Appendix A
Practical technical implementation

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. When the next person retrieves these pages, they are served from the local proxy server instead of from the Internet. This results in faster Web access in most cases. Bandwidth usage can also be reduced. When a proxy server is implemented, the administrator should also be aware that some pages are not cacheable – for example, pages that are the output of server-side scripts. Loading of Web pages is also affected. Where previously the page would begin to load slowly, first showing some text and then displaying the graphics one by one, in a network with a proxy server, there could be a delay when nothing seems to happen, and then the page will load almost at once. The time it takes to load the whole page might take 10 seconds, whereas without a proxy server it may take 30 seconds, but unless this is explained to some impatient users, they may say the proxy server has made things slower. It is usually the task of the network administrator to deal with user perception like these.

1.1 Proxy server products

Squid. Open source Squid is the de facto standard at universities. It is free, reliable, easy to use and can be enhanced with for example content filtering and advertisement blocking. Squid produces logs that can be analysed using software such as Sawmill (commercial) or Awstats, Webalizer (Open Source), both of which produce good graphical reports. In most cases, it is easier to install as part of the distribution than to download it from www.squid-cache.org (most Linux distributions such as Debian, as well as other versions of Unix such as NetBSD and FreeBSD come with Squid). A good Squid configuration guide can be found at squid-docs.sourceforge.net/latest/book-full.html.

Microsoft Proxy server 2.0. Not available for new installations because it has been superseded by Microsoft ISA server and is no longer supported. It is nonetheless used by some institutions, although it should perhaps not be considered for new installations.

Microsoft ISA server. A very good proxy server program, that is arguably too expensive for what it does. However, with academic discounts it may be affordable to some institutions. It produces its own graphical reports, but its log files can also be analysed with popular analyser software such as Sawmill. Administrators at a site with MS ISA Server should spend sufficient time getting the configuration right; otherwise MS ISA Server can itself be a considerable bandwidth user. For example a default installation can easily consume more bandwidth than the site has used before, because popular pages with short expiry dates such as news sites, are continually being refreshed. Therefore it is important to get the pre-fetching settings right, and to configure pre-fetching to take place mainly overnight. ISA Server can also be tied to content filtering products such as WebSense.

Links: www.microsoft.com/isaserver/ and www.isaserver.org/

Other products include:

- Sun ONE Web Proxy Server (formerly iPlanet Web Proxy Server) is a Sun Solaris product. See www.sun.com
- Novell BorderManager. See www.novell.com/products/bordermanager/
- High end products aimed at ISP's: NetCache (www.netapp.com), and CacheFlow (www.bluecoat.com).

1.2 Preventing users from bypassing the proxy server

Techniques for bypassing the proxy server can be found at <http://www.antiproxy.com>. This is useful for administrators to see how their network measures up against these techniques.

Trust. In the layout below, the administrator has to trust that his users will not bypass the proxy server.

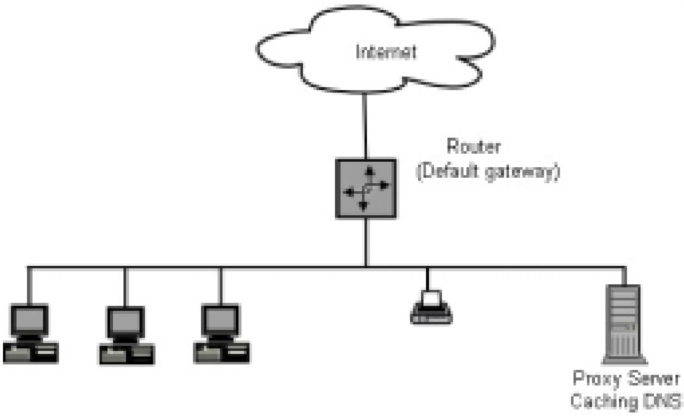

In this case the administrator typically uses one of the following techniques:

- Not giving out the default gateway address through DCHP. However, some network-savvy users who want to bypass the proxy might find or guess the default gateway address.
- Using domain or group policies is very useful for configuring the correct proxy server settings for Internet Explorer on all computers in the domain, but is not very useful for preventing the proxy to be bypassed, because it depends on a user logging on to the NT domain. A user with a Windows 95/98/ME computer can cancel his log-on and then bypass the proxy, and someone who knows a local user password on his Windows NT/2000/XP computer, can log on locally and do the same.
- Begging and fighting with users.

The only way to ensure that proxies cannot be bypassed is by using the correct network layout, by using one of the three techniques described below.

Firewall

A more reliable way to ensure that PCs don't bypass the proxy can be implemented using the firewall. The firewall can be configured to allow only the proxy server through, i.e. to make HTTP requests to the Internet. All other PCs are blocked, as shown in the diagram below.

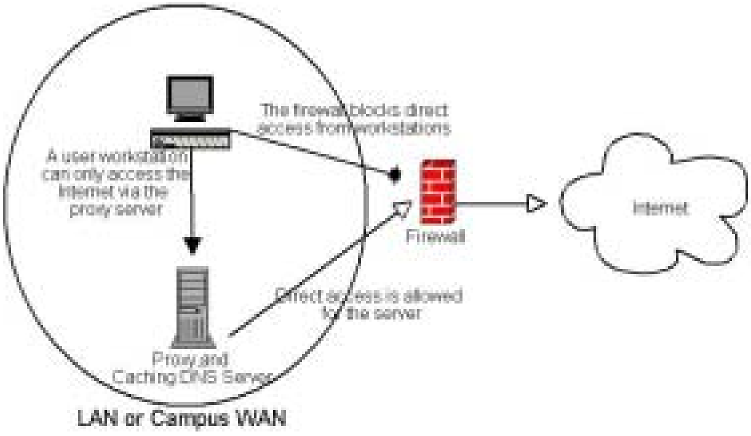
The firewall blocks direct access from workstations
A user workstation can only access the Internet via the proxy server
Firewall
Internet
Direct access is allowed for the server
Proxy and Caching DNS Server
LAN or Campus WAN

Relying on a firewall as in the above diagram may or may not be sufficient, depending on how the firewall is configured. If it only blocks access from the campus LAN to port 80 on Web servers, there will be ways for clever users to find ways around it. Additionally, they will be using other protocols such as Kazaa.

Two network cards

Perhaps the most reliable method is to install two network cards in the proxy server and connect the campus network to the Internet as shown below. In this way, the network layout makes it physically impossible to reach the Internet without going through the proxy server.

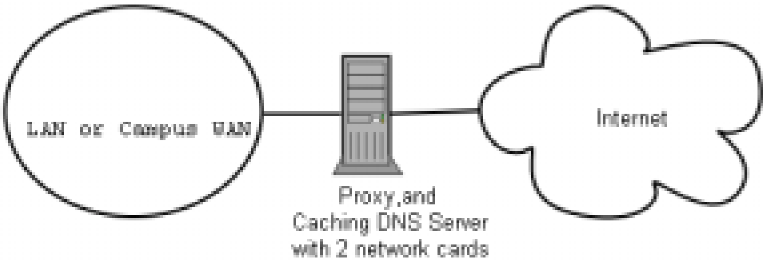

The proxy server in this diagram should not have IP forwarding enabled, unless the administrators knows exactly what they want to let through.

Policy-based routing

One way to prevent bypassing of the proxy using Cisco equipment is with policy routing. The Cisco router transparently directs Web requests to the proxy server. This technique is described in the Makerere University case study. The advantage of this method is that, if the proxy server is down, the policy routes can be temporarily removed, allowing clients to connect directly to the Internet.

1.3 Mirroring a Web site

With permission of the owner or web master of a site, the whole site can be mirrored to a local server overnight (if it is not too large).

This is something that might be considered for important Web sites that are of particular interest to the organization or that are very popular with Web users. This may have some use, but it has some potential pitfalls. For example, if the site that is mirrored contains CGI scripts that require interactive input from the user, this would cause problems. An example is a Web site that require people to register online for a conference. If someone registers online on a mirrored server (and the mirrored script works), the organizers of the site will not have the information that the person registered.

Because mirroring a site may infringe copyright, this technique should only be used with permission of the site concerned.

If the site runs rsync, the site could be mirrored using rsync (this is the best way). See section 14 for more details on rsync.

If the remote Web server is not running rsync, the recommended software to use is a program called wget. It is part of most versions of Unix/Linux. A Windows version can be found at <space.tin.it/computer/hherold>. There is also a Windows installer-based version at <http://homepage.mac.com/shadowboxer/unxutils.exe>.

A script can be set up to run every night on a local Web server and do the following:

- Change directory to the Web server document root: for example, /var/www on Unix, and C:\Inetpub\wwwroot on Windows.
- Mirror the Web site using the command: wget --cache=off –m http://www.python.org

The mirrored Web site will be in a directory www.python.org. The Web server should now be configured to serve the contents of that directory as a name-based virtual host. Set up the local DNS server to fake an entry for this site. For this to work, client PCs should be configured to use the local DNS server(s) as the primary DNS. (This is advisable in any case, because a local caching DNS server speeds up Web response times.)

1.4 Pre-populate the cache using wget

Instead of setting up a mirrored Web site as described in the previous section, a better approach is to populate the proxy cache using an automated process. This method has been described by

J. J. Eksteen and J. P. L. Cloete of the CSIR in Pretoria, South Africa, in a paper entitled “Enhancing International World Wide Web Access in Mozambique Through the Use of Mirroring and Caching Proxies”. In this paper
(available at <http://www.isoc.org/inet97/ans97/cloet.htm>) they describe how the process works:

> “An automatic process retrieves the site's home page and a specified number of extra pages (by recursively following HTML links on the retrieved pages) through the use of a proxy. Instead of writing the retrieved pages onto the local disk, the mirror process discards the retrieved pages. This is done in order to conserve system resources as well as to avoid possible copyright conflicts. By using the proxy as intermediary, the retrieved pages are guaranteed to be in the cache of the proxy as if a client accessed that page. When a client accesses the retrieved page, it is served from the cache and not over the congested international link. This process can be run in off-peak times in order to maximize bandwidth utilization and not to compete with other access activities.”

The following command (scheduled to run at night once every day or week) is all that is needed (repeated for every site that needs pre-populating).

```
wget --proxy-on --cache=off --delete after –m http://www.python.org
```

Explanation:

- -m: Mirrors the entire site. wget starts at www.python.org and follows all hyperlinks, so it downloads all subpages.
- --proxy-on: Ensures that wget makes use of the proxy server. This might not be needed in set-ups where transparent proxying is employed.
- --cache=off: Ensures that fresh content is retrieved from the Internet, and not from the local proxy server.
- --delete after: Deletes the mirrored copy (the mirrored content remains in the proxy cache if there is sufficient disk space, and the proxy server caching parameters are set up correctly.

In addition, wget has many other options, – for example, to supply a password for Web sites that require them.

When using this tool, Squid should be configured with sufficient disk space to contain all the pre-populated sites and more (for normal Squid usage for pages other than the pre-populated ones). Fortunately, disk space is becoming ever cheaper and disk sizes far larger than before. However, this can only be done with a few selected sites. These sites should not be too big for the process to finish before the working day starts, for example, and an eye should be kept on disk space.

1.5 Cache hierarchies

When an organization has more than one proxy server, the proxies can share cached information among them. For example, if a Web page exists in server A's cache, but not in the cache of server B, a user connected via server B might get the cached object from server A via server B. ICP (Inter-Cache Protocol) and Cache Array Routing Protocol (CARP) can share cache information.

CARP is considered the better protocol. Squid supports both protocols, and MS ISA Server supports CARP. See <http://squid-docs.sourceforge.net/latest/html/c2075.html> for more information. This sharing of cached information reduces bandwidth usage in organizations where there are more than one proxy.

1.6 Transparent proxying

A cache server can be configured to cache Web requests transparently. This means that whether the browser has been configured to use the proxy server or not, Web requests will be passed to the proxy, and the proxy will deliver the content to the user. For this to work, user PCs should be prevented from bypassing the proxy server by physical layout, or else Web requests should be redirected to the proxy using Cisco policy routing, as described in the

Makerere case study. The documents below describe how to implement transparent proxying with Squid:
<http://www.squid-cache.org/Doc/FAQ/FAQ-17.html>
<http://en.tldp.org/HOWTO/mini/TransparentProxy-2.html>

1.7 Proxy specifications

On a university campus network, there should be more than one proxy server, both for performance and also for redundancy reasons.

With today's cheaper and larger disks, powerful proxy servers can be built, with 50 GB or more disk space allocated to the cache, for example. Disk performance is important, therefore the fastest SCSI disks would perform best. RAID or mirroring is not recommended.

It is also recommended that a separate disk be dedicated to the cache. For example, one disk could be for the cache, and a second for the operating system and cache logging. Squid is designed to use as much RAM as it can get, because when data is retrieved from RAM it is much faster than when it comes from the hard disk. For a campus network, RAM memory should be 1GB or more:

- Apart from the memory required for the operating system and other applications, Squid requires 10 MB of RAM for every 1 GB of disk cache. Therefore, if there is 50 GB of disk space allocated to caching, Squid will require 500 MB extra memory.
- The machine would also require 128 MB for Linux and 128 MB for X-windows.
- Another 256 MB should be added for other applications and in order that everything can run easily. Nothing increases a machine's performance as much as installing a large amount of memory, because this reduces the need to use the hard disk. Memory is thousands of times faster than a hard disk. Modern operating systems keep frequently accessed data in memory if there is enough RAM available. But they use the page file as an extra memory area when they don't have enough RAM.

1.8 Authentication

In order to establish the identity of Web users, it is necessary to allow them to log on before using the Internet. If this is done, and abuse takes place, it would be possible to determine which user is responsible.

In a network where there is a Windows security context, such as a domain, MS ISA Server, MS Proxy Server and Squid can get a user log-on from the domain. MS ISA Server and MS Proxy Server would have this configuration by setting authentication to take place in the proxy server configuration program. Squid can be configured to use either Unix authentication or Windows domain authentication by using the pam_auth module for Squid, and the authenticate_program directive in squid.conf.

2. DNS caching

Caching-only DNS servers don't host zones (they are not authoritative for any domains) but rather just cache results from queries asked them by clients. Just like a proxy server that caches popular Web pages for a certain time, DNS addresses are cached until their time to live (TTL) expires. This will reduce the amount of DNS traffic on a Wide Area Network (WAN), as the DNS cache may be able to satisfy many of the queries locally.

Of course, client computers must be configured to use the caching-only name server as their DNS server. When all clients use this server as their primary DNS server, it will quickly populate a cache of IP addresses to names, so that previously requested names can quickly be resolved.

DNS servers that are authoritative for a domain also act as cache name-address mappings of hosts resolved by them.

Bind (Named) in Unix. When installed and running, this will act as a caching server (no configuration is necessary). Bind can be installed from a package such as a Debian package or an RPM. Installing from a package is usually the easiest method.

In Debian, type:
apt-get install bind9

Windows NT

To install the DNS service on Windows NT4: select Control Panel > Network > Services > Add > Microsoft DNS server. Insert the Windows NT4 CD when prompted.

Configuring a caching-only server in NT is described in Knowledge Base article 167234. From this article:

> "Simply install DNS and run the Domain Name System Manager. Click on DNS in the menu, select New Server, and type in the IP address of your computer where you have installed DNS.
>
> You now have a caching-only DNS server."

Windows 2000

Install DNS service: Start > Settings > Control Panel > Add/Remove Software > Add/Remove Windows Components > Components > Networking Services > Details > Domain Name System (DNS)

Start the DNS MMC (Start > Programs > Administrative Tools > DNS)

From the Action menu select "Connect To Computer..."

In the Select Target Computer window enable "The following computer:" and enter the name of a DNS server you want to cache.

If there is a . [dot] in the DNS manager (this appears by default), this means that the DNS server thinks it is the root DNS server of the Internet. It is certainly not. Delete the . [dot] for anything to work.

3. Content filtering

Content filtering enables an organization to save bandwidth by blocking access to certain categories of sites, such as MP3 and video download sites, online gambling, shopping, pornography and advertisements. Web access logs of almost any organization show that people spend amazing amounts of time and bandwidth on these kinds of sites, as well as on sites that are innocent but really waste an incredible of time, such as <http://www.newfunpages.com>.

There are many products for content filtering, most are listed in the Google Directory (<http://directory.google.com>) under the following headings:
Computers > Software > Internet > Servers > Proxy > Filtering
Computers > Software > Internet > Servers > Proxy > Filtering > Censorware

Some of these products are intended for home use, and not suitable for universities or other large institutions.

3.1 Open-source products

DansGuardian. This is claimed to be the fastest open-source content-filtering product for Squid. The product itself is free for educational institutions (it is open source but not free for commercial companies). It is also free to download the blacklist once. But because new undesirable sites appear on the Internet daily, it is necessary to keep the blacklist up to date. To keep the list up to date, a paid-for daily or weekly update is necessary. This service is still very cheap (at the time of writing $120 per year for a weekly download, which is far cheaper than commercial products). See <http://www.dansguardian.org> for more information.

SquidGuard. Another fast content-filtering add-on for Squid. It is open source and even the blacklist is free. See <http://www.squidguard.org>.

3.2 Commercial products

Websense <http://www.websense.com>: Websense is a content-filtering solution that integrates with many products, such as MS ISA Server and Cisco Pix Firewall. A list of network products that can be integrated with Websense can be found at <http://www.websense.com/products/about/datasheets/index.cfm>. Websense comes with good reporting tools; for example, in the screenshot below, a manager can see which Web sites the user Jane Doe visited.

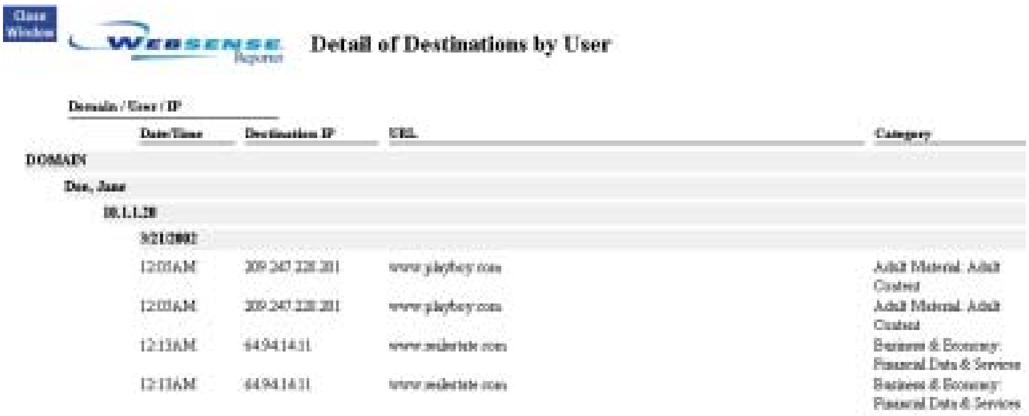

Detail of Destinations by User

Domain / User / IP	Date/Time	Destination IP	URL	Category
DOMAIN				
Doe, Jane				
10.1.1.20				
	9/21/2001			
	12:05AM	209.247.228.201	www.playboy.com	Adult Material: Adult Content
	12:03AM	209.247.228.201	www.playboy.com	Adult Material: Adult Content
	12:13AM	64.94.14.31	www.realestate.com	Business & Economy: Financial Data & Services
	12:13AM	64.94.14.31	www.realestate.com	Business & Economy: Financial Data & Services

N2H2 See <http://www.n2h2.com>. N2H2 is claimed to have the largest database of undesirable Web sites, and N2H2 products can be used in many different systems.

3.3 Integrated products

Netpilot. Equiinet's NetPilot is a server that is sold as an all-in-one proxy server, content filter, mail server, mail virus scanner and firewall. The content filtering makes use of N2H2's list. NetPilot is very easy to use, and all options are configured using a Web front-end. See <http://www.netpilot.com> for more information. Since NetPilot comes as a pre-installed server, organizations wishing to use it must ensure that its hardware specifications are sufficient for the number of users it has to serve. NetPilot is based on Linux, Squid and Postfix. Its advertisement blocking is based on Adzapper (see below).

The screenshot below shows NetPilot's category selection area.

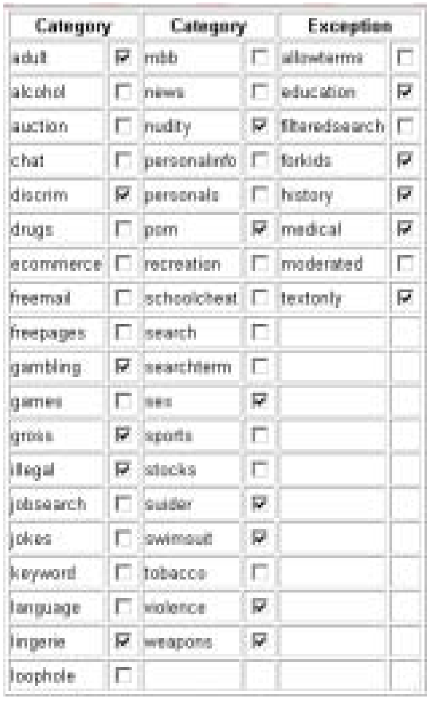

Category		Category		Exception	
adult	☑	mbb	□	allowterms	□
alcohol	□	news	□	education	☑
auction	□	nudity	☑	filteredsearch	□
chat	□	personalinfo	□	forkids	☑
discrim	☑	personals	□	history	☑
drugs	□	porn	☑	medical	☑
ecommerce	□	recreation	□	moderated	□
freemail	□	schoolcheat	□	textonly	☑
freepages	□	search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

SmoothWall. SmoothWall is a Linux-based firewall product that can be installed on any PC server. It is also configurable through a Web front-end, and can be used as a proxy server with content filtering. See <http://www.smoothwall.co.uk> and <http://www.smoothwall.org>.

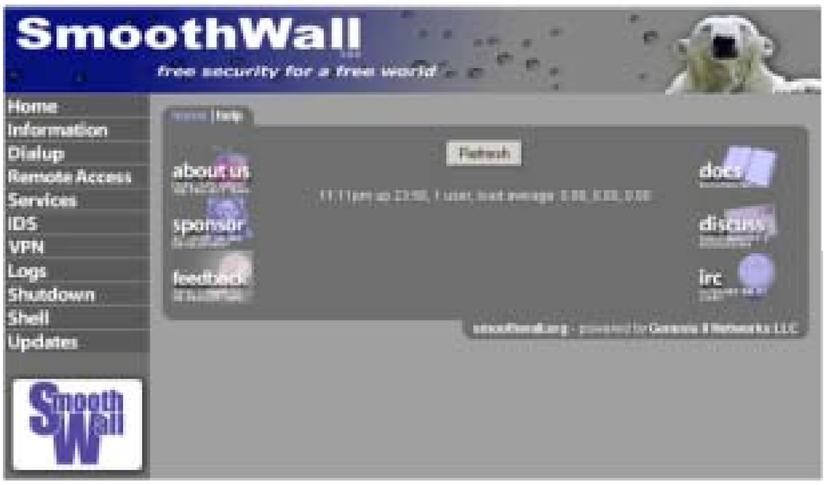

3.4 Advertisement blocking

Addzapper. “This is a redirector for Squid that intercepts advertising (banners, popup windows, flash animations, etc), page counters and some Web bugs (as found). This has both aesthetic and bandwidth benefits.” It is open source, and can be downloaded from <http://adzapper.sourceforge.net>.

Adservers. Another way to redirect Squid is to use the Squid adservers list available at <http://pgl.yoyo.org/adservers>. This file can be saved as /etc/squid.adservers on the proxy server. Then add these lines to squid.conf:

```
acl ads src "/etc/squid.adservers"
http_access deny ads
```

Hosts file. The list at <http://pgl.yoyo.org/adservers> is also available in hosts file format. This list can be added to a PCs hosts file (in c:\winnt\system32\drivers\etc\hosts on Windows NT/2000). It points the hostname of the ad server to 127.0.0.1, which is the local PC. It works because Windows first consults the hosts file, then the DNS to resolve a name. Since the hosts file resolves the name to the localhost, the browser is unable to retrieve the advertisement, and no bandwidth is used. The drawback of this approach is that the hosts file should be updated on all PCs.

Blocking pop-ups advertising in browsers. Many modern browsers (such as Mozilla and Safari) can block pop-up advertising unless the user specifically allow them for a trusted site. This functionality can be added to Internet Explorer by using the Google toolbar

Blocking Windows messenger pop-ups. There is another type of pop-up advertising that sends Windows messenger pop-ups across the Internet to Windows based computers. These can be stopped by disabling the Windows Messenger service.

Other techniques

There are many other ways to block advertising. A comprehensive list of techniques is at <http://www.flourish.org/adremove>.

4. Monitoring

There are a vast number of network monitoring, capturing and analysis tools available. One comprehensive list is maintained at <http://www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html>.

Before using any of them, it is a good idea to pause and decide when and how to apply them. These can be used to account for bandwidth usage, to troubleshoot networks or to identify high bandwidth users. Log files are also analysed to show how the Internet connection is being used.

4.1 Accounting for bandwidth usage

A network administrator should be able to add up the bandwidth used by his proxy, mail and Web servers, using the logs produced by these (with software or server-based analysis). He should then also measure the bandwidth to the Internet by getting data from the router or firewall. For example, in the diagram below, the mail-server and proxy-server logs should be analysed and aggregated using log-file analysis software (see below). The router or firewall traffic should also be analysed (using software that aggregates SNMP counters on the router; see below for packages), and the administrator should expect this to be roughly equal to what is logged in the mail-server and proxy-server logs of the two servers. If it is more, he should be able to explain the difference – perhaps that one of the two servers is acting as a router (forwarding packets) and users in the campus network are using Kazaa.

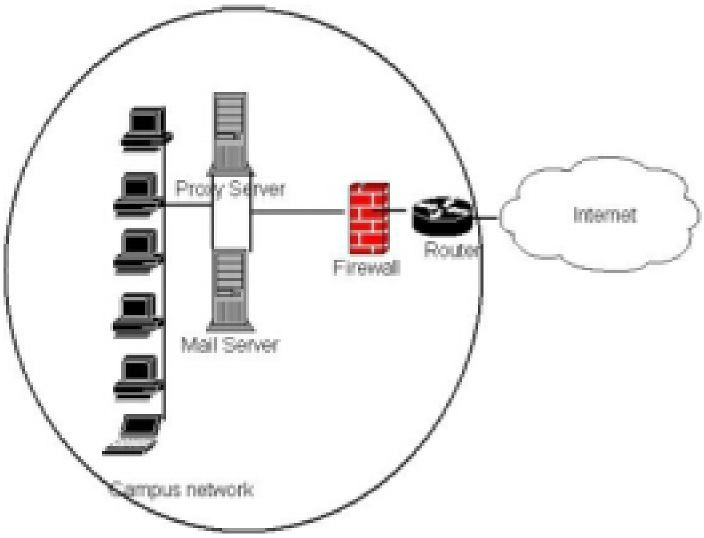

4.2 Router/firewall-based traffic analysis

This type of analysis measures what is going through the routers or firewalls. It is the actual bandwidth usage that is measured here.

MRTG. The Multi-Router Traffic Grapher (MRTG) a very useful tool that combines data collection with graphical trending to monitor the traffic load on network links. It can graph the bandwidth in and out of any SNMP-enabled network device; these include servers, routers, firewalls and switches. It comes with virtually all versions of Unix, and is available for Windows. MRTG's log files do not grow owing to a data consolidation algorithm. See <http://people.ee.ethz.ch/~oetiker/webtools/mrtg>.

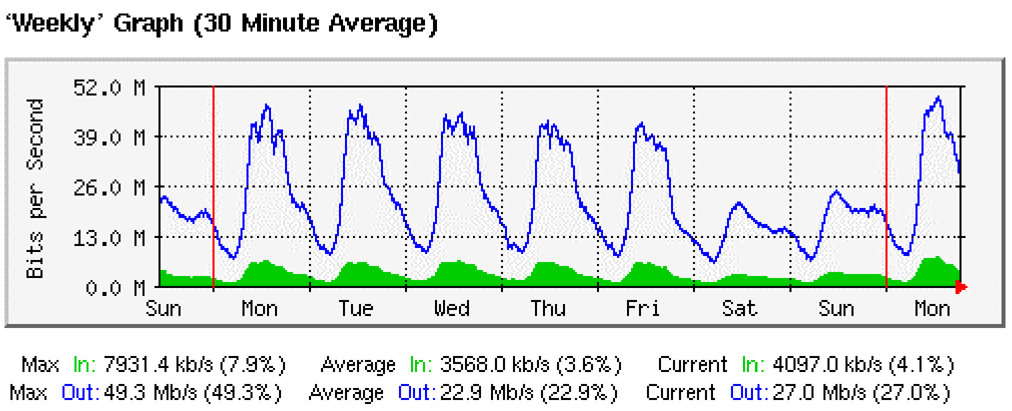
'Weekly' Graph (30 Minute Average)

Max In: 7931.4 kb/s (7.9%) Average In: 3568.0 kb/s (3.6%) Current In: 4097.0 kb/s (4.1%)
Max Out: 49.3 Mb/s (49.3%) Average Out: 22.9 Mb/s (22.9%) Current Out: 27.0 Mb/s (27.0%)

Some more screenshots of other MRTG graphic types can be seen at <http://www.stat.ee.ethz.ch/mrtg>.

If the firewall is Linux iptables based, MRTG can be configured as follows (from <http://techupdate.zdnet.co.uk/story/0,,t481-s2120161,00.html>):

The file /etc/mrtg.cfg should contain the following:

```
WorkDir: /var/www/mrtg
Target[in_packets]: `/usr/local/mrtg/mrtg.sh`
MaxBytes[in_packets]: 1500
Title[in_packets]: Inbound packets
PageTop[in_packets]: <h1>Inbound Packet Stats</h1>
```

The collection script, /usr/local/mrtg/mrtg.sh, looks like this:

```
#!/bin/bash
IPTABLES="/sbin/iptables"
UPTIME="/usr/bin/uptime"
$IPTABLES -nvxL | grep INPUT | awk '{ print $5 }'
$UPTIME | awk '{ print $3, $4, $5 }'
```

MRTG Related links:
<http://people.ee.ethz.ch/~oetiker/webtools/rrdtool>
<http://cricket.sourceforge.net>

RTG

RTG is a high-performance SNMP-statistics monitoring system written in C (C is far better than Perl in terms of performance; MRTG is written in Perl with performance-crucial sections in C). RTG inserts data into a MySQL database where complex queries and reports may be generated. The screenshot below is a weekly bit-rate graph produced by RTG. It is available from <http://rtg.sourceforge.net>, and runs on any version of Unix.

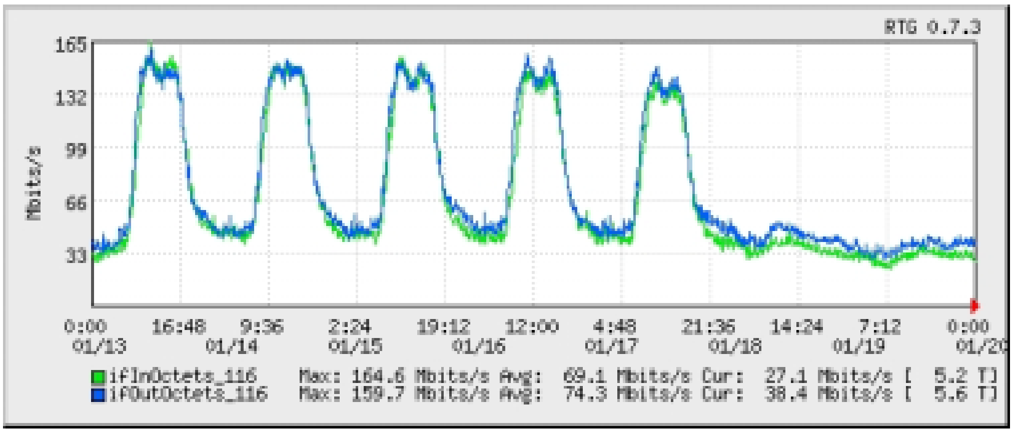

Ntop. If the firewall is Unix-based, Ntop (see below) can give very good graphical statistics of network usage.

Netflow is a feature of Cisco routers that collects network data. See <http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps_wp.htm> for more information about Netflow. Various tools are available to analyse Netflow data.

FlowAnalyzer. A Netflow tool by Cisco that runs as a Java application. See <http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/nfa/nfa_iug/nfa_over.htm>; <http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/nfa/nfa_iug/nfa_mng.htm>.

Cflowd. A network flow analysis tool used for analysing Cisco's NetFlow-enabled switching method. Organizations using Cisco equipment can use it to find out how they use bandwidth. See <http://www.caida.org/tools/measurement/cflowd> for information on Cflowd.

Flowscan. A network analysis and reporting tool that processes IP flows recorded in Cflowd format and reports on what it finds. See <http://net.doit.wisc.edu/~plonka/FlowScan>.

Flavio. A data grapher for Cisco's Netflow; see <http://flavio.sourceforge.net>.

Learnstat: An open-source tool developed at Moratuwa’s Computer Science and Engineering department that analyses traffic flow data collected by NeTraMet and uses RRDTOOL for plotting graphs. See <http://www.ac.lk/LEARNStat>.

Other. Some firewalls and bandwidth-manager products have their own statistical interface, but most will also support SNMP. More information about SNMP is available at <http://www.simpleweb.org>, <http://www.snmplink.org> and at <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/snmp.htm>.

RMON (remote monitoring) is a network-management protocol that allows network information to be gathered at a single workstation. For RMON to work, network devices, such as hubs and switches, must be designed to support it. See <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/rmon.htm> for more information.

(See also Section 5, because some of these quality-of-service or traffic-shaping tools have their own monitoring interfaces.)

Tool comparisons

A good discussion of many of the monitoring tools mentioned above can be found at <http://www.cse.mrt.ac.lk/~pushpa/cssl_paper.pdf>.

4.3 Mail, Web and proxy log file analysis

Looking at usage from the software or server side is done to compare bandwidth usage with what is going through the router/firewall. The router/firewall usage figures will typically be more accurate in terms of bandwidth usage. Server-based logs are then tallied, and they should more or less add up to the amount of bandwidth that went through the router/firewall, or the difference should be explained.

These logs can also be used to determine which Web sites are important, whether some users abuse the system, to what extent e-mail is being used (or if people are using Web-based e-mail), whether the mail server is being abused to send out mass e-mails, whether the proxy server is conserving bandwidth (or accelerating Web access), etc.

Webalizer. An excellent open-source log file analysis program for Unix or Windows. See <http://www.mrunix.net/webalizer>. It adds up all the bandwidth, and also shows the top sites, etc. Much faster than Perl-based analysis programs. Supports Web-server logs, wu-ftpd xferlog, FTP and Squid log formats.

Webalizer usage examples:

webalizer -f Squid
This creates an HTML report with graphics in /var/lib/webalizer. This can be set up to run every day using cron. In Windows the Squid log must be placed in a directory accessible to Windows (the G:drive in this case).

webalizer -F squid -o g:\webalizer g:\access.log
This creates the HTML report in the folder g:\webalizer.

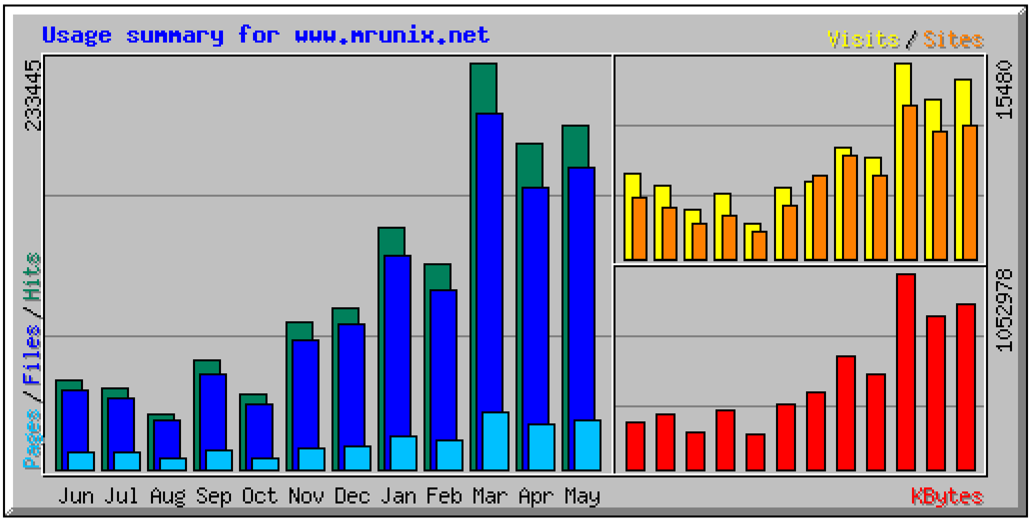

Sawmill. An excellent commercial log file analysis program for Windows or Unix that can analyse a vast range of log files, yet is reasonably priced (currently US$69 for an academic licence). Sawmill has a Web interface with a login screen. Administrators can connect from anywhere and view the analysis if they know the password.

The log files include proxy logs (Squid, ISA, MS Proxy 2), Web (Apache, IIS), mail (Exchange, Postfix, etc.), and firewall logs (Cisco, Linux, etc.). This is also a very responsive company, and they improve the program (adding features, for example) on request. See <http://www.sawmill.net> for more information. The screenshot below is of a Squid log. Sawmill requires a lot of processing power and disk space.

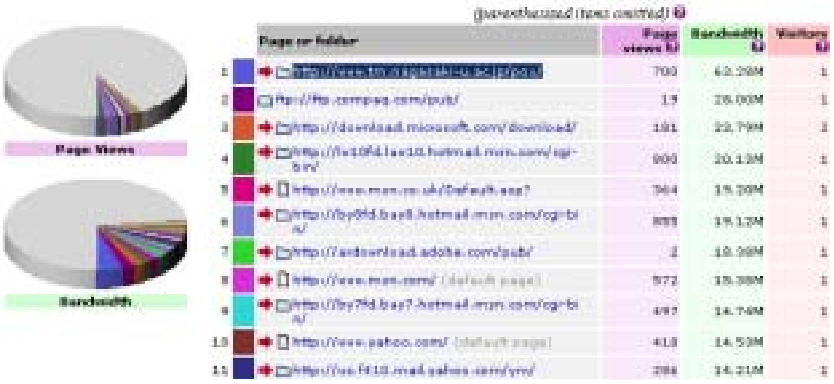

(parenthesized items omitted)

	Page or folder	Page views	Bandwidth	Visitors
1	[illegible]	700	62.29M	1
2	ftp://ftp.compaq.com/pub/	19	28.00M	1
3	http://download.microsoft.com/download/	181	22.79M	2
4	http://lw10fd.law10.hotmail.msn.com/cgi-bin/	900	20.13M	1
5	http://www.msn.co.uk/Default.asp?	364	19.20M	1
6	http://by8fd.bay8.hotmail.msn.com/cgi-bin/	855	19.12M	1
7	http://ardownload.adobe.com/pub/	2	18.38M	1
8	http://www.msn.com/ (default page)	572	15.38M	1
9	http://by7fd.bay7.hotmail.msn.com/cgi-bin/	497	14.74M	1
10	http://www.yahoo.com/ (default page)	418	14.53M	1
11	http://us.f410.mail.yahoo.com/ym/	286	14.21M	1

Awstats. An open-source Perl log file analysis program for Unix or Windows. The advantage of Awstats is that it supports many different log file types: "Apache NCSA combined log files (XLF/ELF) or common (CLF), IIS log files (W3C), WebStar native log files and other Web, proxy, wap or streaming servers log files (but also ftp or mail log files)".

The Windows version requires Perl (for example, Active Perl –
see <http://www.activestate.com>.

Calamaris. A Perl script that analyses Squid proxy logs. It summarizes the log, and also produces a Cache Hit Rate figure. The Cache Hit Rate is a measure of how much Web content was retrieved from the proxy server rather than from the Web. The higher the better, but the hit rate is determined by many factors such as the diversity of Web pages visited (if everyone visits the same pages, all these pages will be cached, and the hit rate will be very high). Normally, the hit rate is between 30 and 60 per cent. To create a summary file from the Squid log, type: cat /var/log/squid/access.log|calamaris > ~/squid.txt

Calamaris is available from <http://cord.de/tools/squid/calamaris>. It can only run on Windows if Perl is installed (see <http://www.activestate.com>).

Calamaris should be used for calculating the Cache Hit Rate, even if another program is used for further analysis.

Analog. A very good open-source log file analysis program for Unix or Windows. Available from <http://www.analog.cx>. A sample analysis report can be viewed at <http://www.statslab.cam.ac.uk/webstats/stats.html>. Analog can be extended using Reportmagic; see <http://www.reportmagic.org>.

4.4 Finding out what is happening on the LAN

If something is causing congestion, it is worth running a network analyser, or even a packet capture tool, to see which PC is causing it. A packet capture tool also enables an administrator to examine the content of the packets, and with luck and patience learn what it is.

Network analysers

IPTraf: A very good tool for finding out in real time what traffic is connecting to which port, etc. Can work in promiscuous mode, and can also log traffic. IPTraf can be used to determine which host on the network is generating traffic at times when there is a general slowness in the network. It comes with all distributions and versions of Unix. Connections can also be sorted by bytes or by packets.

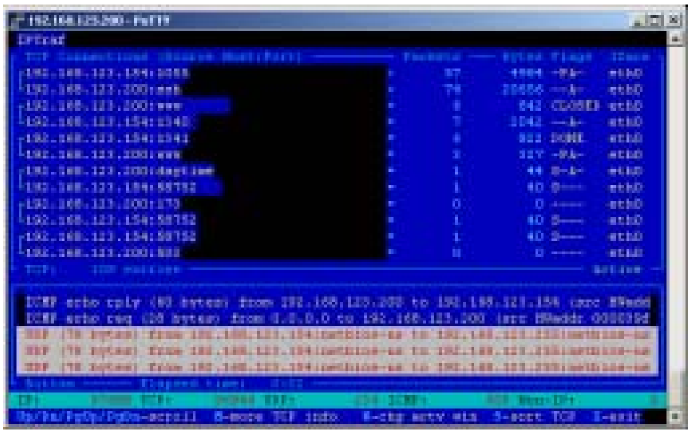

```
IPTraf
 TCP Connections (Source Host:Port) ------------ Packets --- Bytes Flags  Iface
┌192.168.123.154:1055               >      97     4904 -PA-   eth0
└192.168.123.200:ssh                >      78    20656 ---A   eth0
┌192.168.123.200:www                >       8      842 CLOSED eth0
└192.168.123.154:1540               >       7     1042 ---A   eth0
┌192.168.123.154:1541               >       8      922 DONE   eth0
└192.168.123.200:www                >       3      327 -PA-   eth0
┌192.168.123.200:daytime            >       1       44 S-A-   eth0
└192.168.123.184:55752              >       1       40 S---   eth0
┌192.168.123.200:173                >       0        0 ----   eth0
└192.168.123.154:55752              >       1       40 S---   eth0
┌192.168.123.154:55752              >       1       40 S---   eth0
└192.168.123.200:883                >       0        0 ----   eth0
 TCP:    [illegible] entries ------------------------------------------ Active
 ICMP echo rply (60 bytes) from 192.168.123.200 to 192.168.123.154 (src HWadd
 ICMP echo req (28 bytes) from 0.0.0.0 to 192.168.123.200 (src HWaddr 000039d
 UDP (78 bytes) from 192.168.123.154:netbios-ns to 192.168.123.255:netbios-ns
 UDP (78 bytes) from 192.168.123.154:netbios-ns to 192.168.123.255:netbios-ns
 UDP (78 bytes) from 192.168.123.154:netbios-ns to 192.168.123.255:netbios-ns
 Bottom ------ Elapsed time: [illegible]
 IP: [illegible] TCP: [illegible] UDP: [illegible] ICMP: [illegible] Non-IP: [illegible]
 Up/Dn/PgUp/PgDn-scroll  M-more TCP info  W-chg actv win  S-sort TCP  X-exit
```

Network Probe. An excellent commercial network tool for Windows or Unix that enables an administrator to determine easily what is going on. See screenshot below:

It is available from <http://www.objectplanet.com/probe>.

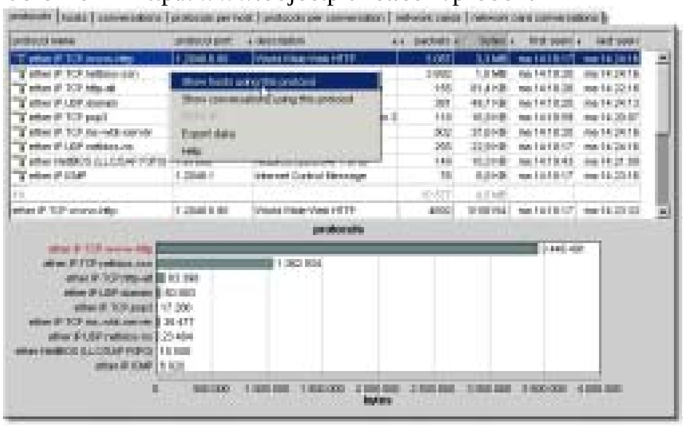

Ntop: A very good statistical network analysis program or network probe that comes with most versions of Linux/Unix. It can also be downloaded from <http://www.ntop.org>, and can be installed on Windows. It has a Web interface, and produces graphical results, as can be seen below:

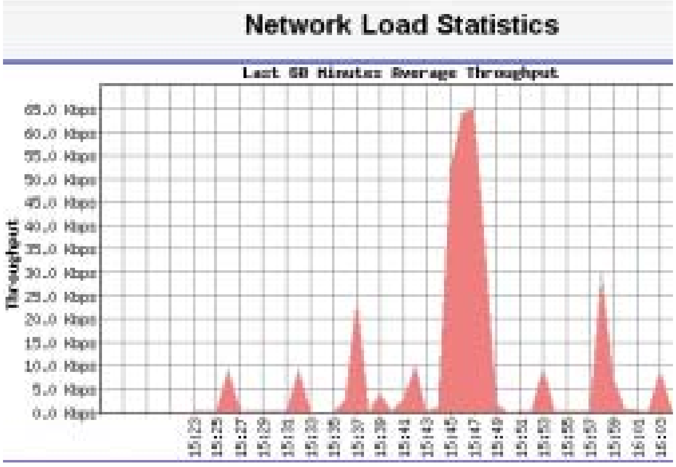

Other free tools that come with the operating system:
Nbtstat (for Windows SMB traffic)
Netstat (Unix and Windows)
tcpdump (Unix)
tcpdump -i eth1 This shows all the packets going through interface eth1.
tcpdump > ~/ntop.log Logs all packets to a file. This file can become very big very quickly.

How to determine whether the LAN is congested

On an overloaded network there are many collisions. On Ethernet, every time there is a collision, the sending machine has to retransmit. Collisions are indicated by a hub/switch with a solidly lit collision LED, which is the first sign that the network may be congested. Another tool that can be used is netstat (available in Unix and Windows). Netstat can be executed on any machine, but it makes more sense to run it on the (busiest) server.

Netstat can be used as follows (in Unix and Windows):

```
netstat -s
```

The result will contain something like this (which are statistics kept by the operating system since boot time):

```
    333914 segments sent out
    73 segments retransmitted
0 bad segments received.
```

(73/333914)*100 is about 0.022%, which is not much. If the percentage of retransmitted packets is more than 5%, the packet loss is high, and something is causing a problem. There can be one of two causes:

- There is congestion (or bad hardware such as a faulty network card on the LAN).
- The packet was dropped at a router owing to congestion on the Internet link.

This information should be interpreted with care, because of the fact that the statistics are from the previous boot time (which may be months ago), and most of the retransmissions may have occurred on a given day owing to circumstances specific to that day. To determine how congested the network is on any given day, one would take a reading in the morning, and another one in the afternoon, and see what percentage of that day's traffic was retransmitted.

Packet capture tools

Analyzer: A network capture program for Windows that enables an administrator to see what is on the network; see <http://analyzer.polito.it>. Since a capture program captures a vast amount of data, it can be hard to interpret.

Ethereal: A Unix-based capture program that comes with most distributions; see <http://ethereal.zing.org>.

5. Bandwidth management

Bandwidth management, prioritization, quality of service (QoS) and traffic shaping are terms that refer to controlling traffic on a network. These terms are related but not exactly equivalent, and this section aims only to be an introduction. Through the use of a bandwidth management product, an administrator can 'shape' traffic in the following ways.

- Make the best use of more than one Internet link, and route traffic through either or both, transparent to the user.
- Prioritize some types of traffic. This is done because some TCP/IP-based applications are much more sensitive to latency (delays) than others. For example, most people would agree that it does not matter if an e-mail is delayed by 1 or two minutes, but that two minutes is a long time to wait for a Web site to load. Voice-over IP, Microsoft's RDP and Citrix are also very sensitive to delays caused by congestion. A conversation over VoIP would break up if the link gets too congested, even if the congestion lasts only a few seconds.
- Provide equal distribution of bandwidth across the various subnets.
- Allocate extra bandwidth for critical applications or user groups.
- Throttle, or limit, bandwidth for certain applications, for example Kazaa, or certain sections of the campus where abuse is known to be taking place.
- Report on how much bandwidth is used by which users, and which protocols they are using.

The QoS or traffic-shaping products mentioned here can enable administrators to provide higher priority to latency-sensitive applications such as voice or Web traffic, and lower priority to mail and FTP traffic, for example. In addition, these products may enable the network team to do accurate bandwidth accounting, and enable them to charge more to departments that use more bandwidth, if charging is considered the right approach for a university. Since some of these products have their own reporting and monitoring tools, they may be able to do many of the functions listed under the monitoring section of this document.

Since certain P2P programs such as Kazaa can change ports if they are blocked, some organizations prefer to use bandwidth-management products to “throttle” them. This means that they are allowed only a small amount of the available bandwidth.

Iproute2 and CBQ

Some of the commercial products below use the routing and traffic control features of the Linux kernel. The programs and techniques used to achieve bandwidth management through Linux are described in the HowTo document at <http://lartc.org/lartc.html>. Apart from the HowTo, there are many resources and a mailing list at this site. With some work, and a cheap PC, a bandwidth manager can be built in-house at low cost. Iproute2 provides Class Based Queuing (CBQ) and Stochastic Fairness Queuing (SFQ).

MikroTik

MikroTik's RouterOS software is based on Linux and Iproute2 that can do bandwidth management. It can be installed on almost any PC. The required specifications are very low:

- At least 32MB of RAM (48MB suggested).
- At least 32MB Flash Drive/HDD.
- Ethernet NE2000 compatible interface.
- 486DX or better processor.

The processors used in expensive routers such as Cisco, are typically less powerful than the Intel 486, and processors are seldom the bottleneck (usually there is not enough bandwidth). This explains why such an old machine can be used. The RouterOS software costs less than US$100. MikroTik also make rack-mountable routers for between US$500 and US$1000. The screenshot and text below is from the MikroTik Web site <http://www.mikrotik.com>.

> “Queuing is a mechanism that controls bandwidth allocation, delay variability, timely delivery, and delivery reliability. MikroTik™ Router supports the following queuing mechanisms: PFIFO, BFIFO, RED, Split and SFQ.”

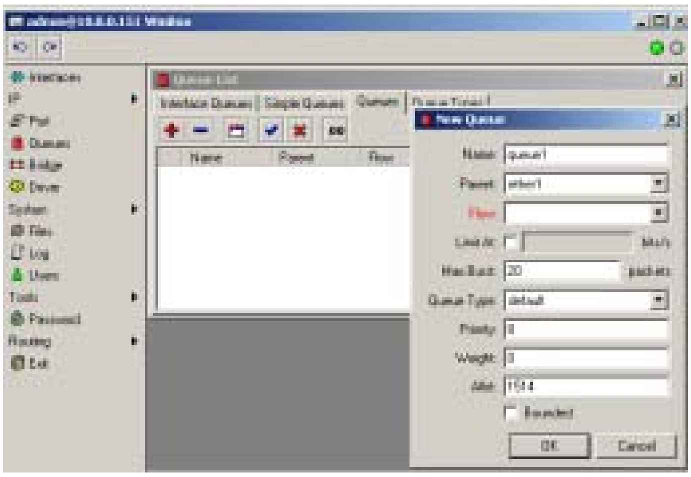

Name: queue1
Parent: ether1
Flow:
Limit At: □ Mbit/s
Max Burst: 20 packets
Queue Type: default
Priority: 8
Weight: 0
Allot: 1514
□ Bounded

Allot

See <http://www.allot.com>. This product is used by the University of Dar es Salaam. Many of the Allot NetEnforcer family of LAN appliances are built on open-source products already discussed (except the graphical interfaces). The following claims from the Allot Web site show that it sees the NetEnforcer as powerful enough to serve at ISP level, which is far more powerful than an organization such as a university typically requires:

> NetEnforcer family of bandwidth management LAN appliances offers best-of-class traffic shaping technology for QoS/SLA enforcement and real-time traffic monitoring on Enterprise and Service Providers networks. Individual models are optimized to support bandwidths from 128 Kbps to 155 Mbps and 1,000 to 10,000 users. NetEnforcer's traffic shaping is complemented and enhanced by additional software modules including NetAccountant (IP accounting), CacheEnforcer® (traffic redirection) and NetBalancer® (server load balancing).
>
> Enterprise customers use Allot products to ensure the performance of business-critical applications such as Citrix, SAP and VoIP, while throttling less critical traffic from P2P and other applications. With Allot, the enterprise network manager can control, monitor and reduce bandwidth consumption on the corporate Internet/WAN link to improve users' productivity and satisfaction as well as contain networking costs.
>
> Service providers (xSPs) use Allot solutions to manage, monitor and invoice their customers according to their service level agreements. With Allot, service providers are able to maximize ROI by managing over-subscription, throttling P2P traffic on upstream and downstream links and delivering tiered services and classes of service.

The diagram below shows Allot's graphical reporting. See the University of Dar es Salaam case study for more screenshots.

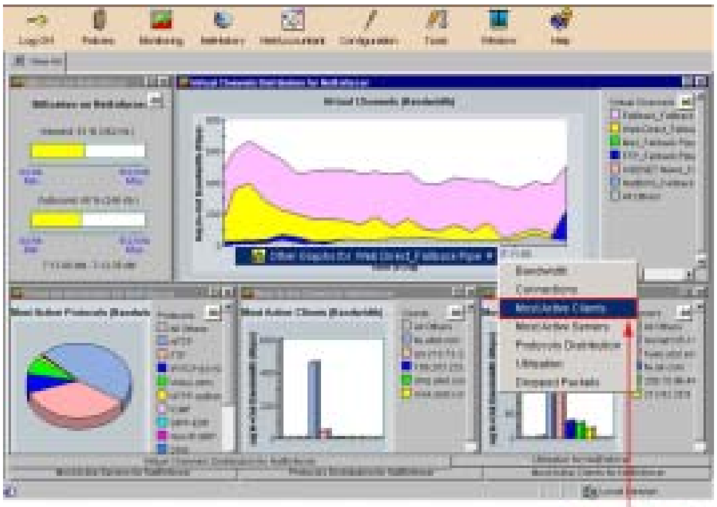

Click here to see your heaviest Web users

Emerging Technologies

The ET/BWMGR – Bandwidth Manager (see <http://www.etinc.com>) is used by Makerere University. Emerging Technologies claim that the ET/BWMGR, which runs on Linux or FreeBSD, performs better at bandwidth management than its competitors. They make a compelling case – see <http://www.etinc.com/index.php?page=bwcompare.htm>. However, this is not an independent test, but marketing literature. However, according to this document, this performance advantage only makes a difference at more than 10 MB/s because they say that "if you ... have bandwidth requirements of 10 Mbps or less, you can 'roll your own' and get all of the power of the ET/BWMGR for the price of a good word processing package". Since most universities in developing countries don't have nearly as much as 10 MB/s, Emerging Technologies seem to recommend using SFQ or CBQ on a regular PC for connections up to 10 Mbps.

The ET/BWMGR appears to be stronger on performance, but less good at reporting and accounting and with a less user-friendly interface than Allot's NetEnforcer.

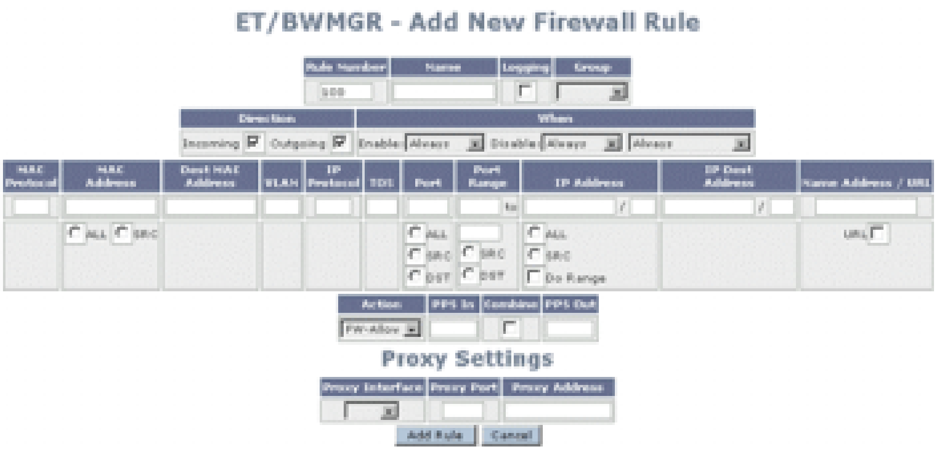

Rule Number: 100

Incoming ☑ Outgoing ☑ Enable: Always Disable: Always Always

Action: FW-Allow

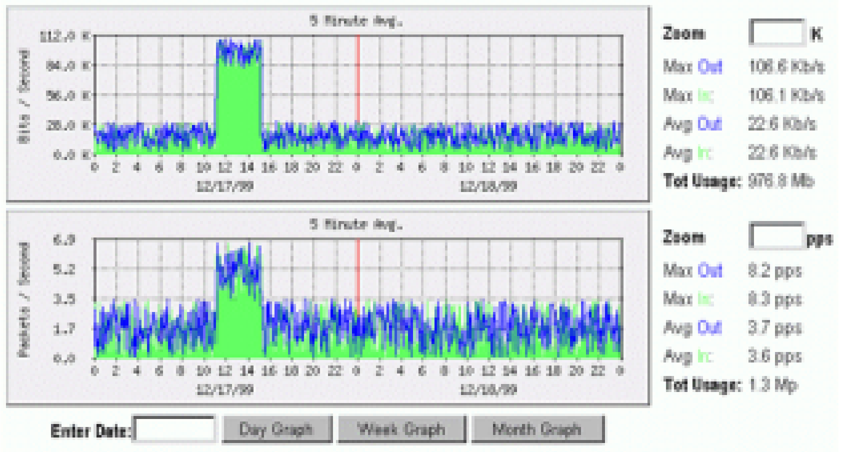

Max Out 106.6 Kb/s
Max In 106.1 Kb/s
Avg Out 22.6 Kb/s
Avg In 22.6 Kb/s
Tot Usage: 976.8 Mb

Max Out 8.2 pps
Max In 8.3 pps
Avg Out 3.7 pps
Avg In 3.6 pps
Tot Usage: 1.3 Mp

Packet Shaper

This is the bandwidth manager with the most memorable name, see <http://www.packeteer.com>. Packet Shaper is the most well known and by far the most expensive. Packeteer has released a number of lower-priced products that are based on Packet Shaper, perhaps because of stiff competition from Emerging Technologies, Allot and open-source products. A seemingly convincing comparison between Packet Shaper and ET/BWMGR shows Packet Shaper in a very bad light in terms of performance.
See <http://lists.insecure.org/lists/security-basics/2003/Jan/0278.html>.
Emerging Technologies also compare the ET/BWMGR with Packet Shaper at <http://www.etinc.com/index.php?page=bwcompare.htm>.
Below is a screenshot of a Packer Shaper graph at the College of Medicine.

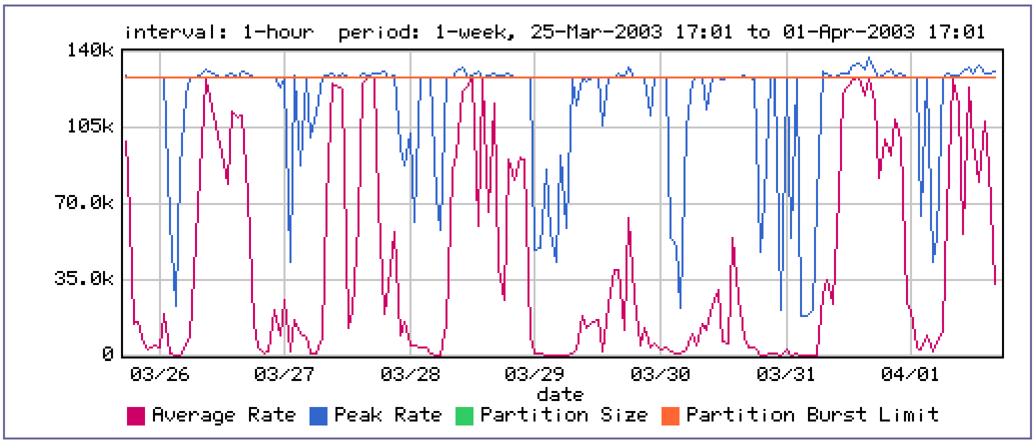

Satelcom MegaPAC

Not a bandwidth manager product as such, the MegaPAC satellite router is mentioned here because it can provide some prioritization. These routers can control how much CIR and BE different networks get. They also have the capacity to prioritize and deprioritize protocols, and to allocate bandwidth by DLCI. Each DLCI can serve a different network, so the equipment can be used to determine how much bandwidth various departments get. Additionally, the VSAT protocol contains features such as TCP acknowledgement spoofing to improve throughput over satellite links (see <http://www.satelcom.com>).

Cisco

Cisco equipment can do a variety of bandwidth-management-related tasks – for example priority, custom and weighted fair queuing. See <http://www.cisco.com> for more details.

Other bandwidth-management products:

- Sun Microsystems: <http://wwws.sun.com/software/bandwidth>
- Lightspeed Systems: <http://www.lightspeedsystems.com>
- NetScreen Technologies: <http://www.netscreen.com>
- Sitara: <http://www.sitaranetworks.com>
- Lucent Access Point QVPN (formerly Xedia): http://www.lucent.com/products

6. Security

Poor security can lead to the following bandwidth-related problems.

- Traffic caused by viruses replicating themselves.
- Traffic generated by people elsewhere on the Internet taking control of a university server and using it as an FTP server for illegal downloads, etc. A poorly secured mail server might be abused for relaying spam to thousands of e-mail addresses.

Apart from the above, hackers might deface a university Web site, gain access to student or financial records, or damage or shut down an essential machine such as a bandwidth manager or a proxy server. Below is are recommendations to ensure the security of a campus network.

6.1 Start with the Top 20

The SANS/FBI Top Twenty List is at <http://www.sans.org/top20>. This list, which consists of the top 20 Windows and Unix exploits, covers most of the threats in these systems and how to protect against them. That is because most threats exploit these vulnerabilities.

6.2 Staying informed

A good system administrator should stay informed about the latest security threats and vulnerabilities. The following Web sites are good starting points:

- **Cert** (<http://www.cert.org>) is a centre for Internet security expertise. The Web site contains a large amount of security-related information. The CERT Advisory Mailing List informs administrators of security threats and how they can be fixed. To subscribe, send an e-mail to <majordomo@cert.org> with "subscribe cert-advisory" (without quotes) in the message body.
- **The Sans Institute** (<http://www.sans.org>) is security-related co-operative research and education organization. It has several newsletters that inform administrators about security threats (see <http://www.sans.org/newsletters>). There is also a lot of security-related information on this site, such as the Top Twenty List, which is essential reading. There is also the SANS NewsBites, a weekly high-level executive summary of the most important news articles that have been published on computer security during the previous week.

More mailing lists, some of them operating-system specific, can be found at <http://directory.google.com/Top/Computers/Security/Mailing_Lists>.

6.3 Passwords

The first thing to be said about password security is that it should be easy for users. If they have to change their password every few days, or remember different passwords for different systems, it is not going to work. If the e-mail password is not the same as the password to log on to the workstation, users will switch to commercial Web-based e-mail as soon as they have a problem. And if an awkward password system is employed, users will share passwords just to get something done. In this way, what is intended to be a highly secure system can become very insecure and counter-productive.

This does not mean that a lax approach to password security is being recommended. Administrative and root passwords should be subject to all the recommended measures – for example, not being susceptible to a dictionary attack, and only known to a few key people. User log-ons should not give the user access to anything that needs to be protected by strong security. Therefore, security does not begin with hassling users with password schemes. It begins with deciding what needs to be secure, putting it on secure servers that users cannot modify, even though they might be able to view it (for example, a Web site), and deciding what should a user not be able to view even though he has logged on (for example, other students' examination results).

6.4 Hardening servers

Most vulnerabilities involve exploiting services or applications, such as a Web server. Hardening involves disabling any unnecessary services to make a server run only those services that are actually used. For example, if a machine is used as a file server, it does not need to be running a Web server. Therefore, the Web server should be stopped, and prevented from starting at boot time.

It also involve setting secure passwords, and disabling the Guest account on Windows. Furthermore, all security-related patches and updates released since the operating system itself was released must be installed/applied.

There are many good hardening guides around. A good way to start is to look for hardening guides on the Sans and Cert Web sites and also at <http://www.securityfocus.com>. The screenshot below shows how a service can be disabled on a Windows 2000 server.

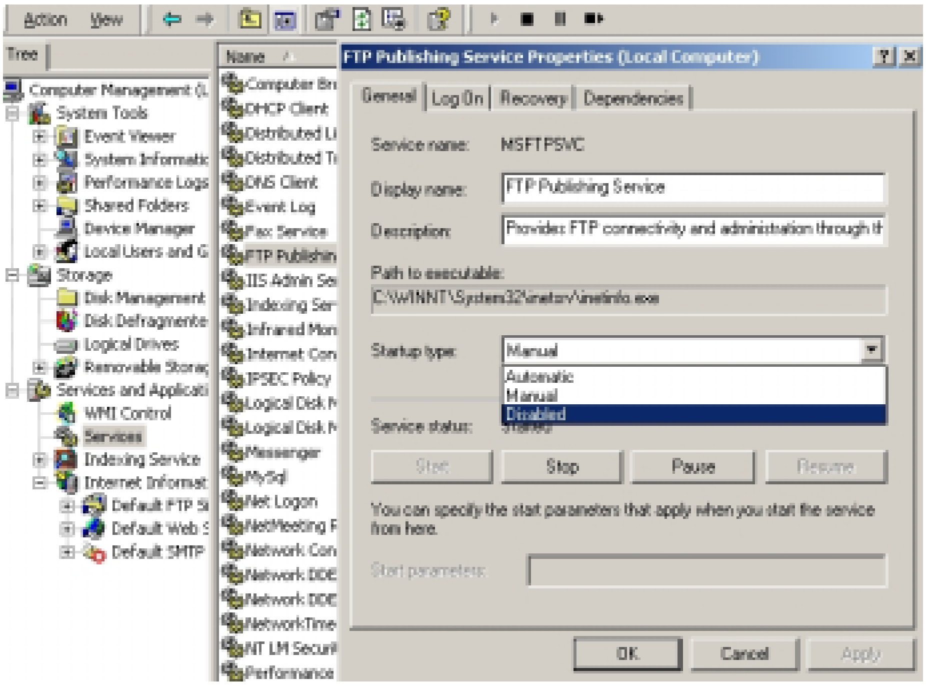

FTP Publishing Service Properties (Local Computer)

Service name: MSFTPSVC

Display name: FTP Publishing Service

Description: Provides FTP connectivity and administration through th

Path to executable:
C:\WINNT\System32\inetsrv\inetinfo.exe

Startup type: Manual

Automatic
Manual
Disabled

Service status: Started

You can specify the start parameters that apply when you start the service from here.

Start parameters:

6.5 Firewalls

These days it is recognized that a network should be "hard on the inside", meaning that everything of importance on the network should be secured rather than relying on a firewall. It is not necessary to get an extremely expensive firewall. But it has become standard practice to install an organizational network behind a firewall, and to assign RFC 1918 addresses to the

internal network. The RFC1918 addresses are
10.0.0.0 - 10.255.255.255 (10/8 prefix)
172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

There are many firewalls that can translate addresses between such a private network and the Internet.

Firewall with open source

A regular PC running Linux, and with two network cards, can act as a firewall/NAT box; see <http://www.netfilter.org> and <http://www.linuxguruz.com/iptables>.

NetPilot

This must be the easiest firewall to configure and run. Although it is aimed at the non-technical user, it has standard firewall functions such as NAT and port-forwarding. It can also act as a DNS and DHCP server; its content-filtering and mail-scanning capabilities have already been mentioned. NetPilot is a transparent proxy. If it is used as a proxy server or mail server, it is important to get the specifications right, or to use more than one. Since Netpilot is used on an Intel-based PC, it has more than enough processing power for the firewall functions, but as mail scanning and proxying is disk and memory intensive, the machine might not perform satisfactorily. See <http://www.netpilot.com> for more information.

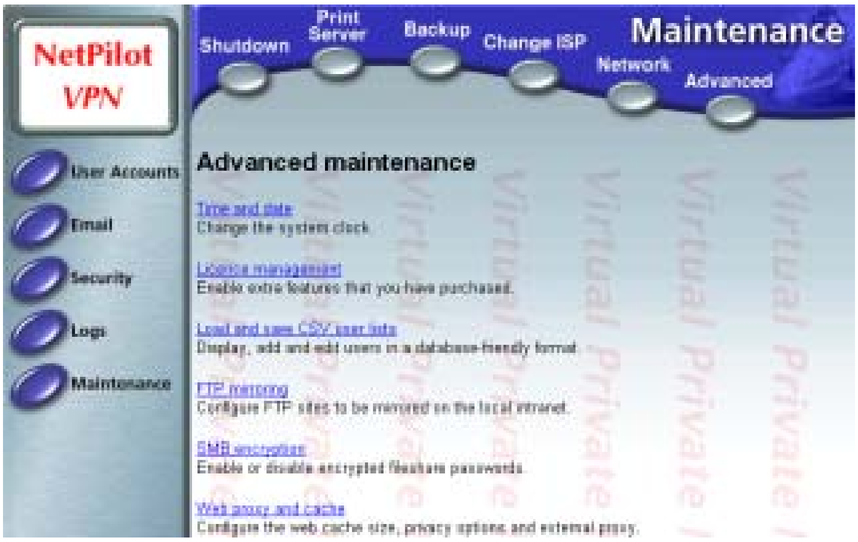

Advanced maintenance

Time and date
Change the system clock

Licence management
Enable extra features that you have purchased.

Load and save CSV user lists
Display, add and edit users in a database-friendly format

FTP mirroring
Configure FTP sites to be mirrored on the local intranet.

SMB encryption
Enable or disable encrypted fileshare passwords

Web proxy and cache
Configure the web cache size, privacy options and external proxy.

Smoothwall

See <http://www.smoothwall.co.uk> and <http://www.smoothwall.org>. A very good, stable Linux firewall that can be installed on any PC-based hardware.

Cisco Pix

The standard against which all other firewalls are measured; see <http://www.cisco.com>.

6.6 Security scans and other tools

Very good vulnerability-scanning tools are available for free. These tools scan for known vulnerabilities. It is a good idea to run these not only from the inside the network but also from outside.

Nessus. A program that runs on Unix, but scans entire networks for both Windows and Unix exploits; see <http://www.nessus.org>. Nessus is famous for staying up to date with the latest exploits; there are daily updates. “The Nessus Project aims to provide to the Internet community a free, powerful, up-to-date and easy to use remote security scanner”. The screenshot below shows the type of HTML report produced by Nessus. When the administrator clicks on the issue regarding a port, it provides details of the threat and how to fix it.

Xscan (Windows). See <http://www.xfocus.org>. A Windows program that scans entire networks for both Windows and Unix exploits.

Nmap. A Unix and Windows port scanner, very useful for scanning a whole network with a single command to find which ports are open. In this way unnecessary services that are running can be identified.

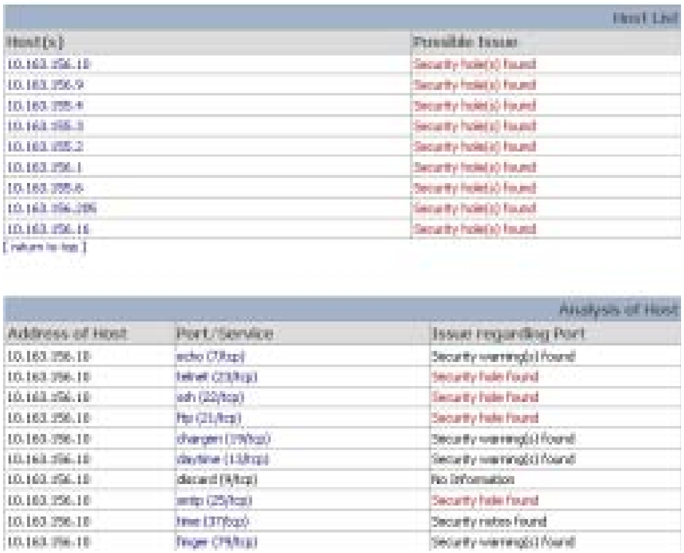

Host List	
Host(s)	Possible Issue
10.163.256.10	Security hole(s) found
10.163.256.9	Security hole(s) found
10.163.255.4	Security hole(s) found
10.163.255.3	Security hole(s) found
10.163.255.2	Security hole(s) found
10.163.256.1	Security hole(s) found
10.163.255.6	Security hole(s) found
10.163.256.205	Security hole(s) found
10.163.256.16	Security hole(s) found

[return to top]

Analysis of Host		
Address of Host	Port/Service	Issue regarding Port
10.163.256.10	echo (7/tcp)	Security warning(s) found
10.163.256.10	telnet (23/tcp)	Security hole found
10.163.256.10	ssh (22/tcp)	Security hole found
10.163.256.10	ftp (21/tcp)	Security hole found
10.163.256.10	chargen (19/tcp)	Security warning(s) found
10.163.256.10	daytime (13/tcp)	Security warning(s) found
10.163.256.10	discard (9/tcp)	No Information
10.163.256.10	smtp (25/tcp)	Security hole found
10.163.256.10	time (37/tcp)	Security notes found
10.163.256.10	finger (79/tcp)	Security warning(s) found

SoftPerfect Network Scanner. This free scanner from <http://www.softperfect.com> easily identifies all the shares in a network. There is a screenshot elsewhere in this document (Section 13.7).

Process to TCP/IP port-mappers

Fport. When a machine has been hacked, the hackers might run their own application on that computer – for example, running a Windows program called ServUDaemon to distribute illegal downloads to other people on the Internet. Fport, available from <http://www.foundstone.com>, can be used to determine which program is listening on which port.

Administrators can schedule the following batch file to run on every important server to e-mail them the result of Fport. In this way, they will know if a new program started running on the server:

```
fport /p >> c:\report.txt
xmail 172.16.1.1 server@mimcom.net gventer@africonnect.com c:\report.txt
```

For this to work, the administrator should install the xmail command line e-mail program, which is available from <http://www.xfocus.org>.

The screenshot below shows how Inetinfo.exe (Internet Information Services) is listening on port 21 (FTP), port 80 (HTTP) and port 443 (HTTPS). If a process other than Inetinfo.exe was listening on one of these ports, it could be an indication that the server has been compromised.

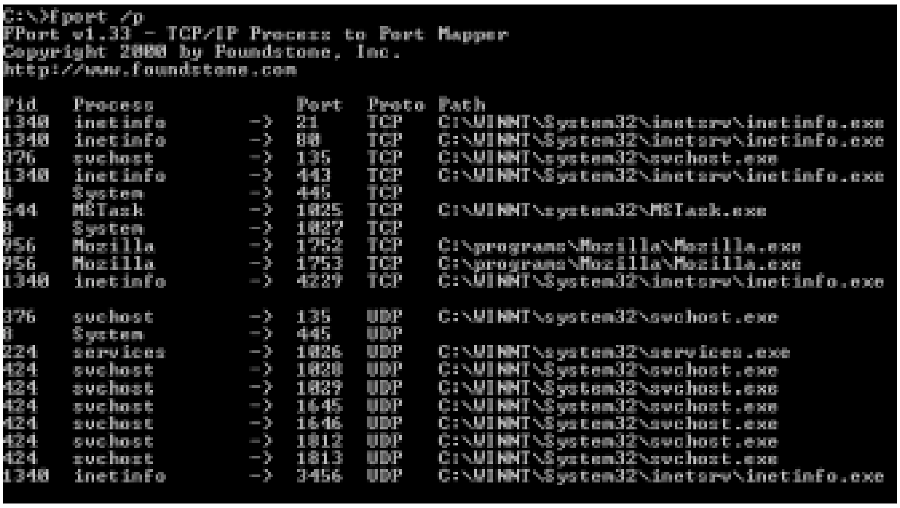

```
C:\>fport /p
FPort v1.33 - TCP/IP Process to Port Mapper
Copyright 2000 by Foundstone, Inc.
http://www.foundstone.com

Pid   Process            Port  Proto Path
1340  inetinfo       ->  21    TCP   C:\WINNT\System32\inetsrv\inetinfo.exe
1340  inetinfo       ->  80    TCP   C:\WINNT\System32\inetsrv\inetinfo.exe
376   svchost        ->  135   TCP   C:\WINNT\system32\svchost.exe
1340  inetinfo       ->  443   TCP   C:\WINNT\System32\inetsrv\inetinfo.exe
8     System         ->  445   TCP
544   MSTask         ->  1025  TCP   C:\WINNT\system32\MSTask.exe
8     System         ->  1027  TCP
956   Mozilla        ->  1752  TCP   C:\programs\Mozilla\Mozilla.exe
956   Mozilla        ->  1753  TCP   C:\programs\Mozilla\Mozilla.exe
1340  inetinfo       ->  4229  TCP   C:\WINNT\System32\inetsrv\inetinfo.exe

376   svchost        ->  135   UDP   C:\WINNT\system32\svchost.exe
8     System         ->  445   UDP
224   services       ->  1026  UDP   C:\WINNT\system32\services.exe
424   svchost        ->  1028  UDP   C:\WINNT\System32\svchost.exe
424   svchost        ->  1029  UDP   C:\WINNT\System32\svchost.exe
424   svchost        ->  1645  UDP   C:\WINNT\System32\svchost.exe
424   svchost        ->  1646  UDP   C:\WINNT\System32\svchost.exe
424   svchost        ->  1812  UDP   C:\WINNT\System32\svchost.exe
424   svchost        ->  1813  UDP   C:\WINNT\System32\svchost.exe
1340  inetinfo       ->  3456  UDP   C:\WINNT\System32\inetsrv\inetinfo.exe
```

lsof

This is the Unix program that may have inspired Fport. It can do the same as Fport and more. For example, to list which applications has which ports open, type: lsof -i:1-1024
To get this list via e-mail: lsof -i:1-1024| mail -s "processes on server" gventer
This can be scheduled to run daily using Cron.

7. Mail and dealing with spam

7.1 Keeping local e-mail local

Users sometimes use e-mail to send a file to the person sitting next to them. If the sender or recipient uses Web-based e-mail, international bandwidth is used to deliver that e-mail, slowing down all other users. It essential to keep all local e-mail local by providing a reliable e-mail system.

7.2 Mail sizes

E-mail is not a protocol designed for distributing large files. Even in countries where there is a lot of bandwidth, it is seldom considered acceptable to send e-mails larger than 1 MB or 2 MB. It is therefore necessary to put a limit on the size of e-mail messages that people can send out, but also to allow users to use other protocols to transmit larger files. Generally speaking, files less up to 10 MB in size should be transmitted via FTP, and larger files transmitted overnight.

7.3 Prevention of spam

The first point of defence is to prevent spammers from obtaining e-mail addresses. This requires user education, as well as prevention measures such as not exposing e-mail addresses to the Web, where this can be avoided. An article entitled “How do spammers harvest e-mail addresses?” describes how spammers get addresses; this article can be found at <http://www.private.org.il/harvest.html>. The first principle is never to reply to spam, or to supply university e-mail addresses to the public Internet, where this can be avoided.

A very good article on avoiding spam is at
<http://www.bris.ac.uk/is/services/computers/nwservices/mail/spam.html>.

If someone’s address must be on a university Web page, his or her e-mail address can be concealed, using the Javascript method described at sites such as these:
<http://dawesbiz.net/spam/#concealing>
<http://www.spywareinfo.com/articles/spam/spambots.html>
<http://www.joemaller.com/js-mailer.shtml>

7.4 Spam scanners

If a spam e-mail has reached a university’s mail server, it has already used bandwidth. Therefore, installing a spam scanner is not going to save any bandwidth. However, to save users’ time, and to prevent untrained users from replying, a spam scanner can be installed.

A good place to scan for spam is at the SMTP server that receives the organization’s mail. The best known spam filter is Spamassassin, which is also used in a variety of other products; see <http://www.spamassassin.org>.

Another program is MailStripper (<http://www.eridani.co.uk/MailStripper), an anti-spam mail scanner with anti-virus capability.

8. Web-based e-mail services

8.1 Why not to use Web-based e-mail services such as Hotmail

(This list was distributed to users on the MIMCOM network)

- They are very wasteful of bandwidth – filling several screens full with colourful graphics instead of just pure text.
- They are used in real time and mail doesn't just come in in the background when bandwidth is available.
- They tend to attract huge quantities of junk mail (which then makes the bandwidth problem worse). After a year or two, unwanted junk mail easily outnumbers real e-mails.
- The mail doesn't get counted as mail within the MIMCom system, and so we can't divert it on to the special bandwidth that we have set up for mail, nor do we scan it for viruses, nor can we do any statistical analysis or take any control of it. And neither can any of the SysOps. It therefore adversely affects those trying to access journals, etc.
- It is not backed up anywhere – unlike your normal mail which should be backed up to tape, etc., from your server.
- The size of the mailboxes is very limited (2MB on Hotmail, and 3MB on Yahoo) and you have to throw away mail quickly. Once your Hotmail box reaches its limit of 2MB then legitimate mails that are sent to you will simply get thrown away.
- Likewise, if you don't use your Hotmail account for a pre-defined length of time it will be automatically closed and mail will be returned to the sender. This means that using it for back-up purposes is difficult as you have to use it regularly to keep it in service.
- It does not form part of the collective memory of your institution.
- It is hard to control acceptable usage under any policy that your site might apply, and is therefore open to abuse (for hate mail, etc.).
- Your institutional e-mail address gives your site a proper name and therefore gives you an institutional identity. You can also work out people's addresses and even publish them. This does not apply to Hotmail, etc. An e-mail from a Hotmail address could be from anyone, and looks far less professional.
- If you have a network policy of using Outlook (or Groupwise) then you can integrate your mail into other local applications. This does not work with Web-based mailers.

8.2 Web-based e-mail server packages

Squirrelmail. Web-based e-mail server software for Unix or Windows.
See <http://www.squirrelmail.org>.

Outlook Web Access. Web-based front-end for Microsoft Exchange. Ships as part of Microsoft Exchange server.

SqWebMail. SqWebMail is a Web CGI client for sending and receiving e-mail using Maildir mailboxes (on Unix systems). Used by Makerere University.
See <http://www.inter7.com/sqwebmail.html>.

SilkyMail. Used by the University of Bristol. See <http://www.cyrusoft.com/silkymail>.

SuSE OpenExchange Server. This is a competitor for Microsoft Exchange and, in addition to Web-based e-mail, it offers Microsoft Outlook compatibility and functions such as scheduling and shared calendar.
See <http://www.suse.de/en/business/products/suse_business/openexchange/index.html>.

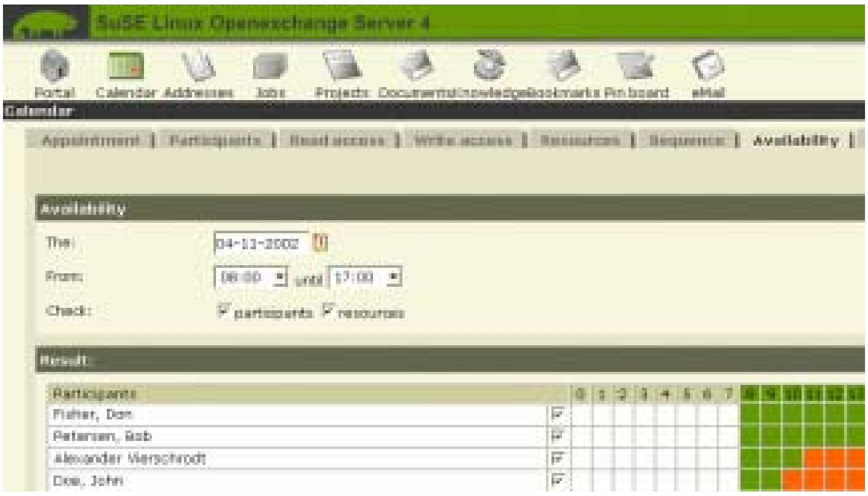
SuSE Linux Openexchange Server 4
Portal
Calendar
Addresses
Jobs
Projects
Documents
Knowledge
Bookmarks
Pin board
eMail
Calendar
Appointment | Participants | Read access | Write access | Resources | Sequence | Availability |

Availability

The: 04-11-2002

From: 08:00 until 17:00

Check: ☑ participants ☑ resources

Result:

Participants		0	1	2	3	4	5	6	7	8	9	10	11	12	1
Fisher, Don	☑									🟩	🟩	🟩	🟩	🟩	🟩
Petersen, Bob	☑									🟩	🟩	🟩	🟩	🟩	🟩
Alexander Vierschrodt	☑									🟩	🟩	🟩	🟧	🟧	🟧
Doe, John	☑									🟩	🟩	🟧	🟧	🟧	🟧

Other

More packages are listed in the Google directory (<http://directory.google.com>) under
Computers > Software > Internet > Clients > Mail > Web-Based
And Computers > Software > Internet > Servers > Mail

8.3 Allowing Web-based e-mail only outside working hours, using Squid

The section below shows how Makerere University uses a Squid acl (access control list) in the squid.conf file to disallow Web-based e-mail during working hours. The file /usr/local/squid/etc/webmail contains a list of Web-based e-mail sites.

```
acl LOUSY_SITES url_regex "/usr/local/squid/etc/webmail"
acl COME_BACK_LATER time M T W H F 09:00-17:00
http_access deny LOUSY_SITES COME_BACK_LATER all
```

Notice that the access lists are placed next to each other to ensure an AND kind of logic. See the Makerere University case study for more details.

9. Anti-virus

9.1 Anti-virus packages for client workstations

When deciding on an anti-virus package, it is important to select a product that can distribute anti-virus updates to clients' computers from a central distribution server. (Unfortunately, central distribution seems to be problematic with Windows 95/98 and ME, and works best in a domain environment.) Central distribution is important because:

- Each workstation PC does not download its updates directly from the Internet; instead, an update is downloaded once, then distributed, which saves bandwidth.
- Administrators do not have to visit every PC to install the update but can do it across the network.
- The update console can indicate which PCs are on which update level; therefore, the administrator can be certain that all PCs are protected.

The products below all have a central distribution function:
Grisoft AV: <http://www.grisoft.com>. An excellent program, and competitively priced.
Sophos AV. <http://www.sophos.co.uk>
Norton AV. <http://www.symantec.com>
McAfee AV. <http://www.mcafee.com>

Reliable AV. <http://www.ravantivirus.com>. Offer discounts for academic institutions.

Other. Many more products are listed in the Google directory (<http://directory.google.com>) under Computers > Security > Anti Virus > Products

9.2 User education

The one thing every user should know is that they should not launch unknown executable files, or open files of unknown origin, especially if these arrive via e-mail. The University of Bristol have comprehensive online user education at the following URLs:
Anti-virus main page: <http://www.bris.ac.uk/is/services/computers/virus>
Anti-virus strategy: <http://www.bristol.ac.uk/is/selfhelp/documentation/av-r1/av-r1.htm>
Anti-virus tutorial: <http://www.bristol.ac.uk/is/services/computers/virus/tutorial>

9.3 Mail server scanners

Some organizations have Microsoft Exchange servers that are hidden from the Internet behind a Unix/Linux SMTP server. Others use only Unix/Linux mail servers. In either case, it is a good idea to use a mail relay server, as shown in the diagram below. If this server is compromised, the users' e-mail messages are still safe on the mail server in the internal network. The anti-virus mail scanner is often installed on the mail relay server.

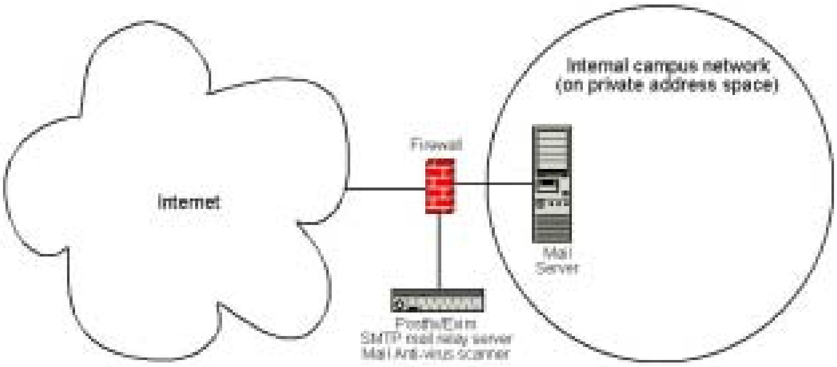

Amavis (<http://www.amavis.org>) is an open-source mail scanner for Unix/Linux. It is free, but uses AV engines such as Sophos which have to be paid for. Amavis comes with most Linux distributions.

Mailscanner (<http://www.sng.ecs.soton.ac.uk/mailscanner>) scans e-mail for viruses and spam. It uses the McAfee anti-virus engine (<http://www.mcafee.com>) and Spamassassin (<http://www.spamassassin.org>).

Reliable AV (<http://www.ravantivirus.com>) has mail scanner software for most mail server products running on most operating systems. Also includes an anti-spam module.

Antigen (<http://www.antigen.com>) makes an anti-virus scanner for Microsoft Exchange that can use a number of anti-virus engines (such as Sophos) to scan.

10. Charging mechanisms

The University of Zululand uses a Squid redirector to enable them to make charges based on usage. There is a small disadvantage to this method, and that is that not all traffic is Web or FTP traffic. However, since the bulk of all traffic is HTTP, HTTPS or FTP, charging only on these protocols represents a fair method of charging. This solution is presented in detail in the University of Zululand case study.

A much less satisfactory solution would be to use Internet Cafe Manager Pro (<http://www.cafemanager.com>). This product measures only the time the user spent on the computer. As bandwidth and not time on the computer is the expensive resource, this is not really a satisfactory product for charging.

Alternative charging mechanisms could also be investigated, such as one of the bandwidth managers described in Section 5, or a bandwidth billing application that ISPs use. The problem with this approach (apart from expensive software) is that it is not as easy to measure and charge for bandwidth at an institution as it is at an ISP. This is because students cannot be charged for local use, whereas all traffic through an ISP is to the Internet and can therefore be charged. Therefore any solution will have to determine which part of a student's usage was traffic to the Internet and exclude local traffic. A list of ISP bandwidth accounting and billing packages can be found at <http://directory.ensim.com/solution.asp?mid=1&sid=4>.

11. TCP/IP factors over a satellite connection

The "long fat pipe network" factors affect TCP/IP performance on networks that have relatively large bandwidth, but a high latency as can be found on satellite networks. Most Internet connections in Africa and other parts of the developing world are via VSAT. Therefore, even if a university gets its connection via an ISP, this section might apply if the ISP's connection is via VSAT. The high latency in satellite networks is due to the long distance to the satellite, which adds about 520 ms to the round-trip time (RTT) (compared to a typical RTT between Europe and the USA of about 140 ms).

The factors to be looked at are long RTT, large bandwidth delay product, and transmission errors.

Generally speaking, operating systems that support modern TCP/IP implementations should be used in a satellite network. These implementations support the RFC 1323 extensions:

- The "window scale" option for supporting large TCP window sizes (larger than 64KB).
- Selective Acknowledgement (SACK) to enable faster recovery from transmission errors.
- Timestamps for calculating appropriate RTT and retransmission timeout values for the link in use.

11.1 Long round-trip time (RTT)

Satellite links have an average RTT of around 520ms. TCP uses the slow-start mechanism at the start of a connection to find the appropriate TCP/IP parameters for that connection. Time spent in the slow-start stage is proportional to the RTT, and for a satellite link it means that TCP stays in slow-start mode for a longer time than would otherwise be the case. This drastically decreases the throughput of short-duration TCP connections. This is can be seen in the way that a small Web site might take surprisingly long to load, but when a large file is transferred acceptable data rates are achieved after a while.

Furthermore, when packets are lost, TCP enters the congestion-control phase, and owing to the higher RTT, remains in this phase for a longer time, thus reducing the throughput of both short- and long-duration TCP connections.

11.2 Large bandwidth-delay product

The amount of data in transit on a link at any point of time is the product of bandwidth and the RTT. Because of the high latency of the satellite link, the bandwidth-delay product is large. TCP/IP allows the remote host to send a certain amount of data 'in advance' without acknowledgement. An acknowledgement is usually required for all incoming data on a TCP/IP connection. However, the remote host is always allowed to send a certain amount of data without acknowledgement, which is important to achieve a good transfer rate on large bandwidth-delay product connections. This amount of data is called the TCP window size. The window size is usually 64KB in modern TCP/IP implementations.

On satellite networks, the value of the bandwidth-delay product is important. To utilize the link fully, the window size of the connection should be equal to the bandwidth-delay product. If the largest window size allowed is 64KB, the maximum theoretical throughput achievable via satellite is (window size) ÷ RTT, or 64KB ÷ 520 ms. This gives a maximum data rate of 123KB/s, which is 984 Kbps, regardless of the fact that the capacity of the link may be much greater.

Each TCP segment header contains a field called ‘advertised window’, which specifies how many additional bytes of data the receiver is prepared to accept. The advertised window is the receiver's current available buffer size. The sender is not allowed to send more bytes than the advertised window. To maximize performance, the sender should set its send buffer size and the receiver should set its receive buffer size to no less than the bandwidth-delay product. This buffer size has a maximum value of 64KB in most modern TCP/IP implementations.

To overcome the problem of TCP/IP stacks from operating systems that don't increase the window size beyond 64KB, a technique known as ‘TCP acknowledgement spoofing’ can be used (see Performance Enhancing Proxy, below).

11.3 Transmission errors

In older TCP/IP implementations, any packet loss is always considered to have been caused by congestion (as opposed to link errors). When this happens, TCP performs congestion avoidance, requiring three duplicate ACKs or slow start in the case of a timeout. Because of the long RTT value, once this congestion-control phase is started, TCP/IP on satellite links will take a longer time to return to the previous throughput level. Therefore errors on a satellite link have a more serious effect on the performance of TCP than over low latency links. To overcome this limitation, mechanisms such as Selective Acknowledgement (SACK) have been developed, which specifies exactly those packets that have been received, allowing the sender to retransmit only those segments that are missing because of link errors. The Microsoft Windows 2000 TCP/IP Implementation Details White Paper states that "Windows 2000 introduces support for an important performance feature known as Selective Acknowledgement (SACK). SACK is especially important for connections using large TCP window sizes."

11.4 Implications for universities

If a site has a 512 Kbps connection to the Internet, the default TCP/IP settings are sufficient, because a 64 KB window size can fill up to 984 Kbps. But if the university has more than 984 Kbps, it might in some cases not get the full bandwidth of the available link due to the “long fat pipe network” factors discussed above. What these factors really imply is that they prevent a single machine from filling the entire bandwidth. This is not a bad thing during the day, because many people are using the bandwidth. But if, for example, there are large scheduled downloads at night, the administrator might want those downloads to make use of the full bandwidth, and the “long fat pipe network” factors might be an obstacle. Administrators might consider taking steps to ensure that the full bandwidth can be achieved by tuning the TCP/IP settings described below.

If a university has implemented a network where all traffic has to go through the proxy (enforced by network layout), then the only machines that make connections to the Internet will be the proxy and mail servers (they make connections on behalf of client computers, as shown below).

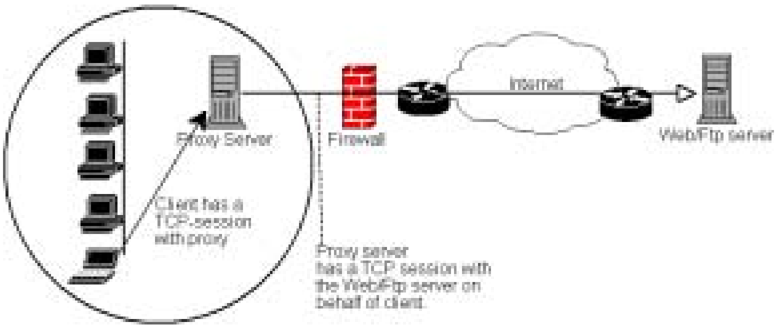
Proxy Server
Firewall
Internet
Web/Ftp server
Client has a TCP-session with proxy
Proxy server has a TCP session with the Web/Ftp server on behalf of client

Therefore, the only computers that should be affected by the “long fat pipe network” factors are the mail and proxy servers. If, however, the proxy server also routes some protocols (other than Web and FTP) to the Internet, workstations on the network using these protocols will be subject to “long fat pipe network” factors. It is possible that upgrading the mail and proxy servers can achieve an improvement in speed by following the steps described at <http://www.psc.edu/networking/perf_tune.html>.

11.5 Performance-enhancing proxy (PEP)

This is described in RFC 3135 (see <http://www.zvon.org/tmRFC/RFC3135/Output>), and would be a proxy server with a large disk cache that has RFC 1323 extensions, among other features. Below is an example of how a PEP might work. The laptop at the campus in the diagram (which is downloading files from the Web server) has a TCP session with the PEP at the ISP. That PEP, and the one at the satellite provider, communicate using a different TCP session or even their own proprietary protocol. The PEP at the satellite provider gets the files from the Web server. In this way, the TCP session is “split”, and thus the link characteristics that affect protocol performance (long fat pipe factors) are overcome (by TCP acknowledgement spoofing, for example). Additionally, the PEP makes use of proxying and pre-fetching to accelerate Web access further.

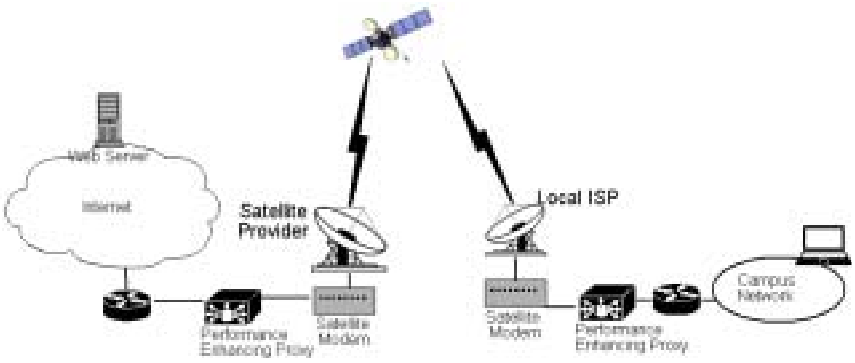

Such a system can be built from scratch using Squid, for example, or purchased ‘off the shelf’. Products to consider are tq®-TELLINET (see <http://www.tellique.de>) and Satelcom MegaPAC (see <http://www.satelcom.com>).

12. Major problem areas

12.1 Hosting a Web site locally

If a university hosts its Web site locally, visitors to the Web site from outside the campus and the rest of the world will compete for bandwidth with the university’s staff; this includes automated access from search engines. A solution is to use split DNS and mirroring. The university mirrors a copy of its Web sites to a server at, say, a European hosting company, and uses split DNS to direct all users from outside the university network to the mirror site, while users on the university network access the same site locally.

12.2 Open proxies

A proxy server should be configured to accept only connections from the university network, not from the rest of the Internet. This is because people elsewhere can connect and use open proxies for a variety of reasons, such as to avoid paying for international bandwidth. The way to configure this is to specify the IP address range of the campus network in the squid.conf file as the only network that can use Squid.

12.3 Open relay hosts

An incorrectly configured mail server will be found by unscrupulous people on the Internet, and be used as a relay host to send bulk e-mail and spam. They do this to prevent getting caught. To test for an open relay host, the following test should be carried out on the mail server (or on the SMTP server that acts as a relay host on the perimeter of the campus network).

Use telnet to open a connection to port 25 of the server in question (with some Windows versions of telnet, it may be necessary to type 'set local_echo' before the text is visible):

```
telnet mail.uzz.ac.zz 25
```

Then, if an interactive command-line conversation can take place (for example, as follows), the server is an open relay host:

```
MAIL FROM: spammer@waste.com
250 OK – mail from <spammer@waste.com>
RCPT TO: innocent@university.ac.zz
250 OK – rcpt to spammer@waste.com
```

Instead, the reply after the first MAIL FROM should be something like:

```
550 Relaying is prohibited.
```

An online tester is available at sites such as <http://www.ordb.org>. There is also information about the problem at this site.

Since bulk e-mailers have automated methods to find such open relay hosts, an institution that does not protect its mail systems is almost guaranteed to be found and abused. Configuring the mail server not to be an open relay consists of specifying the networks and hosts that are allowed to relay mail through them (this should be the IP address range of the campus network).

12.4 Peer-to-peer networking

Bandwidth wastage through peer-to-peer (P2P) file-sharing programs such as Kazaa, Morpheus, WinMX and BearShare can be prevented in the following ways:

- Making it impossible to install new programs. By not giving regular users administrative access to PC workstations, it is possible to prevent the installation of programs such as Kazaa. Many institutions also standardize on a 'desktop build', where they install the required operating system on one PC. They then install all the necessary applications on it, and configure these in an optimal way. The PC is also configured in a way that prevents users from installing new applications. A disk image of this PC is then cloned to all other PCs using software such as Partition Image (see <http://www.partimage.org>) or Drive Image Pro (see <http://www.powerquest.com>).
- From time to time, users may succeed in installing new software or otherwise damaging the software on the computer (causing it to 'hang' often, for example). When this happens, an administrator can simply put the disk image back, causing the operating system and all software on the computer to be exactly as specified.
- Blocking these protocols is not a solution. This is because Kazaa is clever enough to bypass blocked ports. It defaults to port 1214 for the initial connection, but if that is not available it will attempt to use ports 1000 to 4000. If these are blocked, its uses port 80, making it look like Web traffic. For this reason, ISPs don't block it but 'throttle it', using a bandwidth-manager product (see Section 5).
- If rate-limiting is not an option, change the network layout. If the proxy server and mail servers are configured with two network cards (as described in Section 1.2) and these servers are not configured to forward any packets, this would block all P2P traffic. It would also block all other types of traffic (such as Microsoft Netmeeting), and only Web, FTP and mail services will work. However, in low bandwidth networks it may be decided that the advantages outweigh the disadvantages, especially since serious information is available through these services while other protocols are used mostly for informal or recreational activities.

12.5 Programs that install themselves (from the Internet)

There are programs that automatically install, 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. These programs are preventable to some extent by user education and locking down PCs to prevent administrative access for normal users. In other cases, there are software solutions such as Spychecker (<http://www.spychecker.com>) and xp-antispy (<http://www.xp-antispy.de>).

12.6 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 downloads security patches, bug fixes and feature enhancements from the Microsoft Web site. This can consume massive amounts of bandwidth on an expensive Internet link. The two possible approaches to this problem are:

Disable Windows updates on all workstation PCs. The security updates are very important for servers, but whether workstations in a protected private network such as a campus network need them is debatable.

Install a Software Update Server. This is a free program from Microsoft that enables you to download all the updates from Microsoft overnight on to a local server and distribute the updates to client workstations from there. In this way, Windows updates need not use any bandwidth on the Internet link during the day. Unfortunately, all client PCs need to be configured to use the Software Update Server for this to have an effect. This is only a good option for large networks.

Blocking the Windows updates site on the proxy server is not a good solution because the Windows update service (Automatic Updates) keeps retrying more aggressively, and if all workstations do that, it places a heavy load on the proxy server. The extract below is from the proxy log (Squid access log) where this was done by blocking Microsoft's cabinet (.cab) files. Much of the Squid log looks like this:

```
2003.4.2 13:24:17 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:18 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:18 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab HEAD 0
2003.4.2 13:24:19 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:19 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:20 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:21 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:21 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab GET 0
2003.4.2 13:24:21 192.168.1.21 http://windowsupdate.microsoft.com/ident.cab *DENIED* Banned extension .cab HEAD 0
```

12.7 Programs that assume they are on a high bandwidth link

In addition to Windows updates, many other programs and services assume that bandwidth is not a problem, and therefore consume bandwidth for reasons the user might not predict – for example, an anti-virus package that updates itself automatically and directly from the Internet. It is better that these updates are distributed from a local server. The latest versions of Window also have a time service. This keeps the Windows clock accurate by connecting to time servers on the Internet. It is better to install a local time server and distribute accurate time from there.

12.8 Windows traffic on the Internet link

Windows computers communicate with each other via NetBIOS and Server Message Block (SMB). These protocols work on top of TCP/IP or other transport protocols. It is a protocol that works by holding ‘elections’ to determine which computer will be the ‘master browser’. The master browser is a computer that keeps a list of all the computers, shares and printers that you can see in ‘Network Neighbourhood’ or ‘My Network Places’. Information about available shares are also broadcast at regular intervals.

The SMB protocol is designed for LANs and causes problems when the Windows computer is connected to the Internet, because, unless it is blocked, this traffic also tend to spread to the Internet link, wasting the organization's bandwidth. The following steps might be taken to prevent this:

- Block outgoing SMB/NetBIOS traffic on the perimeter router or firewall. In fact, it may save a lot of bandwidth if all access to the Internet is blocked except for the proxy and mail servers. (Only the IP addresses of these servers are allowed to access the Internet directly).
- Install ZoneAlarm on all workstations (not the server). It can be found at <http://download.zonelabs.com/bin/free/1001_cnet_zdnet/zaSetup_37_159.exe>. This program allows the user to determine which applications can make connections to the Internet and which ones cannot. For example, Internet Explorer needs to connect to the Internet, but Windows Explorer does not. ZoneAlarm can block Windows Explorer from doing so.
- Reduce network shares. Ideally, only the file server should have any shares. You can use SoftPerfect Network Scanner (from <http://www.softperfect.com>) to identify all the

shares in your network easily (see screenshot below).

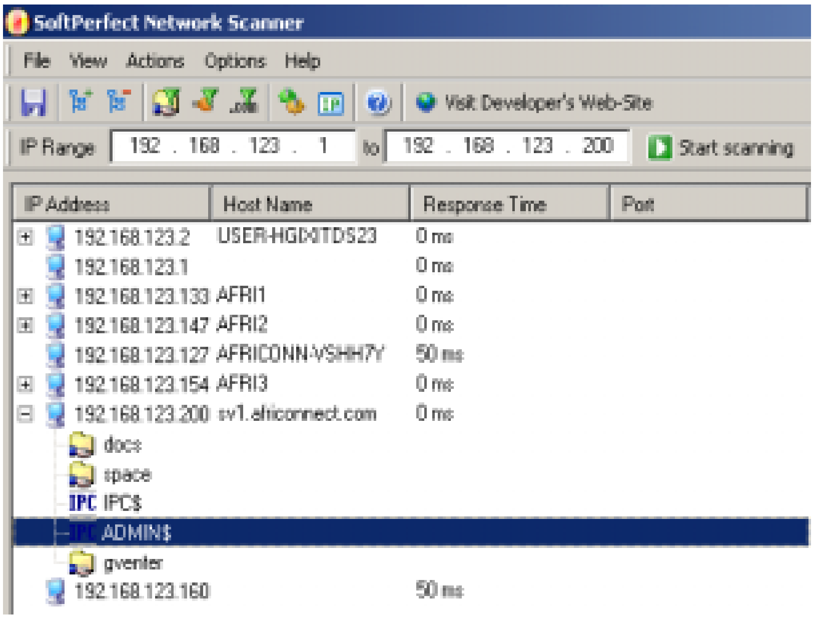

IP Range 192 . 168 . 123 . 1 to 192 . 168 . 123 . 200

IP Address	Host Name	Response Time	Port
192.168.123.2	USER-HGDXITDS23	0 ms	
192.168.123.1		0 ms	
192.168.123.133	AFRI1	0 ms	
192.168.123.147	AFRI2	0 ms	
192.168.123.127	AFRICONN-VSHH7Y	50 ms	
192.168.123.154	AFRI3	0 ms	
192.168.123.200	sv1.africonnect.com	0 ms	
docs			
space			
IPC$			
ADMIN$			
gvenier			
192.168.123.160		50 ms	

12.9 Worms and viruses

Worms and viruses can generate enormous amounts of traffic. The W32/Opaserv worm, for example, is still prevalent, even though it is an old one. It spreads through Windows shares and is detected by other people on the Internet because it attempts to spread further. It is therefore essential that anti-virus protection is installed on all PCs. Furthermore, user education about executing attachments and responding to unsolicited e-mail is essential. The steps given in Section 13.7 about reducing the number of shares is recommended for the prevention of worms and viruses. In fact, it should be a policy that no workstation or server should run unused services. A PC should not have shares unless it is a file server; and a server should not run unnecessary services either. For example, Windows and Unix servers typically run a Web server service by default. This should be disabled if that server has a different function; the fewer services a computer runs, the less there is to exploit.

12.10 E-mail forwarding loops

Occasionally, a single user making a mistake can cause a problem. For example, a user whose university account is configured to forward all mail to her Yahoo account. The user goes on holiday. All e-mails sent to her in her absence are still forwarded to her Yahoo account, which can grow to only 2 MB. When the Yahoo account becomes full, it starts bouncing the e-mails back to the university account, which immediately forwards it back to the Yahoo account. An e-mail loop is formed that might send hundreds of thousands of e-mails back and forth, generating massive traffic and crashing mail servers.

There are features of mail server programs that can recognize loops. These should be turned on by default, but administrators must also take care that they do not turn this feature off by mistake, or install an SMTP forwarder that modifies mail headers in such a way that the mail server does not recognize the mail loop.

12.11 Large downloads

There is no way an effective network can function if the administrator cannot monitor the bandwidth usage of individual users. Therefore it is essential that monitoring software (described in Section 4) is in place. It is also possible to charge according to bandwidth usage, as described in the University of Zululand case study, because sometimes large bandwidth usage can be legitimate (e.g. ISO images of software CDs).

A user may start several simultaneous downloads, or download large files such as 650MB ISO images. In this way, a single user can use up most of the bandwidth. The solutions to this kind of problem lie in training, offline downloading, and monitoring (including real-time monitoring). Offline downloading can be implemented in at least two ways:

- At the University of Moratuwa, a system was implemented using URL redirection:

  > Users accessing FTP: URLs are served by their program, which presents them with a directory listing in which each file has two links, one for normal downloading, and the other for offline downloading. If the offline link is selected, the specified file is queued for later download and the user notified by e-mail when the download is complete. The system keeps a cache of recently downloaded files, and retrieves such files immediately when requested again. The download queue is sorted by file size. Therefore, small files are downloaded first. As some bandwidth is allocated to this system even during peak hours, users requesting small files may receive them within minutes, sometimes even faster than an online download.

- Another approach would be to create a Web interface where users enter the URL of the file they want to download. This is then downloaded overnight using a cron job or scheduled task. This system would only work for users who are not impatient, and are familiar with what file sizes would be problematic for download during the working day.

12.12 Sending large files

Remote uploads

When users need to transfer large files to collaborators elsewhere on the Internet, they should be shown how to schedule the upload them. In Windows, an upload to a remote FTP server can be done using an FTP script file, which is a text file containing FTP commands, similar to the following (which is saved as c:\ ftpscript.txt):

```
open ftp.ed.ac.uk
gventer
mysecretword
delete data.zip
binary
put data.zip
quit
```

This can executed as:

```
ftp -s:c:\ftpscript.txt
```

On Windows NT, 2000 and XP computers, the command

```
ftp -s:c:\ftpscript.txt
```

can be saved into a file such as transfer.cmd, and scheduled to run at night using the Scheduled Tasks (Start > Settings > Control Panel > Scheduled Tasks). In Unix the same can be achieved, using “at” or “cron”.

Users sending each other files

Users often need to send each other large files. It is a waste of bandwidth to send these via the Internet if the recipient is local. A file share should be created on the local Windows or Samba or Novell server, where a user can put the large file for others to access.

Alternatively, a Web front-end can be written for a local Web server to accept a large file and place it in a download area. After uploading it to the Web server, the user receives a URL for the file. He can then give that URL to his local or international collaborators, and when they access that URL they can download it. This is what the University of Bristol has done with their FLUFF system:

> The University offers a facility for the upload of large files (FLUFF) available from <http://www.bristol.ac.uk/fluff>. These files can then be accessed by anyone who has been given their location.

The advantage of this approach is that users can give external users access to their files, whereas the file share method above can work only for users on the campus. A system like this can easily be implemented as a CGI script using Python and Apache.

13. Split DNS and a mirrored server

The aim of split DNS or ‘split horizon’ is to present a different view of your domain to the inside and outside worlds. There is more than one way to do split DNS, but the recommended way for security reasons is to have two separate ‘internal’ and ‘external’ content DNS servers, each with different databases.

Split DNS can enable clients from a campus network to resolve IP addresses for the campus domain to local RFC1918 IP addresses, while the rest of the Internet resolves the same names to different IP addresses. This is achieved by having two zones on two different DNS servers for the same domain.

One of the zones is used by internal network clients and the other by users on the Internet. For example, in the network below the user on the Makerere campus gets <http://www.makerere.ac.ug> resolved to 172.16.16.21, whereas the user elsewhere on the Internet gets it resolved to 195.171.16.13.

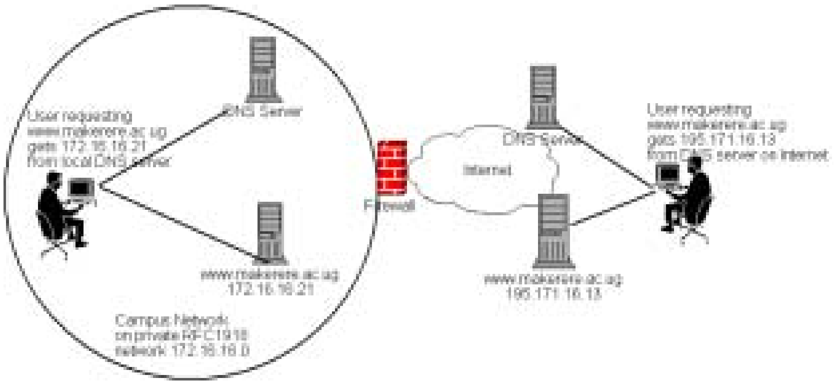

The DNS server on the campus in the above diagram has a zone file for <makerere.ac.ug> and is configured as if it is authoritative for that domain. In addition, it serves as the DNS caching server for the Makerere campus, and all computers on the campus are configured to use it as their DNS server. The DNS records for the campus DNS server would look like this:

```
makerere.ac.ug
www          CNAME          webserver.makerere.ac.ug
ftp          CNAME          ftpserver.makerere.ac.ug
mail         CNAME          exchange.makerere.ac.ug
mailserver   A       172.16.16.21
webserver    A       172.16.16.21
ftpserver    A       172.16.16.21
```

But there is another DNS server on the Internet that is actually authoritative for the <makerere.ac.ug> domain.

The DNS records for this external zone would look like this:

```
makerere.ac.ug
www          A       195.171.16.13
ftp          A       195.171.16.13
mail         A       16.132.33.21
    MX       mail.makerere.ac.ug
```

Split DNS is not dependent on using RFC 1918 addresses. An African ISP might, for example, host Web sites on behalf of a university but also mirror those same Web sites in Europe. Whenever clients of that ISP access the Web site, it gets the IP address at the African ISP, because that way the traffic stays in the same country, while if visitors from other countries access that Web site, they get the IP address of the mirrored Web server in Europe. In this way, international visitors do not congest the ISP's satellite connection when visiting the university's Web site.

This is becoming an attractive solution, because Web hosting close to the Internet backbone has become very cheap because of the numerous hosting companies.

Mirroring

An attractive mirroring program is rsync. It is very bandwidth-efficient because it copies only the differences in files that have actually changed, compressed and through secure shell for security (new files are also transferred, of course).

For example, if the administrator changes a file on the Web server at the campus and scheduled mirroring takes place overnight, rsync does not transfer the new version of the file; only the parts of the file that are needed to make the two copies identical are transferred.

rsync comes with most versions of Unix. Its Web site is at
<http://samba.anu.edu.au/rsync>.
There is also an rsync 'how to' at <http://everythinglinux.org/rsync>.

14. Bandwidth testing

A good test of bandwidth is to download a file of reasonable size (between 1 MB and 2 MB, for example), and time how long it takes. If a 1 MB file can be downloaded in 60 seconds, the data rate is 1024 KB/60 seconds, or 17.1 KB/sec. Since there are 8 bits in a byte, it follows that the links from which 17.1 KB/sec was achieved is equivalent to 136.8 Kbps – a little better than ISDN. Windows 'Scheduled Tasks' or Unix 'cron' or 'at' can be used to schedule a download test in the middle of the night, using FTP or wget.

Here is an example of an upload test to see if the full uplink bandwidth can be achieved (the uplink is the link to the Internet, which is typically smaller than the downlink). A 3561 KB file was uploaded from a MIMCOM site to a server on the Internet to test if the 32 Kbps uplink could be fully utilized. This was done during the night to limit interference from other bandwidth usage.

Run number	Time taken (seconds)	Data Rate (KB/s)	Bit rate (Kbps)	Utilization of available capacity (%)
1	1082.5	3.37	26.96	84.3
2	1087.4	3.35	26.8	83.8
3	1054.9	3.46	27.68	86.5

For comparison, a download on a 64 Kbps leased line typically achieves a data rate of 7.36 KB/s, which is 58.88 Kbps or a link utilization of 92%. The difference between the leased line and the satellite network may be due to other traffic on the link and the higher error rate on a satellite connection. Therefore, the 84% utilization obtained is not a bad result.

Other tests that can be used include the Web-based test at
<http://www.bandwidthspeedtest.com> and MicroTik's free Windows-based bandwidth-testing utility at <http://www.microtik.com/download.html>.

Testing consumes a lot of bandwidth and should be done only occasionally, otherwise it will cause problems for other users.

15. Authentication

Many of the optimization strategies described in this document rely on authentication. Large organizations need a system whereby all users log on to use the computer systems, and to access the Internet and e-mail. In order to prevent widespread confusion and support problems, this log-on should be unified. The main options are:

- Microsoft Active Directory
- Windows NT Domain
- Novell NDS
- OpenLDAP
- Using Sun NIS with PCNetlink.

If a university wants to use Novell or Microsoft for a user database but use a Unix program such as Squid, Apache, or any Unix e-mail system, PAM (pluggable authentication modules) can be used to authenticate against the Microsoft or Novell databases, rather than a separate Unix user database having to be maintained. A good description of PAM can be found at <http://www.linuxgeek.net/index.pl/authentication>. In this case, users can access their e-mail messages, use the Internet or access a password-protected Web page on Apache, using their Windows or Novell password.

If it is decided that a Unix user database should be used for all authentication (for example, because there is already a Unix user list for e-mail), there are two possibilities:

Use Samba as the PDC or BDC for the Windows computers (Windows 2000 and XP can also log on to an NT4/Samba domain), and then synchronize the Samba and Unix passwords. See <http://freebooks.by.ru/view/SambaIn24h/ch16-02.htm> for more details of how this can be done automatically. There is also a Web-based tool to allow users to change their passwords at <http://changepassword.sourceforge.net>.

Gina is the Windows authentication DLL. There is a Gina replacement called pGina (that replaces the regular Windows Gina) in order to allow Windows users to authenticate against a Unix directory such as OpenLDAP. See <http://pgina.xpasystems.com> and <http://www.openldap.org> for more details. If this is implemented, the Windows computers can authenticate against the Unix user directory.

Sun’s Solaris PCNetLink (which is now free) is an all-in-one solution and enables Windows users to authenticate to a Sun NIS directory. To Windows NT, 2000 and XP users, this is like using a Windows NT4 domain. The server also supplies file and print services.
See <http://www.sun.com/solutions/interoperability/netlink/pcnetlink1_2/index.html>.

16. Network layout

16.1 Allowing only certain protocols

If it is decided that only certain protocols are to be allowed (as discussed in the main document), the network design should be as shown in Section 1.2, forcing all traffic through the mail and proxy servers, either by using two network cards in each proxy and mail server (one connected to the campus side, and one to the Internet side), or by using policy routing as described in the Makerere case study.

16.2 Diagrams

In order to understand a network properly, and to design it for security and performance, it is necessary to draw a network diagram.

The network diagram should enable anyone to understand how network packets flow from the user's PC to the server on the Internet. These diagrams can be detailed for the technical team, but simpler versions can be made to explain things to management.

There are several packages for creating network diagrams:

- Dia: An open-source drawing program that is available for Unix and Windows, available from <http://www.lysator.liu.se/~alla/dia>. The diagrams in this document were all drawn using Dia.
- MS Visio. See <http://www.microsoft.com>

16.3 Network design

There are guides to network design freely available on the Internet – for example, the Cisco guide at <http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2002.htm>. The disadvantage of a vendor guide like this is that creates the impression that everything should be done with Cisco equipment in order to maximize their sales. The possibility of using a powerful PC server with two network cards as a proxy/firewall/router is not mentioned, even though this is a common technique. The network that a Cisco engineer designs will be different from the network that a server administrator designs. In designing a network, it is useful to be aware of all the options.