Effectiveness of Internet Filtering Software Products

Prepared for

NetAlert and the Australian Broadcasting Authority

Paul Greenfield, Peter Rickwood, Huu Cuong Tran

Mathematical and Information Sciences

September 2001

CONTENTS

1 Introduction

This report presents the findings of a study commissioned by the Australian Broadcasting Authority (ABA) and NetAlert into the effectiveness of a number of Internet content filtering products. The study examined both how easy the products were to install and use, and how effectively they filtered Internet content.

The products under evaluation all attempt to effectively filter the Internet, blocking access to 'undesirable' content, such as pornography or racist propaganda, and letting all other content pass through untouched. In reality, this is an impossible goal as the Internet is just too big and dynamic, and all products will pass through some content they should have blocked and block some content that should have passed through.

This report answers the following questions for most of the approved filtering products currently on the IIA list:

- Is it easy to install, configure, use and update?
- Is it easy to disable or bypass?
- How well does it stop access to undesirable content?
- Does it stop access to desirable content as well?
- Can it effectively track access?

The products under test include both home-based and server-based filters, and include products using a variety of filtering techniques, including white lists, black lists and content analysis. Not all products on the IIA list were evaluated, as some of the products have been withdrawn from sale or their vendors did not reply to requests for evaluation copies.

1.1 How this report is organised

The report builds on the earlier reports commissioned by NetAlert, adding the results of the studies into usability and effectiveness. There are two main sections.

An overview of blocking and filtering

This section gives a general overview of Internet filtering and blocking technologies, including their strengths and weaknesses.

An appraisal of filtering products

This section assesses each product that was submitted for evaluation, looking at:

1. A summary of the capabilities of each product and a subjective assessment of how easy the product is to install, configure, and use.
2. An assessment of how effectively the product filters Web content, looking at both how well it blocks access to 'undesirable' sites, and how many 'desirable' sites are blocked as a side effect.

2 An Overview of Blocking and Filtering

2.1 The nature of the Internet

The Internet links together computers from all over the world into a single, seamless network, with little regard to geographical and political[1] boundaries. Computers on the other side of the world are just as easy to access as ones in the same city, perhaps just a little slower in responding because of the distances involved.

The Internet is single network, but it is one constructed by linking together a large number of smaller, and autonomous, networks. The Internet is not owned, controlled or managed by any single organisation or government.

Individuals generally access the Internet by subscribing to a service provided by an ISP. This service usually includes at least two things:

- Access to the Internet from their computer. Users first connect to their ISP, via a dial-up modem or DSL/cable link, and their computer is then 'on' the Internet and they can access everything the Internet has to offer.
- Optional Web site hosting. Subscribers can create their own Web pages and have them hosted on their ISP's servers. These pages are then available to anyone, anywhere in the world, and at any time.

Businesses access to the Internet in a similar manner, often using permanent ISP connections that link their internal networks and servers into the Internet. Larger businesses will often run their own e-mail and Web servers, only using the ISP as a pathway into the Internet.

Publishing on the Internet is remarkably easy and inexpensive, in contrast to more traditional media such as newspapers, books and television. Printing is moderately expensive and published material has to be physically distributed, making it easier to control and regulate. Radio and television are broadcast media with low distribution costs, but need expensive equipment and scarce (and tightly regulated) spectrum space. In contrast, publishing on the Web can cost extremely little, often this is a free service included as part of the end-user packages offered by ISPs. Governments can easily regulate and control traditional media, but controlling the Internet is much harder because of its scale and global nature. In most parts of the world, anyone can easily publish on the Web, without having to have their material vetted or approved. Publishing is just a matter of copying files to a Web server, and this simple act makes them available to the world.

2.2 Internet Content

The Internet is just a network of computers, and can carry anything that can be converted into digital form and passed from computer to computer, including text, pictures, sound and video. This content can be published on a Web site, sent via an e-mail, posted to a newsgroup, discussed in a chat room or transferred as files.

[1] Apart from the very few countries that rigidly filter Internet traffic to protect the security and stability of their political systems and society.

All that the ISP's and carriers who run the Internet 'backbone' see are packets of data, with source and destination addresses. Video or text, Web pages or e-mail, all look the same to the routers that are switching the packets on their way to their final destination. It is only when the data packets arrive at their final destination that they are reassembled and can then be interpreted as a video stream or photograph.

The Internet is more than just the Web, and filtering products have to be able to address many possible sources of content if they are to be really effective. These sources include:

World Wide Web	The Web is the most commonly used method of accessing Internet information. The client software is generally referred to as a *Web browser*, or simply a *browser*. Information published on the Web can be found through the use of search engines such as Google and AltaVista. These engines run 'Web crawlers' that travel the Internet by following links, and build up searchable indexes.
Newsgroups	There are thousands of newsgroups worldwide, covering almost every conceivable topic. Users need to become subscribers to newsgroups to be able to access the news, and any subscriber can post material to a news group. Newsgroups started off as public e-mail discussion forums but are also used as data repositories. The 'alt.erotica' family of newsgroups contains very large amounts of erotic material, including text, photographs and video. The 'newsnet' Internet mechanism broadcasts megabytes of 'news' everyday on a very wide range of topics. ISPs can take selected parts of this 'feed' and make it available to their customers. There are also public newsnet repositories available on the Web.
ftp sites	The file transfer protocol, *ftp*, is used to transfer a file from one computer to another. A file can be any sort of content, such as a spreadsheet, a report, or a picture.
Chat rooms	Chat rooms are discussion groups that people can enter and leave at any point in time. Some chat rooms have restricted membership, but others are public. It is commonplace for chat room visitors to adopt a pseudonym.

Table 1 - Some common Internet content sources

2.3 What is Blocking and Filtering

The terms 'blocking' and 'filtering' are often used as synonyms for the technologies that prevent access to particular types or specific pieces of the content that is available on the Internet. This report will use 'blocking' to refer to the techniques used within routers that stop Internet traffic based on its addresses, and 'filtering' to those techniques that stop access to content based on its content. The term 'filtering' will also be used as a general term for both blocking and filtering.

While it is technically feasible to block access to all 'undesirable' Internet content, no Internet blocking or filtering scheme will ever be 100% effective, or resist a determined and informed attacker, but many of them will be perfectly adequate in normal use.

Truly effective Internet blocking and filtering technologies have to address all the possible ways that Internet content can be distributed. Filtered Internet services may have to block access to file transfer and chat rooms, as well as denying access to known 'undesirable' Web sites. A completely 'safe' Internet may well be a very restricted Internet, especially when new types of content and new distribution technologies emerge.

Although most attention is normally paid to the ability of filtering technologies to block access to pornography and other 'undesirable' content, many of the commercial products also provide for the blocking of a wide range of other content, such as sport and other hobbies. These products are intended for use in work environments where the management does not want their employees accessing the Internet for non-work related activities. Such products allow administrators to select the categories of content that should be blocked.

2.4 Approaches to Content Filtering

There are three basic approaches to filtering Internet content:

- Allowing through known 'good' content (inclusion filtering)
- Blocking known 'bad' content (exclusion filtering)
- Examining content and blocking when it fails acceptability tests

Many filtering products are based on lists of Web sites that are supplied by their vendor. These lists are expensive to produce, as they have to be compiled by having people examine and classify Internet content, and as a result these lists are often closely held proprietary information. The secret nature of these lists can make it difficult to know just what content is being blocked and for what purpose.

These lists also reflect the values of the organisations and people who compile them, and may not reflect the values of Australian society as a whole. Some Internet activists[2] complain that commercial filtering products reflect US-based conservative and religious values, and as such may not reflect the more liberal values held by Australian society. Cultures differ considerably in their concepts of 'acceptable' content and filtering products really have to customise their lists to meet local cultural norms. Most products allow their lists to be customised by parents and administrators, by blocking or allowing specific sites, or specific categories.

Inclusion Filtering (white lists)

Products based on inclusion filtering only allow access to a relatively small number of sites that are known to be 'acceptable', and block access to the rest of the Internet – a 'guilty until proven innocent' approach to filtering. The product vendor or some other body trawls the Internet and comes up with a list of sites that meet their criteria of 'acceptability'. These sites are put on a 'white list' of sites that can be accessed.

[2] www.peacefire.org

Attempts to access any other Internet content are blocked. This type of filtering can be 100% effective as only known 'good' sites can be accessed.

The main problem faced by inclusion filtering products is the size of the Internet. The very scale of the Internet means that almost all Internet content will be blocked by default, regardless of whether it would be regarded as 'acceptable' if it was examined. The publicly indexable Web contained an estimated 800 million pages as of February 1999, encompassing about 15 terabytes of information or about 6 terabytes of text after removing HTML tags, comments, and extra whitespace[3]. No inclusion-based filtering product can hope to cover anything more than a very small percentage of the entire Internet. One well known publicly available inclusion list is Yahooligans[4], which contains in the order of 3000 Web addresses, which represents a tiny (and shrinking) part of the entire World Wide Web.

Exclusion Filtering

Exclusion filtering is based on *black lists* (or *blocking lists*) of known objectionable sites and is a more common form of filtering than inclusion filtering. Exclusion filtering adopts the 'innocent until proven guilty' philosophy and allows through all content it does not know to be unacceptable. It has the advantage of giving access to the whole Internet, rather than restricting users to a 'walled garden'.

The advantages of exclusion filtering come at the cost of allowing through much 'unacceptable' content. Sites and content that have not yet been classified by the product vendor pass through the filter, and the sheer scale of the Internet means that much 'undesirable' content could pass through unscathed. These products rely on their black lists being comprehensive and up-to-date. Some black lists, such as N2H2's *Bess* list, are large – over 130,000 pages covering 40,000 servers[5]. Vendors of exclusion filtering products cannot guarantee that their users will not come across 'undesirable' content, but they can block much such material, particularly if it is coming from well known and established sites.

Content filtering

The last approach is to allow access to the entire Internet but to examine the content retrieved before allowing it through to the user. These filtering products will look for certain 'key words' in Web pages or for other characteristics that are supposed to indicate dubious content, such as graphics with large amounts of 'flesh tones'. Content that fails to meet the acceptability tests will be blocked, regardless of whether it is a Rubens painting or a Penthouse centrefold.

Content filtering is appealing because it dynamically classifies incoming content as it arrives. Vendors do not have to manually examine large numbers of Web sites, and users do not have to constantly update lists of acceptable or unacceptable sites. The problem it faces is that accurately determining whether content should be allowed through is a very difficult computing task.

---

[3] "Accessibility and Distribution of Information on the Web", Lawrence and Giles, *Nature*, 400, 107-109. Also see http://wwwmetrics.com/ .

[4] http://www.yahooligans.com .

[5] http://www.n2h2.com/customer/support/old_docs/choicenet_bess.html . This paper is a well written reference on content filtering.

Combined filtering

Commercial filtering products often combine these techniques to improve their effectiveness. For example, a product may primarily use a black list but also use key word filtering to limit access to undesirable sites that have not yet been classified. This approach allows access to the broader Internet while still preventing access to some undesirable content. It also inherits the disadvantages of both techniques, including the need for large exclusion lists and the inadvertent blocking of useful content.

2.5 Content-based Filtering

Products that use content-based filtering techniques examine incoming content and outgoing requests to determine if they appear to be 'unacceptable'. These products employ a variety of methods such as looking for key words, analysing images and looking for 'known' characteristics of 'undesirable' Web pages.

Key word Filtering

Products using key word filtering scan Internet content as it is being loaded into a user's computer and look for words that are included in a black list. A page is blocked if it contains any of the words in the block list. Filtering products also often check requests before they are sent out to prevent users from using search engines to find sites that may contain 'undesirable' content but not included in the product's black lists.

Key word filtering can be very efficient and so is suitable for older, less powerful, personal computers.

Commercial pornographic sites need to be easily found by their potential customers and the search engines they use to surf the Web. One common mechanism by such sites is to add a number of likely search keywords to their Web pages, even if they are not actually displayed. These words will be picked up by the Web crawlers, added to the indexes used by their search engines and can then be found by Web surfers. These same keywords can also easily be picked up by the content filters and used to deny access to Internet content.

There are several problems with key word filtering technologies:

- They only check *text*, and cannot block objectionable pictures that are not accompanied by (in)appropriate text. This could be a particular problem for pornographic content, as Russian or Japanese sexually explicit photographs look much the same as Australian or US pornography but may not come with any helpful English key words.
- They have to be able to distinguish the 'acceptability' of a word from its context. Early key word scanning products had a reputation for being simplistic, blocking words regardless of how they were being used, and unnecessarily blocking access to desirable content as a result. The classic example is the term *breast cancer*, which would be picked up by a key word filter looking for the word *breast*, resulting in blocking the entire site[6]. Another

[6] The text in this very report, if put on the Internet, would likewise be blocked by such a simplistic technique, simply because it contains the words *breast*, *sex*, and *pornography*.

problem arises when black-listed words are contained inside other words – for instance a key word filter with *sex* in its black list may block documents containing the word *Middlesex.* Attempting to overcome these problems, for example by only blocking a page if it contains a certain number of trigger words, can partially overcome these problems, but it is not obvious how to decide on this number[7].

The discussion of keyword filtering so far has focussed on the problems inherent in classifying pornographic content, but, as far as categories go, this may be the *least* problematic category. Many filtering products allow the user to select different categories which should be blocked, but trying to classify a site as sex-education or gambling or drugs based on keywords alone may be more problematic than classifying pornography, because the words occurring in them are more likely to be closer to normal unobjectionable English.

Phrase Filtering

Phrase filtering is a more sophisticated extension of keyword filtering. Phrase filtering does not consider words in isolation, but as part of a phrase. This allows for more fine-grained classification, as it would allow one to consider the phrases *huge breasts*, and *breast cancer* in their respective contexts. While this approach might be expected to do better than keyword filtering alone, it still has many of the associated problems (such as deciding how many objectionable phrases are required before a page is blocked, and being useless for non-English sites), and, in addition, has the added difficulty of having to enumerate all the different phrases that are considered objectionable.

Profile filtering

Several companies have introduced products that filter Internet content based on the *characteristics* of the received content. Vendors tend to be circumspect on how these products work, but some of the features they look for include the ratio of pictures-to-text and links to other known 'undesirable' sites.

Profile analysis can be computationally intensive and result in an unacceptable slow down in perceived Internet access times. Product vendors have countered this problem by performing the analysis in the background, after the page has been displayed to the user, and adding it to the black list if it fails the acceptability tests. Subsequent attempts to retrieve similar content from the same site will be blocked. Content-based filtering is often used in conjunction with other methods, such as URL filtering, to evaluate content that is not already on a black list.

Image analysis filtering

Some filtering products examine images as they are delivered to a user. This approach tries to determine if incoming content contains images of naked bodies, often looking for large amounts of 'skin tones' and on the analysis of images themselves. It is computationally intensive and a difficult task, and computers will invariably experience difficulty in distinguishing between art and pornography, between a

[7] Would it make sense, for example, for this number to vary depending on the amount of text on the Web site? If so, one needs to decide on a mathematical relationship between the amount of text on the page, and the threshold number. If not, one runs the risk of classifying Web sites with lots of text as pornography based on the presence of only a small number of words.

Rubens painting and a Penthouse centrefold, or even between pornography and pictures of the family at the beach.

The time taken to examine image content may also prove to be a problem, as filtering software has to be basically transparent to its users and not degrade the Internet browsing experience. As a result, some of these products perform their image analysis after displaying the content to the user. Pages that are found to contain 'undesirable' images are then added to the black list and will not be available in the future.

2.6 Source-based filtering

Products using source-based filtering techniques block requests going to or responses coming from known 'unacceptable' sites. The blocking is based on the addresses contained in the request or response, using either URLs in the messages or IP addresses in the data packets.

Packet filtering

All content and other information travels across the Internet as IP data packets. Each packet has the IP address of where it is going to (its destination), as well as the IP address of where it has come from (its source). Packet filtering mechanisms examine the source IP address of every packet and block them if they are coming from banned sites.

Packet filtering was examined in some detail in a 1998 CSIRO report on Content Filtering[8]. Packet filtering can be implemented in the routers[9] that provide access to the Internet, such as those used by ISPs to link their customers to the Internet backbone. These routers already have tables that say how to switch packets based on their addresses and packet filtering simply sends packets arriving from blocked sites to oblivion rather than on to the requesting user. Top-of-the-range routers can implement packet filtering without any performance degradation as they are designed to look at packets and switch them at very high speeds.

The main problems with packet filtering are its granularity and the possible impact on ISPs. Packet filtering is particularly coarse as it works at the level of IP addresses. Each IP address represents a computer, not a Web site and filtering an Internet site by using its IP number may block a large number of legitimate sites hosted on the same computer as well. There have been examples of large public Web sites, such as Geocities and OzEmail being blocked because they also hosted a small number of Web pages that contained content that was considered unacceptable.

ISPs will also run into practical difficulties if packet filtering is to be implemented on a large scale. The routers used by ISPs to link into the Internet only have limited space available to hold routing lists and holding large numbers of blocked IP addresses will rapidly exhaust their capacity. Routers are a critical piece of the Internet infrastructure and ISPs need them to be very stable if they are going to provide a reliable service to their customers. Having to constantly update routing lists in order to block new sites will risk introducing instability, and degrade performance and reliability.

---

[8] Available from http://www.cmis.csiro.au/Reports/filtering.pdf .

[9] Routers are special purpose computers that look at packets and steer them from their source to their requested destination.

Packet filtering is not widely used as the basis of filtering products because of these problems, nor is it recommended.

URL Filtering

The most common, and effective, form of source-based filtering is based on URLs, the human-readable address of a Web page. URLs provide more fine-grained control than packet filtering as they name individual Web pages rather whole computer systems. URL filtering can be used with both inclusion (white lists) and exclusion (black lists) techniques.

Practical URL filtering relies on lists of acceptable or unacceptable sites and parts of sites, not on lists of individual pages. A single Web site, such as www.playboy.com, will contain hundreds or perhaps thousands of individual pages, each with their own URL. The hierarchical (left-to-right) structure of Web page names allows filter vendors to block access to entire sites or just parts of sites by specifying partial URLs.

All the pages hosted on the Playboy site will normally have addresses that start with www.playboy.com and this is all that needs to be specified in a filtering list to block access to the entire site, including both www.playboy.com/June2001/centrefold.html and http://www.playboy.com/May2001/centrefold.html.

These partial names can be specified at any level, allowing access to some parts of a site while still blocking access to other parts. For example, a vendor could block access to www.bigpond.com.au/users/~ahacker without blocking access to other users with personal home pages hosted on the Big Pond Web server.

Partial URL matching relies on the use of recognisable 'domain names', such as *www.playboy.com,* within URLs. These domain names are very dynamic and a site can have many domain names, and get new ones allocated quickly. A site, such as Playboy, could use *www.playboy.com* to refer to their main Web server but use additional servers with their own domain names, such as ww2.playboy.com and ww3.playboy.com, to supply most of their actual content. Some large Web sites, such as HotMail, DejaNews and EBay have a thousand or more servers behind a single domain name. This use of multiple servers is done to increase capacity, not to evade filtering, but all of these potential names or IP addresses may still have to appear in filter lists for blocking to be effective.

Domain names, such as *www.playboy.com,* only exist for the convenience of people and their use is entirely optional. Each computer system connected to the Internet is really known by its 'IP address', a 32-bit number commonly written down in 'dot notation' such as *206.251.29.10*. IP addresses are perfectly acceptable within URLs in place of domain names. For example, the URLs *http://www.playboy.com/centrefolds/* and *http://206.251.29.10/centrefolds/* are completely equivalent and interchangeable. The result of this is that filters have to be able to block both forms of a URL, so adding to the size of their black lists or increasing the filtering delays if domain names have to be translated to IP addresses before they can be checked.

Computer systems on the Internet do not have to have domain names at all, and this is commonly the case when a high capacity Web site is built up from a number of smaller but equivalent servers. Each one of these Web servers can normally be accessed individually, causing problems for filtering products. This issue is discussed in section 3.1.

2.7 Auditing

Auditing and tracking access is an important feature of many Internet filtering products. Rather than just blocking access to 'unacceptable' content or sites, they also securely record the attempted access for later review by parents or managers.

2.8 Where can Content Filtering occur?

Filtering can take place on a user's personal computer, on a corporate server, at an ISP or corporate server, or on a third-party system.

Filtering on the user's personal computer

Most of the commonly used filtering products are designed to run on the user's personal computer. These filtering products can use any of the approaches discussed previously, and using a combination of black lists and some content filtering is quite common. This class of products is illustrated in Figure 1.

Some products are set up so that they periodically request the latest list of blocked sites from the software vendor. Others require the user to update the list periodically from the supplier of the filtering software.

Filtering on the user's personal computer has the advantage of relatively plentiful resources, such as processor time. This makes it feasible to implement more computationally expensive techniques such as content-based filtering.

This class of filtering products may also be the least reliable, as the filtering software is running on an insecure system in an insecure environment. Once the filtering software is disabled or bypassed, the whole Internet, in all its diversity, is available.

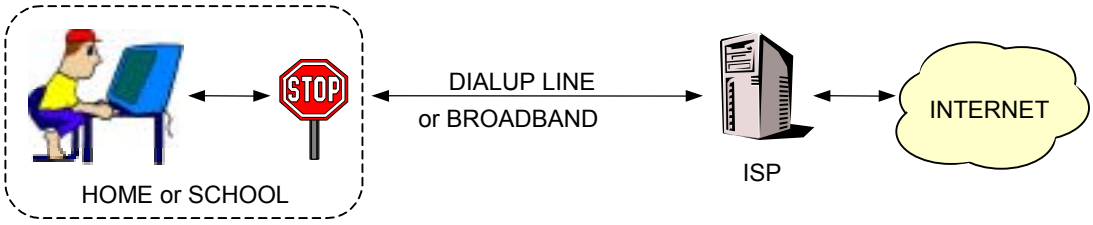

Figure 1 - Content filtering on the user's computer

Filtering at the ISP or corporate server

Filtering products can also be installed and run by an ISP or IT department, as illustrated in Figure 2. Any filtering technique can be used, but the shared nature of the ISPs computers means that only highly efficient mechanisms, such as checking requested URLs against black lists, are really viable.

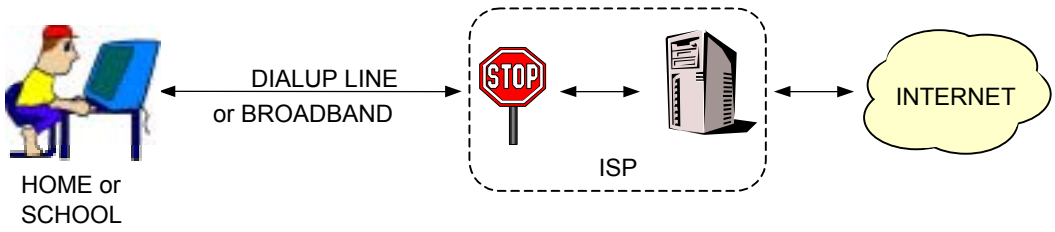

Figure 2 - Content filtering by the ISP

System administrators will normally get lists of blocked sites from the vendor of the filtering software, and will ensure that updates are installed regularly.

Server-based filtering can be very secure because the user only ever sees a filtered view of the Internet and the filter is running on a secure system.

Filtering by a third party

In this situation, user requests are passed through the ISP directly to a nominated third party, where the request is checked against a filter list. For this to be effective, the end user's browser must be configured to point to the third party's Website for any requests and it should not be possible to access the Internet without going through the third party's site. This is illustrated in Figure 3.

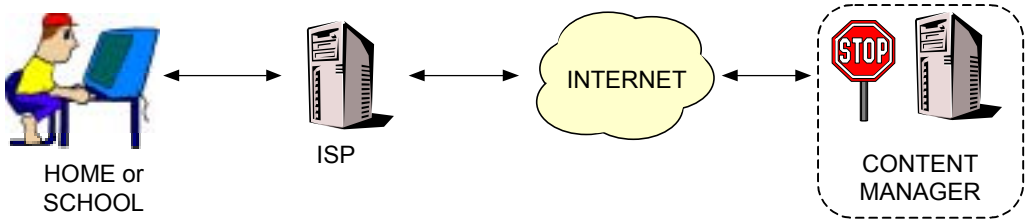

Figure 3 - Content filtering by a third party.

This category of product requires that special software, possibly including specially modified browsers, to be loaded onto the user's machine. This software works with the Content Manager's Web server, giving the third party complete control of the user's Internet experience.

2.9 Server-based Content filtering

Content filtering may be carried out by ISPs and organisations using several server-based technologies, which differ in their cost, effectiveness and potential damage to the experience of other Internet users. Server-based filtering has the advantage that it is the most secure and hard to bypass. The user does not need to load special software on to their home computers, and all their access to the Internet has to pass through the filters.

The primary disadvantages of server-based filtering come from the scale of the task they face. Home filtering products are only working for one user and can afford to spend considerable (in a technical sense) time on checking requests and Web content. User response times will not be adversely impacted by a filtering product taking 0.1 second to examine content as it passes through. The same 0.1 second of processor

usage would be unacceptable to an ISP serving hundreds or thousands of concurrent users.

Proxy servers

The most common server-based filtering technology is based on proxy servers. Proxy servers sit in the path between the user and the Internet and can examine all requests and returned content on their way through. Figure 4 shows a proxy server filtering Web and ftp requests. All clients must go through this proxy server to be able to access the Internet 'proper'. Clients may be required to configure their software to 'point to' this proxy server to be able to access Web pages and ftp files, although newer 'transparent proxies' can be used to avoid this onerous administrative task, at the risk of lower availability and reduced performance.

Proxy servers are general-purpose computers that act as a 'proxy' for a real server elsewhere in the Internet. Their main roles are to act as security gateways, checking traffic going in and out, and to improve Internet access times by keeping copies of frequently accessed information ('caching'). Proxy servers can also act as filters, using basically the same, efficient, mechanism used to implement caching. The proxy can get a Web request and quickly look up the URL in a list of allowed or blocked sites and pages.

Proxy servers can be selective about what they block, and can be configured to block or permit access to a range of Internet-based services, not just Web pages.

On the down side, proxy servers are designed to improve security or access times, not to filter content. In particular, they may not be proof against simple countermeasures, such as using non-standard 'port' numbers for accessing a Web site.

Although proxy servers are general purpose, and can be used to support any filtering technique, in practice they are limited to the highly efficient techniques such as black list filtering based on partial URLs. Proxy servers have to handle large numbers of requests every second and will not normally have enough processor time available to run slow, content-based, filtering tools.

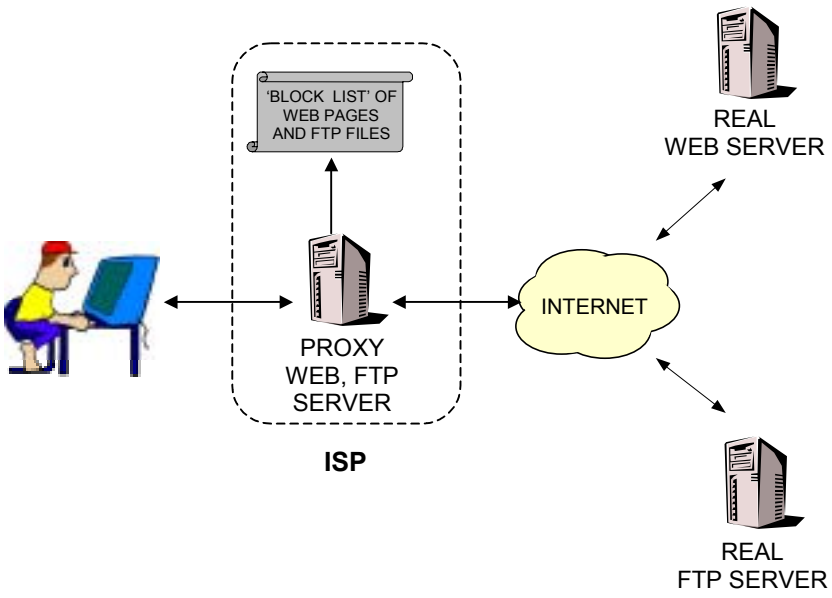

Figure 4 - Content filtering by an ISP using a proxy filter

ISPs can also use specialised caching equipment to filter Internet access. These specialised devices are effectively proxy servers that have been optimised to improve performance and reduce network costs through caching. They are faster than general-purpose proxy servers but are even more limited in their ability to filter.

Differentiated services

Filtering gets in the way of Internet use in many ways. ISPs can provide filtered services, using proxies or specialised caches, but this does involve considerable costs in equipment, operations and administration. Filtered services also tend to deny access to services they don't know about – better safe than sorry. This restricts user's ability to make use of new Internet services and new technologies.

ISPs often provide their clients with both filtered and unfiltered service to overcome these problems. The optional filtered service provides a 'safe' but restricted environment to some of their customers, and the unfiltered service provide fast, low-cost access to everyone else. Making the filtered service optional will reduce costs because it will only ever be used by a subset of customers and so can make use of smaller filtering computers.

Differentiated services are normally provided through the use of multiple access points of multiple accounts. An ISP could offer two dial-in numbers, one for the filtered service and one for unrestricted access. Parents would set up the filtered dial-up for their children, possibly using the unrestricted service themselves and protecting it with a secure password. Some products use a single dial-in number but assign the use to filtered or restricted services based on their account name. Some of the home-based filtering products also offer more or less filtering, based on password-protected accounts.

Figure 5 illustrates shows how an ISP can provide differentiated services for different user groups using different phone numbers.

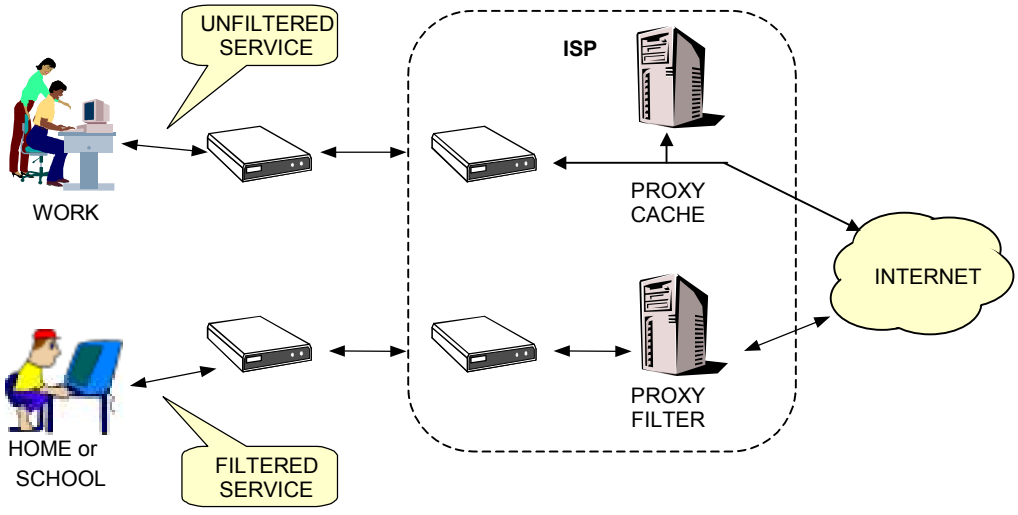

Figure 5 - An ISP can provide differentiated services using different dial-in lines

3 Countermeasures and Other Complications

Client-side filtering products have always had the disadvantage that they are somewhat optional. The user (or their parents or managers) has to choose to install the filtering product and it can be by-passed given sufficient effort and ingenuity[10]. Commercial client-side filtering products are designed to be hard to disable, making it unlikely that they can be by-passed without considerable effort and without this being noticed. The underlying connection to the Internet though is unfiltered and if the user does succeed in disabling the client-side filter they get full access to the wild and untamed world of the Internet.

ISP-based filtering has always had difficulties with performance, cost and possible impact on other Internet users but it has had the advantage of being fairly effective and much harder to avoid. The ISP filter is directly in the path between the user and the Internet and any content coming to the user is already filtered. Users cannot disable this filtering because it is taking place at a secure and remote site.

All filtering technologies are fallible, and the more effective they are, the more they risk intruding on general Internet usage. Products have to strike a balance between filtering out undesirable content, and allowing access to (possibly unknown) useful content. The 'white list' products are the most 'effective' because they are the most restrictive and constrain users to a very small part of the Internet.

Filtering products are susceptible to a large number of attacks and face an equal number of problems caused by the nature of the Internet and Web technology. The Internet world is actually quite complex, even if it seems simple on the surface, and filtering products have to deal with this complexity if they are to be effective.

3.1 URLs and IP addresses

The Internet works with IP addresses. Domain names, such as www.playboy.com, only exist so that people haven't got to remember unwieldy numbers, and in fact, the friendly 'domain names' used in URLs are actually translated to numeric IP addresses before they are used to access content. Users can always use the numeric form of a URL, and this may well deceive a filter looking just for the friendly domain name form. For example, For example, the URLs http://www.playboy.com/centrefolds/ and http://206.251.29.10/centrefolds/ are completely equivalent and interchangeable and filters have to be able to block both forms.

Computer systems on the Internet do not have to have domain names at all, and this is commonly the case when a high volume Web site is built using a number of equivalent servers. The user sends a request to what appears to be a single computer system but their request is actually forwarded on to one of the many computers servicing the site, as shown in Figure 6. Each of these computers has its own IP address and may be directly accessible from the Internet, and capable of returning content on request. Filtering products may block access to the main IP address but users may be able to access the same content by going directly to one of the real servers. Filters that effectively blocked access to content coming from such a site would have to block each possible IP address, a difficult task and a moving target.

[10] Or with the assistance of anti-censoring groups such as Peacefire.

The owners of these Web sites are not trying to make it harder to filter out their content – it is simply a consequence of the technologies they use to handle their very large numbers of requests.

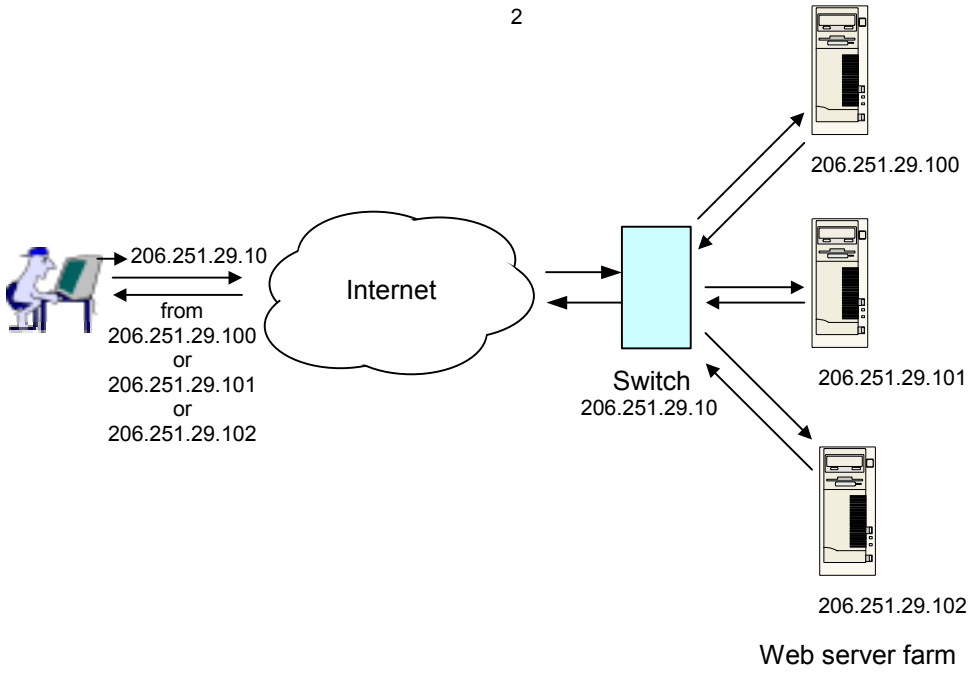

Figure 6 - High volume Web server

A Web page is also not a single monolithic entity, unlike a printed page, but instead is composed of a number of independent components, each with their own URL, that are fetched separately and independently by the browser. Each of these components is directly accessible through its URL and so may also be a candidate for filtering. For example, a filter may block access to *http://www.playboy.com/* but this may not stop direct access to pictures that are used on the Playboy home page, such as *http://www.playboy.com/centrefolds/dec99.gif*.

As discussed in section 2.6, the normal approach to this problem is to block on domain names or partial URLs, such as *www.playboy.com*, within URLs. These domain names are very dynamic and a site can have many domain names, and get new ones allocated quickly. A site, such as Playboy, could use *www.playboy.com* to refer to their main Web server but use additional servers with their own domain names, such as *ww2.playboy.com* and *ww3.playboy.com*, to supply most of their actual content. This use of multiple servers is done to increase capacity, not to evade filtering, but all of these potential names may still have to appear in filter lists for blocking to be effective.

3.2 Filtering the Web or the Internet?

Much attention is paid to filtering Web pages but 'undesirable' content can be found in many places on the Internet, including newsgroups and file servers. Some of the more tightly filtered Internet services, such as some of those designed for the educational market, resolve this problem by completely blocking access to all Internet services other than the Web and e-mail. This approach is certainly safe, but would be

unacceptable for the general Australian community and so these other sources may have to be filtered as well. Web-based news readers, such as DejaNews, can be filtered through the use of URL black lists, although expensive pattern-based filtering might be necessary to deal with URLs as complex as:

http://x17.dejanews.com/fetch.xb?newsgroup=alt.sex.bad.stuff&context=9714053

Proxy servers have to be able to separate out Web traffic from other Internet traffic as it goes past so that it can be sent through to the filtering engine for its consideration. This separation is usually done simply by trapping requests being sent to port 80, the standard port used for Web traffic. Web servers can actually be easily configured to use any port number at all and Web traffic going to other than port 80 is unlikely to be filtered. A site operator wishing to make filtering more difficult could just using another port and include this port number in their URLs, for example htttp://www.porn.com:6969/ [11]. Browsers can use any port number with equal ease, and users going to a site by clicking on a link would not even be aware that an unusual port was being used. It should be noted that currently such schemes, while posing no technical hurdles, are very rare.

An emerging problem with filtering Web traffic through the use of server-side filters is the rapidly increasing use of the Web's protocol (HTTP) and port (80) for other purposes, such as e-commerce and 'Web Services'. Filtering all HTTP traffic could result in degraded performance for major applications, rather than just slowing down interactive Web browsing.

3.3 Blocking useful content

The major problem faced by the developers of filtering products is the sheer size of the Internet, and the difficulties this creates in building up lists of 'unacceptable' sites. Although the actual techniques used by filter vendors is proprietary information, it seems that most of them build their black lists by crawling the Web (in the same way as the search engines build their indexes), and adding pages containing suspect material to 'interim' black lists. These interim lists are supposed to be checked by real people before the sites make it onto the published black lists. Anti-censorship groups, such as Peacefire, have expressed strong doubts about the effectiveness of this manual review, given the number of innocuous sites that have been blocked by major filter vendors.

Even if this 'trawl and review' process works perfectly, vendors will still run into problems of scale, given the hundreds of millions of Web pages and their constant churn, and many sites containing 'undesirable' content will still be accessible. To remedy this problem, some filter vendors couple black lists with content examination techniques, providing a second way of detecting 'undesirable' content.

Content filtering is a difficult problem. Even text-based filtering requires some ability to determine context (and meaning) for words they discover. Early products were infamous for simplistic filtering, with the blocking of 'breast' cancer content being the most quoted example. Filtering products have improved since those early days but the task is still very difficult and moderately high error rates can be expected. Filtering

[11] The '6969' in the URL is the port number to be used by the browser when opening a connection to the Web server.

out non-textual information, such as photographs or video, is much more difficult and problematic.

Filtering errors will either result in 'undesirable' content coming through, or in 'desirable' content being blocked. Black list products, even with back-up content analysis, suffer from both of these problems. White list products should, in theory, not allow any 'undesirable' content through but this is achieved at the expense of denying access to almost all of the Internet.

3.4 Tunnelling

ISP-based filtering has one major advantage over home-based filters – it is much harder to bypass as the Internet feed itself is filtered. Nothing can be done on the PC, apart from using another connection, that will let the user see content that has been filtered out at the ISP – as long as the blocked Web sites are being accessed directly. However, ISP-based filtering *can* be bypassed by accessing blocked sites indirectly, either by using redirectors or tunnelling.

Redirectors can be set up as a publicly available service accessible from anywhere on the Internet. Web users can send their browsing requests to these remote proxies, and possibly bypass any filtering imposed by their local ISP. The ISP's filters would have to scan the entire URL, not just the first part of it to detect that the request was really for a blocked Web page, or block all access to the redirector. For example, a user could send off the request 'http:/open-redirector.org/fetch-www.playboy.com' to a redirector. Their local ISP's filter may just look at the first part of the URL 'open-redirector.org' and pass it through, not realising its real destination. The actual destination could even be encrypted, making filtering more difficult.

Tunnelling raises similar, but much more difficult problems. Tunnelling was originally designed to support Virtual Private Networks (VPNs). VPNs allow an organisation to create a secure private network that runs on top of the public network, linking far flung parts of an organisation without the expense of leasing data lines and building a private network.

Virtual private networks use 'tunnelling' technology to carry private network data across the public network. Tunnelling works by wrapping the IP packet to be sent inside another IP packet, giving the wrapper the address of an appropriate gateway computer and sending it out over the Internet. When the wrapped packet reaches the gateway, the inner packet is extracted and sent on for delivery within a local or private network. Figure 7 shows how tunnelling can be used to secure send packets from one part of an organisation's network to another part, over the public and insecure Internet. Virtual Private Networks are well supported in commercial operating systems, such as Microsoft Windows.

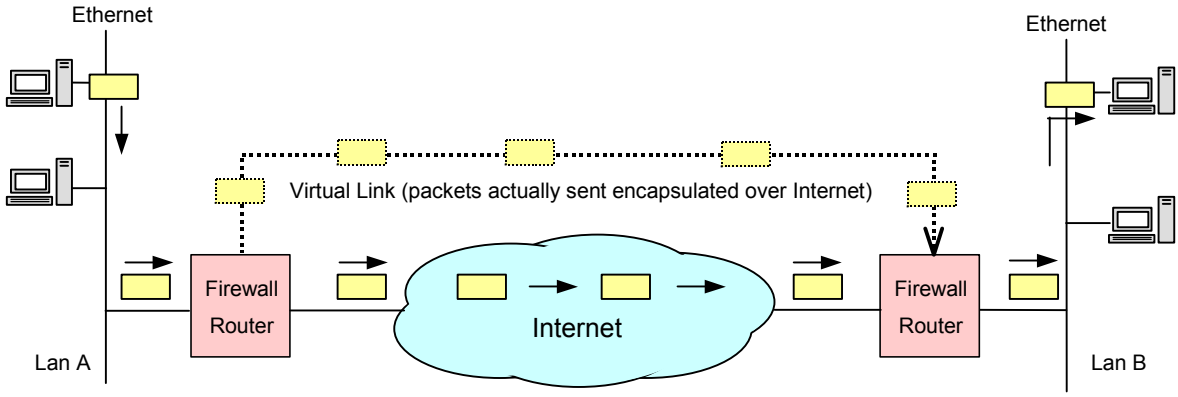

Figure 7 – Use of tunnelling in a Virtual Private Network

Tunnelling and VPNs are also used to support 'roaming'. Roaming allows a user anywhere in the world to access their home network over the public Internet . Travellers can access computers and file servers in their own company back home from the other side of the world, simply by connecting to a local ISP where they happen to be and tunnelling through to their own network. For example, a CSIRO user can access their files and mail in Sydney from New York, simply by dialling a New York ISP and tunnelling across the Internet back to Sydney to a gateway in Australia. All they have to pay for is the local connection in New York, rather than for expensive international phone calls. This use of VPNs and tunnelling is becoming more common as companies set up tunnelling gateways. There are also standards emerging that will make such roaming commonplace. People will soon expect to be able to be anywhere in the world and connect back to their home network using a cheap local Internet connection.

The impact of roaming on filtering is that it can easily be used to by-pass ISP-based filters. An Australian user can establish an account with a US-based ISP and connect to the Internet from there, only using their local ISP to gain access to the tunnel. The local ISP just sees encrypted tunnelling traffic and cannot determine what the traffic means. The user could be accessing a gambling or pornography site anywhere in the world but the local ISP has no idea that this is happening and so cannot filter any such traffic.

The only way of blocking undesirable content coming through an encrypted tunnel is to block tunnelling altogether or to block selected tunnel portals using their IP address and port numbers. Both of these alternatives are quite undesirable and have the potential to cause considerable collateral damage. Blocking tunnelling will block legitimate Internet traffic as well as possible undesirable content, and the ISP and Internet user community will find this unacceptable in general.

Client-side filtering products will still work with tunnelling, as any undesirable data has to be decrypted before it is given to the Web browser to display. ISPs who offer filtered Internet access may have to block tunnelling in their filtered service, otherwise it could be used to by-pass the filtering.

4 Evaluation of filtering products

An initial call for filtering products for evaluation was issued by the National Office for the Information Economy (noIE) and the Internet Industry Association (IIA) in October 1999. The products submitted at this time formed the basis of the products listed in the *Internet Industry Codes of Practice, Schedule 1, Approved Filters.*

This list of 'approved' filtering products was extended through a series of ongoing evaluation reports commissioned by NetAlert and delivered by CSIRO from May 2000 to February 2001. There were 24 products on the list submitted to NetAlert as of the last report in February 2001.

The first stage of this evaluation project was to contact the vendors or local agents for all of the products on the list and request that they provide access to their products for testing. Those vendors who did not reply to this initial request were followed up, sometimes several times, by both CSIRO and the ABA. At the end of this process, we had 14 products available to test. Some of the remaining products had been withdrawn from sale, but most of the others were not tested simply because the vendor did not reply to our requests. The products actually tested, and the reasons for not testing the others, can be found in Table 2.

4.1 Products evaluated

All of the filtering products currently on the list of 'approved' filters are shown in Table 2. Some of these listed products have been withdrawn from sale or have been replaced. Products that were not made available for test, and were consequently not tested, are shaded in the table.

Product	Comment	Date listed	Type	Type Tested	Mode tested
AOL Parental Control	Multiple services, tailored to different age groups	Dec 99	Filtered service	Service	Under 12's 13-15 and 16-17
Arlington Browser	Replacement for KidSafe Browser. This is a customised browser that can be configured to only allow access to specified sites. **No list of allowed or disallowed sites came with the product.**		Client	Client	Site lists
BAIR Filtering System	No reply from vendor	Dec 99	Server	Not tested	
CSM Proxy Server	No reply from vendor	Dec 99	Both client & server	Not tested	

Product	Comment	Date listed	Type	Type Tested	Mode tested
Cyber Patrol	No response from vendor after initial contact. Tests carried out on evaluation copy of product downloaded from vendor Web site.	Dec 99	Both client & server	Client	Site lists + keywords
Cyber Sentinel	Product relies solely on keyword filtering.	Dec 99	Both client & server	Client	Keywords only
CYBERsitter		May 00	Both client & server	Client	Site lists + keywords
Eyeguard	Product relies solely on image analysis.	Dec 99	Both client & server	Client & Server	Images only
Genesis	Manufacturer requested that Genesis be removed from the Schedule.	Dec 99	Server	Not tested	
IFilter	This was a service run by ISeek using the N2H2 filter. This service has been withdrawn and ISeek requested that the N2H2 filter be tested instead.	Dec 99	Filtered service	Not tested	
I-Gear		Dec 99	Server	Server	Site lists + content analysis of unknown sites
Internet Sheriff	Tested using service provided by vendor.	Dec 99	Server	Server	Site lists
InterScan WebManager	No reply from vendor	May 00	Server	Not tested	
KahooTZ	Vendor requested that the product be taken off the IIA list.	Dec 99	Client	Not tested	
KidSafe Browser	Replaced by Arlington Custom Browser	May 00	Client	Not tested	
Kidz.net	Withdrawn from market and may be re-launched. Decision made to defer testing.	Dec 99	Client	Not tested	
N2H2	Added to replace ISeek at the request of the vendor		Client, server & service	Client	Site lists
Net Nanny		Dec 99	Client	Client	Site lists

Product	Comment	Date listed	Type	Type Tested	Mode tested
Norton Internet Security	Filtering supported only in the Family Edition	Aug 00	Client	Client	Site lists
Smart Filter		Aug 00	Server	Server	Site lists
SuperScout	No reply from vendor	Nov 00	Server	Not tested	
SurfWatch	No reply from vendor	Dec 99	Client, server and toolkit	Not tested	
Too C.O.O.L.	White list of sites and proprietary browser	Dec 99	Client	Client	White list of sites
Websense	No reply from vendor	Dec 99	Server	Not tested	
X-Stop		May 00	Both client & server	Client	Site lists

Table 2- Evaluated Filtering Products

4.2 Evaluation methodology

The general approach of the research was to answer the following questions for each of the scheduled filters:

- Is it easy to install, configure, use and update?
- Is it easy to disable or bypass?
- How well does it stop access to undesirable content?
- Does it stop access to desirable content as well?
- Can it effectively track access?

Every available product on the list was installed and tested to see how effectively it carried out the task of filtering and tracking Internet access.

The assessment of ease of use was largely subjective, based on experiences in installing, using and de-installing the products under test. Relative performance scores were given for some product attributes, such as ease of installation. Other capabilities were just checked and given a yes/no score. No attempt was made to combine these subjective ratings to give an overall product score.

Filtering effectiveness was tested by installing the product under test and then attempting to access all of the Web pages on our standard test list. This list includes 895 sites covering 27 content categories, and includes both sites that could be expected to be blocked and sites that should be passed through.

4.3 Testing Ease of Installation and Use

Home-based products

All home-based products were tested for ease of installation and use. Each product was installed, configured, used and its filtering lists updated (where appropriate). The products were evaluated and scored according to the following criteria:

- Ease of installation and deinstallation. How easy is it for an inexperienced user to install and uninstall the product?
- Ease of configuration. How easy is it to configure the product, including setting up classes of users and customised filtering criteria if appropriate.
- Impact on system stability. Did the installation damage the system in any way? Does the product uninstall cleanly?
- Ease of use/transparency. Is the product easy to use? Does it interfere with normal Internet access? Does it slow down Internet access and degrade performance?
- Fetching/installing upgrades to filtering lists. How easy is it to fetch and install updated filtering lists (where appropriate)? Can the process be automated and how easily?
- Accuracy and usefulness of documentation. Is the documentation accurate and complete? Does it provide solutions to common problems that might be encountered during installation and use?

Server-based products

The same basic usability tests were carried put on the server-based products, although we placed less emphasis on ease of use as these products will be installed by people with considerable technical skills, and a more complex installation and configuration process is quite acceptable.

All server-based products were also evaluated for ease of installation and use as perceived by the end user. Some of these products required special PC configuration settings and they also have the potential to degrade Internet access performance.

4.4 Verifying Claimed Capabilities

Basic tests were run on every product to ensure that they meets the claims made by their vendors or agents. The actual tests run in this phase varied from product to product and included:

- Does it block access to Web sites based on their URL? Does it block both the domain name and IP address forms of the URL?
- Does it block searches that include ‘undesirable’ words?
- Does it block content based on included text? Does it try to determine context to avoid ‘collateral blocking’ (such as passing Middlesex and breast cancer)?
- Does it block access to chat rooms, newsgroups or e-mail?

- Does any claimed configurability actually work? Can the system actually be configured for multiple users with different levels of access?
- Does any claimed tracking mechanism actually work? Can it be configured? Are tracking logs held securely?

4.5 Determining Effectiveness

The effectiveness tests aim to determine how well a filtering product does its job. A truly effective filter meets the following criteria:

1. It blocks all undesirable Internet content.
2. It passes all other Internet content through untouched.
3. It is not easily bypassed or disabled.
4. It securely tracks all attempts to access undesirable content.

The diversity and size of the Internet, and the ingenuity of users, means that no product will ever fully meet these goals.

The first two criteria were tested using carefully constructed lists of Web sites, containing both sites with content that could be expected to be blocked and sites with content that should pass through untouched. This standard test list was developed with the assistance of the ABA and NetAlert and covered the following types of sites and content:

Pornography/erotica	Art/Photography
Nudism	Glamour/Lingerie models
Swimsuit models	Sex Education
Contraception	Abortion
Sexual Health	Medical/Health
Gay Rights/Politics	Politics
Drug Education	Drug policy
Free Speech	Filtering Information
Racist/supremacist/nazi/hate	Cults
Drug Advocacy	Macabre/Gross content
Bomb-making/terrorism/…	History of facism/racism
Anti-racism/hate	Atheism/anti-church
Profanity	Anarchy/revolutionary/
Sex laws/issues/…	Redirectors

These categories were chosen to include content that could be expected to be blocked by any effective filter, such as pornography, and innocuous or educational sites that could be blocked accidentally if the filtering was not suitably discerning. The list of

sites does not include any that are included in the list of notified sites sent out to registered filter vendors by the ABA.

We aimed for about 30 distinct and randomly chosen Web pages in each category, and always had at least 20 pages. These numbers were chosen after consultation with CSIRO statisticians who work in the field of experimental design. The final list has 895 sites covering 24 categories. The only exception is 'redirectors' but this category is only included as a test of whether products can be easily bypassed by use of a redirector.

The tests also considered whether the filtering products were easily bypassed or disabled. Peacefire[12], an organisation opposed to Internet censorship, have released an application that attempts to disable many of the commonly used filters and we tried this out on all the products under test.

Filters that block using URL black lists have to take both domain name and IP address forms of the URLs into account (e.g. both www.XXXX.com/pictures.html and 180.123.456.82/pictures.html forms). We carried out tests on each product to determine if it blocked both forms of some selected URLs.

In order to reduce the amount of effort needed to test products against our test list, we automated the testing process as much as possible. Our primary tool was a test driver that automatically fetched all of the pages in the list and recorded the result of the request and the retrieved content. This driver worked well for most products under test but had to be modified to handle a few products that directly worked with a specified browser. Two products, EyeGuard and CyberSentinel, were not automatable using our test tools and had to be tested manually, and as a result were not evaluated against the complete site list.

The test results represent the state of the Internet at the moment in time that the specific test was run. All the sites/pages on the list were chosen because they were stable and normally available. Despite this, some sites were unavailable at times. This problem was exacerbated by the time it took to run all the tests on all products, and some sites were simply unavailable while some tests were being run. The impact of this is that some products returned 'failed' rather than 'passed' or 'blocked' at times.

4.6 Test configurations

All client product were tested on a Dell Pentium (166 MHz processor) system with 64MB of RAM, running Windows Me. This 'low-end' computer was deliberately chosen to exacerbate any performance penalty imposed by the filtering products, and because such systems are still in common use in homes and schools. Microsoft Internet Explorer 5 was used as the browser when the product under test did not come with its own Internet access software.

In most cases, we installed server-side products on a Dell Pentium II system (dual 400MHz processors) running Windows 2000 Server and Microsoft Internet Security and Acceleration Server as the proxy server. Some products were tested against a service provided by the products vendor and in these cases we did not carry out installation or configurability tests.

[12] www.peacefire.org

Most filtering products provide their users with the ability to configure aspects of their behaviour, including what content categories are filtered out. Where possible, we tested products on their default settings, assuming that most home users would just install the product and run it 'out of the box', rather than spend the time needed to understand and configure it. The exception to this rule is for the server-based products, where we assumed that the installation and configuration task would be performed by a technically skilled person. We configured these products after installation to match the filtering profile common to home-based filters, and the configuration used is specified in the relevant product evaluation.

4.7 Notes on test results

Care and diligence has been taken to obtain an accurate understanding of each product's capabilities and shortcomings. However, the necessarily limited time available to install, test, document, and de-install each product makes it possible that some products' features (or shortcomings) may not have been fully appreciated. Consider, for example, testing each product's effect on performance and stability. Stability and performance of a computer are notoriously hard to test, as they can be affected by numerous transient factors that have nothing do to with any installed filtering product. While efforts have been made to minimise such effects, it is impossible to discount them altogether. Similarly, while we are confident that, on the whole, we have produced a fair and accurate assessment of each product, it is *possible* that the product has some feature or behaviour that we have not noticed or understood correctly, and so it is impossible to guarantee that each assessment is 100% accurate. Future versions of this report will aim to rectify any such errors, if they occur. Vendors who feel that some feature of their product has been overlooked should contact us.

We have only documented and assessed the observable behaviour of each product. We have *not* made extensive efforts to discover how each product works internally, although we may speculate on this at times. Instead, we just describe how each product behaves from an end-user point of view, and assume that any sophisticated or intelligent filtering technology will result in the product being more effective on our standard tests.

5 Product Evaluations

All of the products available for test were installed if necessary and then evaluated for both usability and effectiveness. The assessment for each product includes a description, an assessment of its usability and the results of the effectiveness tests. The following information is provided for every product:

Product type

Is the product intended for installation on a home computer or on a server run by an ISP or large organisation?

Product description

This is taken from information supplied by the manufacturer or local distributor, based on public information available from Websites, or from material supplied along with the supplied evaluation software. In some cases, discussions were held with the distributors or manufacturers for purposes of clarification.

System requirements

This describes the type of computer environment that is needed to run the software. This is important in the case of user-based filtering products, since many home and school users may not have modern and powerful computing equipment.

Ease of installation and de-installation

This is regarded as an important parameter for end user products in particular, since it is likely that the filtering software will be loaded by parents or teachers who may not be technically skilled. Some end-user products are provided on CD-ROMs, while others are down-loaded over the Internet. Server-based products may require the installation of matching end-user software.

Ease of use

Once the filtering software has been set up, it is important that it be easy to use. It must be straightforward to change from one user profile to another and easy to over-ride the filter when necessary.

Configurability

Some degree of configuration is required for most filtering products, be it setting up profiles for different members of the family, creating different accounts with an ISP for different family members, or customising a filter list. Since this task is likely to be carried out by a person who may not be technologically skilled, it is important that the product can be configured easily.

How updates are handled

We look at how users of the filtering software update their filter lists and other operational databases, and how frequently these databases are updated.

We also look at any mechanisms provided to let users tailor blocking lists to meet their own needs.

Performance and Stability

We note whether the product has any noticeable effect on system performance and/or stability.

Documentation and Support

We assess the adequacy of product documentation, and the nature of other support available to users of the product.

Notes

If important features or limitations of the product are not covered under any of the other headings, they are noted here.

Price

We quote a retail price where possible. Prices for consumer products are quite elastic, and the price per copy paid by a large ISP may well be considerably less than retail. Prices are generally not quoted for server-side products as these are not widely advertised. All prices are quoted in A$ where possible.

Usability Assessment

For ease of reference, in addition to discussing each product's capabilities, we provide a table giving numeric or yes/no scores for key product attributes. The numeric scores are out of a maximum of 10, with higher numbers indicating a better performance. The entries in that table are:

Easy to install – is the product easy to install for someone with little or no technical knowledge?

Easy to de-install – is the product easy to de-install for someone with little or no technical knowledge?

Simple configuration – is the product easy to configure? For example, is there an intuitive user interface for changing program settings and choosing program options?

Multiple users – if the product supports multiple users, how well is this managed and how easy is this to set up?

Flexible – how flexible is the product in allowing users to configure what is filtered and not filtered? Is it flexible enough to permit access at certain times of the day only, or only for a certain number of hours per week?

Breadth – does the product cover Web pages only, or does it cover other Web services such as email, chat, ftp, and news?

Impact on system performance – does the filtering software noticeably slow down system performance and Internet access? The effect on Internet access speed was assessed by downloading 1k and 11k Web pages from a remote site, measuring download times (over 100 separate downloads), and comparing this to download times on the same machine without any filtering software installed.

Impact on system stability – does the software cause the computer the freeze, crash, or have any other impact on system stability?

Easy to use – is the product easy to use? Does it have an intuitive and simple user interface?

Side effects – does the product interfere with any other activities that the computer might be used for? A score of 10 indicates that the product has no effect on other activities. A score of 0 indicates it has a major impact.

Impact on other users – is it possible to still give some people unfiltered access? How are such users affected by the filtering software? A score of 10 indicates that it has no impact on other users.

Intrusive change – does the software require any major change in the way in which users access the Internet? Is the user required to use a specific Web browser or email program? A score of 10 indicates that users can continue using the Internet without change.

Ease of updating – if the program works with lists of allowed and/or disallowed sites, how easy is it to obtain the latest version of these lists?

Automatic updates (yes/no) – is the software capable of automatically updating its blocked/allowed site lists?

Installation instructions – does the software come with sufficient installation instructions[13]?

Configuration instructions – does the software come with sufficient documentation (either online or on paper) to configure the system?

Troubleshooting and support – how good are the product's troubleshooting and support resources?

Basic URL blocking (yes/no) – is the product capable of blocking URL's (such as, for example http://www.playboy.com)?

IP blocking (yes/no) – is the product capable of blocking the IP address version of a URL (such as, for example, http://129.94.242.46)

Blocking searches (yes/no) – is the product capable of preventing the user *searching* for objectionable content?

Blocking on text (yes/no) – does the product filter content based on the text that occurs in the page?

Context sensitive (yes/no) – does the software do any clever content based filtering that somehow takes into account the context of any page content.

Multiple users and access (yes/no) – is it possible to configure the product for use by multiple people, with varying levels of access.

Add to blocking list (yes/no) – does the product allow the owner of the software to specify sites that should be filtered?

Allow sites through (yes/no) – does the product allow the owner of the software to specify sites that should not be filtered?

Tracking all access (yes/no) – is the product capable of tracking Internet access and reporting or logging attempts to access filtered sites?

[13] Note that if the program's installation is so simple that it requires no installation instructions, it can still obtain maximum marks in this category even if it comes with little or no documentation.

Blocking chat rooms (yes/no) – does the product block Internet chat rooms?

Blocking newsgroups (yes/no) – is the product designed to be capable of blocking access to newsgroups?

Blocking email (yes/no) – is the product capable of controlling the receipt and sending of emails?

Effectiveness Evaluation

The effectiveness evaluation consists of the results of accessing the standard list of 895 sites through the filter under test. These results are presented as a chart showing the percentage of sites blocked or passed in each of the 24 categories. The passed/blocked numbers often do not add up to 100% because tests only reflect the state if the Internet at the time the test was run. Sites come on–line and go away again, leading unavoidably to access failures at times.

5.1 AOL Parental Control (AOL version 6.0)

America Online Inc.
http://www.aol.com/ *or* http://www.aol.com.au

Product Type

ISP filtered service, based on inclusion and exclusion lists depending on the user's profile.

Product Description

AOL provides the Parental Control facility to their members as part of their subscription.

Parental controls are based on user profiles. AOL allows the master account to set restrictions or limitations on the Internet services and content that can be accessed by sub-accounts controlled by the same user.

AOL lets users assign sub-accounts to one of four pre-defined categories:

- Adult (no filtering)
- Mature Teen ('black' list)
- Young Teen ('black' list)
- Kids Only ('white' list)

Each category has pre-set restrictions; some of them can be modified by the master account. Mail, Chat room, Instant Message, Web, Newsgroups, and downloading, are blocked or filtered according to the category settings.

System Requirements

No additional requirement above that required for accessing the AOL service.

Ease of Installation and De-Installation

The system is straightforward to install and de-install.

Ease of Use

The system is quite easy to use. However, the user is required to use AOL's own proprietary software to access the Internet. Users can substitute some non-AOL software, such as email programs, but these are less appealing because they are not integrated with the rest of the AOL software world, and hence can be more difficult to set-up, configure and use. While the user is not *forced* into using AOL's software for Web access, *and* ICQ, *and* IM, *and* email, there is a disincentive in not doing so. Some users will not care about this, and may even welcome the simplicity that comes with a fully integrated platform, but others will still want to use software they have become accustomed to using and may be unwilling to convert just to gain access to a filtered Internet service.

Configurability

AOL Parental Control System comes with four pre-defined user categories, each with its own pre-set restrictions; some of which can be modified by the master account

user. Mail, Chat room, Instant Message, Web, Newsgroups, and downloading are all filtered or blocked according to the category settings.

These in-built categories make it easy for parents to limit access to Internet content and services that they consider unsuitable. This ease of use comes at the cost of limiting the amount of control and flexibility available to parents when configuring exactly what their children are allowed to do and see. Parents have no control, for example, over what categories of sites should be considered objectionable for each age group, being limited to what AOL considers acceptable and appropriate. Users can allow or deny access to particular sites, and this facility is quite easy to use.

How updates are handled

AOL handles its filtering lists internally and the user never has to be concerned with updating lists.

Performance and Stability

On our test machine, there was a noticeable deterioration in network access speed but this could just have been the result of having to access the Internet via the AOL server rather than directly.

Documentation and Support

The online documentation, and online (and telephone) support provided by AOL are excellent.

Local support is available from AOL by calling 1300 734 357, or by email on AussieHelp@aol.com . Live online support is also available 7 days a week on special help/support chat sites set up by AOL.

Response to filtering violation

Attempts to access objectionable material are blocked. No other action is taken.

Price

$A 24.95/month subscription to the AOL service for unlimited hours. $A 8/month for 4 hours access/month.

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	9
Multiple Users	9
Flexible	8
Breadth	9
Detrimental to system performance?	No

Detrimental to system stability?	No
Ease of Use	
Easy to use	9
Side effects	10
Impact on other users	10
Intrusive change	6[14]
List updates	
Ease of updating	10
Automatic updates?	Yes
Documentation	
Installation instructions	10
Configuration instructions	10
Troubleshooting and support	10
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	No[15]
Blocking on text	No
Context sensitive	Yes
Multiple users & access?	Yes
Add to blocking list?	No
Allow sites through?	No
Tracking all access?	No
Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	Yes

Filtering Effectiveness

We tested AOL at each of the pre-defined user categories. As expected, the filtering intended for younger users blocked more Internet content. The 'under 12' category uses a white list to restrict access to a small 'safe' subset of the Internet. The other categories (13-15 and 16-17) use black lists and provide access to much more of the Web, at the cost of allowing through some content that could be considered undesirable. The older teenager category blocks less of the available 'mature' content, such as information on sexual health and sex education.

[14] AOL is a full-functional network, not just a filter or a simple Internet gateway. AOL comes with an integrated set of tools and using other, less proprietary tools may be difficult at times. Internet filtering is only possible using the AOL browser.

[15] On the 'young child' setting, search engines are not included in the 'white' list, so this does limits Web searches for people in this category.

The AOL service was also quite effective at blocking anonymous Web access and the use of redirectors.

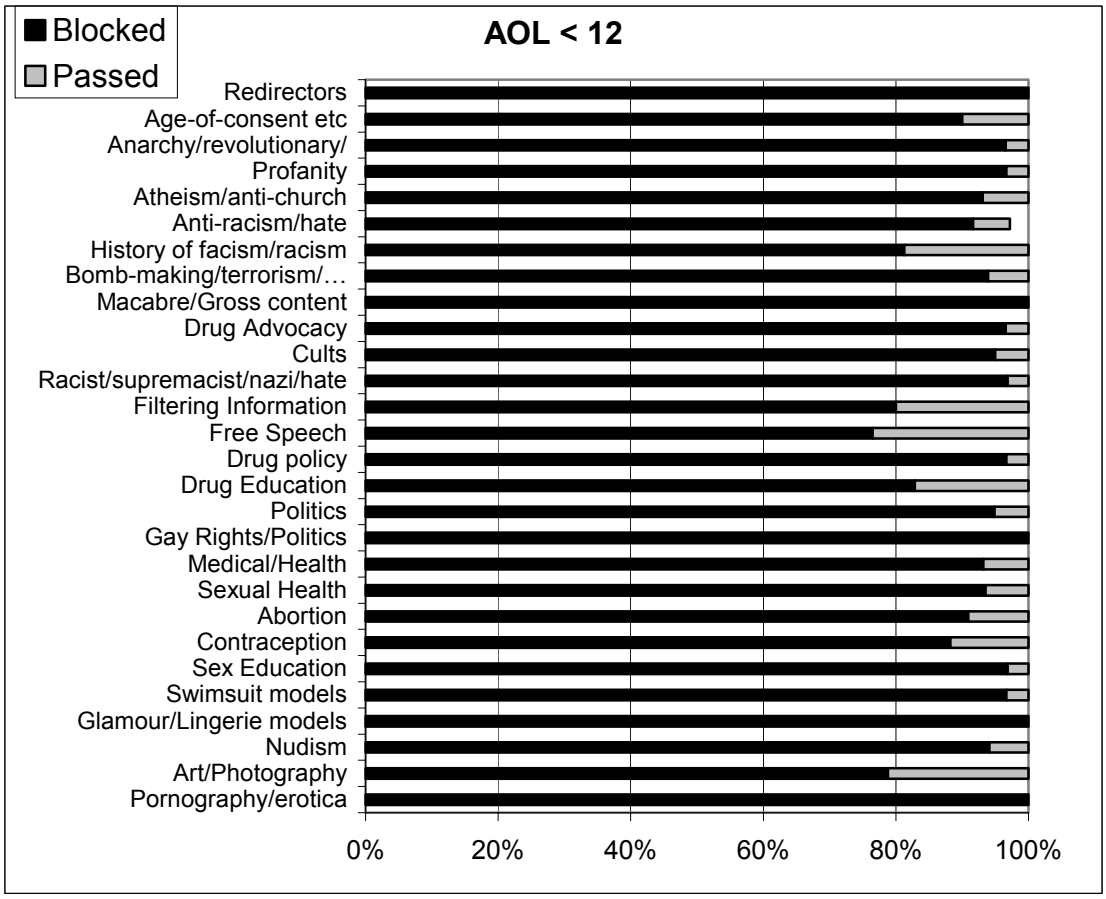

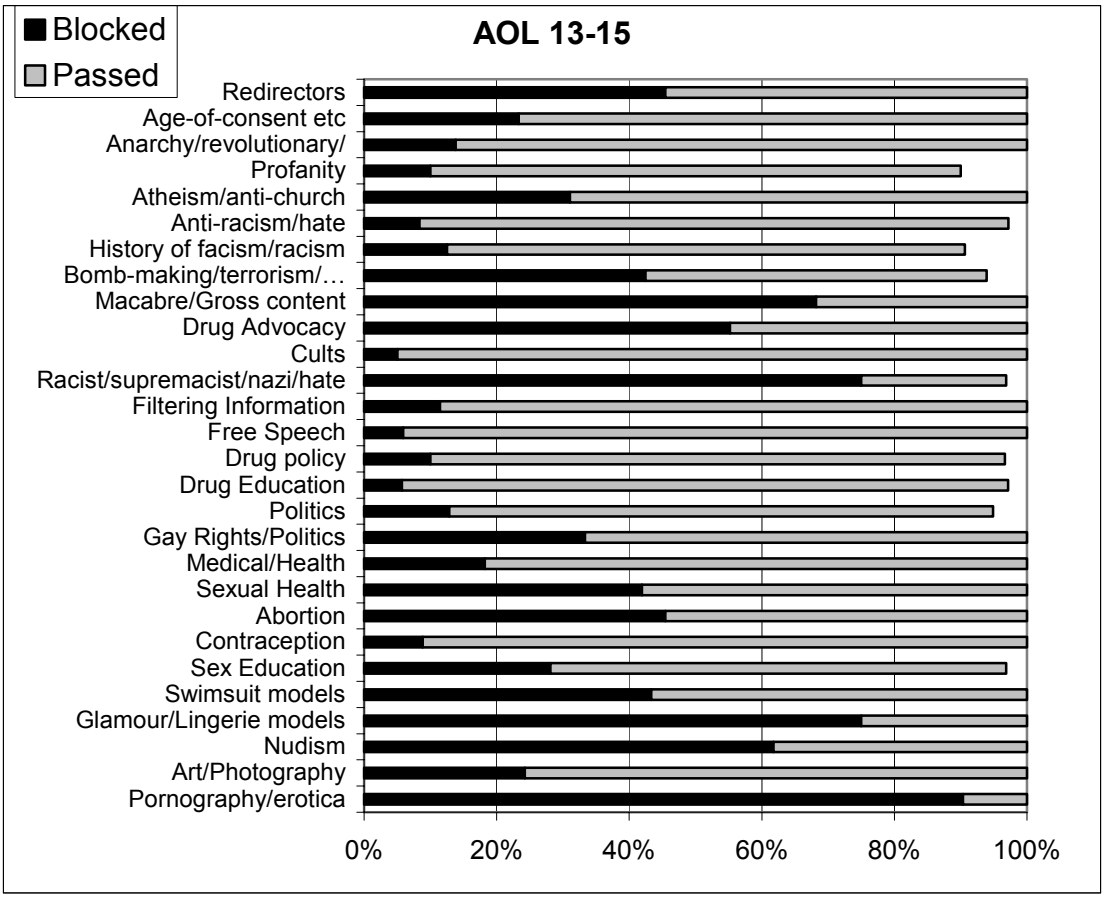
AOL 13-15
Blocked
Passed
Redirectors
Age-of-consent etc
Anarchy/revolutionary/
Profanity
Atheism/anti-church
Anti-racism/hate
History of facism/racism
Bomb-making/terrorism/…
Macabre/Gross content
Drug Advocacy
Cults
Racist/supremacist/nazi/hate
Filtering Information
Free Speech
Drug policy
Drug Education
Politics
Gay Rights/Politics
Medical/Health
Sexual Health
Abortion
Contraception
Sex Education
Swimsuit models
Glamour/Lingerie models
Nudism
Art/Photography
Pornography/erotica
0%
20%
40%
60%
80%
100%

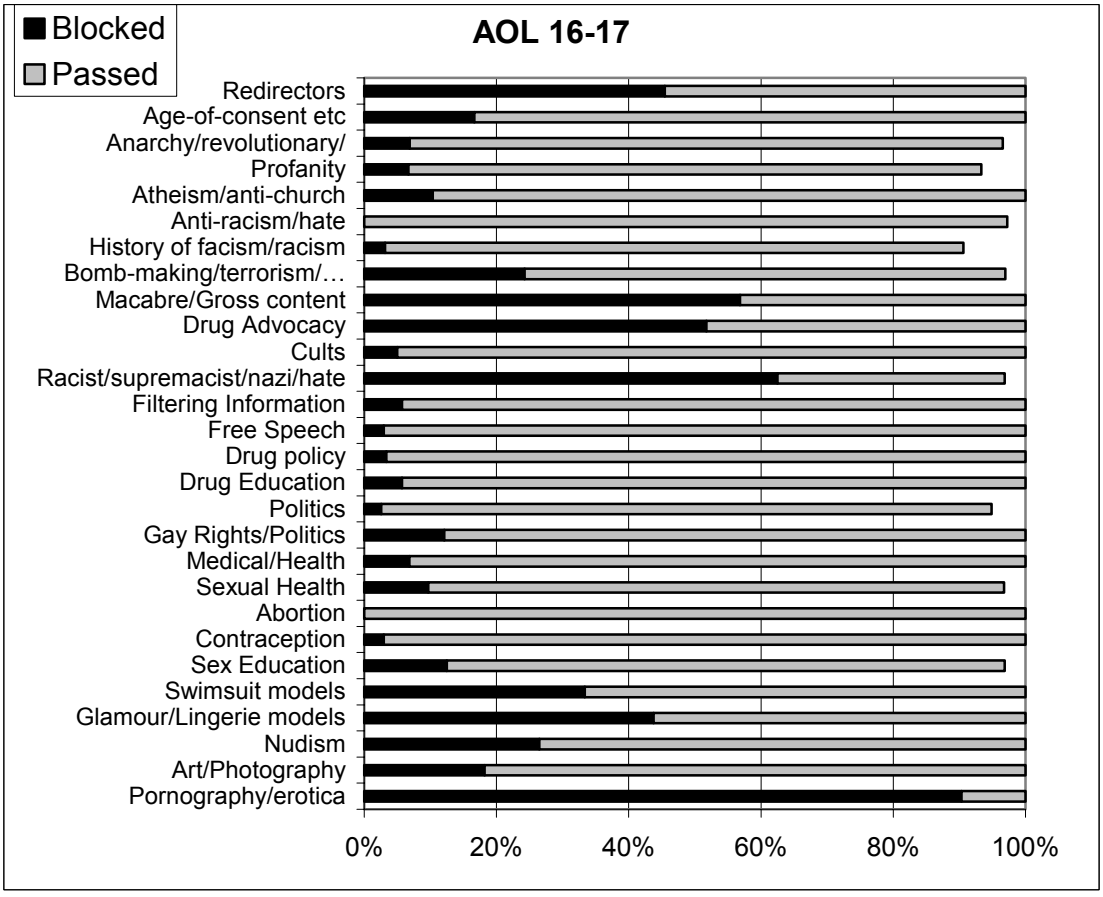
AOL 16-17
Blocked
Passed
Redirectors
Age-of-consent etc
Anarchy/revolutionary/
Profanity
Atheism/anti-church
Anti-racism/hate
History of facism/racism
Bomb-making/terrorism/…
Macabre/Gross content
Drug Advocacy
Cults
Racist/supremacist/nazi/hate
Filtering Information
Free Speech
Drug policy
Drug Education
Politics
Gay Rights/Politics
Medical/Health
Sexual Health
Abortion
Contraception
Sex Education
Swimsuit models
Glamour/Lingerie models
Nudism
Art/Photography
Pornography/erotica
0%
20%
40%
60%
80%
100%

5.2 Arlington Custom Browser

Arlington Technology P/L
http://www.arlington.com.au

Product Type

An end-user product.

Product Description

Arlington Browser is a customised browser that runs on a client machine. The browser is meant to replace standard Web browsers such as Netscape. The difference between the Arlington browser and a normal browser is that the Arlington browser has built-in options for controlling Internet access.

System Requirements

Windows 95 or later, and Internet Explorer 5.0 or later.

Ease of Installation and De-Installation

Very easy.

Ease of Use

The product is very easy to use. The user simply accesses the Web through the Arlington browser as they would through any other Web browser.

Using the Arlington browser in its restricted mode changes the way the user interacts with the Internet – by forcing the user to use a custom browser and preventing access to other programs that may be used to access the Internet. Some users may find this restricted environment unacceptable.

Configurability

The Arlington browser can be configured to give unrestricted access to the Internet, or can restrict the user to specifically sanctioned parts of the Internet. On its most controlled setting, the Arlington browser locks access to the computer, only letting users interact with the browser, so stopping them from using another, unrestricted, browser.

In restricted mode, Arlington browser works on the basis of a white list of 'allowed' sites that is set up by the user. Initially, after installation, these lists are near empty[16] and the user is expected to add sites to them as required.

Arlington browser also has several other capabilities, such as disallowing the download of files with particular extensions (only *.zip, *.exe, *.mp3, and *.com extensions are recognised), and disallowing the launching of particular programs (such as standard Windows programs for running streaming video or sound). The browser is also capable of logging Internet access.

Users can be provided with a password to give them unrestricted access to the Internet.

[16] There are half a dozen sites included in predefined lists for instructional purposes.

How updates are handled

The company does not provide lists of allowed or blocked Web sites; instead the user has full control of what Web sites are allowed.

Performance and Stability

There was no noticeable effect on stability or performance on the client machine.

Documentation and Support

Support is provided by email only (support@arlington.com.au). The trial version of the software we used to evaluate this product did not come with any paper documentation, but installation was straightforward, and, once installed, the browser did include access to help pages.

Response to filtering violation

Attempts to access objectionable material are blocked. Logging of Internet access can also be enabled. No other action is taken.

Price

$US 69 for a licence (from Arlington Web page)

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	10
Multiple Users	0
Flexible	6
Breadth	7
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	5
List updates	
Ease of updating	NA
Automatic updates?	NA

Documentation	
Installation instructions	10
Configuration instructions	10
Troubleshooting and support	7
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes[17]
Blocking searches	No[18]
Blocking on text	No
Context sensitive	No
Multiple users & access?	No
Add to blocking list?	No
Allow sites through?	Yes
Tracking all access?	Yes
Blocking chat rooms?	No
Blocking newsgroups?	No
Blocking email?	No

Filtering Effectiveness

Arlington Browser was not tested for effectiveness, as it does not come with a default list of acceptable Web sites. It relies on the purchaser of the software to specify which sites should be allowed.

[17] Since the user must specifically add 'allowable' sites, these can be automatically blocked.
[18] Of course, searches can be blocked simply by blocking all search engines.

5.3 Cyber Patrol 5.0

SurfControl Inc.
http://www.cyberpatrol.com or http://www.surfcontrol.com

CyberPatrol is now part of the SurfControl family of filtering products. CyberPatrol comes in an end-user (client) product, intended for home and educational use. There is also a proxy-server version of CyberPatrol. In addition, SurfControl offers another product, SuperScout, targeted towards businesses rather than schools or home users but this was not supplied for testing.

Product Type

Client-side filtering product.

Product Description

CyberPatrol is a client-side product designed to be installed on a home (or library or office) computer. Different users on the same computer can be given their individual profiles that determine how and when they can access the Internet. Each user simply logs on with their own password, and their pre-defined filtering settings are then applied.

CyberPatrol works on the basis of 'black' lists of sites. CyberPatrol has an internal list of recognised Web sites, each of which is assigned to a content category. Each CyberPatrol user has a customisable profile that defines which categories they are allowed to access. Separate facilities allow configuration of access to email, chat, and Internet news.

The different content categories defined by CyberPatrol are:

- Violence/profanity
- Partial Nudity
- Full nudity
- Sexual acts/text
- Gross depictions/text
- Intolerance
- Satanic or cult
- Drugs/drug culture
- Militant/extremist
- Sex education
- Questionable/illegal or gambling
- Alcohol and Tobacco

By default, *all* categories are blocked. The user has to explicitly modify these default settings if they want different filtering behaviour. The user is also free to create their own lists of sites that should be allowed or blocked.

The product was tested on its default settings (i.e. all categories blocked, no user specified allow/deny lists).

There are also configurations settings that restrict access to email, Internet chat rooms and newsgroups. It is possible to limit Internet access times to particular hours of the day, or to set a limit on daily/weekly usage. Each of these settings can be different for each individual user.

Finally, CyberPatrol allows the user to enter specific words that they consider objectionable, and should be taken into account when filtering Internet access. For example, CyberPatrol can be configured to block all Web sites with 'sex' in the URL, but to allow through URL's containing 'Wessex'. Similar facilities exist for news, email and Internet chat. Note that, for Web sites, this keyword filtering does not apply to text that is displayed on the sites, only to the URL of the site.

System Requirements

Cyber Patrol is available for Windows 95 or later, and Macintosh 7.1 or later, although the latest version (5.0) is not yet available for Macintosh.

Ease of Installation and De-Installation

Installation is straightforward.

The product provides a de-installation option from the control panel. De-installation is very easy.

Ease of Use

Once installed and configured, it is very easy to use. The user simply accesses the Internet as they normally would.

Configurability

The software provides a password protected control panel where options are provided for configuration. Users can:

- configure the categories for access filtering.
- enable/disable Internet access or automatic downloading.
- update lists to specifically allow or deny access to particular sites.
- configure local allow/block list based.
- configure list and time for Internet Relay Chat access.
- enter personal information to prohibit its transmission (phone numbers, credit card numbers, address, etc).
- set up access control for news groups, ftp, games and other application.
- change passwords and create/delete users.

The configuration, while adequate, is not as intuitive as it could be. It doesn't guide the user as much as it could, and its interface for doing things that should be obvious (such as adding new users and modifying their profiles) could be improved.

The product covers more than simple Web page filtering, and is quite flexible in allowing different filtering configurations for different users. Categorizing Web pages into separate categories is especially helpful in allowing users to alter filtering settings if they choose to do so.

How updates are handled

Updates to the list of blocked or allowed sites can be performed automatically at set times, or at the request of the user. The procedure for doing either of these is straightforward.

Performance and Stability

There was no noticeable effect on stability or performance.

Documentation and Support

Web-based support and documentation are available. Local phone support is available. The software comes with online configuration support and help.

Response to filtering violation

Attempts to access objectionable material are blocked. No other action is taken.

Price

$A 65.77 (from Dymocks catalogue), $US 49.95 (from vendor Web site)

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	8
Multiple Users	10
Flexible	8
Breadth	8
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	Yes
Documentation	
Installation instructions	10
Configuration instructions	7

Troubleshooting and support	8
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	No
Blocking on text	Yes[19]
Context sensitive	No[20]
Multiple users & access?	Yes
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	No
Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	Yes

Filtering Effectiveness

The product was tested on its default settings (i.e. all categories blocked, no user specified allow/deny lists). These settings blocked access to the following categories of sites:

> *Violence/profanity, Partial Nudity, Full nudity, Sexual acts/text, Gross depictions/text, Intolerance, Satanic or cult, Drugs/drug culture, Militant/extremist, Sex education, Questionable/illegal or gambling, Alcohol and Tobacco*

CyberPatrol did not block access to redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector, indicating that the returned page was blocked or the outgoing URL was analysed.

[19] As mentioned, the content of Web pages is not checked for the presence of unacceptable words, but the tool does allow for limited blocking based on text in URL's, and in emails, chat rooms and newsgroups.
[20] We do not consider the limited text based filtering provided to be a useful context-based filtering scheme.

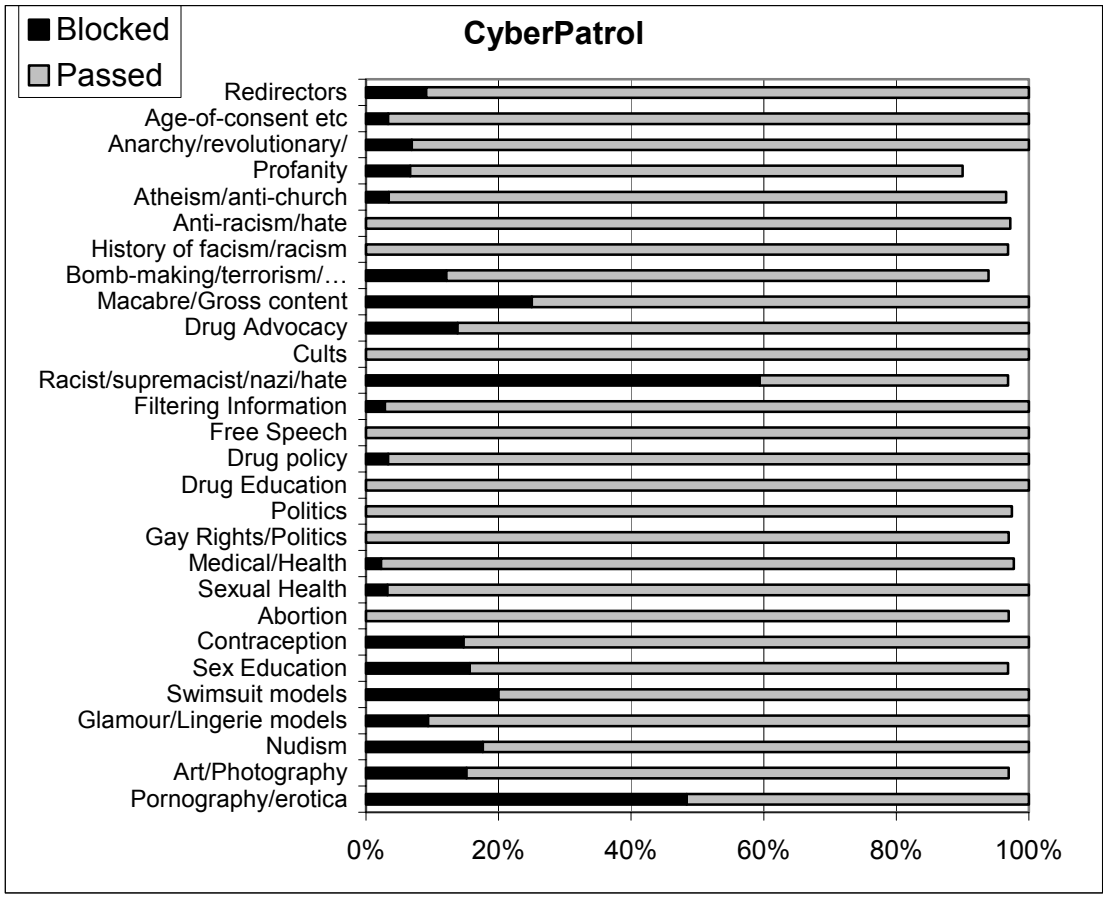
Blocked
Passed
CyberPatrol
Redirectors
Age-of-consent etc
Anarchy/revolutionary/
Profanity
Atheism/anti-church
Anti-racism/hate
History of facism/racism
Bomb-making/terrorism/…
Macabre/Gross content
Drug Advocacy
Cults
Racist/supremacist/nazi/hate
Filtering Information
Free Speech
Drug policy
Drug Education
Politics
Gay Rights/Politics
Medical/Health
Sexual Health
Abortion
Contraception
Sex Education
Swimsuit models
Glamour/Lingerie models
Nudism
Art/Photography
Pornography/erotica
0%
20%
40%
60%
80%
100%

5.4 Cyber Sentinel 2.0

Security Software Systems, Inc
http://securitysoft.com/

Security Software Systems offer several other filtering products in addition to CyberSentinel. Only CyberSentinel is on the list of approved filtering products.

Product Type

Client-side filtering product.

Product Description

CyberSentinel is an unusual product in that it uses neither inclusion nor exclusion lists. Instead, the product works primarily by analysing words that are displayed on the screen. The product does not distinguish between objectionable words displayed on a Web page, or objectionable words entered in an Internet chat session, or even objectionable words displayed in a word processor. This makes the tool very effective in blocking explicit content (whether it be stored on the Web or on your home computer). On the down side, because the product uses a simple word/phrase matching technique, it does not exercise any discretion about what is blocked. If a document (Web page, word document, *any* document) contains an objectionable word or phrase, it will be blocked. This behaviour could be considered intrusive and unacceptable by some users.

The product also supports conventional list-based filtering but these lists have to be built up by the user and none are supplied with the product.

System Requirements

Windows 95 or later.

Ease of Installation and De-Installation

Installation and de-installation are both straightforward. The one trouble we had with installation was that the CyberSentinal CD-ROM did not automatically launch the installation process.

Ease of Use

Once installed and configured, it is very easy to use. The user simply goes about their business as usual.

The pervasiveness of the product can be intrusive, and the fact that it relies so heavily on text displayed on the screen means that it does not prevent access to pornographic pictures with embedded text or without accompanying English text.

Configuring the product to block anything other than pornography is quite difficult. While some key words and phrases occur almost exclusively in a pornographic context, it is harder to think of other words and phrases for subjects such as drugs and bomb-making that would not disallow access to much legitimate material as well.

Configurability

The product has quite good coverage, but is not very flexible. For example, the product allows you to allow, or deny, access to Web services (http), Internet news (nntp), email (smtp, pop), file transfer (ftp), and more, but you have no other option other than to ban *all* access or to allow *all* access. If access is allowed, you must rely on the programs analysis of on-screen text.

The product does not allow different profiles for different users. The filtering product *can* be turned off, but there is no way to have separate setting for different users.

The product automatically logs attempts to access objectionable content by saving a screen shot of the offending document to disk, and can be additionally configured to send an email to a specified guardian. The product can also be configured to perform several different actions when objectionable content is detected, such as freezing the computer (until the administrator password is entered), or warning the user, or shutting down the application that is displaying the objectionable content.

How updates are handled

Since the product does not work on the basis of vendor-supplied allowed/denied lists, it doesn't need any sort of update facility.

Performance and Stability

There was no noticeable effect on stability or performance.

Documentation and Support

Support is available through Banksia Software in Australia. Phone and email support is available (9424 2580 and support@banksiasoftware.com.au). The product also comes with an online manual, but this manual is far from comprehensive in its coverage.

Response to filtering violation

CyberSentinel can be configured to lock access to the computer (until the administrator password is entered), or simply warn the user about an access violation. Logging of access violations also takes place in both cases.

Notes

The on-screen word detection that CyberSentinel employs to enforce its filtering policy does produce some strange results. For example, one nudist site we came across has an entry page which had the message "This site contains No pornography. It contains no pictures of male or female genitals…", and this page was blocked by CyberSentinel, presumably because the words 'pornography' and 'genitals' occurred on the screen.

The filter runs in the background, and new content can be displayed for a second or two before it is blocked.

Price

$US 34.95 (from vendor Web site)

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	9
Easy to uninstall	10
Ease of Configuration	
Simple configuration	9
Multiple Users	0
Flexible	5
Breadth	10
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	6
Impact on other users	7
Intrusive change	10
List updates	
Ease of updating	NA
Automatic updates?	NA
Documentation	
Installation instructions	10
Configuration instructions	10
Troubleshooting and support	8
Claims/Capabilities	
Basic URL blocking	Yes[21]
IP blocking	No
Blocking searches	No
Blocking on text	Yes
Context sensitive	No[22]
Multiple users & access?	No
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	Yes

[21] Site list filtering is supported but the lists have to be built up by the user.
[22] We do not consider the text and phrase-based filtering that CyberSentinel employs to be context sensitive in the proper sense. It is instead a limited extension of simple key-word matching.

Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	Yes

Filtering Effectiveness

CyberSentinel was difficult to test using our automated tools and had to be tested manually. As a result we only tested the Pornography category completely and sampled 1 in 5 of all the other sites on our standard list.

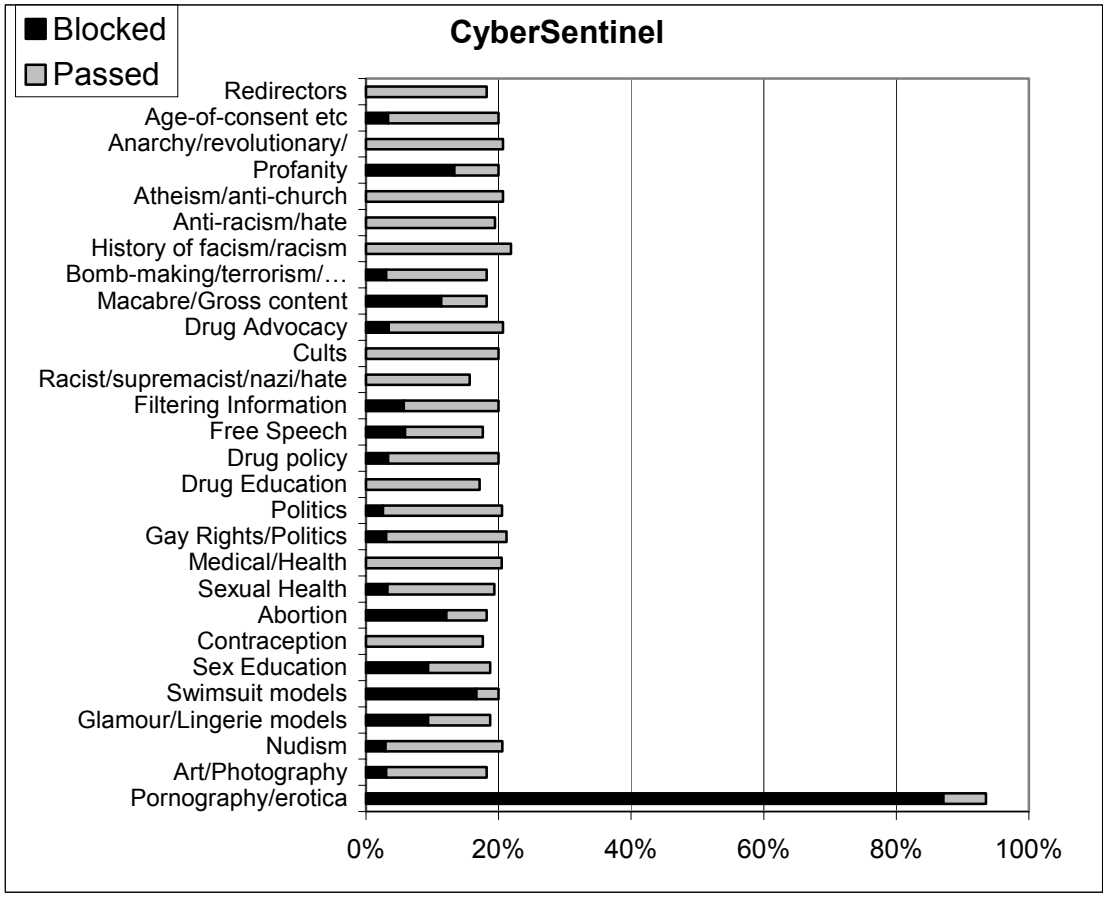

5.5 CyberSitter 2001

Solid Oak Software, Inc.
http://www.cybersitter.com

Product Type

Client-side filtering product. There is also a proxy-server version of the product. Only the client product was tested.

Product Description

CyberSitter primarily uses exclusion lists of banned sites – black lists. Known sites are classified into a large number of different categories, and the user can select which categories should be allowed and which blocked.

Users are free to add sites to either the exclusion lists, or 'white' inclusion lists (which are initially empty) of sites that should be allowed.

CyberSitter also attempts to examine and modify unclassified Web pages to filter out objectionable content. See the notes and effectiveness evaluation sections for more comments on this dynamic filtering technique.

System Requirements

Windows 95 or later.

Ease of Installation and De-Installation

Installation and de-installation are both straightforward.

Ease of Use

Once installed and configured, the product is very easy to use. The user simply goes about their business as usual and, if they wish to use the Internet, must identify themselves and log in using their individual password.

Configurability

CyberSitter is very configurable. The user has a lot of control over what should be filtered. CyberSitter recognises 30 categories of Web content, and the user can specify which categories of Web site are filtered. One major drawback is that the product does not support multiple users, so users wishing to allow varying levels of access for children of different ages would need to reconfigure the settings every time a different child wanted to access the Internet.

In addition to providing a lot of flexibility in choosing which Web categories to block, CyberSitter also has good coverage of other Internet services, such as Instant Messaging, ICQ (Internet chat), ftp, news, and email. The user is also able to limit Internet access times, and to log attempts to access material that does not pass the filtering policy in effect.

The configuration options available allow flexible configuration, and provide good filtering coverage of most Web services. Despite the flexibility of configuration, it is still easy to configure the product, although a little more documentation and/or online help would make this even easier.

How updates are handled

The product can be set up to download updates automatically, or the user can request updates whenever they wish by clicking the appropriate buttons.

Performance and Stability

There was no noticeable effect on stability or performance.

Documentation and Support

Free email technical support 7 days a week on support@cybersitter.com. Phone support in the US only. Comes with good online documentation and help. There is a newsgroup set up for CyberSitter questions (which you can subscribe to at www.cybersitter.com), and on-line Web forms you can fill out if you have a question to ask and do not have email.

Response to filtering violation

Attempts to access objectionable material are blocked. Access violations are also logged.

Notes

CyberSitter is unusual in that it uses both a black list of sites (split up into separate categories) and content analysis to filter out objectionable Internet content.

When CyberSitter does not recognise a Web site as one in its black list, it examines the page, looking for objectionable words and/or phrases. These words are either just deleted if they are found or they can trigger truncation of the page. CyberSitter regularly made such drastic changes to Web pages that they were either not displayable by our Web browser, or severely truncated. Despite truncation, some undesirable content made it through at times.

Price

$US 39.95 for 1 year (from vendor Web site)

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	9
Multiple Users	0
Flexible	10
Breadth	10
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	

Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	10
Configuration instructions	8
Troubleshooting and support	9
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	No[23]
Blocking searches	Yes[24]
Blocking on text	Yes
Context sensitive	No[25]
Multiple users & access?	No
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	Yes
Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	Yes

Filtering Effectiveness

CyberSitter was tested on its default 'out-of-the-box' settings.

CyberSitter's performs dynamic content analysis and attempts to alter or truncate Web pages that contain offensive content. This truncation can be minor (offensive words censored out of the page), or major (the page is altered so markedly that it is not viewable). In the first analysis below, where a truncation has occurred, and that truncation results in a significant portion of some Web page being un-viewable, we count the site as having been blocked, even though part of the page may still be viewable. This is not an ideal solution but it allows easier comparison with other products.

[23] In our limited testing of this feature, we entered the IP address version of several URL's that the product blocked. The product failed to block these equivalent IP address versions, but in many cases, the text content analysis resulted in the page being censored or truncated anyway.
[24] Searches are redirected to the CyberSitter search engine.
[25] We do not consider the text and phrase-based filtering that CyberSitter employs to be context sensitive in the proper sense.

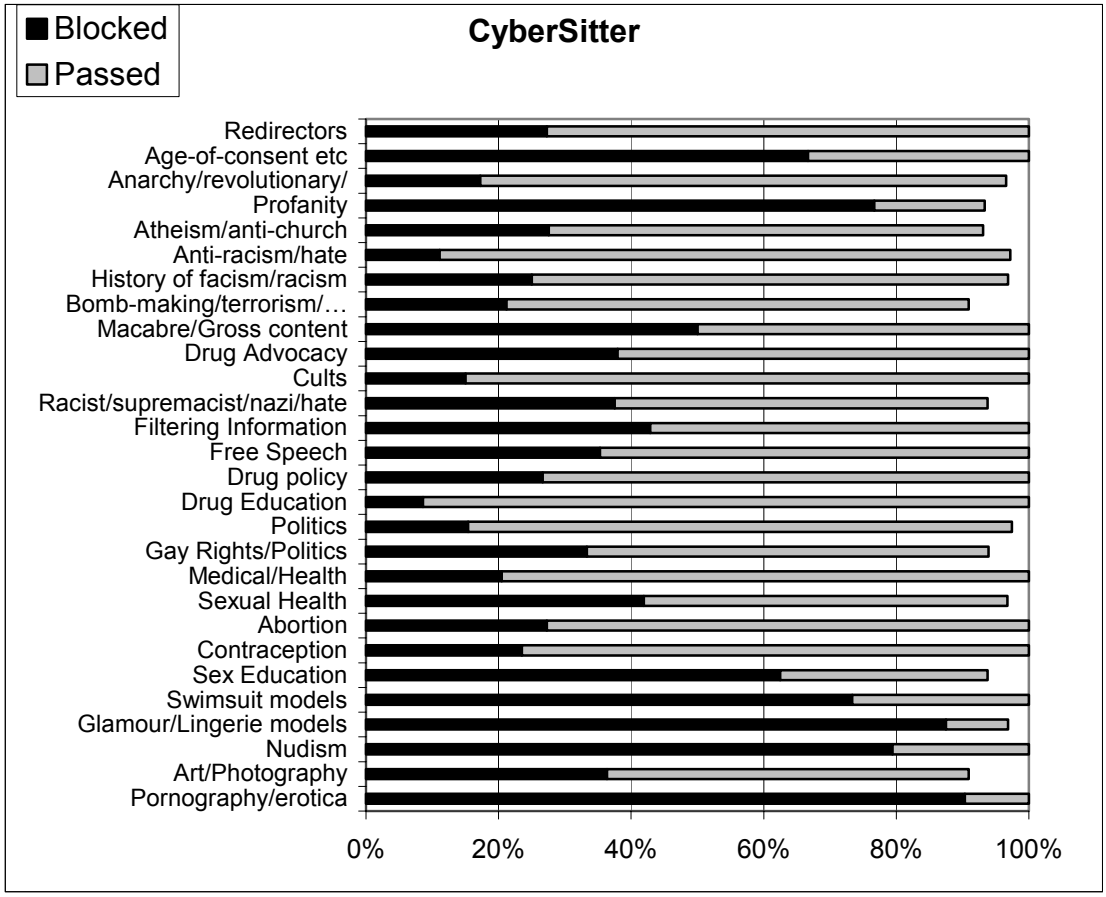

CyberSitter blocked access to two of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector, indicating that the returned page was blocked or the outgoing URL was analysed.

The next analysis chart shows the impact of the dynamic analysis and the resulting truncation of Web content.

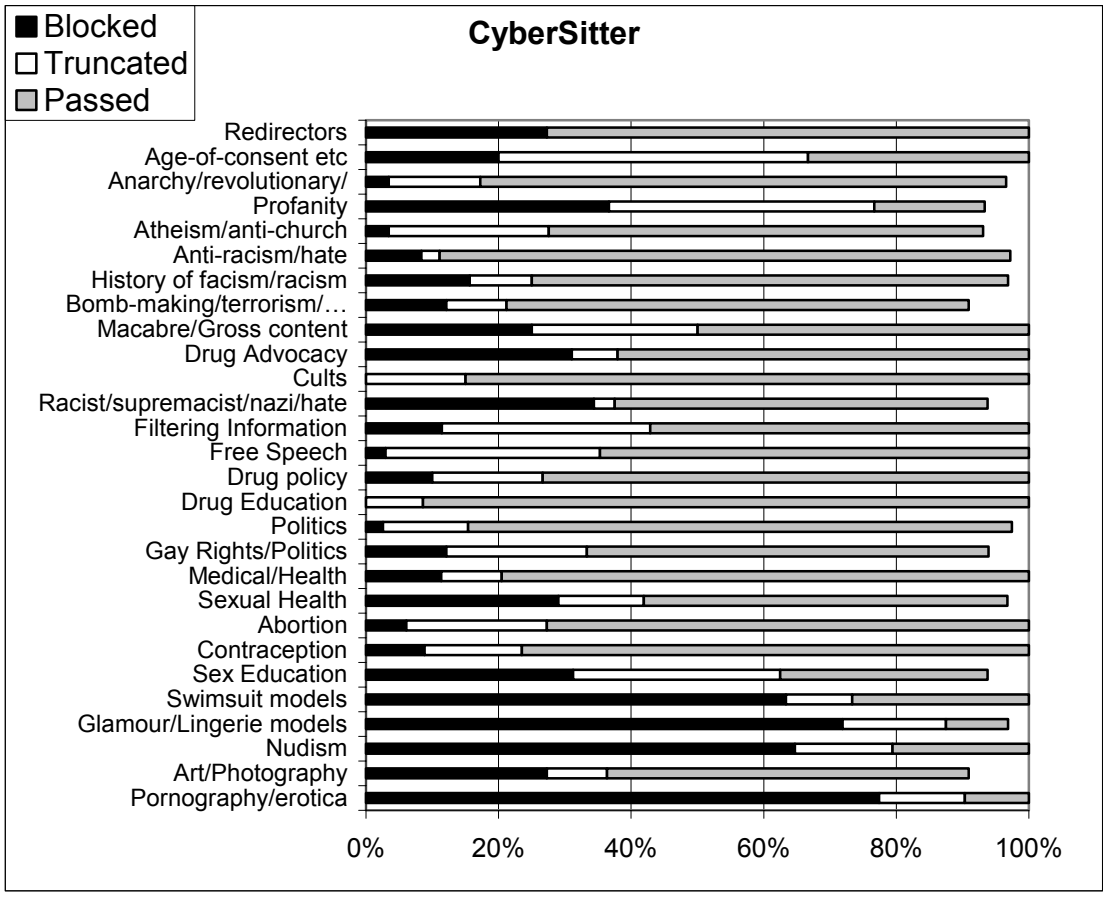
CyberSitter
Blocked
Truncated
Passed
Redirectors
Age-of-consent etc
Anarchy/revolutionary/
Profanity
Atheism/anti-church
Anti-racism/hate
History of facism/racism
Bomb-making/terrorism/…
Macabre/Gross content
Drug Advocacy
Cults
Racist/supremacist/nazi/hate
Filtering Information
Free Speech
Drug policy
Drug Education
Politics
Gay Rights/Politics
Medical/Health
Sexual Health
Abortion
Contraception
Sex Education
Swimsuit models
Glamour/Lingerie models
Nudism
Art/Photography
Pornography/erotica
0%
20%
40%
60%
80%
100%

5.6 Eyeguard

EYE-t Technology
http://www.eye-t.com/

Product Type

An end-user product based on image analysis. Eyeguard has separate client and server components, but both of these components can be configured on a single machine if required. Each machine that is to have its access filtered needs to have an Eyeguard client installed, and access to an Eyeguard server.

Product Description

Eyeguard is an unusual product. It uses image analysis to check images as they are being displayed, and alerts the user if they are considered to be pornographic. It is the only such product tested in this report.

Once installed, Eyeguard attempts to detect sexually explicit images displayed on the screen from *any* source (Internet, e-mail, image files, CD ROMs). The PC can be configured to remain unusable until the supervisor unlocks it with the relevant password if an unacceptable image is detected. Alternatively, it can act as a silent watchdog, detecting the unacceptable images silently, leaving the user unaware. In both cases, the offending images are stored in a central log.

System Requirements

486DX4/100 or higher, Windows 95 or later, 5MB free disk space, 8MB RAM.

Ease of Installation and De-Installation

The installation is straightforward, although necessarily more complicated than a freestanding end-user system, because both an Eyeguard Server and an Eyeguard client must be installed. We had trouble de-installing the product; after selecting de-install from the Control Panel, the product failed to deinstall, unregister, and clean up properly.

Ease of Use

Once installed and configured, the product is easy to use and unobtrusive. The user goes about their business as normal.

Configurability

The product provides a control panel for configuration that is password protected. The authorised user can enable/disable the software, change the way in which the images are interpreted, review the history log and change the password etc. The layout and design of the user interface for configuration, in our opinion, results in some of the configuration functions being unnecessarily complex and difficult to find.

Eyeguard is based on analysing images, and the user can modify a large number of settings that affect how this analysis is done. For example, the user can modify the minimum image area that is to be analysed, and various other setting that govern how sensitive the image analyser is to contrast, skin-tones, and other image qualities. In addition, the user can alter settings that determine how sensitive Eyeguard is – the

more sensitive settings increase the chance of detecting pornographic material, but also increase the chance of incorrectly identifying non-pornographic material as pornographic. Documentation is provided on how to configure Eyeguard, and what affect each tuneable setting has. Despite this, the technical nature of parameter tuning makes configuration of the product much less accessible to the average non-technical user. Anyone wishing to actually configure the product would likely need to engage in extensive trial-and-error tuning.

All testing was carried out with the default settings.

Eyeguard only attempts to block pornographic content, and seems to do this mainly by looking for skin-tones in pictures displayed on the screen. As a result, if the user wishes to block pornographic content, similar pictures of a family at the beach are likely to be blocked as well. There is little that the user can do about this.

How updates are handled

There is no filter list to update.

Performance and Stability

There was no noticeable effect on stability. The product does consume processor resources, but seems fairly efficient in this regard. In our normal testing scenario, with an Internet browser and 1 or 2 other programs running, there was no noticeable degradation in performance.

Documentation and Support

There is online support from 8am-midnight (AESDT) at www.eye-t.com. There is also free email support at eg-tech@eye-t.com. The Eyeguard Web site also contains FAQ's and support documentation. The product itself comes with good documentation.

Response to filtering violation

Eye-Guard can be configured to lock access to the computer entirely (until it is unlocked by an administrator), or to simply warn the user than an access violation has been detected. Access violations are logged.

Notes

Out experience of testing the product suggested to us that, on the default settings at least, skin-tones were the primary trigger for detecting pornographic content. Many black & white nude photographs went undetected by Eyeguard, and unusual lighting, colours, and shadows could throw Eyeguard off. In addition, many pictures with fleshy tones (pictures of faces, and even a couple of desert scenes with skin-like colours) were wrongly classified as pornography.

Price

295 British Pounds for a site licence (from vendor Web site)

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	9
Easy to uninstall	3
Ease of Configuration	
Simple configuration	5
Multiple Users	10
Flexible	NA[26]
Breadth	NA
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	NA
Automatic updates?	NA
Documentation	
Installation instructions	10
Configuration instructions	7
Troubleshooting and support	9
Claims/Capabilities	
Basic URL blocking	NA
IP blocking	NA
Blocking searches	No
Blocking on text	No
Context sensitive	No
Multiple users & access?	Yes
Add to blocking list?	NA
Allow sites through?	No
Tracking all access?	Yes
Blocking chat rooms?	NA

[26] We do not feel it is fair to score this product in comparison with others, as it is so different from other products in the way that it works that we believe a numeric comparison would be misleading.

Blocking newsgroups?	NA
Blocking email?	NA

Filtering Effectiveness

We have not performed a complete test of Eyeguard because it was incompatible with our automated testing tools. Instead, we tested the product against all of the more graphic and sexually oriented categories where it was more likely to be effective. We did not seek out non-sexual graphics that were likely to confuse the product, although the art/photography group did contain some such images.

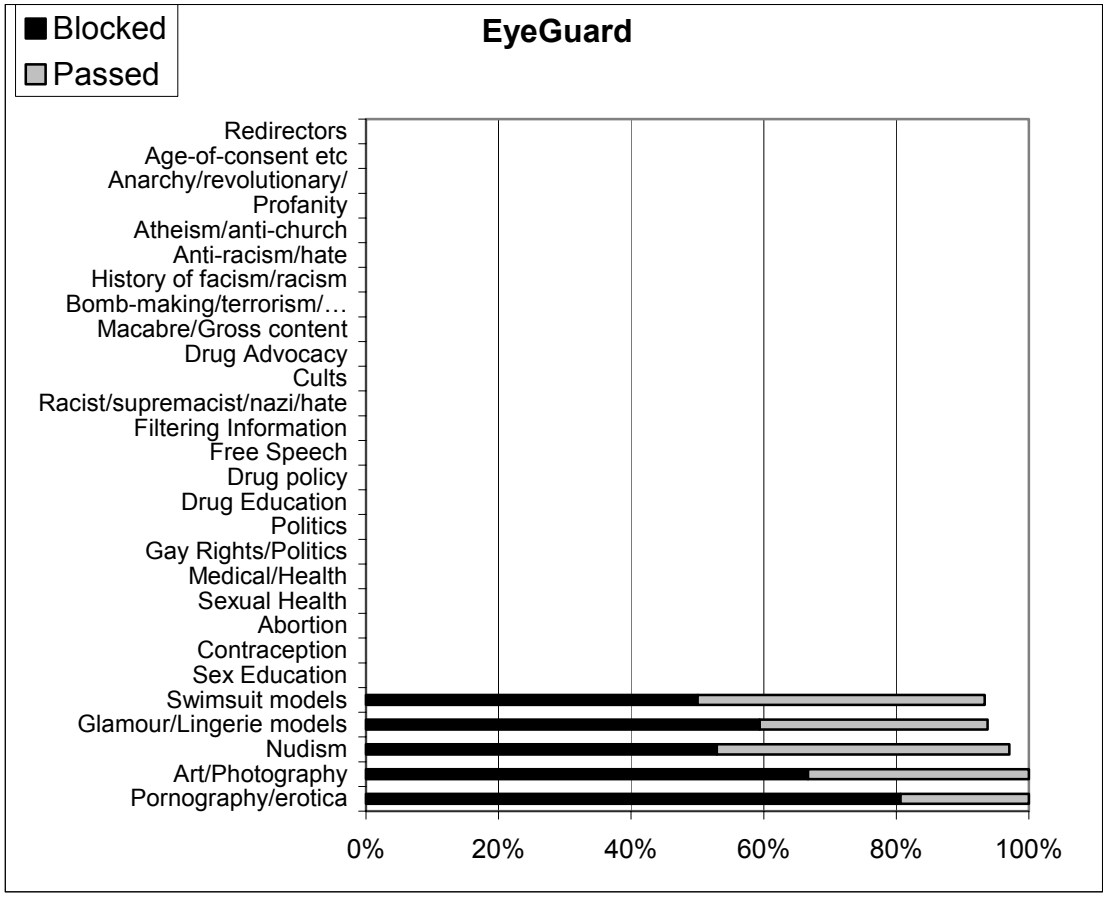

In testing Eyeguard, we varied our normal testing regime to more fully test Eyeguard's capabilities. Some pages (including many pornography sites) on the list are 'entrance' pages that identify the site and give the user links to gain access to the real content in the site. Often these entrance pages contain no offensive material, or no images, but provide links to other Web pages that do. In these situations, to more fully test Eyeguard's effectiveness, we followed these links to a page where graphics were viewable, if such links were available.

Another phenomenon with EyeGuard is that it mostly did not object to the 'thumbnail' graphics commonly found on many pornography sites. Clicking on these

thumbnails led to larger versions of the same images that were blocked. These thumbnail images can still be quite explicit.

5.7 Internet Sheriff

Tel.net Media Pty Ltd
http://www.telnetmedia.com.au/

Product Type

A proxy server based filtering product.

Product Description

Internet Sheriff is a server-side product that filters out different categories of Web content. The system classifies online content into categories such as: *Anarchy and Extremism, Adult Sex,* and *Drugs legal.* The administrator of the system is then free to choose which categories should be blocked.

Our testing was performed by using a remote computer as a Web proxy. This remote computer was owned (and set-up by) TelnetMedia Inc, so we were not able to assess how easy it is to install or configure the product. We requested that the remote computer be set up to block pornographic & extremist sites only. As a result, TelnetMedia set up a proxy-server with the following categories blocked: *Drugs illegal, Vilification and intolerance, Antisocial and offensive, Adult sex* (including *Adult sex related*, and *Sex non-explicit*). Other adult content categories, such as *Adult Nudity,* were allowed.

Internet Sheriff combines exclusion filtering and content based filtering. TelnetMedia maintains a list of 'known' sites, each of which is assigned to a particular content category. In addition, sites that are not classified are subject to content-based filtering on keywords and phrases. TelnetMedia's Web site also claims that 'unknown' sites are classified by comparing them with 'models' that have been build up to classify Web content.

System Requirements

ISP/Web-proxy capable environment.

Ease of Installation and De-Installation

Not applicable. The proxy-server was set up by TelnetMedia.

Ease of Use

Users only need to change their Web-browser's proxy settings and then they can work as normal.

Configurability

We did not configure the server, so we were unable to test its configurability. The product filters Web access only (http/https), but the company plans to extend filtering capabilities to email (SMTP) as well in future versions.

How updates are handled

Internet Sheriff is a server-based product, so updates are effectively automatic for clients. Since we did not have access to the server-side product, we could not investigate the manner in which it obtains list updates.

Performance and Stability

We could not assess the performance and stability of the product on the server end because it was located at an external site. We did notice a slight slowdown in Internet access speed on our client machine, but the nature of the proxy setup (a third party administered proxy at a remote site) made it impossible to tell how much of this was due to merely to the fact that Internet access had to go through a Web-proxy, and how much was due to the filtering software installed on that proxy.

Documentation and Support

We did not obtain a copy of the product's installation and configuration documentation, as the product was installed and configured by TelnetMedia. TelnetMedia's Web site (www.telnetmedia.com.au) lists an email support contact (info@isheriff.com.au).

Price

Negotiable.

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	Not tested
Easy to uninstall	Not tested
Ease of Configuration	
Simple configuration	Not tested
Multiple Users	Not tested
Flexible	Not tested
Breadth	5[27]
Detrimental to system performance?	Not tested
Detrimental to system stability?	Not tested
Ease of Use	
Easy to use	Not tested
Side effects	Not tested
Impact on other users	Not tested
Intrusive change	Not tested
List updates	
Ease of updating	Not tested
Automatic updates?	Not tested
Documentation	
Installation instructions	Not tested

[27] Web filtering only.

Configuration instructions	Not tested
Troubleshooting and support	Not tested
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	Yes
Blocking on text	Yes
Context sensitive	Not tested
Multiple users & access?	Not tested
Add to blocking list?	Not tested
Allow sites through?	Not tested
Tracking all access?	Not tested
Blocking chat rooms?	Not tested
Blocking newsgroups?	Not tested
Blocking email?	No

Filtering Effectiveness

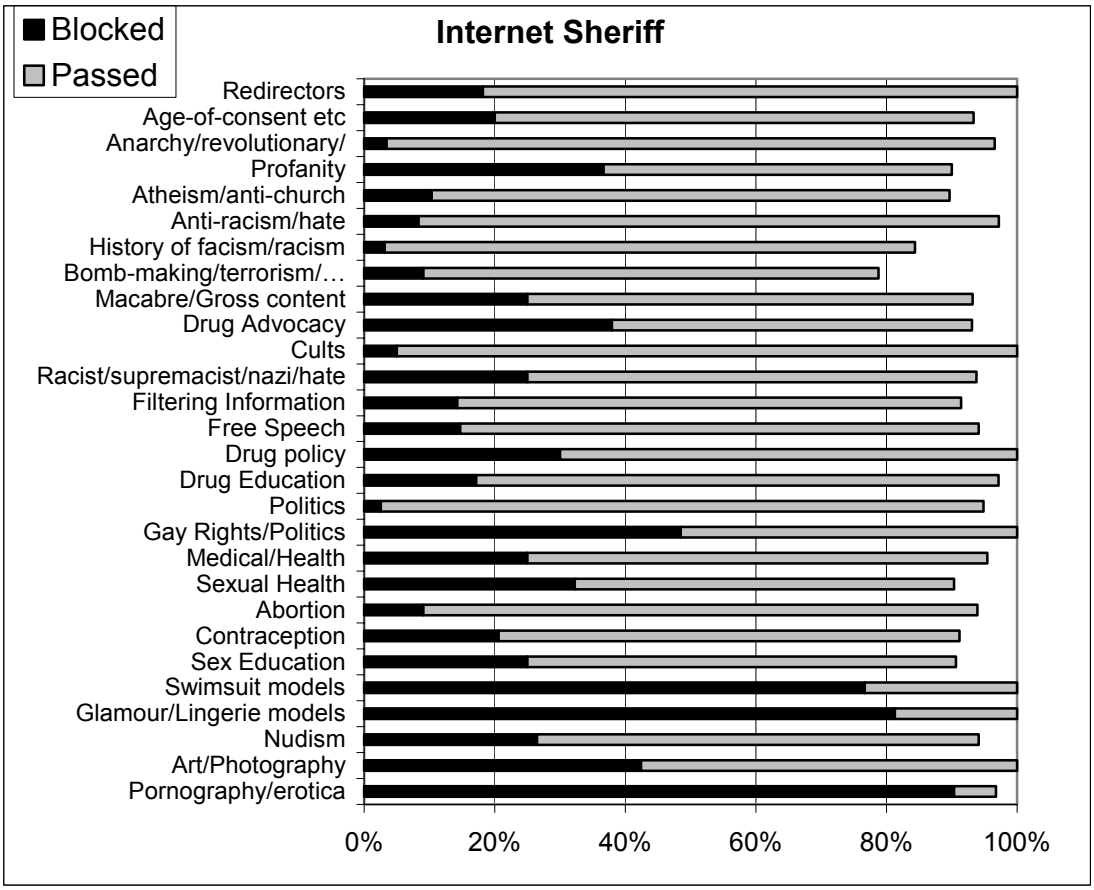

Internet Sheriff did not block access to most of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector, indicating that the returned page was blocked or the outgoing URL was analysed.

5.8 I-Gear 3.5

Symantec
http://enterprisesecurity.symantec.com/

Product Type

A server-based system, employing both exclusion and inclusion lists.

Product Description

I-Gear is a server-side product for filtering Web (http/https) traffic. It has other capabilities but Web traffic is its principal focus. I-Gear recognises 27 different categories of Web site content. Some Web sites may fall into more than one category. Examples of Web-site categories are *Crime* and *Sex-Acts*. A list of known Web sites is maintained by Symantec, and each 'known' Web site is assigned to one or more of these categories.

In addition, to cover Web sites that are not in any list, I-Gear incorporates a Dynamic Document Review (DDR) feature, where html is examined in real time. Incoming text is scored on a points system, using key words and phrases that have both positive and negative weights. The resultant points score for a document is compared with a threshold to determine whether the document should be blocked. This decision is based on words and phrases in the document, *and* on the categories chosen to be blocked. So, for example, if the user configures I-Gear to filter the *Drugs-Advocacy* category, then the phrase 'great pot' may score highly, whereas if this category is allowed, this phrase may attract no score.

System Requirements

Windows NT Server 4.0 with Service Pack 3 or later, OR

Sun Solaris 2.x or later, OR

Red Hat Linux 5.2 or later.

Ease of Installation and De-Installation

The product was installed on a Windows 2000 Server. The installation was very easy, despite the fact that server-side software usually involves more detailed installation and configuration than client-side products.

Ease of Use

The product is straightforward to use from both a client and server-administrator point of view. The interface is clean and intuitive.

Configurability

The product is highly configurable, and the associated administration features are good. Various classes of users or groups of users can be created by the administrator, with differing degrees of access.

Grades of access control range through unfiltered, audit only, black list filtering, white list filtering, local files only, and locked. In addition blocking of download files by their extension can be used to stop any of the standard extensions such as gif, jpg, exe,

as well as proprietary extensions. Access can be set up according to a pre-defined weekly timetable. The reports that can be generated are thorough.

In our testing, we configured the product to block the following categories of sites: *Crime, Drugs/Advocacy, Drugs/Non-medical, Sex/Acts, SexEd/Sexuality, Violence.* Categories not blocked were: *Alcohol/Tobacco, E/Games, E/Sport, Finance, Gambling, Interactive/Chat, Interactive/Mail, Job Search, News, Occult/New Age, Prescription/Medicine, Real Estate, Religion, Sex/Attire, Sex/Nudity, Sex/Personals, SexEd/Advanced, SexEd/Basic, Vehicles, Weapons.*

The Dynamic Document Review (DDR) facility comes pre-configured, and the administrator does not need to configure this at all. If the administrator of the product does wish to configure the DDR facility, configuration tools are provided for adding words or phrases, or adjusting the sensitivity of the dynamic filtering mechanism.

How updates are handled

Clients need not be aware of list updating, since they access the Internet via the proxy-server. List updating on this proxy-server is very easy (handled by a single button click). The tool can be configured easily to handle these updates automatically.

Performance and Stability

There was no noticeable effect on stability or performance on the client machine. Of course, proxy-based filtering does impose demands on the proxy machine that does the filtering, but we did not perform testing on software's effect on the proxy-machine. Instead, we ran the proxy on a lightly loaded machine handling no other network access.

Documentation and Support

The product comes with good documentation, describing the installation, configuration, and behaviour of the component. Online help is incorporated into the configuration tools.

The Symantec support Web page for I-Gear (http://www.symantec.com/techsupp/ent/i-gear), contains good support options, from online troubleshooting guides, to contact options for Web support, to phone support[28].

Response to filtering violation

Attempts to access objectionable material are blocked. The product also can log Web access. No other action is taken.

Price

Negotiable.

[28] Phone support attracts a support fee.

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	9
Multiple Users	10
Flexible	10
Breadth	7
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	10
Configuration instructions	10
Troubleshooting and support	10
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	Yes
Blocking on text	Yes
Context sensitive	Yes
Multiple users & access?	Yes
Add to blocking list?	Yes
Allow sites through?	Yes

Tracking all access?	Yes
Blocking chat rooms?	Yes[29]
Blocking newsgroups?	No
Blocking email?	No

Filtering Effectiveness

In our first test, we configured the product to block the following categories of sites: *Crime, Drugs/Advocacy, Drugs/Non-medical, Sex/Acts, SexEd/Sexuality, Sex/Nudity, Violence.* Categories not blocked were: *Alcohol/Tobacco, E/Games, E/Sport, Finance, Gambling, Interactive/Chat, Interactive/Mail, Job Search, News, Occult/New Age, Prescription/Medicine, Real Estate, Religion, Sex/Attire, Sex/Personals, SexEd/Advanced, SexEd/Basic, Vehicles, Weapons.*

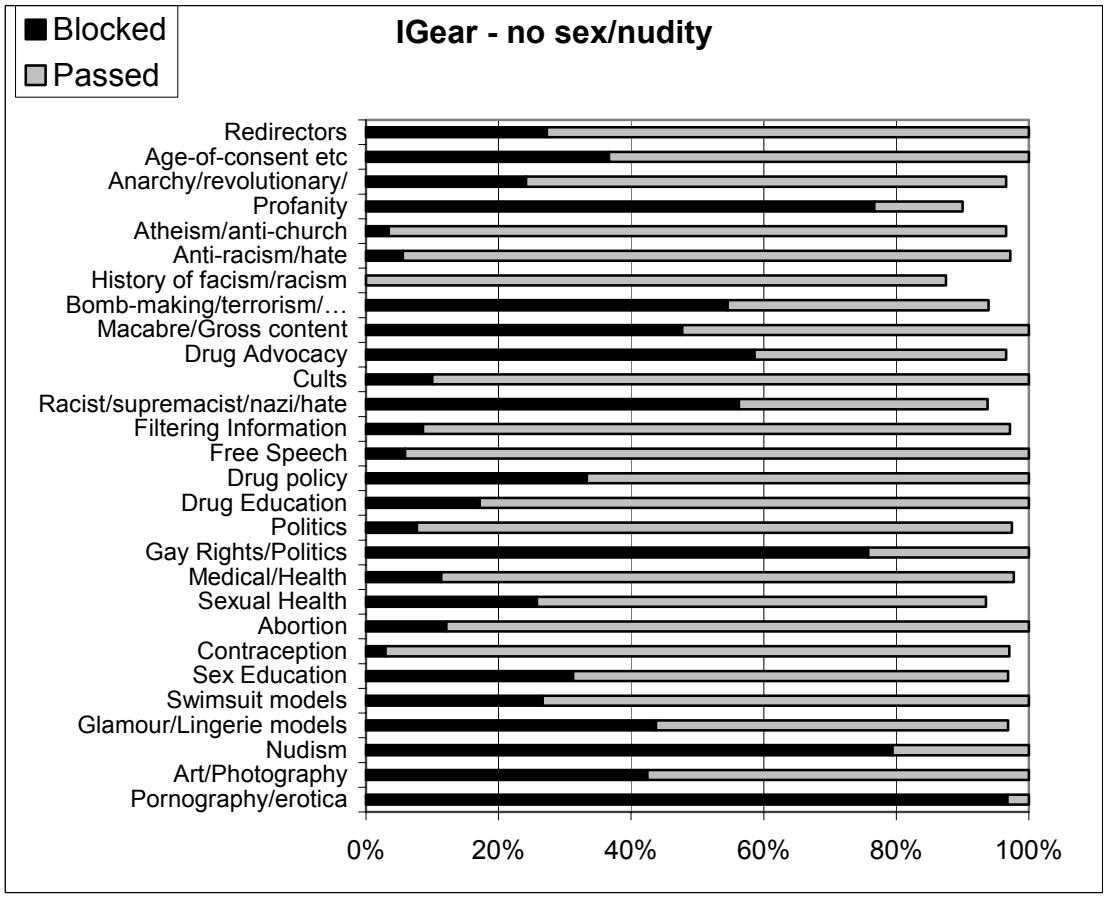

We then repeated the test allowing *sex/nudity* through unscathed. The main impact of this change was to allow access to nudism sites and to more art/photography content.

[29] I-Gear has a *Chat* category that can be blocked or filtered, but, because it is primarily focussed on filtering http traffic, it does not include facilities for blocking and monitoring other chat mediums such as Instant Messaging or ICQ.

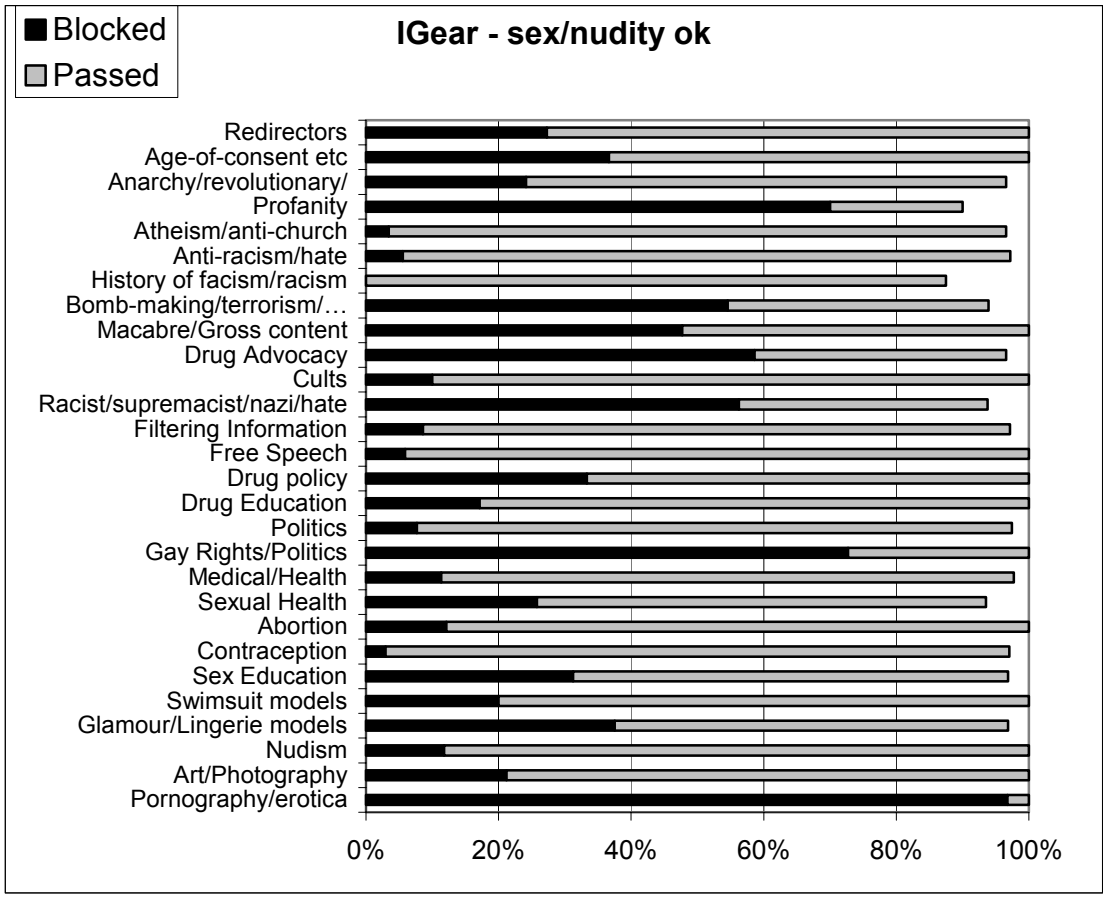

IGear did not block access to most of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector, indicating that the returned page was blocked or the outgoing URL was analysed.

5.9 N2H2

N2H2
http://www.n2h2.com/

Product Type

N2H2 makes client-side (end-user) and server-side filtering products, as well as offering a server-side filtering service. We evaluated only the client-side product.

Product Description

N2H2's Internet client-side filtering product works based on banned/allowed categories of Web content. The user installs the software and, after configuration, accesses the Internet as normal.

N2H2 was not on the original IIA list of approved filtering products, but it was the filtering engine behind the registered ISeek service. ISeek have withdrawn this service and requested that we test N2H2 instead.

System Requirements

Windows 95 or later.

Ease of Installation and De-Installation

Installation and de-installation are straightforward.

Ease of Use

The product is very easy to use, once installed and configured. The user simply accesses the Web as they normally would. Each user has a specific profile, and needs to enter their password to access the Internet.

Users can be granted unrestricted access, and simply have to enter the appropriate password to be granted this access.

Configurability

N2H2 allows four different filtering levels:

1) No Filtering
2) Minimal Filtering
3) Typical School Filtering
4) Maximum Filtering

Each Internet user can be assigned a different user ID and a different filtering level.

Beyond setting the filtering policy for each user, no other configuration of the filtering policy is possible.

How updates are handled

Updates are handled automatically.

Performance and Stability

We noticed a marked slow-down in the speed of Internet access when using N2H2. The product had no effect on system stability.

Documentation and Support

Installation and configuration of N2H2 is very straightforward, and the minimal online documentation on these topics is more than adequate.

The N2H2 Web site (http://www.n2h2.com), contains solutions to common installation and configuration problems. Phone support is also available for a fee. No email support.

Response to filtering violation

Attempts to access objectionable material are blocked. No other action is taken.

Price

$US 39.95 (from vendor Web site)

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	10
Multiple Users	10
Flexible	5
Breadth	5
Detrimental to system performance?	Yes
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	10
Configuration instructions	10

Troubleshooting and support	8
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	No[30]
Blocking on text	No
Context sensitive	No
Multiple users & access?	Yes
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	No
Blocking chat rooms?	No
Blocking newsgroups?	No
Blocking email?	No

Filtering Effectiveness

The nature of our automated testing framework did not allow us to test N2H2 on settings other than its 'maximum filtering' setting. We present results below for that setting only.

N2H2 blocked access to most of the redirectors and sites that support anonymous Web access. It also blocked access to a pornographic site accessed through a redirector while allowing access to an innocuous site, indicating that the returned page was blocked or the outgoing URL was analysed.

[30] The filtering policy with maximal filtering does block search engines, however.

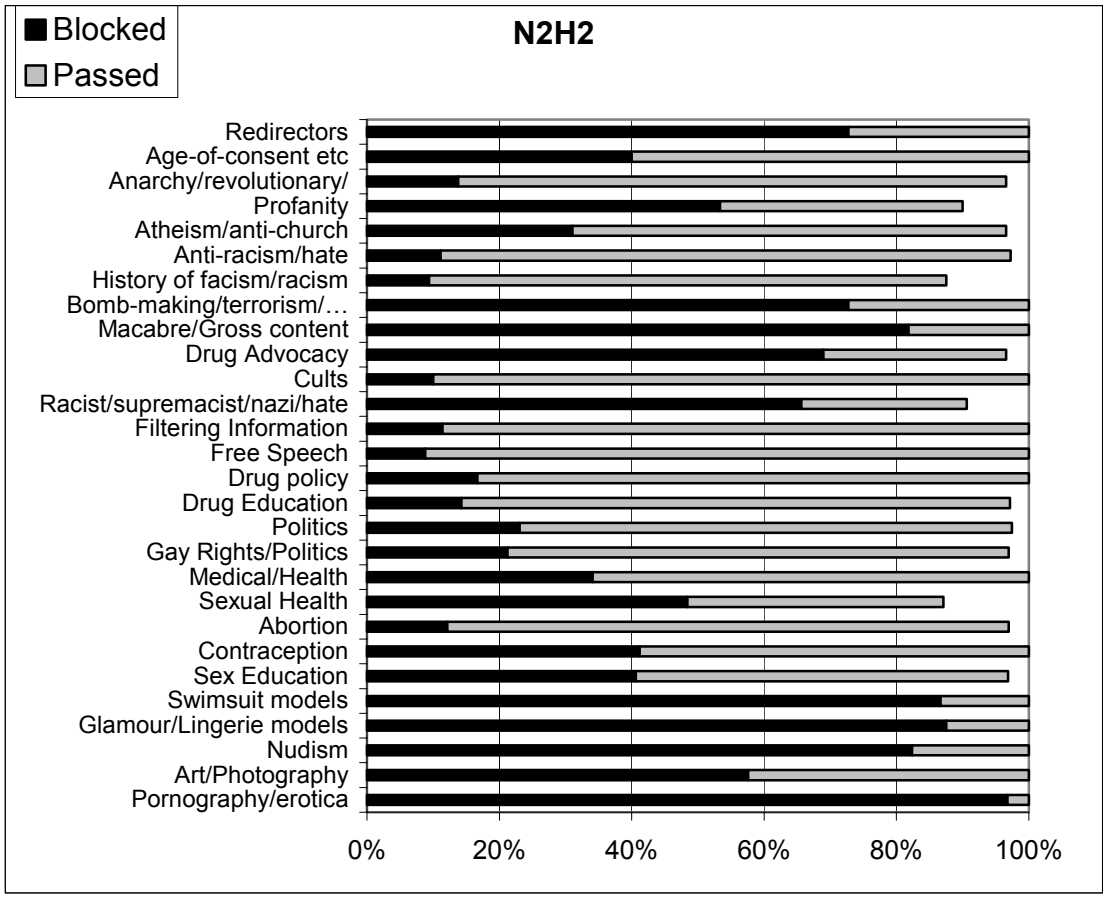
Blocked
Passed
N2H2
Redirectors
Age-of-consent etc
Anarchy/revolutionary/
Profanity
Atheism/anti-church
Anti-racism/hate
History of facism/racism
Bomb-making/terrorism/…
Macabre/Gross content
Drug Advocacy
Cults
Racist/supremacist/nazi/hate
Filtering Information
Free Speech
Drug policy
Drug Education
Politics
Gay Rights/Politics
Medical/Health
Sexual Health
Abortion
Contraception
Sex Education
Swimsuit models
Glamour/Lingerie models
Nudism
Art/Photography
Pornography/erotica
0%
20%
40%
60%
80%
100%

5.10 Net Nanny 4.0

Net Nanny Software International Inc
http://www.netnanny.com/

Product Type

An end-user product.

Product Description

NetNanny is an end-user product that controls access to Web sites using both "black" and "white" lists. Access to other Internet services, such as Internet chat or news, can also be somewhat restricted.

System Requirements

Windows 95 or later.

Ease of Installation and De-Installation

Very easy.

Ease of Use

The product is very easy to use, once installed and configured. The user simply accesses the Web as they normally would. Each user has a specific profile, and needs to enter their password to access the Internet. NetNanny prompts for this password when the user first accesses the Internet.

Users can be granted unrestricted access, and simply have to enter the appropriate password to be granted this access.

Configurability

NetNanny comes with a predefined "black" list of Web sites. Unlike other filtering software vendors, this list is not encrypted and the user is able to browse through the list of blocked sites.

The user can easily add to the list of blocked sites, and, in addition, can add sites to a "white" list, so that they are never blocked. Both these operations are easy to perform.

NetNanny can be configured to limit Internet access times, and can record filtering policy violations. It can also prevent Web browsers from downloading or displaying particular file types – for example, the user can stop common image files from being displayed.

NetNanny supports different profiles for different users. Each user has a separate password and a separate filtering policy can be applied for each user.

The major shortcoming of NetNanny's administration is its "black" list is an "all-or-nothing" approach. Web sites are not broken down into categories. Instead, the user must decide whether all sites in the "black" list should be blocked, or none of them. It would be useful if the list of banned sites was further divided into separate lists, where each list contained Web sites on some common theme.

We found the configuration set-up of NertNanny a little confusing and hard to navigate through. A nice graphical interface is provided, but the layout of the different

menus, and the way in which configuration options must be accessed is not as intuitive as it could be.

How updates are handled

Updates can be requested by clicking a button. This obtains the latest list of banned sites from the NetNanny homepage. NetNanny can also be configured to obtain such updates automatically.

Performance and Stability

We noticed a marked slow-down in the speed of our Internet connection when using NetNanny. This slow-down was noted consistently across 3 separate installs of the product, and was substantial.

We also found that NetNanny seemed to clash with the AntiVirus software installed on our test machine, although this behaviour was not always repeatable.

Documentation and Support

NetNanny has associated documentation that describes the installation and configuration of the product. Frequently asked questions about how to install, configure, and use NetNanny also appear on the company Website.

Support is via email (support@netnanny.com) only.

Response to filtering violation

Attempts to access objectionable material are blocked. Access violations are also logged.

Price

$A 79.95 (from Dymocks catalogue), US$ 39.95 (from vendor Web site).

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	7
Multiple Users	10
Flexible	8
Breadth	7
Detrimental to system performance?	Yes
Detrimental to system stability?	Yes
Ease of Use	
Easy to use	10
Side effects	10

Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	10
Configuration instructions	9
Troubleshooting and support	9
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	No
Blocking on text	No[31]
Context sensitive	No
Multiple users & access?	Yes
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	Yes
Blocking chat rooms?	No
Blocking newsgroups?	No
Blocking email?	No

[31] NetNanny does have an option which replaces obscene words with '#'s. The rest of the page that contained these obscene words is still displayed.

Filtering Effectiveness

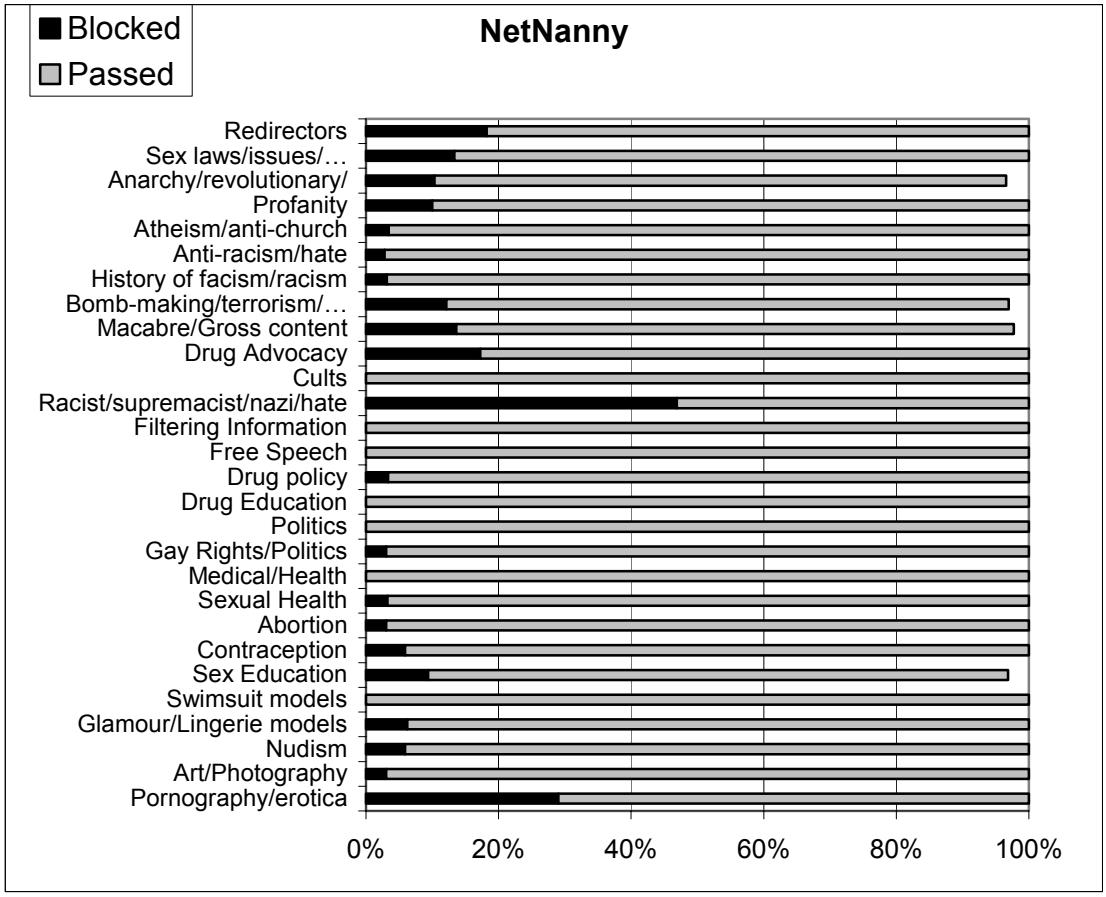

NetNanny allowed access to most of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector, indicating that the returned page was blocked or the outgoing URL was analysed.

5.11 Norton Internet Security 3.0 Family Edition

Symantec
http://www.symantec.com

Product Type

An end-user product.

Product Description

Norton Internet Security is a product that controls access to Web sites by means of both "black" and "white" lists. Access to all other Internet services, such as news, ftp, and instant messaging, can also be blocked.

Internet filtering is just part of a complete Internet Administration package from Norton that attempts to protect the user from malicious programs and Internet content. The product is, in effect, a security and network management software package that includes a capability to filter Internet content. We assess the product only from an Internet filtering perspective.

System Requirements

Windows 95 or later and Internet Explorer 4.01 or later.

Ease of Installation and De-Installation

Installation is fairly long, and not trivial. This is partly due to the fact that the software is much more than just Internet filtering software. De-installation is straightforward.

Ease of Use

The product is very easy to use, once installed and configured. The user simply accesses the Web as they normally would. Each user has a specific profile, and needs to enter their password to access the Internet.

Users can be granted unrestricted access, and simply have to enter the appropriate password to be granted this access.

Configurability

Norton Internet Security is quite complicated for a non-technical person to configure. This is due partly to the fact that it is designed to be both broad *and* flexible. In consequence, Norton Internet Security allows you a lot of control over the actual filtering policy in effect, but also requires effort on the user's part to understand all the myriad configuration options being offered.

Norton is so configurable that we will not go into all the possible configuration details here. As an example of the configurability of Norton Internet Security, consider that it is possible for individual users to have different access policies, and it is possible to simultaneously customise access policies for different programs (for example, you can configure your Web-browser with a different access policy to your email client). In addition, you are free to customise access via all the well-known Web protocols (IM, ICQ, HTTP, HTTPS, NNTP, FTP, etc).

The Web filtering that Norton provides is implemented by means of 'black' lists of banned sites. Sites are grouped into separate categories, and the user can select to block or allow different categories. The user can also enter individual sites that they wish to block or allow.

How updates are handled

Updates can be requested by clicking a button. This obtains, over the Internet, the latest list of banned sites. Updates can also be obtained automatically by configuring the product appropriately.

Performance and Stability

There was no noticeable effect on performance or stability.

Documentation and Support

Symantec's Web site has a good database of user documentation, and answers to frequently asked questions. The product also ships with built-in help pages, and a thorough manual. The manual is quite technical in nature (discussing the basics of Internet plumbing, such as TCP/IP). Telephone support is available (for an additional fee per call).

Response to filtering violation

Attempts to access objectionable material are blocked. Limited logging of access violations (number of violations, most common violation type, etc), is also possible.

Notes

Norton Internet Security is designed as a complete solution to both Internet filtering and Internet security/privacy, and, as a result, can control network access to the computer on which it is installed. When we installed the software, we found that access to our LAN was cut off until we configured the product appropriately to allow local network traffic.

Price

$A 117.30 (Australian retail price). Note that this product offers a complete Internet security package and does much more than filtering.

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	9
Easy to uninstall	10
Ease of Configuration	
Simple configuration	7
Multiple Users	10
Flexible	10
Breadth	10

Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	9
Side effects	8
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	10
Configuration instructions	10
Troubleshooting and support	10
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	Yes
Blocking on text	No
Context sensitive	No
Multiple users & access?	Yes
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	No
Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	Yes

Filtering Effectiveness

We tested the product on its default settings. By default, the following categories were blocked: *Adult Humour; Alcohol-Tobacco, Anonymous Proxies, Crime, Drugs/Advocacy, Drugs/Non-medical, Gambling, Intolerance, Occult/New Age, Sex/Acts, Sex/Attire, Sex/Nudity, Sex/Personals, Sex Education/Advanced, Sex Education/Basic, Sex Education/Sexuality, Violence, and Weapons.*

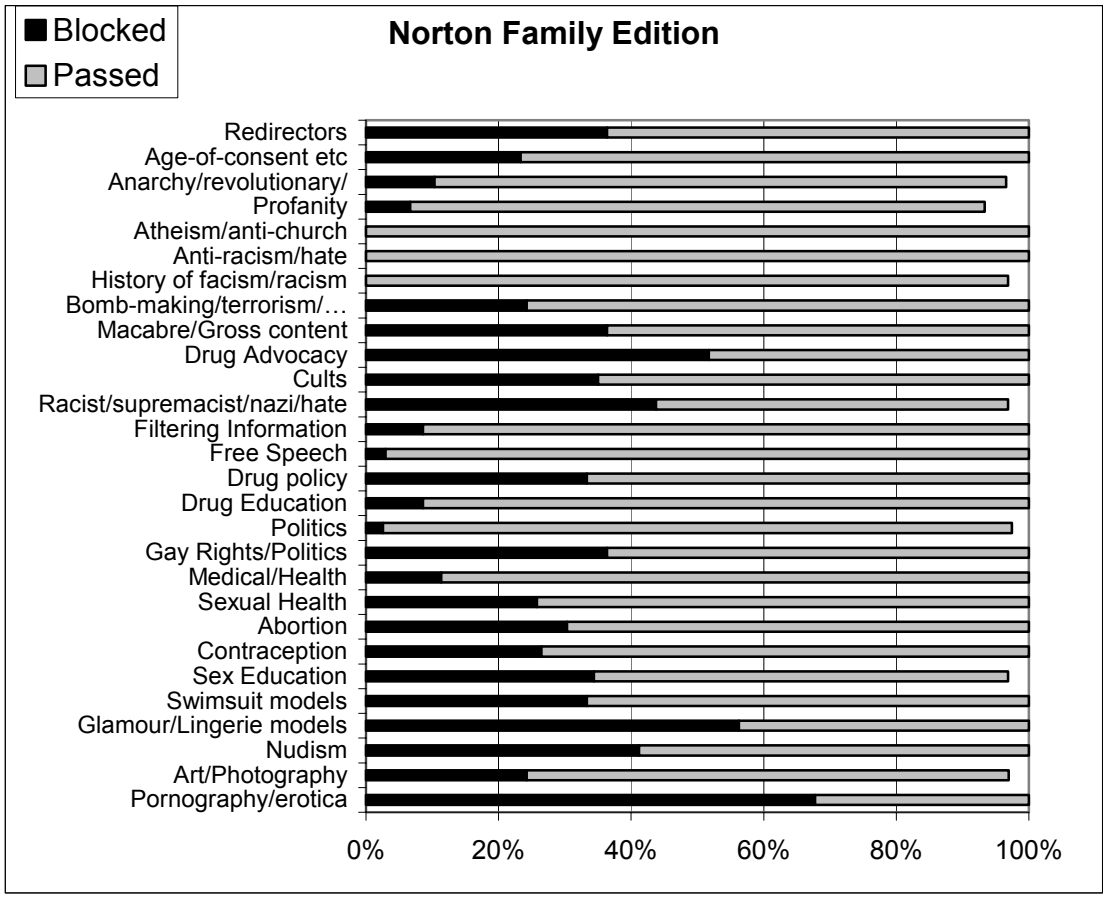

Norton Family Edition blocked access to some of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector, indicating that the returned page was blocked or the outgoing URL was analysed.

5.12 Smart Filter 3.0

Secure Computing
http://www.securecomputing.com

Product Type

A server-side product based on black lists.

Product Description

Smart Filter can be installed by an ISP to provide its subscribers with filtered access to the Internet. It can also be installed by a business on its Internet gateway to provide users within the company with filtered access to the Internet.

System Requirements

SmartFilter supports Microsoft, Solaris, and Linux platforms. Depending on the platform, various other proxying software must also be installed. For our testing purposes, we tested the product on a computer running Microsoft Server 2000, using Microsoft ISA server as a proxy server.

Ease of Installation and De-Installation

If Microsoft ISA server is already configured and running on the computer in question, setting up SmartFilter is easy. There were, however, a number of ISA Server/SmartFilter interoperability issues that can cause problems if either is not configured properly. SmartFilter did not work "out-of-the-box" with a default install of Microsoft ISA Server, and we needed to go through the SmartFilter manual and release notes carefully to get the two programs working together. Since the product is a server side product, it is reasonable to assume that skilled technical staff will have little trouble installing and integrating SmartFilter.

We cannot comment on other software/hardware configurations.

Ease of Use

The product is very easy to use, once installed and configured. Client computers, once configured to access the Web-proxy on which Smart-Filter is installed, do not see any change to the way they access the Internet.

Configurability

Configuration of SmartFilter is somewhat complicated. The way in which users, groups, and access policies are administered (and integrated) is quite complicated.

Administration of filtering policies and lists is relatively straightforward.

Access to Web sites is restricted based on "black" lists of Web sites. Each list contains sites with a common theme. There are 30 content categories.

How updates are handled

Updates are handled automatically.

Performance and Stability

There was no noticeable effect on performance or stability. Of course, proxy-based filtering necessarily places extra demands on the Web proxy that is performing the filtering, but our limited testing was not intended to test for such proxy bottlenecks.

Documentation and Support

The product comes with complete documentation, and built-in help pages. Support is available on the Web, via email, or via phone. We received prompt responses from support requests made via email.

Response to filtering violation

Attempts to access objectionable material are blocked. A complete activity logging system is built into SmartFilter.

Price

Negotiable.

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	6
Easy to uninstall	10
Ease of Configuration	
Simple configuration	6
Multiple Users	10
Flexible	9
Breadth	8
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	9
Configuration instructions	10
Troubleshooting and support	10

Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	Yes[32]
Blocking on text	No
Context sensitive	No
Multiple users & access?	Yes
Add to blocking list?	Yes
Allow sites through?	Yes
Tracking all access?	Yes
Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	No

Filtering Effectiveness

In our testing, we configured SmartFilter to block the following categories: *Anonymiser/Translator, Criminal Skills, Drugs, Extreme, Hate Speech, Sex.* We did not block access to categories of interest to business, such as sport or stock trading.

SmartFilter blocked access to most of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector but passed through a request for an innocuous site, indicating that the returned page was blocked or the outgoing URL was analysed.

[32] SmartFilter can be configured to block searches by disallowing particular search terms.

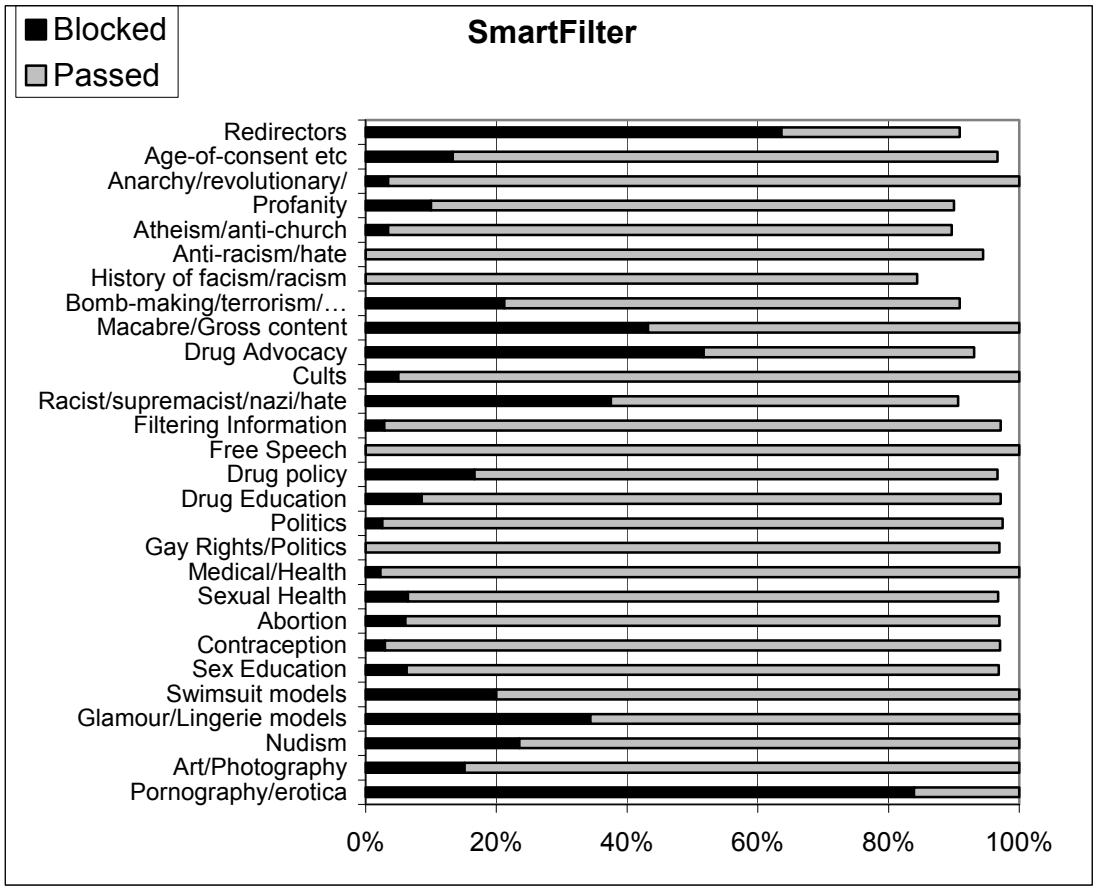
Blocked
Passed
SmartFilter
Redirectors
Age-of-consent etc
Anarchy/revolutionary/
Profanity
Atheism/anti-church
Anti-racism/hate
History of facism/racism
Bomb-making/terrorism/…
Macabre/Gross content
Drug Advocacy
Cults
Racist/supremacist/nazi/hate
Filtering Information
Free Speech
Drug policy
Drug Education
Politics
Gay Rights/Politics
Medical/Health
Sexual Health
Abortion
Contraception
Sex Education
Swimsuit models
Glamour/Lingerie models
Nudism
Art/Photography
Pornography/erotica
0%
20%
40%
60%
80%
100%

5.13 Too C.O.O.L.

Jordan/Chans Co., LLC (a subsidiary of Software 2010 LLC)
http://www.just2cool.com/

Product Type

An end user service, based on inclusion filtering ("white list") and a monitored chat room.

Product Description

Too C.O.O.L. has a proprietary browser designed for children to surf the net without possibility of being exposed to undesirable Web sites, or to undesirable people in chat rooms. In addition, it has proprietary content such as comic books, Horoscope, Joke of the Day etc. An e-mail service will be available in the future. It is designed for children aged between 7 and 13.

A monitored chat room is provided, where the user can navigate through a graphical environment, as well as communicating with other children.

System Requirements

Windows 95 or later.

Ease of Installation and De-Installation

Installation and de-installation are both trivial.

Ease of Use

The product is quite easy to use, but does require that the user learn to use the custom browser and the associated chat software.

Configurability

There is nothing for the user to configure, because the list of allowable sites is automatically maintained.

How updates are handled

Updates are handled automatically at the TooCOOL server.

Performance and Stability

Performance testing was not performed, as the TooCOOL custom browser hampered our attempts to automatically time HTTP request times. The software had no effect on stability.

Documentation and Support

Technical support is available via email at tech@jordanchans.com, or via a US phone number.

Response to filtering violation

Attempts to access objectionable material are blocked.

Notes

The TooCOOL CD that we were supplied with had the unacceptable restriction that it had to be in the CD drive at all times when the TooCOOL browser was being used.

Price

No price given by vendor.

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	10
Ease of Configuration	
Simple configuration	NA
Multiple Users	NA
Flexible	NA
Breadth	NA
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	0
Intrusive change	0
List updates	
Ease of updating	10
Automatic updates?	10
Documentation	
Installation instructions	10
Configuration instructions	NA
Troubleshooting and support	7
Claims/Capabilities	
Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	NA
Blocking on text	No
Context sensitive	No
Multiple users & access?	No

Add to blocking list?	No
Allow sites through?	No
Tracking all access?	No
Blocking chat rooms?	Yes
Blocking newsgroups?	Yes
Blocking email?	No

Filtering Effectiveness

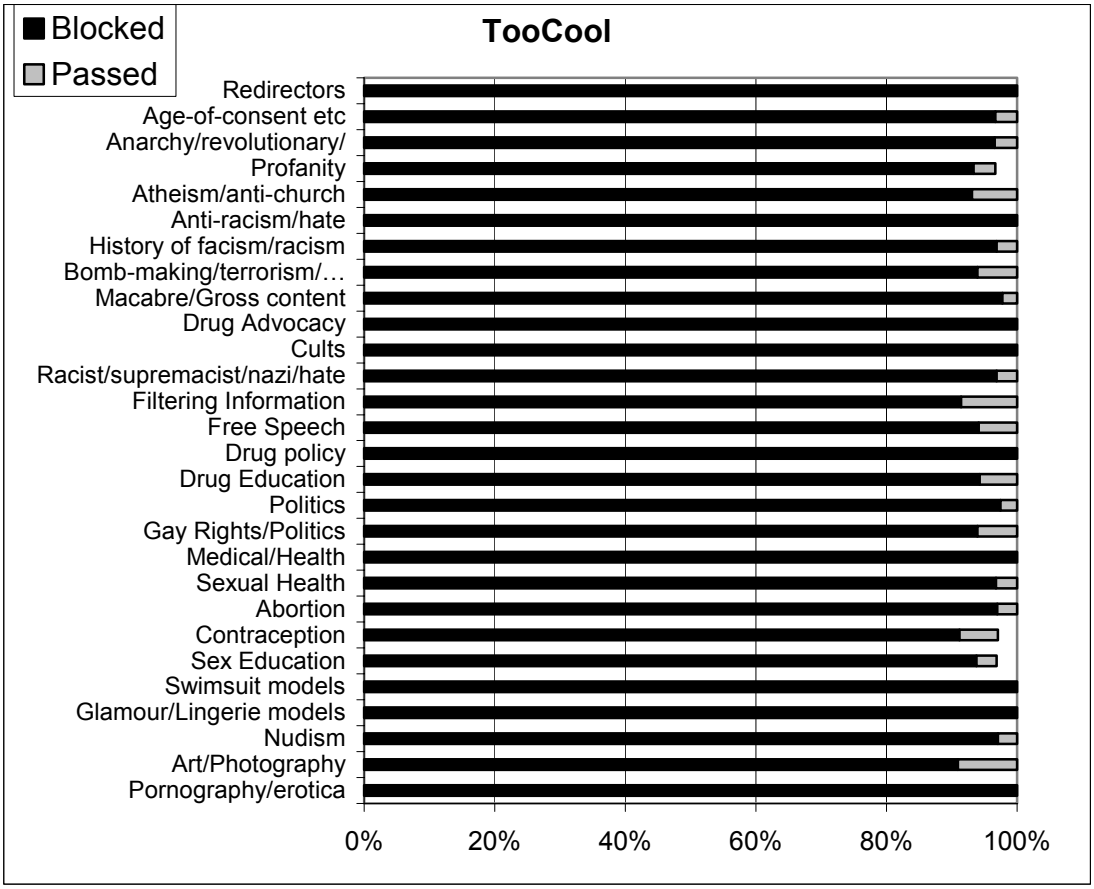

TooCOOL blocked access to all of the redirectors and sites that support anonymous Web access.

5.14 X-Stop 3.04

8e6 Technologies
http://www.xstop.com/

Product Type

X-Stop is an end-user filtering product. 8e6 Technologies (the makers of X-Stop) also make server-side filtering products, but these are not evaluated here.

Product Description

X-Stop is a filtering product with an emphasis on pornography filtering. Once installed, it runs in the background and monitors Internet access.

X-Stop has two filtering mechanisms. Firstly, it has list of banned Web sites that are automatically blocked. Secondly, it censors offensive words entered into the computer, preventing the user from entering them into Web searches or emails. This has the side effect that the user also cannot type any such words into any document or other application.

System Requirements

Windows 95 or later.

Ease of Installation and De-Installation

Installation is easy. De-installation is conceptually straightforward, but we did encounter problems when de-installing the product – the de-installation process froze before finishing. This may be due to the operating system our test machine is running (Windows Millennium Edition) -- 8e6 technologies currently does not recommend the product for computers running Windows Me, or Windows 2000.

Ease of Use

The product is very easy to use. Once installed, users need not be aware that it is running at all.

Configurability

Apart from having the ability to add/remove words from an offensive word list, and add/remove URL's from the banned Web site list, there is nothing that the user can configure.

X-Stop filters Web traffic only.

How updates are handled

The user can request an update of the banned site list by selecting a menu item.

Performance and Stability

No noticeable degradation in performance was observed. Apart from the problems associated with de-installation (see ***Ease of Installation and De-installation***), stability was not affected.

Documentation and Support

The 8e6 support Web site (www.8e6technologies.com/techsupport/prod_client.html) provides information. The product itself comes with no online documentation, but is simple enough that little documentation is needed. Email support is available at support@8e6technologies.com. Phone support is available via a US telephone number.

Response to filtering violation

Attempts to access objectionable material are blocked. No other action is taken.

Notes

The *Peacefire (www.peacefire.com)* disabling software was able to disable X-Stop.

Price

$US 60/year for client-side product (from vendor Web site).

Usability Assessment

Category	Score or Result
Ease of Installation	
Easy to install	10
Easy to uninstall	7
Ease of Configuration	
Simple configuration	10
Multiple Users	0
Flexible	3
Breadth	4
Detrimental to system performance?	No
Detrimental to system stability?	No
Ease of Use	
Easy to use	10
Side effects	10
Impact on other users	10
Intrusive change	10
List updates	
Ease of updating	10
Automatic updates?	0
Documentation	
Installation instructions	10
Configuration instructions	8
Troubleshooting and support	8
Claims/Capabilities	

Basic URL blocking	Yes
IP blocking	Yes
Blocking searches	Yes
Blocking on text	No
Context sensitive	No
Multiple users & access?	No
Add to blocking list?	Yes
Allow sites through?	No
Tracking all access?	No
Blocking chat rooms?	No
Blocking newsgroups?	No
Blocking email?	No

Filtering Effectiveness

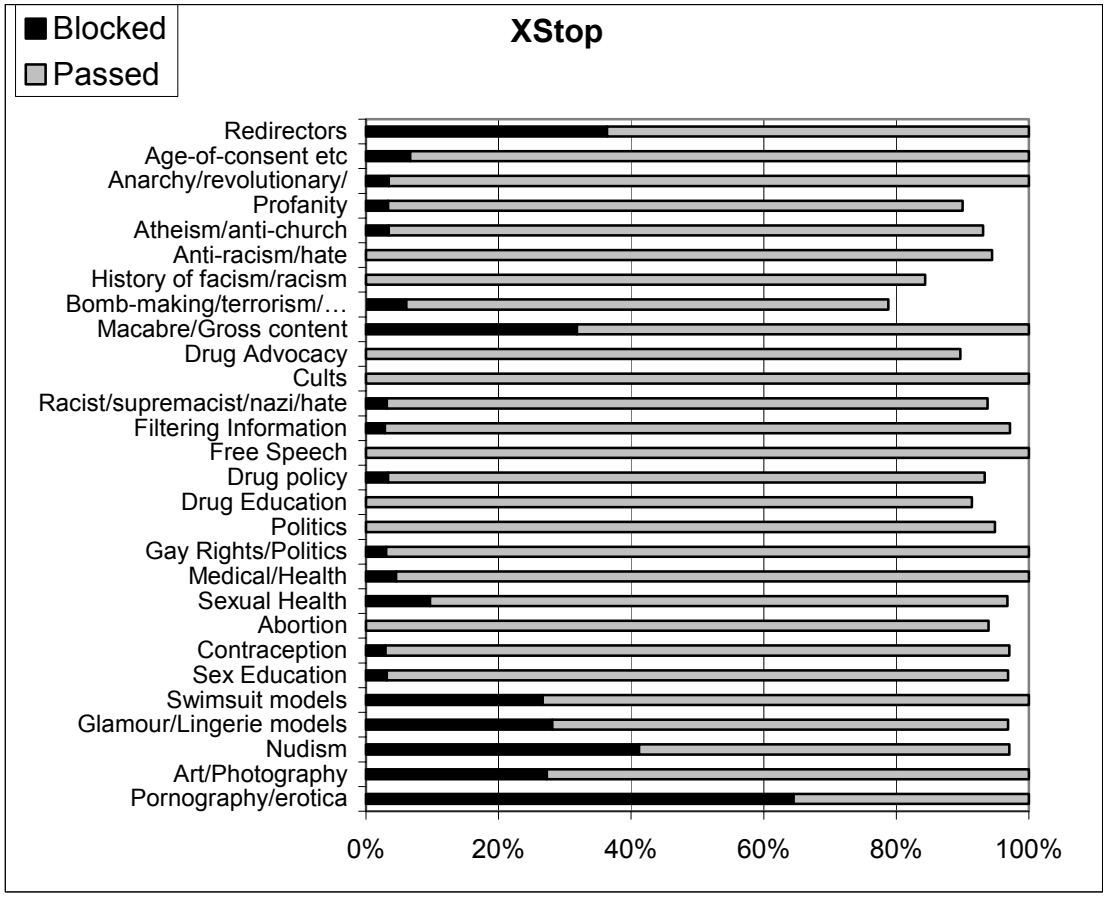

XStop blocked access to some of the redirectors and sites that support anonymous Web access. It blocked access to a pornographic site accessed through a redirector but

passed through a request for an innocuous site, indicating that the returned page was blocked or the outgoing URL was analysed.