4. Network optimization review

4.1 Web caching

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

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

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

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

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

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

4.2 DNS caching

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

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

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

4.3 Content filtering

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

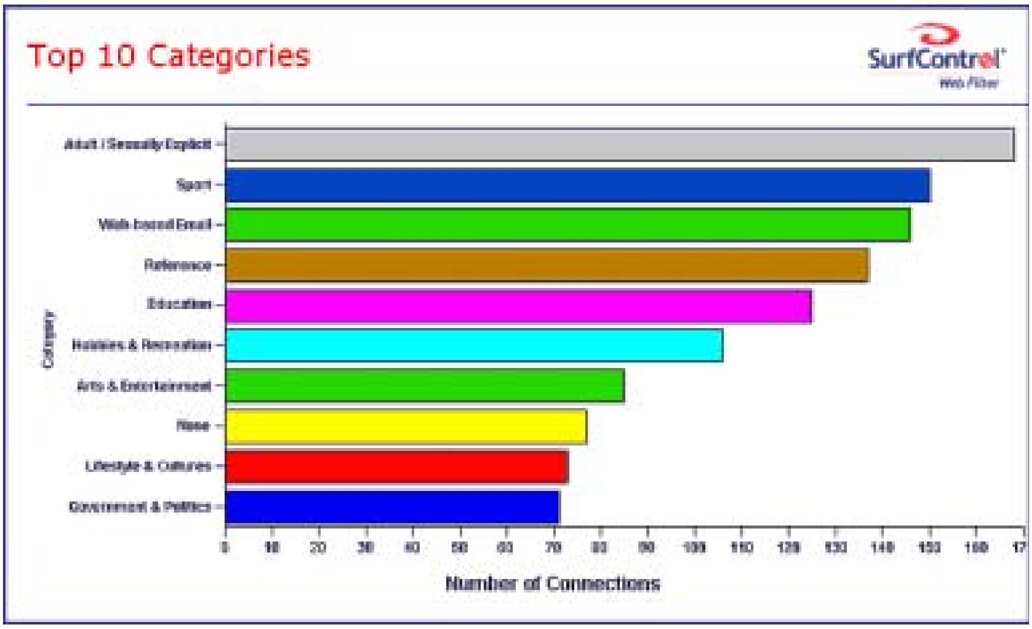

Content filtering can be used to prevent users from accessing certain categories of Web sites or from downloading certain types of files. The programs that can be used for content filtering are typically implemented along with the proxy server, and the measures for ensuring that the proxy server is not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

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

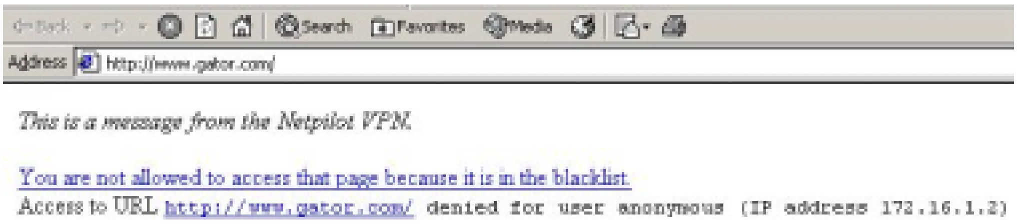

This is a message from the Netpilot VPN.

You are not allowed to access that page because it is in the blacklist.

Access to URL http://www.gator.com/ denied for user anonymous (IP address 172.16.1.2)

The main aims of content filtering are:

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

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

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

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

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

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

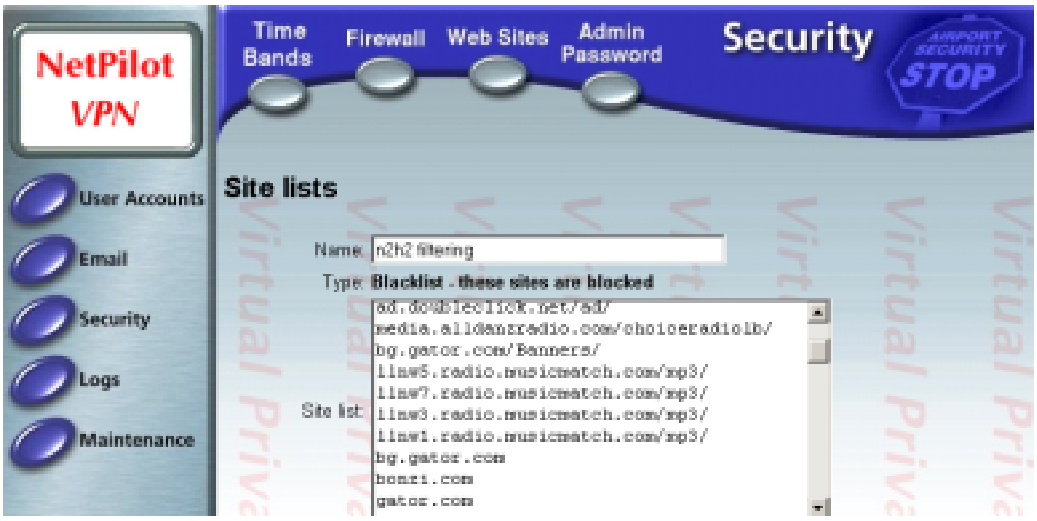

Site lists

Name: n2h2 filtering

Type: Blacklist - these sites are blocked

Site list:

```
ad.doubleclick.net/ad/
media.alldanzradio.com/choiceradiolb/
bg.gator.com/Banners/
liaw5.radio.musicmatch.com/mp3/
liaw7.radio.musicmatch.com/mp3/
liaw3.radio.musicmatch.com/mp3/
liaw1.radio.musicmatch.com/mp3/
bg.gator.com
bonzi.com
gator.com
```

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

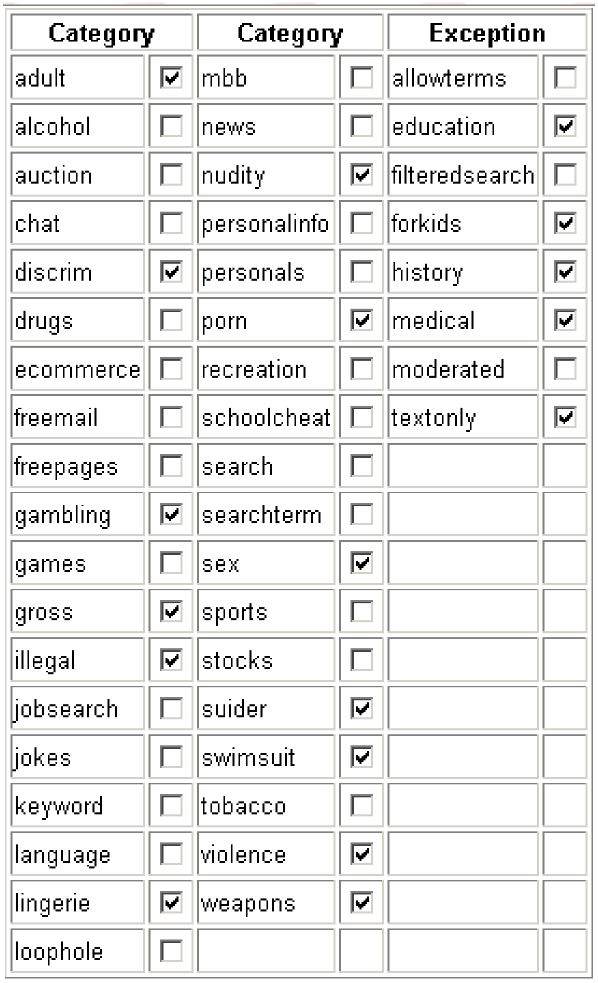

Category		Category		Exception	
adult	☑	mbb	□	allowterms	□
alcohol	□	news	□	education	☑
auction	□	nudity	☑	filteredsearch	□
chat	□	personalinfo	□	forkids	☑
discrim	☑	personals	□	history	☑
drugs	□	porn	☑	medical	☑
ecommerce	□	recreation	□	moderated	□
freemail	□	schoolcheat	□	textonly	☑
freepages	□	search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

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

4.4 Monitoring

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

4.5 Bandwidth management

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

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

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

4.6 Security

Apart from being important in its own right, security is also important because there are people on the Internet who 'hack' into a vulnerable machine and then use it for sharing things like illegal copies of software, music and movies. These people hide the fact that they have compromised a machine because they want to keep using it for their own purposes. Illegally shared files are often called ‘warez’, and participants use IRC to tell each other where to download the files from. Kazaa, BearShare, LimeWire, IRC’s download facility and FTP are used to download these files. Once a host has been compromised, a university or other organization may find that its bandwidth is completely consumed by these activities, leaving nothing for its intended use.

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

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

The following security measures should be considered:

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

See Appendix A for more detail on these recommendations.

4.7 Dealing with spam

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

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

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

Another form of spam avoidance now appearing is to list e-mail addresses on Web sites in a form different from their correct appearance. For example, avoiding the symbol “@” and giving <jsmith at abc.ac.uk> instead of <jsmith@abc.ac.uk>. This is because some companies that sell lists of e-mail addresses to spammers have software that searches the Web and collects e-mail addresses from Web sites, mailing-list archives, etc.

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

4.8 Web-based e-mail services

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

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

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

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

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

4.9 Anti-virus software

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

The main way in which viruses are spread is though the Microsoft Outlook and Microsoft Outlook Express e-mail programs. This is because these applications have capabilities that are tightly integrated with the Windows operating system, which unfortunately makes them ‘insecure by design’ and very commonly targeted. This has led some institutions to actively discourage or prohibit the use of Microsoft Outlook, and to require users to use other e-mail client software such as Eudora or Mozilla Mail. Bristol University is one institution that takes this approach, and they provide no support for Outlook.

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

4.10 Major problem areas

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

4.11 Training

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

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

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