per active user (i.e. someone who is online and using the Internet) is too slow, while 3 Kbps per active user is generally sufficient (this can be increased outside working hours).

If users can be limited to a certain amount of bandwidth, deciding how much bandwidth is needed can be done simply by multiplying the number of active users during peak time with the bandwidth that was deemed sufficient. Therefore, if there are 1000 active users during peak time, the university would need 3000 Kbps.

The figure of 3 Kbps was arrived at by Moratuwa through experimentation and deciding what is acceptable.

6.2.1 Avoid 'flat topping'

The graphic below shows a highly congested link at Malawi College of Medicine. The college has a 128 Kbps link to the Internet, and the peak rate is 128 Kbps throughout the working day. When the graph 'flat tops' like this, it means that some users have to wait a very long time for anything to happen, and a bandwidth upgrade is needed (unless the congestion is being caused by unnecessary traffic). It also means that bandwidth is being wasted, as applications such as Web browsers have to keep retrying to do the same thing.

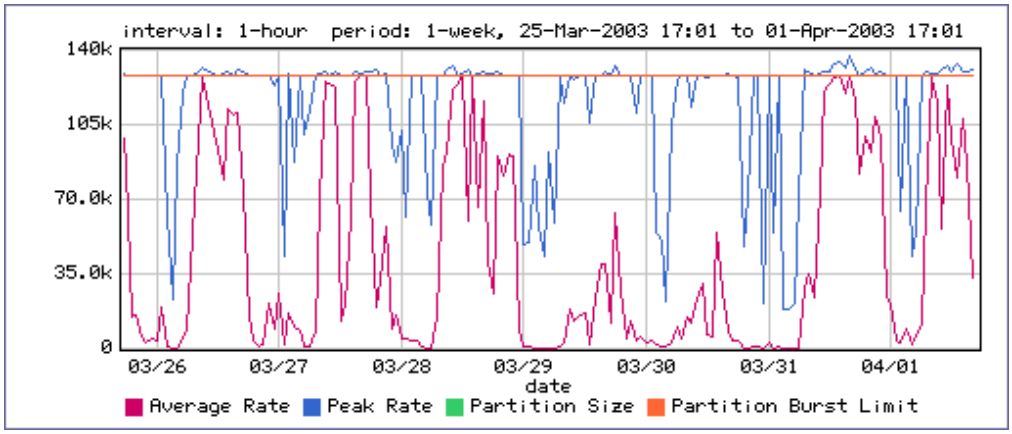

Keeping graphs like this one (or the raw data that produced it) will enable the network administrator to know how congested the connection is and how it changes over time.

6.2.2 Be within the CIR 80% of the time

In the previous example, the College had a 128 Kbps link to the Internet, without the possibility of BE. In a shared network, it works a bit differently.

In a shared network, it is a good idea to be within the CIR 80% of the time (during the working day). For example, if a network has a CIR of 768 Kbps and a BE of 1 Mbps (and if the network team has properly dealt with abuse and unnecessary traffic), the network could use around 768 Kbps 80% of the time, and up to 1 Mbps for the remaining 20% of the time. This would be a network that is making good use of the available bandwidth, without 'flat topping'. In the example below, the usage seldom exceeds 100%; when it does, it means that the network is using BE. This network has enough bandwidth, because it is using BE less than 20% of the working day.

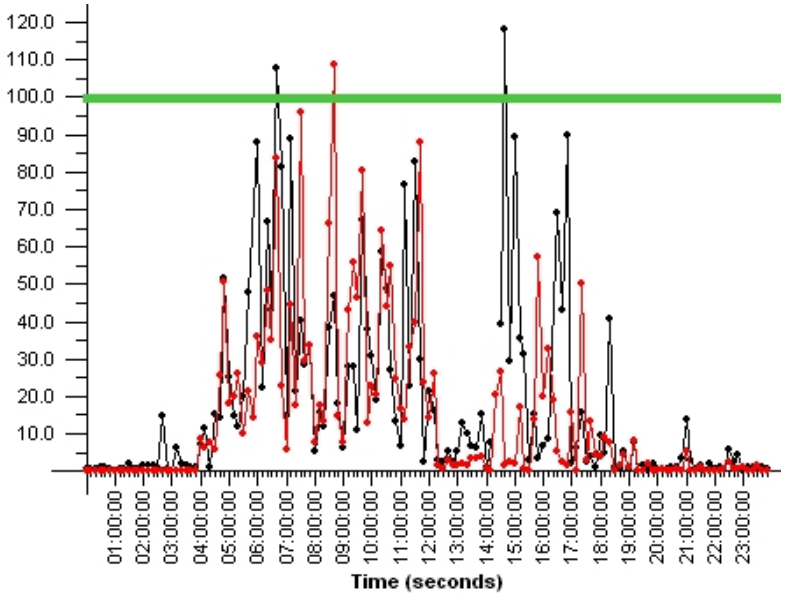

7. Network design

Management could ask the network team to present a diagram (such as the example below) that indicates how network packets get to a Web server on the Internet and back, and showing the position of the proxy servers, for example.

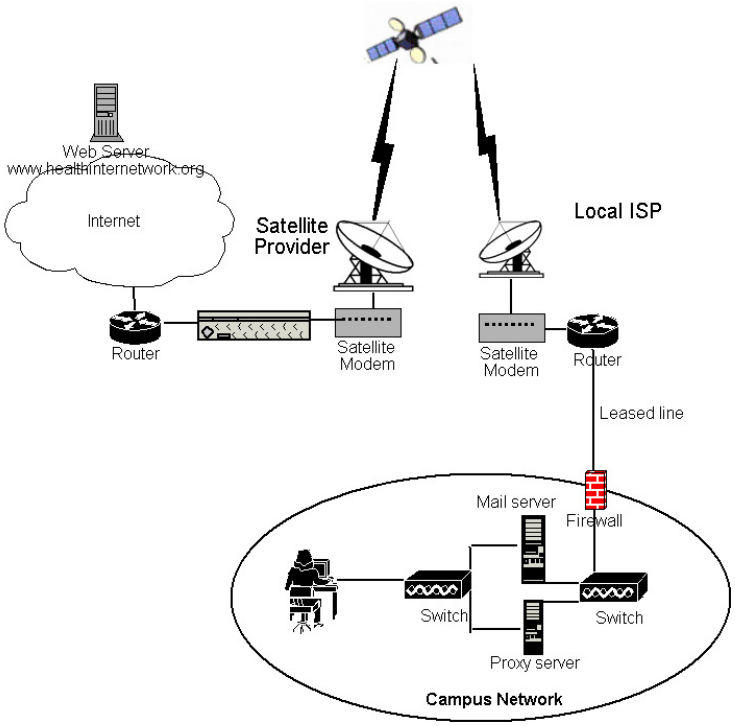

It is also useful to know about the various protocols used on the Internet, and to decide whether all of them should be allowed, or only some.

Web (HTTP) traffic remains the biggest user of Internet traffic (more than Mail, FTP or entertainment-related protocols such as IRC and Kazaa). Web traffic is also increasing its share against protocols such as FTP and Mail because file downloads are increasingly being done using HTTP, and e-mails are being sent and read using the HTTP protocol (through Web-based e-mail services such as Hotmail). Almost everything that is worthwhile for academic purposes is available via HTTP and FTP. E-mail has become essential for personal, administrative and academic communications. HTTPS is essential for secure communications such as online banking and shopping.

A university might decide that, in view of the expense and scarcity of their bandwidth, they will allow users access only to HTTP, HTTPS, FTP and Mail, since most other protocols are either entertainment-related, or are related to system administration and therefore not relevant to most users. Examples of the latter include RDP, pcAnywhere, VNC and Secure Shell.

But this decision must consider other protocols essential to someone's work. For example, an external collaborator might make certain resources available using a different protocol. (This should be discouraged if possible.) If this is the case, an alternative method of access should be devised for the department concerned.

Certain network layouts, discussed in Appendix A, would make it completely impossible to use protocols such as Kazaa, and to only use HTTP, HTTPS, FTP and Mail.

8. Usage policies

In order to maintain control, organizations should get their users to sign a usage policy agreement before they can use the Internet. This policy should describe appropriate use: for example, a user would agree not to use the university's computers for illegal activities or mass e-mailing.

Examples of usage policies appear in Appendix C.

Internet usage policies are essential for any large organization, and require not only a good policy document but also attention from a high-level ICT directorate, because policy implementation cannot be dependent upon the personality of the network administrator but must involve higher-level management.

If a policy is not enforced, it is not worth having. Successful implementation and enforcement of usage policies requires a network to which users have to log on. Authentication is discussed in the next section.

Management needs to think about how to deal with abuse. It is wise to make users aware that their usage is monitored, but also that their privacy not invaded. Where management becomes aware of abuse, a general warning is suggested, rather than an example being made of the first offender; only when warnings are not heeded should action be taken. Over-zealous enforcement has its own dangers: for example, users might begin to think they live in a police state and switch to Web-based e-mail because they think that their e-mail messages are being read; or users who can, may make dial-up connections to the Internet in order to bypass monitoring, but dial-up is expensive, and also thwarts security and anti-virus measures. Pragmatism is the best approach. For example, lists of the top sites that have been visited can be posted on notice boards to indicate that the organization is aware of the activities that users are undertaking.

9. Authentication

9.1 Reasons for authentication

Management would normally want an authentication system to be implemented. This should preferably be a single system, e.g. not different log-ons for e-mail, PC access or the Internet. Many network optimization strategies rely on authentication – for example, for the enforcement of usage policies (to identify users who break the rules described in the policy).

It is recommended that all users sign the policy document before they are given an account to log on to the network. It is not a good idea to let users use the Internet without logging on, because there is no way of knowing why large amounts of bandwidth are being used, why the connection is slow at certain times, and who is abusing the system.

Activities such as hacking into other systems are also less likely if everyone has to log on.

Authentication improves security, enables administrators to trace problems to specific users, and allows the option of charging for usage.

9.2 Authentication systems

Different authentication systems are discussed in Appendix A; there is a technical description of the various systems that enable a single log-on across all systems. This is necessary because most users can deal with only one password; password administration otherwise becomes too difficult. Modern operating systems such as Windows XP, Windows 2000, Windows NT and Linux can all be used to authenticate from a single user directory. While the log-on of older operating systems such as Windows 95, 98 and ME can easily be bypassed, it is possible to enforce a log-on for these systems before users access the Internet. When users open their Web browser and try to access the Internet, they are first presented with a log-on dialog box; they cannot access the Internet before entering their username and password.

On these systems, it might be preferable in a multi-user environment for the institution's Web-based e-mail service to be used rather than an e-mail client. This is because Windows 95, 98 and ME don't support multi-user profiles properly, and therefore cannot keep different users' e-mail separate. Additionally, using Web-based e-mail on these systems allows for authentication to take place in the Web browser instead of at operating-system level.

9.3 Password creep

Password creep is when users share passwords to get around difficulties caused by a strict but ill-conceived security policy. Where a desire for security introduces multiple barriers to users, such as separate passwords for different applications, or regular changes to them, users will start to freely exchange passwords just to get to the Internet or to get their work done. In this way, what was intended to be a highly secure system can become very insecure and the measures become counter-productive. When this occurs, management must recognize that the password policy has failed and needs to be revised. These principles should be followed:

- Password security should be easy for users.
- A single authentication scheme should be implemented. If an e-mail password is not the same as the password used to log on to the workstation or network, users will switch to commercial Web-based e-mail when they have a problem.
- User log-ons should be equivalent where possible – there should be no reason why users would want to find out what someone else's log-on is because that log-on enables them to do more. User log-ons should not give the user access to anything that needs to be protected by strong security.

If these principles are applied, there should not be an incentive for 'password creep'.

10. Managing the IT department

Just as there should be a policy requiring users to obey certain rules in order to use the Internet, there should also be a document outlining the tasks and aims of the IT department. Institutions recently connected to the Internet can easily fall into the trap of having a network administrator deciding everything – whether or not to have content filtering or whether or not to produce statistics and analyse log files, for example. Typically, administrators would do these things only if they are particularly diligent or if they feel like it.

Because they are the 'experts', network administrators can easily say that something cannot be done or can make technical excuses for why something is wrong, when neither are necessarily the case.

Furthermore, the IT team is usually the largest consumer of bandwidth. They download things like service packs and other software, and even 650 MB ISO images. They have greater access to the network than anyone else, and they know how things work. They are more Internet savvy, and are as likely as anyone else to abuse bandwidth.

10.1 Getting the IT department to do what a university needs

A structure should be established that gets the IT department to do what a university needs. To this end, the ICT structure at Makerere University is suggested:

The Makerere University Network (MAKNET) is implemented and supported by the Directorate for Information and Communication Technology Support (DICTS), which is also responsible for IT-related projects.

DICTS, in turn, reports to the ICT Steering Committee (ICTSC). This Committee takes high-level and budgetary decisions about the direction of ICT aims and projects. The Director of DICTS is the secretary to this committee. There is also an Architecture Working Group (AWG), which is composed of representatives from all faculty-level academic units as well as administrative units. The aim is to 'provide a forum for the development and continuous review of the University's information architecture, ensuring that it conforms to the common vision of the end users'.

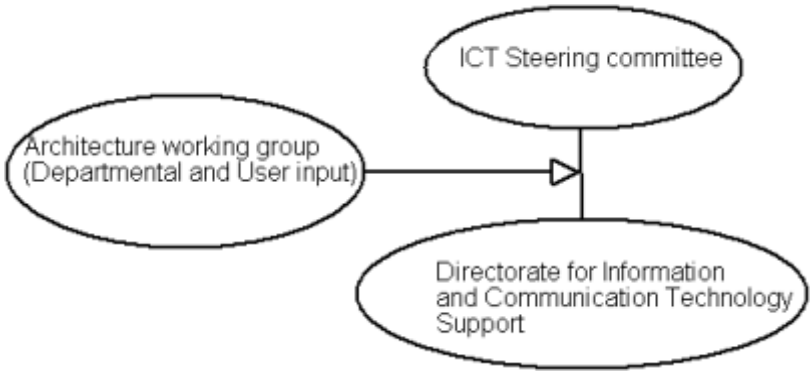

The ICT structure at Makerere frees the network technicians from making policy decisions and enables them to work towards the goals set by this structure. Another attraction of this structure is that departments get formal high-level input in the form of the AWG, where they can say what they need without having to talk only to a technician.

10.2 Retaining good staff

It is often difficult for governments and universities to compete with the private sector for good ICT staff. A university might see its best staff continually being lured away by better salaries in the private sector, even though it is in the business of training people for the benefit of the

country; but it needs a stable network and therefore continuity in staffing.

One way to deal with this problem is the semi-commercial approach taken by the University of Dar es Salaam. Their University Computing Centre Ltd. is a company wholly owned by the University of Dar es Salaam. It supplies computing services and Internet access to the University community in Tanzania, and is business-oriented and self-financing. Being self-financing, it can pay its staff market rates. It also allows the University to benefit from a pool of skills and talent that are developed through working for other institutions and companies. But an inherent danger is that the University's students and researchers might receive lower priority in terms of bandwidth and service than other, more demanding, commercial customers.

A better way, perhaps, is for a university simply to pay competitive rates for such staff.

10.3 Division of work

In any network, there are user issues that need to be addressed, and this is often done by a support engineer who attends to the users' PCs. There is also a central role – staff who look after the network and who are not distracted from doing essential work, such as installing a proper backup system, by individual users' problems. The Malawi College of Medicine case study provides an example of a network where important central functions were neglected because the support team were too busy attending to user issues.

11. Using free software

The principles of free software are explained at Web sites such as <http://www.gnu.org/philosophy/free-sw.html> and <http://www.fsfeurope.org>.
Free Software is often referred to as Open Source. These terms are generally (and in this document) used interchangeably, although there are debates about which name should prevail. The freedom that Free Software refers to is being allowed to read and learn from the source code, and to modify the software as needed.

Free Software, or Open Source should not be confused with Freeware. Freeware is software that is free, but the source code is not available. It is usually free because the programmer has abandoned the project. Useful Open Source programs will continue to improve and be supported if the original programmer(s) abandon the project, because the source code is available

The free-software movement is growing fast, and becoming more successful. Universities can save money by using Open Source software where it makes sense, not only because it is usually free but also because the source code can be modified, as was done at Moratuwa and Zululand. In this regard, see also Makerere University's 'Make or Buy Policy' (Appendix D).

Free software is available for the following reasons:

- An organization such as a university may write a piece of software for its own purposes and, because it is not a commercial company, release it in the public domain. This enables programmers from other institutions to improve the software, which also benefits the original institution.
- Programmers may write free software and make it available (including the source code) to the wider world. When it becomes popular, they are able to make money by adapting the software on request. For example, if a company that is using the software needs a specific feature, it can pay the programmers to add that feature.
- Software is written for dedicated devices. Hewlett-Packard, for example, supports the Samba project (see below), because they use Samba in their high-end print-server devices. They sell hardware but use free software, and therefore have an interest in developing that free-software project. Similarly, the free Linux operating system is used in a multitude of commercial communications, routing and firewall products. These companies support the development of Linux with financial contributions.

As with commercial software, there are good and bad Open Source products available. However, free-software projects typically do have very good support. An institution can commit to free software as a policy, or use only the free software that works for them. However, an organization that never uses free software is likely to be wasting money. Notable free-software products include the Apache Web server (which runs 60% of all Web sites on the Internet), MySQL (which provides the database power for most database-driven Web sites), the Linux and BSD operating systems, Samba (which enables organizations to use a Unix or Linux server as a file and print server and domain controller for a domain of Windows computers), as well as the free C, Python, Perl and PHP programming languages, and the OpenOffice office suite.

A very large number of Free Software projects are hosted (and available for download) at SourceForge - http://sf.net

12. Connection options

12.1 Technologies

Options for connecting to the Internet are discussed in Appendix B. These are VSAT, wireless and leased line. Where a country has an undersea cable connection to the Internet, the organization should attempt to get a leased line connection to this infrastructure.

12.2 Peering

University networks are often among the largest computer networks in a country. As such, they might decide to start initiatives to keep local traffic local. This can be achieved by setting up an Internet Exchange Point (IXP). For example, if a university (which already has its own connection to the Internet) makes a connection to the largest ISP in the country, in order that traffic between the university and customers of the ISP can be routed directly within the country, that is called a peering arrangement, and helps to keep traffic between the university and customers of the ISP within the country. Such an initiative could be extended to make connections to other ISPs and universities within the country.

This arrangement requires some investment, but can save the country money, because less money needs to be paid for bandwidth obtained from foreign satellite companies, and local traffic will reach its destination much faster.

Peering also makes creating local software mirrors more useful, because there are more potential users that can access large downloads without using international bandwidth. For example, the Sri Lanka LEARN network, which is a network of educational institutions that have their own network, has its own mirror sites of popular software download sites such as GNU, Simtel, MySQL and Tucows, as well as of Linux distributions.

These links also give local access to the local online newspaper and business sites, and local e-mail remains within the country.

Unfortunately, peering requires co-operation and trust, and these might be difficult to get right. The link and quote below, illustrates this problem:
<http://www.balancingact-africa.com/news/back/balancing-act_156.html>

> The major issue is one of trust. You need to be able to work with your competitors and in some countries this level of trust has not yet been established. As Brian Longwe told a recent workshop at the Southern African Internet Forum: "Getting any IX/peering arrangement off the ground is 10% technical work and 90% socio-political engineering." He also pointed out the importance of getting ("written") regulatory support. Setting up a local IXP is neither costly nor difficult.

Of course, the link between the institution and the ISP must be cable or wireless. It makes no sense to use a satellite connection for such a link.

Appendix B contains more information on peering.

12.3 Academic networks

Where universities and other educational institutions can make local network connections with each other, they can pool resources and make a joint connection to the Internet. This enables them to negotiate a better deal, because they will be buying more bandwidth: buying in bulk is cheaper. It also enables them to keep traffic between these universities local, and to share certain resources, such as mirrors of software download sites, as described in the previous section.

Such co-operation exists in many countries – for example, the LEARN network in Sri Lanka (<http://www.ac.lk>) and the Janet network in the UK (<http://www.ja.net>). These networks are discussed in more detail in the Moratuwa and Bristol case studies.

13. Glossary

Note: The letters q.v. after a term indicate that a separate entry can be found in the glossary for that term.

Analogue
A connection is said to be analogue if it makes use of continuously variable signals. Regular telephone lines are analogue.

ADSL
Asymmetric Digital Subscriber Line. A form of DSL (q.v.) in which the bandwidth available for downstream connection is significantly larger than for upstream connection.

Asymmetric connection
An Internet connection where the inbound bandwidth to your computer is more than the outbound bandwidth. This is efficient, because for web browsing the web request (a click on a hyperlink, for example, that goes out) consumes far less bandwidth than the resulting web page that appears in your browser

Authentication
Identifying a user as having access rights to a computer system by letting him supply a valid username and password

Bandwidth
The amount of data that can be sent through a connection. Usually measured in bits per second (bps). A full A4 page of English text is about 16,000 bits. (*See also* **Data rate**.)

Burst Excess (BE)
The bandwidth that a network gets from a shared system over and above its CIR (q.v.). The amount of BE a network gets depends on how much the other customers of an ISP (q.v.) are using at that moment. If none of them are using any bandwidth, the full BE will be available.

Cache
An area of memory or disk space holding recently accessed data for quick retrieval.

CD-ROM
Compact Disk-Read Only Memory. A storage medium for digital data, CD-ROMs can hold up to 650 MB.

Committed Information Rate (CIR)
The guaranteed amount of bandwidth that a network should get. (*See also* **Burst Excess**.)

Data rate
The capacity of network connections is measured in Kbps (Kilobits per second) or Mbps (Megabits per second). This is an international convention. The data rate (of a file download, for example) is measured in Kilobytes per second. All other computer data sizes use bytes as the basic unit. File sizes, for example, are measured in Kilobytes or Megabytes. There are 8 bits in a byte, 1024 bytes in a Kilobyte, and 1024 Kilobytes in a Megabyte.

Digital
A connection is digital if it makes use of distinct signals that indicate either 0 or 1. All digital data consists of many 0's and 1's that are combined to make up any kind of data. Even graphics and voice calls can be transmitted in this way.

DNS (Domain Name Service)
While computers use TCP/IP addresses to communicate, humans prefer names like www.google.com. So if someone wants to connect to the server at www.google.com, the domain name service matches that name with a TCP/IP address to enable the connection to take place.

Downlink
The link from the Internet. When users click on a hyperlink, this click is transmitted to the Internet via the uplink. The Web page that consequently loads in the browser is transmitted from the Internet via the downlink.

DSL
Digital Subscriber Line. DSL is a set of protocols for high-speed data communication over existing copper telephone lines by utilizing unused frequencies. DSL can allow voice and high-speed data to be sent simultaneously over the same line. Because the service is 'always

available', end-users don't need to dial in or wait for call set-up.

Ethernet
The standard protocol for communication on a LAN. When connecting to the Internet, Ethernet is used to communicate with the router. The router communicates with other routers via Ethernet or other protocols.

Firewall
Software that either protects a whole network, or sometimes just a single computer from unauthorised access from a network such as the Internet.

Free Software
Free software is software that can be freely used, modified, and redistributed with only one restriction: any redistributed version of the software must be distributed with the original terms of free use, modification, and distribution. Free software may be packaged and distributed for a fee; the "free" refers to the ability to reuse it, modified or unmodified, as part of another software package. As part of the ability to modify, users of free software may also have access to and study the source code. (See also **Open Source)**

FTP
File Transfer protocol. Used for transferring files between computers via the Internet.

HTTP
HyperText Transfer Protocol. A protocol for the exchange of HTML documents (Web pages), as well as other things.

HTTPS
HyperText Transmission Protocol, Secure. Used for exchanging secure HTML documents (Web pages), for example for on-line shopping or banking, where credit card information might be used.

ICT
Information and Communications Technology. Telephone lines, mobile phones, computers, networks and the Internet.

Internet
The worldwide network of networks based on the TCP/IP protocol. The Internet is not an online service and has no real central 'hub'. Rather, it is a collection of millions of networks, online services, and single-user components.

Internet Exchange Point (IXP)
A place where the networks of different ISPs are connected to permit the exchange of Internet traffic. Large academic or governmental networks typically also connect to IXPs. (*See also* **Peering**.)

Intranet
Internal systems, based on Internet technology, designed to connect the members of a specific closed-user group. An Intranet is a private Internet: a private network, usually a LAN or WAN, that enables the use of Internet-based applications in a secure and private environment. As on the public Internet, Intranets can host Web servers, FTP servers, and any other IP-based services.

IMAP
Internet Message Access Protocol. A mail server protocol that allows you to store all your messages and any changes to them on the server and/or on your computer's hard disk

Internet backbone
The fast, high bandwidth part of the Internet that connects the economically more powerful countries. It is not a precise term, but indicates the fibre connected part of the Internet where bandwidth is not a major issue.

IP Address
A number such as 216.239.51.100 that is used to address a networked device or computer on the Internet. IP addresses are necessary for routers to carry network traffic to their intended destination.

IRC
Internet Relay Chat. A system that enables participants from all over the world to communicate with each other in real time via the Internet. It is normally used for discussions between people with similar interests.

ISO image
The complete contents of a CD-ROM, packed into one file. CD-writer software can be used to create the image and to write the image on to a blank CD.

ISP
Internet Service Provider. A company that provides a connection to the Internet.

Janet (Joint Academic Network)
A network connecting UK academic institutions with each other and the Internet.

Javascript
A programming or scripting language that adds logic (such as calculations) to a web page.
LAN
Local Area Network. A short-distance network used to link a group of computers together within a building. LANs are typically limited to distances of less than 500 metres and provide low-cost, high-bandwidth networking capabilities within a small geographical area.
Latency
On satellite networks, the signal has to travel a long way (35 000 km) to the satellite, and then back again. This delay is called latency. It can also be experienced in a fibre connection halfway around the world, but is particularly high and unavoidable with satellite connections.
LEARN network
Lanka Academic and Research Network. A network in Sri Lanka that connects academic institutions such as universities to each other and to the Internet.
Link speed
See **Data rate**
Long fat pipe network
A network connection that has relatively high bandwidth, but high latency, such as a satellite network. The design of TCP/IP (especially earlier implementations), interprets long delays as congestion, and therefore if it takes long to receive an acknowledgement packet, the sender reduced the rate of sending data. This is the wrong approach for a "long fat pipe network" that is not congested, and causes lower throughput than is possible.
Mailing list
A way of distributing information via email on a topic of mutual interest to all who are subscibed. A mailing list can either consist of messages sent by a single entity (for example an announcement list), or all subscribers may participate by sending their own views or questions.
Mirror
In order to save bandwidth, large software download facilities are often recreated or mirrored on servers in other countries in order to save international bandwidth.
Operating system
The low-level software that handles the interface to peripheral hardware, schedules tasks, allocates storage, and presents a default interface to the user when no application program is running. Examples include Windows, Unix, Mac OS and NetWare.
Open proxies
Proxy servers that will accept connections from anywhere, and are therefore abused by individuals to make sure their activities on the Internet cannot be traced.
Open relay hosts
Mail servers that will accept connections from anywhere, and are therefore abused by commercial interests on the Internet to distribute spam.
Open Source
In general, open source refers to any program whose source code is made available for use or modification as users or other developers see fit. (Historically, the makers of proprietary software have generally not made source code available.) Open source software is usually developed as a public collaboration and made freely available. (See also **Free Software**)
Packet
Any data (such as an e-mail) that is to be transmitted over a network, is divided into parts, each of which becomes a packet. The packet also contains the necessary addressing information to get it to its destination.
PDF
Portable Document Format. A document format for which free reader software exists for any operating system. An author can distribute a PDF with the knowledge that the recipient should be able to read the document without having to pay for any software. Another reason for distributing a document in PDF format is that the recipient cannot modify the document.
Peer-to-Peer Networking (P2P)
On the Internet, peer-to-peer is a type of transient Internet network that allows a group of computer users with the same networking program to connect with each other and directly access files from one another's hard drives. Napster and Gnutella are examples of this kind of peer-to-peer software. Corporations are looking at the advantages of using P2P as a way for employees to share files without the expense involved in maintaining a centralized server and as a way for businesses to exchange information with each other directly.
Peering
When two or more organizations, companies or ISPs that have their own connections to the Internet create network links between them, it is a peering arrangement. The purpose is to keep traffic between them direct and local, rather than have it routed via third parties.

POP
Post Office Protocol. Makes available client-server e-mail messaging. Messages sent to you are stored on servers at your Internet service provider until you connect and retrieve them. When you use POP to access your e-mail, your messages are downloaded to your computer and removed from the mail server when you retrieve them.

Postscript
A printing language that can also be used to create PostScript documents. As a document format it is similar to PDF in that is used for the same reasons – universal reader software and the reader cannot modify the document.

Proxy server
A server used for caching previously accessed Web pages and files to prevent them from being retrieved from the Internet multiple times. (*See also* **Cache**.)

Router
A device which forwards packets between networks.

Script
A simple program creatable by users that performs a simple task, such as to download certain files every morning at 2 a.m.

Serial port
A connector at the back of the computer for attaching external devices.

Server
A computer that handles requests for data, e-mail, file transfers and other network services from other computers (clients).

Spam
Unsolicited commercial e-mail.

Split DNS or "split horizon"
Using different pointers on internal and external DNS servers in order to present a different view of the organization's domain to the inside and outside worlds. For example, if a university uses an ISP, then all users of that ISP including the campus users will be directed to an address on the campus network, while the rest of the world is directed to a server on the Internet backbone in Europe.

Spyware
A small program that gets installed without the user's knowledge, and that send details about your browsing habits to the company that benefits from this knowledge.

TCP/IP
Transmission Control Protocol/Internet Protocol. The basic communication language or protocol of the Internet. It can also be used as a communication protocol in intranets (q.v.).

Unix
An operating system. There are many versions, distributions and makes of Unix. The most well known are Sun Solaris, Linux, IBM AIX, and BSD. Some are free, while others are very expensive.

Uplink
The link to the Internet, which is sometimes smaller than the downlink. When a user clicks on a hyperlink, this click is transmitted to the Internet via the uplink. The web page that consequently loads in his browser is transmitted from the Internet via the downlink. Communication from the local computer to the Internet goes through the uplink and communication from the Internet to the local computer goes through the downlink.

UPS
Uninterruptible power supply. A device with a strong battery that supplies power to a server or other equipment, and keeps doing so for a short while if the power fails. Modern UPSs connect to the computer's serial port and provide information such as the battery time remaining, allowing the computer to be shut down gracefully before complete loss of power.

Virus
A malicious computer program that replicates itself to other programs.

VSAT
Very Small Aperture Terminal. An earth station, used for the reliable transmission of data, video, or voice via geo-stationary satellite, with a relatively small dish-antenna (often 2.4m or 3.8m in diameter).

WAN
Wide Area Network. Two or more local area networks (see LAN) joined together over any geographical distance.

Worm
A kind of computer virus that only replicates itself. It does not really harm a computer, but the cumulative effect of its replication can use up vast amounts of bandwidth.