



Internet Filtering Software Tests:

Barracuda, CyberPatrol, FilterGate, & WebSense

Sarah Houghton-Jan, Digital Futures Senior Librarian

Original report submitted February 4, 2008

Revised report submitted April 2, 2008

Executive Summary and Background Information

The San José Public Library was asked by the City Council to test various Internet filtering service options for implementation in the Library's public use computers, with a focus on filtering "web sites that contain child pornography or material that is obscene." Councilmember Pete Constant proposed, in his memorandum to the City council Rules Committee dated October 18, 2007, *Attachment G "Proposed City Internet Access Policy,"* that all computers with Internet access use filtering technology. Specifically, the proposed policy states:

"The Library uses filtering technology on all computers with Internet access. Patrons 17 years of age or older are given a choice of an Internet session with a basic filter or one that has additional filtering. The intent of the basic filter is to block web sites that contain child pornography or material that is obscene. The intent of the additional filtering is to block web sites that contain material that is harmful for minors."¹

San Jose Public Library staff explored the Internet filtering market by reading the extensive research and white papers on the topic conducted in the last decade, as well as speaking with nearly three dozen different companies that offer an Internet filtering product, in order to gain an understanding of their product's strengths from their sales and technical staff. We attempted to find a service that only blocks images, specifically, as defined in the proposed policy, images that are obscene and harmful to minors. We were able to identify products that would allow us to choose to functionally block all images of all types on all web sites. We were also able to identify products that allowed for general filtering by keyword and web site address (URL) in many categories, including categories with varying references to adult content, sexual content, etc. We were not able, however, to find any product on the market that successfully allows filtering only of images that are classified as obscene and harmful to minors. Filtering expert Lori Ayre's research holds up our findings of what the Internet filtering market currently offers:

"No filter, however, actually limits its categories to obscene material and child pornography because the current definition of obscenity doesn't work on the Internet." (Ayre, "Filtering and Filter Software," p. 52)

Our research of the market showed that the offerings of today's filtering market is not much different than in 2004, the year of Ayre's report. There are no existent filters that will filter out only obscene and harmful images. Given that we could not fulfill that aspect of the original proposal because the technology simply doesn't exist to do so, we originally tested three filters, and subsequently one additional filter upon Councilmember Constant's request, with various features, granularity, and functionality in an attempt to determine whether, as has been asserted, content filtering technology has improved over the last decade to the extent that over-blocking is minimal

¹ According to California Penal Code Section 311, "obscene matter" is "matter, taken as a whole, that to the average person, applying contemporary statewide standards, appeals to the prurient interest, that, taken as a whole, depicts or describes sexual conduct in a patently offensive way, and that, taken as a whole, lacks serious literary, artistic, political, or scientific value." California Penal Code Section 313 defines "harmful matter" as "matter, taken as a whole, which to the average person, applying contemporary statewide standards, appeals to the prurient interest, and is matter which, taken as a whole, depicts or describes in a patently offensive way sexual conduct and which, taken as a whole, lacks serious literary, artistic, political, or scientific value for minors."

and has little effect on patron research. A second goal of the library research was to learn about the current state of content filtering software's ability to block materials that are harmful to minors.

How Filters Work

Content filters today are powerful and full of features. Filters today have artificial content recognition that help to evaluate content on a more granular level – a single image, a single search result, a single web page. However, filters still lack the ability to successfully evaluate and determine the actual content and context of web pages, including text, still images, video, and more. As a result, filter performance is highly dependent on the programs' artificial content recognition, administrative human intervention, chosen settings, and features.

Network-Based and Stand-Alone Options

There are two major categories of filtering products: network-based and stand-alone. Network-based filters are installed on one central server and individual computers' settings are controlled by the settings on the server. Stand-alone filters are installed on each computer individually and the settings only control that computer. Both categories of products have individual filters that are more or less powerful or complex than others and both have their merits, which is why we tested two network-based filters (WebSense and Barracuda) and two stand-alone products (CyberPatrol and FilterGate).

Filtering by URL or Keyword

Most software now on the market works by filtering based on URLs (web site address) and/or filtering based on content (trigger words, phrases, etc).

- Products that filter based on URLs typically use a search engine (Google in most cases) and run searches for trigger words, like "live sex chat rooms." The list of results from that search is then pared down by removing educational and government sites (done only by removing sites with .edu and .gov suffixes, missing many educational and government sites that choose to be a .net or .org, for example). The remaining sites, generally the top 100 - 500, are then blacklisted on the "trigger URL" list. Some companies stop the process there, while others will have a staff member spot-check for errors, a process whose quality varies greatly from company to company. When the filtering program is in use on a computer, each Internet search result or direct entry of a web address is scanned against the list before results are displayed.
- Products that filter based on content analyze web pages as they are requested by the user, looking for trigger keywords and sometimes phrases as well as other factors such as banner ads, number of links and images, etc. An artificial intelligence software program then looks for a substantive formula of the various criteria and classifies the web page as allowed or blocked.

Blocking (What the User Sees)

Using one or both of these methods, companies build up lists of trigger URLs and/or keywords that they deem should be filtered. When content is blocked, users see a "blocked" message that states, in varying degrees of detail depending on the flexibility of the product, what was blocked, why, and how/if it can be unblocked. Some filters allow for a "warning and bypass" message on the screen, either requiring a simple click-through or a password to get to the content that was blocked.

When access to a filtered page or resource is attempted, some systems will filter out only the triggering content (e.g. only blocking those images on the results page that are triggers) but still allowing the non-triggering content on the page, while other systems will filter out/block the entire page, hiding everything on that page from view, not just the triggering content. Other systems allow you to see references to trigger content on search results pages, but will not let you click on the result to get to the actual page/resource.

Blocking by File Type

A small number of filters allow one to block specific file types – such as video file types (.avi), audio (.mp3), or still images (.jpg). Unfortunately, as previously noted, these programs do not allow you to successfully designate the blocking of those file only for images that are classified as obscene and harmful to minors. It is also impossible to create an exhaustive catalog of all file extensions for a particular file type and expect to block that file type successfully. For example, adult web sites frequently embed their images in another file type (like Flash or even PDF), getting around the blocking of the filters. As a result, if the library wanted to try to block **only** images that are obscene and harmful, it would have to block **all** images due to the limitations of the existing technology.

Some filtering systems block only that one URL (specific web page) when trigger content is found, while others are more broad in their blocking and will block an entire domain (the entire web site: for example, Craigslist or eBay) based on one user or one page with trigger content. Still others are even broader and block anything hosted on that Internet Protocol (IP) address (numerous domain names share a single IP address; for servers that host multiple sites, blocking by IP can result in gross over-blocking).

Classification of URLs and Keywords

One of the challenges to successful filtering in libraries is how web pages are classified in the filtering system – that content is evaluated for the user by automated systems and sometimes IT or clerical subcontractors, not by trained information professionals like librarians. Lori Bowen Ayre sums it up accurately when she writes:

“Ironically, librarians - professionals trained to catalog and evaluate content - subcontract their cataloging job to Internet filter companies when they install a filter. Unlike librarians, the subcontractors are not information professionals, they typically use automated methods to classify the 3 billion web pages on the Internet.” (Ayre, *Internet Filtering Options Analysis: An Interim Report*)

Automated methods result in faster classification, thereby raising the number of “cataloged” sites and the product’s perceived value for the company, but also results in less accurate classification, specifically in more resources being falsely blocked.

Filtering software companies do not tell their customers, in detail, the types of things or what specific sites they block in each category. No examples are given and no information beyond a one or two sentence description is offered. Because companies ferociously protect their list of categorized sites and their process for categorizing, there is no way of obtaining a list of sites that are blocked in certain categories, as that is considered a trade secret and vital to their continued business interests. The subscribers are asked to make global decisions that will affect users’ ability to access content based on these brief descriptions. There is no way to know exactly what sites, or types of sites, are included in the “Illegal or Questionable” or “Tasteless” categories, for example.

All studies of Internet filters show over-blocking and under-blocking. No product is perfect. Lori Bowen Ayre writes:

“All filters overblock. All filters underblock. No filter is 100% accurate because no one agrees on what being 100% accurate is.” (Ayre, “Filtering and Filter Software,” p. 36)

Ayre writes of the desire on libraries’ parts for filters to create more specific “child pornography” categories, something not offered by filtering companies now:

“[F]iltering companies are free to devise filters based on language that works for their target audience – parents, employers and schools. Therefore, you’ll never see a category of web sites defined as “harmful matters” or “child pornography.” Some take the plunge and define web sites as “obscene” but how closely those web sites match the legal definition is anyone’s guess. And since none of the companies release the list of web sites on their radar and the category into which they’ve been placed, the end user has no way of knowing whether the “obscene” sites include some Constitutionally protected sites or not.” (Ayre, *Internet Filtering Options Analysis: An Interim Report*)

Most filters allow for the library or the vendor to apply additional whitelists (sites to always allow) and blacklists (sites to always block) in addition to the vendor’s database of URLs and/or keywords. Some vendors require that any addition to either list be approved by them, while others will allow the local library to apply the change directly. Over time, with the addition of whitelists and blacklists as the library staff and users come across sites that have been categorized incorrectly or not categorized at all, the library is able to build a more effective filter for local needs. This site-by-site method, however, is time consuming and can never cover the ever-growing number of sites on the web.

Until more advanced classification and categorization methods are developed, either through Artificial Intelligence (AI) or human intervention, filters will find difficulty in maintaining accurate categorization without over- or under-blocking, and the market will continue to yearn for effective and accurate “harmful matters” or “child pornography” categories.

Test Description

In our original test, four workstations of various configurations were set up by the library, with the involvement of the City Information Technology Department. As part of our planning for the test, library staff met with Vijay Sammeta (Deputy Director of San José Information Technology Department) on January 14th to review our testing process and set-up. One workstation was set up without any filtering installed and three different filtering programs were also tested: CyberPatrol, FilterGate, and WebSense. Upon the subsequent request two months later by Councilmember Constant, the library, once again with the involvement of Vijay Sammeta, set up a duplicate network and workstations to mimic our original tests and tested one additional filtering program: Barracuda.

Each program offers different options for content filtering, without a one-to-one correlation of settings between programs. However, every effort was made to set up consistent filtering levels on each machine to filter only content of an adult sexual nature. Professional best practices, per the

two paramount filtering reports by the Kaiser Family Foundation and Lori Bowen Ayre, recommend that the filters be set to their lowest setting; in other words, being very specific about the categories one wishes to filter and not choosing every category by default and/or choosing lower levels of intensity within the filtering software.

CyberPatrol was set up to filter *Adult/Sexually Explicit* and *Glamour & Intimate Apparel* content, as well as *Remote Proxies* (well-documented sources for adult content sites). FilterGate's *AdultFilter* option was enabled. WebSense was set up to filter *Adult Material* (including *Adult Content, Lingerie & Swimsuits, Nudity, and Sex*), *Illegal or Questionable sites* (redirect sources for adult content sites), *Information Technology* (including *Proxy Avoidance* and *URL Translation Sites*, also sources for adult content sites). Barracuda was set up to filter the *Sexual* category (including *Adult, Intimate Apparel & Swimsuit, and Porn*) as well as one category of the *Communication & Technology* category (*Proxies*).

While the programs tested do offer the option of whitelists and blacklists, that was not an option we were able to employ during our tests as the content of those lists is built up over time by the local staff to meet the local needs and requirements of the community. Libraries who have had filters installed for a long time can sometimes have substantial whitelists and blacklists that are an overlay on the filter's own database of blocked and/or allowed sites. If the library were to implement filtering, we would anticipate the build-up of these types of list over time.

A set of 135 test questions and scenarios were written based on the existing literature about filtering and staff suggestions of real information requests they have received from their users. The questions/scenarios were broken into the following categories:

- general keyword searches (for both “content of an adult sexual nature” and “content *not* of an adult sexual nature”) in three different web search engines
- direct URL access to a variety of types of sites and content
- image searches (“content of an adult sexual nature” and “content *not* of an adult sexual nature”) in three different image search engines
- email text and photo attachments through several different webmail providers
- RSS feed content access
- searches in the online library catalog, and searches in our proprietary subscription databases

The test questions/scenarios do not represent a scientific random sampling of all information requests or searches. A conscious effort was made to include searches and scenarios that the filters should be able to handle fairly easily as well as attempts to find information that might be incorrectly blocked or attempts to find and view materials that are harmful to minors. No attempt was made to find or view materials, such as child pornography, that are illegal.

For the original tests, four teams of two senior librarians each, with representation from San José Public Library and the San José State University Library, were designated to test the 135 questions and scenarios on each of the three original filters, with an unfiltered computer as a control. For the subsequent Barracuda test, the Digital Futures Senior Librarian conducted the testing with City Information Technology representative, Vijay Sammeta, present for some of the testing. Data was recorded and submitted to the Digital Futures Senior Librarian for central review and processing.

General Findings

Below is the average accuracy percentage in each content category for all four filters combined to show a general sense of how effective these filters were in the various categories. The accuracy rate represents the success of the filter in blocking the content it should block and/or letting through the content it should let through. The perfect score for each category would be 100%.

The success in filtering out content is higher, particularly in keyword searches, than the ability to correctly allow content through that should not be filtered. In other words, the trend is toward over-blocking. The accuracy rates for correctly filtering the non-text and non-standard-text content (images, email attachment images, and RSS feeds) is lower. The accuracy rates for the library's proprietary catalog and databases are on par with the accuracy rates for keyword searching and direct URL access.

Average Filter Accuracy (margin of error +/- 5%)

Type of Content Tested	Accuracy Percentage
Content of an Adult Sexual Nature – direct URL access	87%
Content of an Adult Sexual Nature – keyword searches	81%
Content not of an Adult Sexual Nature – direct URL access	86%
Content not of an Adult Sexual Nature – keyword searches	69%
Image Searches	44%
Email Attachments	25%
RSS Feeds	48%
Library Catalog Searches	75%
Library Database Searches	88%

Reading through the results of all of the major published Internet filtering studies conducted from 2001-2008 (listed at the end of this report), which predominantly tested traditional text-based content such as direct URL access and keyword searching, one will note that our findings are extremely similar to the other studies' findings. In fact, the average accuracy rating of all of the various studies cited is 78.56%. The comparable sections of our informal study (keyword searching, direct URL access, RSS feeds, catalog and database searches) yielded very similar results: an average accuracy of 76.29%, a difference of only 2.27%.

We did, however, experience a much lower success rates for non-traditional and rapidly growing web content in various formats, including images. Only one published study directly addresses the success of image searching, the *Expert Report* by Dr. Paul Resnick for North Central Regional Library District. He found a 48% rate of accuracy in blocking trigger images (images the filter is meant to catch). We tested both images that the filter should catch as well as images that the filter should let through, in both image search engine keyword searching and image email attachments. Our results for image search engine keyword searching, which is the section most comparable to Dr. Resnick's study, yielded an average accuracy of 44%—nearly identical to Dr. Resnick's findings. If you include image email attachments (something Dr. Resnick did not test), our study's findings go down to an average accuracy rating of 34.5%, still not that far off from Dr. Resnick's findings.

In all four filters tested, image filtering had a low rate of accuracy. Many images of an adult sexual nature were displayed on web pages accessed by the testers, and additionally the image search results pages and most of those images' full-size versions and/or parent sites could be accessed as well.

Because of the ability of image search engines (like Google Images and Yahoo Image Search) to display thumbnails which often aren't treated as "real" images by the filtering programs, image filtering is a problem for the filtering software's AI. Images of an adult sexual nature from image search engines, pages with images of an adult sexual nature but "fake" innocent text, or images of an adult sexual nature posted to social sites like Craigslist were consistently displayed in all four filter tests. Additionally, clicking on the search engine results pages' links to "cached" versions of webpages allowed access to those webpages and their images, even though their main entries on the results page were blocked. There were many work-arounds discovered by our testers that allowed access to the very material that the filtering systems were attempting to block. At the same time, many sites without images of an adult sexual nature, or even entire search results pages, were blocked, such as the medical site WebMD or search results pages for a search for "Parents and Friends of Lesbians and Gays."

For two of the four filters tested, over-blocking of text content was a serious problem. Based on our test results, it is apparent that the artificial content recognition in all four filters is heavily reliant on URL and single-word black lists, and not so much on phrases or overall contextual content of a site. As a result, much over-blocking occurs. Numerous searches for content that is not of an adult sexual nature were blocked (e.g. the search results pages were entirely blocked, or various credible results blocked). Direct URL access to sites without content of an adult sexual nature were blocked incorrectly as well, such as VictimsOfPornography.org (a support group for victims of pornography) and Lesbian.org (a lesbian support site).

The same was found, though to a lesser extent, in a small study conducted by the Kaiser Family Foundation: "See No Evil: How Internet Filters Affect the Search for Online Health Information."

"At the least restrictive or intermediate configurations, the filters tested do not block a substantial proportion of general health information sites (1.4%); however, at the most restrictive configuration, one in four health sites are blocked....Even at their least restrictive settings, filters could have a modest impact on those seeking information on sexual health issues; on average, filters incorrectly blocked about one in ten sites on safe sex, condoms, or health issues pertaining to gays." (Kaiser Family Foundation, *See No Evil*)

Blocking of terms of an adult sexual nature across filters and search engines was highly inconsistent. Only one out of the fifteen terms of an adult sexual nature that the testers searched on was blocked in all three search engines in all four filters. The keyword searches that are blocked vary from search engine to search engine, showing inconsistency in the methods by which content is blocked. The more popular sites/engines filtered more out, demonstrating that certain tools may have received more attention from the filtering software developers. In other words, depending on which search tool you happen to use, you will get more or less access to content that the filter is trying to block.

Workarounds to "fool" the filter were also easily successful in every test filter. For example, you could get around the filter's parameters by searching for "pron" instead of "porn," using plural word forms, searching for acronyms instead of the actual institution's name, or getting out to an adult site through a seemingly innocent "portal" site (like Linkbase.org) to get around the filters, clicking on the thumbnail images or "cached" versions of webpages, or using a site like Peacefire.org whose sole purpose is to provide users with a one-click workaround for filtering systems.

The filtering programs' artificial content recognition does not handle non-English language words well, completely allowing Spanish-language terms, including slang, searches and their results, while blocking the English translation of the same term. This is a problem for two chief reasons. First, in our multicultural community many languages are spoken and searches are conducted in numerous languages. Second, with dominantly-English language search engines indexing more and more non-English content, results with Spanish language trigger words would not be caught, thereby allowing more sites with content of an adult sexual nature to be incorrectly displayed.

None of the four filtering programs successfully filtered out emails with content of an adult sexual nature. RSS feeds, however, were blocked appropriately in only one of the four filters.

Filter-Specific Findings

CyberPatrol

CyberPatrol allows for a rather granular level of filtering, but the restrictiveness and lack of description for the settings would make precise and effective configuration difficult. Through all of the various searches and scenarios CyberPatrol allowed fewer images of an adult sexual nature, but also over-blocked quite a bit (compare the first row of accuracy statistics below - the accuracy for "content not of an adult sexual nature" is lower in both categories).

In all image search engines, image filtering was unsuccessful. Many images of an adult sexual nature got past the filters and many images that did not include adult sexual content, and even entire searches, were blocked. Additionally, for most image thumbnails (even those that were deemed "adult" and blocked by the filtering software), if you clicked on the originating site or the blank thumbnail image you could still get through to see the full size image on its original web page. Questionable sites, like a Craigslist posting with innocuous text but a graphic adult photograph, are allowed. Keyword searching results in general inconsistencies in what is and isn't blocked (e.g. "women's asses" is allowed but "Shakespeare and sex" isn't).

Keyword searching within the library's proprietary resources also met with some challenges; for example:

- a search for "orgasm" in the Health and Wellness Resource Center database was blocked
- a search for "vagina" in the World Book Encyclopedia online was blocked

Numerous sites that do not contain content of an adult sexual nature are being blocked as well, both through keyword searching and direct URL access, including:

- WebMD
- the American Urological Association site
- VictimsOfPornography.org
- Univision.com
- DirtyPicturesBand.com (a rock band site with no adult content)
- Amazon and Google Book Search item pages (including the Amazon item page for an album by the band The Cure entitled "Pornography")

Entire domains also appear to be blocked if even one post on one sub-domain contains something of an adult sexual nature (e.g. the entire site, SlideShare, which is a PowerPoint slideshow sharing site, was blocked because of one slideshow discussing sexual positions).

CyberPatrol Accuracy (margin of error +/- 5%)

Type of Content Tested	Accuracy Percentage
Content of an Adult Sexual Nature – direct URL access	87%
Content of an Adult Sexual Nature – keyword searches	96%
Content not of an Adult Sexual Nature – direct URL access	73%
Content not of an Adult Sexual Nature – keyword searches	65%
Image Searches	44%
Email Attachments	25%
RSS Feeds	25%
Library Catalog Searches	75%
Library Database Searches	50%

FilterGate

Because FilterGate allows only for general blocking with their AdultFilter, and does not allow for specific subject-based filtering, many sites without any content of an adult sexual nature are blocked. This rough approach to filtering would not offer us the functionality requested. Most image searches were allowed, and the thumbnails of images, both content of an adult sexual nature and not, were displayed fully and not filtered appropriately.

If a “filtered-out” image of an adult sexual nature appears as a result on a page, the entire results page is blocked, blocking access to content without material of an adult sexual nature. Keyword searching results in general inconsistencies in what is and isn’t blocked (e.g. “big penises” is allowed but “Parents and Friends of Lesbians and Gays” isn’t). Blocking is inconsistent as well: “parents and lesbians” is blocked while “parents and gays” is allowed, “Parents and Friends of Lesbians and Gays” is blocked while “PFLAG” is allowed. Keyword searching within our proprietary resources also met with some challenges; for example, the following searches were *not* allowed in the library’s online catalog:

- lesbianism
- how to build a pipe bomb
- sexual positions

Numerous sites without any content of an adult sexual nature are being blocked as well, including:

- TheSmokingGun.com
- Lesbian.org (a gay/lesbian support site)
- the Wikipedia entry for *Hustler* Magazine
- a World War II history web site
- a UK breast cancer information site
- entire blogs are blocked because one of the many posts discussed something “adult”

FilterGate Accuracy (margin of error +/- 5%)

Type of Content Tested	Accuracy Percentage
Content of an Adult Sexual Nature – direct URL access	93%
Content of an Adult Sexual Nature – keyword searches	74%
Content not of an Adult Sexual Nature – direct URL access	82%
Content not of an Adult Sexual Nature – keyword searches	41%
Image Searches	36%
Email Attachments	25%
RSS Feeds	100%
Library Catalog Searches	25%
Library Database Searches	100%

WebSense

There is more under-blocking than over-blocking in WebSense. This is vastly different from Filtergate and CyberPatrol, which over-blocked, perhaps because of the more granular nature of the filtering categories in WebSense and the increasing dependence on keyword filtering instead of just URL filtering. All image searches were allowed in all search engines, with individual images being erased/blocked on the results page instead. Over-blocking occurred, as in the case of National Geographic images of beavers being blocked. Consistently, however, images of an adult sexual nature still got through the filters and were displayed for nearly every search in their thumbnail format and it was often possible to click on the thumbnail image, even if it was erased, and still get access to the originating web site and larger version of the image. Below are examples of some of the image searches that resulted in numerous instances of graphic content being displayed on the search results page directly and/or allowing click-through access to the original web site and image:

- anal sex pictures
- huge breasts
- rape photos
- Spanish term “cojones”
- Spanish term “putas”

All keyword searches were allowed, but individual results for some searches were blocked, sometimes inappropriately, such as some of the results for searches for:

- how to be a good lover
- gay sex
- *Hustler*
- vibrators

Keyword searching for text results in general inconsistencies in what is and isn't blocked. For example:

- Yahoo's directory of adult sex chat sites is not blocked
- some very graphic search results were viewable through a search for “violent sex site”
- some very graphic search results were viewable through a search for “porn videos”
- Some very graphic search results were viewable through a search for “animal sex photos”

Library catalog and database searches, in this case, were completely successful.

WebSense Accuracy (margin of error +/- 5%)

Type of Content Tested	Accuracy Percentage
Content of an Adult Sexual Nature – direct URL access	87%
Content of an Adult Sexual Nature – keyword searches	78%
Content not of an Adult Sexual Nature – direct URL access	100%
Content not of an Adult Sexual Nature – keyword searches	82%
Image Searches	33%
Email Attachments	25%
RSS Feeds	33%
Library Catalog Searches	100%
Library Database Searches	100%

Barracuda

There is more under-blocking than over-blocking in Barracuda, as in WebSense. All image searches were allowed in all search engines, with no individual images being erased or blocked. All images were displayed, period. The same occurred with image email attachments – everything was displayed. Over-blocking occurred, as in the case of PFLAG.org being blocked. As with the image searching in all other filters, clicking on the thumbnail format of images, or clicking on cached versions of web pages, allowed full access to content of an adult sexual nature.

Below are examples of some of the image searches that resulted in numerous instances of graphic content being displayed on the search results page directly and sometimes also allowing click-through access to the original web site and image(s):

- anal sex pictures
- rape photos
- normal erection
- Spanish term “cojones”
- Spanish term “putas”

All keyword searches were allowed, but individual results for some searches were blocked, sometimes inappropriately, such as some of the results for searches for:

- Breast enlargement surgery
- Parents and Friends of Lesbians and Gays
- *Hustler*
- vibrators

Keyword searching for text results in general inconsistencies in what is and isn't blocked. For example:

- Hustler.com was blocked but HustlerLingerie.com was allowed
- PFLAG.org, the national organization's webpage, was blocked but all of the state and international chapters' websites are accessible
- a page about building a potato gun on hubpages.com and a page about building a flying saucer on beyondweird.com were both blocked incorrectly

- Examples of sites that are allowed incorrectly: AnimalSex.es, PornXTube.net, WildWebCamGirls.com, XXXChatters.com, Adultcyberdating.org, Cruel-Rape.com, and BestExtremeVideos.com/Forced-Fuckers.html and FuckingDickHead.com
- some very graphic search results were viewable through a search for “sex chat rooms”
- some very graphic search results were viewable through a search for “huge breasts”

Numerous sites that do not contain content of an adult sexual nature are being blocked as well, both through keyword searching and direct URL access, including:

- ImplantInfo.com (a site with a wealth of medical information about breast implants)
- PFLAG.org
- A Gay.com article on queer sexuality and another on “Our Trans Children”
- A Nazi history article
- *Hustler’s* homepage
- Lesbian.org (a gay/lesbian support site)
- SexHelp.com

Entire domains also appear to be blocked if even one page on one sub-domain contains something of an adult sexual nature (e.g. the entire site, Squidoo, which is a site that allows users to create “lenses” which result in topical webpage with links to various resources, was completely blocked but it is unclear why.

Library catalog and database searches, in this case, were completely successful.

Barracuda Accuracy (margin of error +/- 5%)

Type of Content Tested	Accuracy Percentage
Content of an Adult Sexual Nature – direct URL access	78%
Content of an Adult Sexual Nature – keyword searches	74%
Content not of an Adult Sexual Nature – direct URL access	90%
Content not of an Adult Sexual Nature – keyword searches	87%
Image Searches	64%
Email Attachments	25%
RSS Feeds	33%
Library Catalog Searches	100%
Library Database Searches	100%

Conclusion

Despite the fact that our test was geared toward filtering out only content of an adult sexual nature, other text and image content that was not of an adult sexual nature was filtered out as a consequence. The filters we tested falsely blocked many valuable web pages and other online resources, on subjects ranging from war and genocide to safer sex and public health. No filter was reliably able to distinguish text or image content including obscenity, child pornography, or “harmful to minors” material from other, legal content. As a result, each filter blocked a wide range of constitutionally protected content in its attempt to block other content. Other, published studies cited in the References section have consistently shown that the more successful the filter is at

blocking the content it wishes to block, the more unsuccessful it is at letting constitutionally protected (i.e., neither illegal nor harmful to minors) content through. This was the case in our test as well.

Because the filtering programs are looking for particular trigger words and URLs, the filtering of images is highly problematic. The only existent way to filter images is based on the words surrounding them – either in the text around an image on the web page, image file names, or alternative text tags (text that is read out loud when a screen readers is used to access the web site, usually in the case of a blind user). There is no artificial content recognition that can evaluate the actual content and context of an image and determine whether or not it falls into a specific category, or contains a particular type of image.

As such, in order to even attempt to block adult images of an adult sexual nature, the library would have to choose to block whole categories of content (e.g. “Adult-Sexual”) including both text and images, and/or block all images on all websites entirely. The result would be that both images and text, not to mention access to entire web sites or web pages, would be blocked—not just images of an adult sexual nature. As our tests show, filtering technology is ill-equipped to deal with newer and non-text and non-standard-text content, such as image results on image search engine pages, image email attachments, RSS feeds, and non-English content.

Our results show that the effectiveness of content filtering either in blocking materials harmful to minors or in allowing access to information including images that is not harmful to minors has not changed significantly in recent years.

References

- Ayre, Lori Bowen. "Filtering and Filter Software." *Library Technology Reports*. American Library Association, March-April 2004.
- Ayre, Lori Bowen. "Infopeople Project How-To Guides: Filtering the Internet." *InFoPeople Project*. September 19, 2002. <http://infopeople.org/resources/filtering/index.html> (Accessed 04/02/08).
- Ayre, Lori Bowen. "Internet Filtering Options Analysis: An Interim Report." *InFoPeople Project*. May 2001. http://statelibrary.dcr.state.nc.us/hottopic/cipa/InternetFilter_Rev1.pdf (Accessed 04/02/08).
- Brunessaux, Sylvie et al. *Report for the European Commission: Review of Currently Available COTS Filtering Tools*. European Commission. 2001. <http://np1.net-protect.org/en/results3.htm> (Accessed 04/02/08).
- Consumer Reports. "Digital Chaperones for Kids." 2001. <http://web.archive.org/web/20010310234724/http://www.consumerreports.org/Special/ConsumerInterest/Reports/0103fil0.html> (Accessed 04/02/08).
- Consumer Reports. "Filtering Software: Better But Still Fallible." June 2005. <http://www.consumerreports.org/cro/electronics-computers/resource-center/Internet-filtering-software-605/overview/index.htm> (Accessed 04/02/08).
- Edelmen, Ben. *Sites Blocked by Internet Filtering Programs: Expert Report for Multnomah County Public Library et al. vs. United States of America et al.* Cambridge, MA: Ben Edelman, 2002.
- eTesting Labs. *Corporate Content Filtering Performance and Effectiveness Testing Websense Enterprise v4.3*. WebSense. 2002. <http://web.archive.org/web/20030406232751/www.websense.com/whyqualitymatters/etestinglabs-fullreport.pdf> (Accessed 04/02/08).
- eTesting Labs. *Updated Web Content Software Filtering Comparison Study*. Department of Justice: October 2001. <http://web.archive.org/web/20030727105727/http://veritest.com/clients/reports/usdoj/usdoj.pdf> (Accessed 04/02/08).
- Finnell, Cory for the Certus Consulting Group. *Internet Filtering Accuracy Review*. Department of Justice. 2001. http://filteringfacts.files.wordpress.com/2007/11/cipa_trial_finnell_ex_report.pdf (Accessed 04/02/08).
- Greenfield, Paul and Peter Rickwood and Huu Cuong Tran. *Effectiveness of Internet Filtering Software Products*. Australian Broadcasting Authority. 2001. <http://www.acma.gov.au/webwr/aba/newspubs/documents/filtereffectiveness.pdf> (Accessed 04/02/08).

- Haselton, Bennet. *Report on the Accuracy Rate of FortiGuard*. American Civil Liberties Union. 2007. http://filteringfacts.files.wordpress.com/2007/11/bradburn_haselton_report.pdf (Accessed 04/02/08).
- Heins, Marjorie, Christina Cho, and Ariel Feldman. *Internet Filters: A Public Policy Report*, 2nd Edition. Brennan Center for Justice, NYU School of Law. 2006. <http://www.fepproject.org/policyreports/filters2.pdf> (Accessed 04/02/08).
- Janes, Dr. Joseph. Expert report of Dr. Joseph Janes. American Civil Liberties Union. 2001. <http://www.aclu.org/FilesPDFs/janesreport.pdf> (Accessed 04/02/08).
- Kaiser Family Foundation. "See No Evil: How Internet Filters Affect the Search for Online Health Information." Kaiser Family Foundation. December 12, 2002. <http://www.kff.org/entmedia/3294-index.cfm> (Accessed 04/02/08).
- Markkula Center for Applied Ethics. *Access, Internet, and public libraries – the effectiveness of filtering software; recommendations*. Santa Clara University. 2007. <http://www.scu.edu/ethics/practicing/focusareas/technology/libraryaccess/> (Accessed 04/02/08).
- National Research Council. "Youth, Pornography, and the Internet." National Academy of Sciences. 2002. <http://www.nap.edu/openbook.php?isbn=0309082749> (Accessed 04/02/08).
- Net Protect. *Report on the evaluation of the final version of the NetProtect Product*. 2004. <http://www.net-protect.org/en/EADS-WP5-D5.2-v2.0.pdf> (Accessed 04/02/08).
- Online Policy Group and the Electronic Freedom Foundation. *Internet Blocking in Public Schools: A Study on Internet Access in Educational Institutions*. San Francisco, CA: Online Policy Group, June 2003. http://www.onlinepolicy.org/access/blocking/net_block_report/net_block_report.pdf (Accessed 04/02/08).
- Resnick, Paul. *Expert Report*. North Central Regional Library District. 2008. http://filteringfacts.files.wordpress.com/2008/02/bradburn_04_05_08_resnick_report.pdf (Accessed 04/02/08).
- Stark, Philip B. *Expert Report*. Department of Justice. 2006. http://filteringfacts.files.wordpress.com/2007/11/copa_trial_stark_report.pdf (Accessed 04/02/08).
- Untangle. *Deep Throat Fight Club Open Testing of Porn Filters*. 2008. http://www.untangle.com/index.php?option=com_content&task=view&id=283&Itemid=1122 (Accessed 04/14/08).
- Veritest. *Websense: Web Filtering Effectiveness Study*. WebSense. 2006. <http://www.lionbridge.com/NR/rdonlyres/websensecontentfilte7fmspvtstryjhojtsecqomzmiriqoefctif.pdf> (Accessed 04/02/08).

Relevant Court Cases

American Civil Liberties Union vs. Gonzalez.

<http://www.paed.uscourts.gov/documents/opinions/07D0346P.pdf> (Accessed 04/02/08).

American Civil Liberties Union v. Miller. <http://www.aclu.org/news/n062097b.html> (Accessed 04/02/08).

American Civil Liberties Union v. Reno II. <http://www.aclu.org/news/2000/n062200b.html> (Accessed 04/02/08).

American Library Association v. Pataki. <http://www.aclu.org/news/nycdahome.html> (Accessed 04/02/08).

American Library Association v. U.S. Department of Justice and Reno v. American Civil Liberties Union. <http://www.ciec.org/> (Accessed 04/02/08).

Mainstream Loudoun v. Board of Trustees of Loudoun County Library.

<http://loudoun.net/mainstream/Library/Internet.htm> (Accessed 04/02/08).

Preliminary Injunction Against Child Online Protection Act and Judge Lowell Reed's Decision.

<http://www.aclu.org/features/f101698a.html> (Accessed 04/02/08).

United States vs. American Library Association (CIPA).

<http://www.supremecourtus.gov/opinions/02pdf/02-361.pdf> (Accessed 04/02/08).

Filtering Studies and Their Findings

Date	Title	Source	Summarized Conclusions
2008	Deep Throat Fight Club Open Testing of Porn Filters	Untangle	<ul style="list-style-type: none"> • Fortinet 97.7% accuracy blocking trigger websites • Watchguard 97.3% accuracy blocking trigger websites • Websense 97.0% accuracy blocking trigger websites • SonicWall 96.1% accuracy blocking trigger websites • Barracuda 94.0% accuracy blocking trigger websites • Average of 99% accuracy allowing non-trigger sites
2008	Expert Report	Dr. Paul Resnick (for North Central Regional Library District)	<ul style="list-style-type: none"> • 93.1% accuracy blocking trigger websites • 48% accuracy blocking trigger images
2007	Report on the Accuracy Rate of FortiGuard	Bennet Haselton (for the ACLU)	<ul style="list-style-type: none"> • 88.1% overall accuracy on .com sites • 76.4% overall accuracy on .org sites
2006	Expert Report	Philip B. Stark (for the DOJ)	<ul style="list-style-type: none"> • 87.2%-98.6% accuracy blocking “sexually explicit materials” • 67.2%-87.1% accuracy allowing “non-sexually explicit materials”
2006	Websense: Web Filtering Effectiveness Study	Veritest (for Websense)	<ul style="list-style-type: none"> • WebSense: 85% overall accuracy • SmartFilter: 68% overall accuracy • SurfControl: 74% overall accuracy

2004	Report on the evaluation of the final version of the NetProtect Product	Net-Protect.org	<ul style="list-style-type: none"> • Surf-mate: 85% accuracy blocking trigger content and 89% accuracy allowing non-trigger content • CyberPatrol: 44% accuracy blocking trigger content and 95% accuracy allowing non-trigger content • Net Nanny: 18% accuracy blocking trigger content and 97% accuracy allowing non-trigger content • CYBERSitter: 24% accuracy blocking trigger content and 97% accuracy allowing non-trigger content • Cyber Snoop: 3% accuracy blocking trigger content and 99% accuracy allowing non-trigger content • NetProtect 2: 96% accuracy blocking trigger content and 83% accuracy allowing non-trigger content
2003	Internet Blocking in Public Schools	Online Policy Group	<ul style="list-style-type: none"> • School curriculum materials accessed with filters set to least restrictive settings: 95-99.5% accuracy • School curriculum materials accessed with filters set to most restrictive settings: 30% accuracy
2002	Corporate Content Filtering Performance and Effectiveness Testing Websense Enterprise v4.3	eTesting Labs (for Websense)	<ul style="list-style-type: none"> • SuperScout: 90% accuracy blocking “adult” materials • SmartFilter: 90% accuracy blocking “adult” materials • WebSense: 95% correct accuracy blocking “adult” materials
2002	No Evil: How Internet Filters Affect the Search for Health Information	Kaiser Family Foundation	<ul style="list-style-type: none"> • 98.6% accuracy in accessing health information on least restrictive settings • 95% accuracy in accessing health information on intermediate restrictive settings • 76% accuracy in accessing health information on most restrictive settings
2001	Expert report of Dr. Joseph Janes	Dr. Joseph Janes (for the ACLU)	<ul style="list-style-type: none"> • 34.3% accuracy in allowing non-trigger content

2001	Internet Filtering Accuracy Review	Cory Finnell for the Certus Consulting Group (for the DOJ)	<ul style="list-style-type: none"> • CyberPatrol: 92.01%-95.31% overall accuracy • Websense: 89.97%-94.75% overall accuracy • Bess: 93.08%-91.64% overall accuracy
2001	Updated Web Content Software Filtering Comparison Study	eTesting Labs (for the DOJ)	<ul style="list-style-type: none"> • 92% average accuracy of four filters in blocking “objectionable” content • 96% average accuracy of four filters in allowing non-trigger content
2001	Digital Chaperones for Kids	Consumer Reports	<ul style="list-style-type: none"> • Cybersitter 2000: 78% accuracy blocking “objectionable” content • Internet Guard Dog: 70% accuracy blocking “objectionable” content • AOL's Young Teen Control: 63% accuracy blocking “objectionable” content • CyberPatrol: 77% accuracy blocking “objectionable” content • NetNanny: 48% accuracy blocking “objectionable” content • NIS Family Edition: 80% accuracy blocking “objectionable” content
2001	Effectiveness of Internet Filtering Software Products	Paul Greenfield, Peter Rickwood, and Huu Cuong Tran (for the Australian Broadcasting Authority)	<ul style="list-style-type: none"> • N2H2 (now Bess), set to “maximum filtering,” was reported as the most effective filter tested in this study • 95% accuracy blocking the “pornography/erotica” category • 75% accuracy blocking the “bomb-making/terrorism” category • 65% accuracy blocking the “racist/supremacist/Nazi/hate” category • 40% accuracy allowing non-trigger content in the “art/photography” category • 60% accuracy allowing non-trigger content in the “sex education” category • 70% accuracy allowing non-trigger content in the “atheism/anti-church” category • 80% accuracy allowing non-trigger content in the “gay rights/politics” category • 85% accuracy allowing non-trigger content in the “drug education” category

2001	Report for the European Commission: Review of Currently Available COTS Filtering Tools	Sylvie Brunessaux et al.	<ul style="list-style-type: none">• Average of the 10 filters tested• 67% accuracy blocking trigger sites in English• 52% accuracy blocking trigger sites in five languages• 91% accuracy allowing non-trigger content
------	--	--------------------------	---