HISTORY AND DEVELOPMENT OF FILTERS

Internet content filters have been available since the mid 90s. The first filters to appear on the market relied largely on keyword blocking, a simplistic and ineffective way to filter content.

Keyword blocking occurs when the searcher uses a word on the filter's long list of forbidden words, and that word is "disappeared." For example, if *breast* was on the keyword list, any search for *breast cancer* would simply be conducted as *cancer,* as if the user had not typed the word *breast* at all.

These simplistic filters even "disappeared" words from the content of a page resulting in pages that made no sense or that stated something quite different from the author's intent. One such incident reported by Peacefire resulted in a filter changing a sentence on a website from "the Catholic Church opposes homosexual marriage" to "the Catholic Church opposes marriage."[3]

URL: Uniform resource locator

Other problems with keyword blocking included rendering URLs inaccessible and altering e-mail and chat messages. The filter companies decided a better approach was to block the entire page, not just the word.

CyberSitter and many other filter products soon began blocking entire Web pages when a keyword or key phrase was encountered. This blocking method prevented some of the embarrassment associated with changing the meaning of sentences on a Web page, but the result was that entire pages were lost because of the presence of a single word.

These early filters were designed for parents to use in the home where concerns about overblocking content were not high.[4]

Although most of these early products are still available and still marketed to parents, several new factors affect the filtering marketplace:

- New markets for filters have emerged.
- Different technologies for filtering have been developed.
- Filters once designed for home use have expanded their customer base to include schools, libraries, and businesses.

The customer base for Internet filters—once limited to parents seeking blocking and monitoring software—has expanded to include churches, public schools, private schools, libraries, businesses, Internet service providers, and even entire countries.

As a result of this expansion, the function of Internet filters has changed. Where the early filters were designed to block all content deemed inappropriate for children, many of today's filter companies are devising new and creative ways to categorize the entire Internet thus providing their customers with the ability to block a broader range of material.

The importance and problem of content categories

Most filters allow for the filter administrator to control, to a large extent, what is blocked and what is allowed. In most cases, however, the decision about what to block and what to allow is made by category, not by domain, website, URL, or page.

The filter company decides which domains or pages fall in any particular category. The specific websites within each filter's categories are not usually made public or even available to libraries on a nondisclosure basis. Rather than publish lists of URLs contained in each content category, filter companies describe each category and sometimes provide examples of pages that belong in it.

In devising a blocking strategy, the filter administrator has only the category descriptions to go on. The administrator can never be certain which sites were chosen to be blocked. For example, although the pornography category may be selected for blocking, no one can be sure Lolita's private webcam is blocked but the Williams College "Sex and Sexuality" pages are not.

That category lists are not publicly available is one of the primary complaints lodged against filter companies. Filter companies argue that their category lists or content databases, as this collection of categorized websites is called, is a major component of what their customers pay for (and their competitors would benefit from).

Free speech advocates argue that the lists should—at the least—be transparent and customizable so that pages that appear in the wrong category can be immediately corrected (or moved to a different category).

A few filters are available that allow administrators to view the category lists and their associated URLs. These products are often based on open-source products such as Squidgard or Dan's Guardian.

Network-based filters with viewable content databases include: Squidgard, Dan's Guardian, Netpure, EngageIP, IF-2K, Corporate Guardian, CyberSetting, and Netsweeper.

How an individual site ends up in any given category is part of the proprietary process devised by each filtering company. The success of the filter largely depends on how accurate the customer believes the classification process is and how useful the categories are. As a result, the content categories often shed light on who the filter company is marketing their product to and what they understand their customers are trying to accomplish with the filter.

No filter is designed exclusively for libraries despite the unique needs of libraries. Filters are designed for parents, schools, and business and have been influenced largely by religious groups and employers whose filter requirements are different from a library's. The influence of these marketplace pressures has changed filters dramatically over the years, particularly in how categories of content are defined.

Influence of faith-based organizations on filter categories

In attempting to serve their religious constituencies, filter companies have added categories of content that meet the needs of people sharing a certain religious or moral point of view. Consider the following Websense categories:

Religion

- Traditional
- Nontraditional

Abortion advocacy

- Pro life
- Pro choice[5]

Saudi Arabia uses Websense for "preserv[ing its] Islamic values, filtering the Internet content to prevent the materials that contradict [its] beliefs or may influence [its] culture."[6] In her article, "Internet Filters: The Religious Connection,"[7] Nancy Willard of the Center for Safe and Responsible Use of the Internet, describes the link among largely Christian organizations and several of the most popular filtering companies including Symantec's I-Gear, N2H2's Bess, 8e6Technologies' X-Stop, Solid Oak Software's CyberSitter, and others. She suggests that many of the content categories users can choose to block have been added to address the views of these faith-based groups.

Here are some examples of categories likely to have arisen from demands by faith-based organizations:

Symantec/I-Gear:

Occult and New Age—Sites dedicated to occult and New Age topics including but not limited to astrology, crystals, fortunetelling, psychic powers, tarot cards, palm reading, numerology, UFOs, witchcraft, and Satanism.

Sex education and sexuality—Sites dealing with topics in human sexuality. Includes sexual technique, sexual orientation, cross-dressing, transvestites, transgenders, multiple-partner relationships, and other related issues.

8e6 Technologies/X-Stop:

Alternative journals—Sites for nonmainstream periodicals, information on self-awareness, spirituality, healing arts, holistic living, junk culture, fringe media, art perspectives, and so on.

Anarchy—Sites contain information regarding militias, weapons, anti-government groups, terrorism, overthrowing of the government, killing methods, and so on.

Cult—Sites promoting cult or gothic subject matter, use of mind control, paranoia, fear, and any other type of psychological control or manipulation.

Lifestyle—Information promoting adultery, swinging lifestyles, and same gender or transgendered relationships.

Influence of businesses on filter categories

Religious groups aren't the only ones who have influenced the content categories found in today's filters. An even larger number of categories have been developed to address employers' desire to prevent their employees from engaging in nonwork-related activities at work.

The Internet provides a potential playground for employees to use while enjoying a high-speed connection to the Internet, which makes employers nervous.

Most office workers today require Internet access to do their jobs effectively. Many employees would prefer to conduct personal Internet business while using that high-speed connection. Internet filters allow employers to restrict access to pages employees can visit and to monitor what they do.

Filter companies cite studies that state 37.1% of employees surf the Web constantly at work, 31.9% do it a few times a day, 21.3% do it a few times a week, and only 9.7% said they never surf at work.

Survey at www.vault.com.

Employers also have become more interested in Internet filters because of the explosion of the number of sexual harassment cases. Not willing to trust their employees to understand the difference between fun and potentially litigious conduct, employers are installing filters, which limit what can be viewed on the Internet and what can be e-mailed from work, to avoid being held liable for the conduct of their employees.

Productivity categories are categories designed to address the problem of employees using the Internet for nonwork-related Internet research or to play games or engage in activities seen by employers as nonproductive. Filtering products designed for the business market generally include many categories that address both content (usually sexually explicit material) and productivity concerns.

These filters, such as the example that follows, attempt to provide a category for every website on the Internet. The goal is to empower customers to configure the filter to suit their needs.

DynaComm i:filter categories

Adult

Advertising sources

Business and consumer products/services

Business conferences, online training, and distance education

Charitable and nonprofit organizations

Chat rooms, forums, and online communities

Complaint sites

Education organizations and institutions

E-mail hosts

Employment and jobs

Internet service providers

Law and legal services

News and weather nonmonitored sites

Personals, dating, and personal websites

Political

Portals and search engines

Professional organizations

Recreational drugs and drug paraphernalia

Religion and spirituality

Reproductive health and sexuality

Entertainment

Financial services

Gambling

Government and military

Hacker and cracker activities/information

Health and medical

Information resource

Shareware and freeware

Shopping

Sports and hobbies

Terrorism

Travel and tourism

Web cams and video-diaries

Web hosting sites

Libraries have benefited from the growth of business customers into the filter market. Unlike the home users, business users require some of the advanced features that libraries also need.

For example, the ability to define numerous filter profiles that can be applied to individual PCs, groups of PCs, or individual users, or groups of users gives the library more flexibility in controlling the way filtering can be done in their library. The filters often allow for blocking pages or just warning that a page many not be appropriate (and then allowing the end user to access the site anyway).

The business filters tend to provide many reports for the administrator including details about which users visit which pages. And many of them allow administrators to apply different Internet-use restrictions based on time of day.

Business requirements caused additional override options to be introduced. Typically, systems staff, or floor staff, have numerous ways to override a blocked page on the spot, rather than having to contact the filter company to ask for a change to add the erroneously blocked site to the 'always allow' list, also known as the white list—a customizable list of sites that are always exempted from filtering.

Business class filters also often permit any staff person to override the blocked page at the user's computer and allow the administrators to recategorize the Web page, rename category names, and add and populate new categories.

The ability to customize the default block page is another development that arises from the application of filters beyond home and school use. Most filters now allow for administrators to redirect the browser to an HTML page they designed rather than displaying a generic "Access to this site is not allowed" message.

HTML: Hypertext markup language

The ability to customize the block page provides libraries with an important opportunity to empower patrons with information the patrons need to challenge the decision to block the page, to be advised of any recourse for avoiding the block, and to learn more about the library's Internet use policy (IUP).

Bandwidth and protocol-based categories

Libraries sometimes use filters to restrict what patrons can do on library computers including: using chat and instant messenger programs, downloading files, playing games, gambling, using peer-to-peer file-sharing programs, and more. Bandwidth- and protocol-based categories can be used to limit many of these activities.

Many Web pages rely on certain types of protocols to function. For example, to download content, the FTP protocol is required. To participate in IRC, the

IRC protocol is required. To log in to another server, the telnet protocol is required. Many filters can be configured to prevent certain protocols from being used.

FTP: File Transfer Protocol

IRC: Internet relay chat

Bandwidth-intensive activities users engage in over the Internet include online chatting, playing audio and video files, playing online games, and participating in videoconferences. Filters can prevent users from accessing pages with chat rooms or MP3 or movie files to download. Another approach is to block certain file types from being read by the browser, which effectively prevents people from conducting these bandwidth intensive activities.

Blocking images but not text

Most filters today are designed to block entire pages, not just the images on the page. The filter companies evaluate the content on the page and then categorize those pages. When the filter administrator chooses a category to block, all the pages in that category are blocked, not just the images.

Blocking specific file types (.mp3, gif, jpg, midi) or disallowing certain protocols (FTP, telnet) is most often an 'always allowed' or 'never allowed' prospect. Limiting the bandwidth activities or the protocols within specific content categories isn't usually possible.

For example, preventing images from being displayed when a page is categorized as 'pornography' might be more useful than blocking access to the entire page including text and images. This less-restrictive approach would comply with CIPA, which only requires libraries to prevent access to visual depictions of certain types of sexually explicit material.

Turning off images for all websites, in every category, however, would not be desirable. To achieve the desired results, the product must allow the filter administrator to turn off images within a category. Only a small number of filters offer this feature.

www.libraryfiltering.org provides a list of products that answered 'yes' to the question 'Admin can choose to block images but not text within a category.'

Some libraries have created their own add-on program to block images within a content category. Tacoma Public Library uses Surfcontrol in combination with a script its technical staff wrote to block images and graphics on any pages identified as inappropriate by the filter. Its librarians say this creative solution brings them into compliance with CIPA while reducing the amount of content being blocked because no text is blocked in the library, only certain images.

Finding and cataloging websites

Every filter company has devised its own strategy for finding and classifying Web content. Business-oriented filters with content, bandwidth, and productivity categories must find and then catalog a larger percentage of Internet content than a simpler product targeted at home users. Products for home and school use have developed classification schemes that can be used to limit access to Internet content based on the age of the Internet user.

A classic example is We-Blocker, a free product designed exclusively for school and home use. It has six categories: porn, adult, violence, hate speech, drugs, gambling, and weapons. For We-Blocker to identify sites that belong in each of its six categories, it can run fairly straightforward searches to locate the URLs that will come up highest on the search engines hit lists.

A company such as Websense, on the other hand, makes finer distinctions including between traditional and nontraditional religions. And CyberPatrol distinguishes between partial nudity and full nudity. These companies' products probably can't rely on simple searches to find and easily categorize pages into their proper category. With these more complex products, the first job is to locate the sites that need to be categorized.

Help from customers

One technique vendors use to locate websites is to let the customer find them. Many filter companies have a place on their website where customers or potential customers can go to see how a Web page will be categorized. Visitors can enter the URL into the online form and information about how that filter classifies the site is returned.

Allegedly, these online forms are made available so potential customers can see how a site will be categorized by the filter. But these forms also are a useful way for filter vendors to learn about sites that their customers are likely to be interested in.

In most cases, these forms allow the visitor to suggest a category for the site if it hasn't yet been categorized. Or if it has been categorized, the visitor can suggest a change. Any suggestion by a visitor for how to categorize a Web page will likely be based on a more thorough evaluation of the content than the automated tools the filter company relies on.

Current methods of filtering

To filter Internet content, one of two methods is generally used: pass-through or pass-by technologies. The most common approach to filtering is a pass-through method.

With a pass-through method, the requested page is first passed through a device, such as a proxy server, where the URL is looked up in a database that indicates whether that page will be allowed. If it is not allowed, the request never goes out over the Web (saving bandwidth) and the end user is sent a block page (see Figure 1) instead of the requested Web page. These types of filters are generally referred to as URL filters.

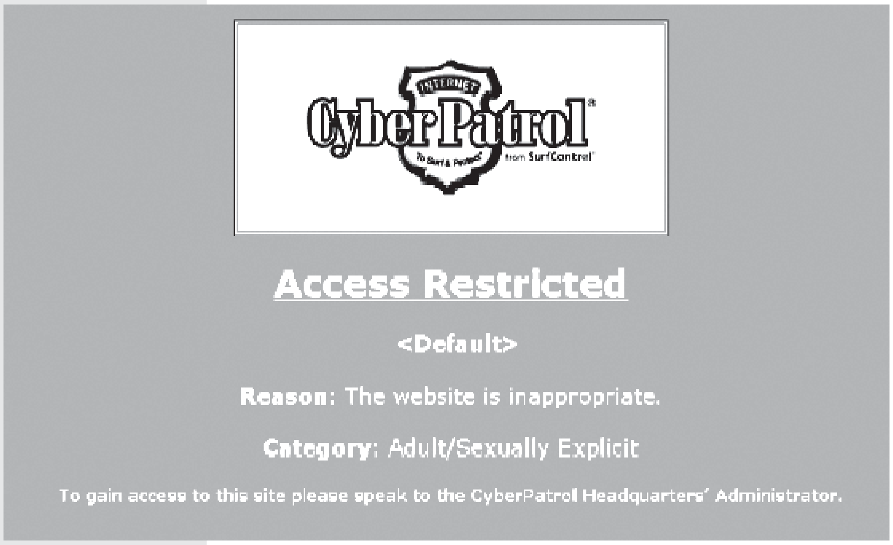

Access Restricted

<Default>

Reason: The website is inappropriate.

Category: Adult/Sexually Explicit

To gain access to this site please speak to the CyberPatrol Headquarters' Administrator.

Figure 1. CyberPatrol Default Block Page.

The other method for filtering Internet traffic is a pass-by approach. These filters analyze, or sniff, each Internet packet as it is retrieved from the

Internet. They determine on-the-fly whether access to the page is allowed. These types of filters are usually referred to as content filters.

On-the-fly: At the moment the page is requested by the user.

URL filters

URL filters rely on populating a list, or database, with URLs. Each URL is associated with one or more categories.

When the Internet user selects a Web page to visit, the URL filter checks to see if that URL is in the database. If the URL is in the databse, the user will be allowed to view the page only if the category the page falls in is allowed by the filter administrator. If the page is not found in the database, it will be blocked or allowed depending on how the filter is set to address pages not yet categorized by the filter company.

Because a URL filter's success relies on the ability to look up a URL in a database, these databases must be continuously updated with new URLs.

Users of URL filters generally enter into a subscription agreement with the filter company to receive updates to their database. Because of this reliance on large lists of URLs, URL filters must continuously seek sites their customers are likely to encounter, categorize them, and place them in the database.

Filter companies may well use popular search tools to increase the chance that they will find the same websites their users will find. The average person using Google, for example, might look through the first 20 to 30 listings from any search they conduct.

They won't likely scroll through 100 to 200 matches looking for what they want rather than conducting a new search. Therefore, some filter companies probably continuously search the Web using tools such as Google or Alta Vista, attempting to find the sites their users are most likely to find. Using the hits returned by the search engine, they are able to collect the most popular URLs and quickly file them in the most suitable category of their content database.

Here's a simplification of how filters find and classify Web pages:

1. Search: shocking sex acts
2. Remove any from domains ending in .edu or .gov
3. Classify top 500 hits in pornography category
4. Spot check for errors

The way filter companies actually find and classify Web pages is more complicated. Many filter companies have designed their own special search tools designed specifically to locate content in their target content categories. And they have developed sophisticated programs for conducting the searches as well as weeding the pages that don't fit.

These types of filters are called URL filters because they rely on comparing the URL accessed by the user with the URLs contained in the filter's database. When the filtering software finds a match, it looks at which content category the URL was found in.

If the category is a blocked category, the end user will be shown the default block page instead of the requested Web page. If the category is an allowed category for that user, the browser will complete the request and the page will display (see Figure 2).

If the URL filter being used is one simple block list, rather than lists of URLs broken into content categories, the filter simply has to check for the presence of the URL in the block list. If it is there, the default block page is displayed. If the URL isn't on the block list, the page is retrieved.

Content filters

Another way filters work is by analyzing the content of the page on-the-fly. That is, instead of precategorizing URLs, only the URLs retrieved in response to the search are categorized.

The browser retrieves the page but does not display it. First the filter analyzes the page to determine what category it should be classified into. Like the URL filter, it will then present the end user with a block page or the requested page, depending on whether the content category it was classified into was selected for blocking.

Every company doing content filtering has developed some kind of proprietary technology for quickly analyzing content on the page. To be effective, the analysis must be quick enough so as not to delay the retrieval process.

The software engineering that goes into these content analyzers has become more sophisticated than its early predecessor—keyword blocking. Although some analysis of words in the URL and on the page is part of the process, other steps are involved in evaluating the page and placing it into a content category.

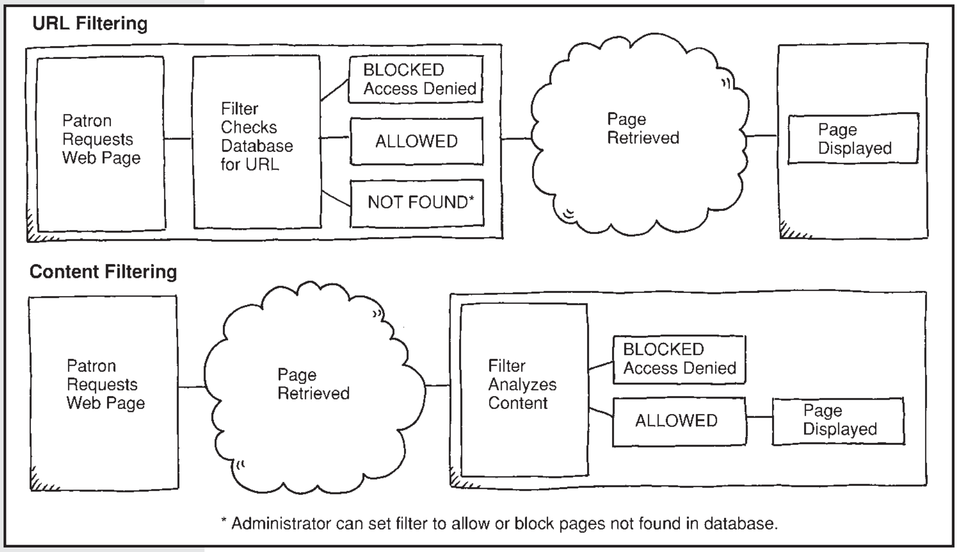

Figure 2. How URL and Content Filtering Works.

Combination products

Because of the time involved in conducting the analysis step, some content filters incorporate a URL database component. Sites that have been processed by the artificial content recognition engine (Figure 3), for example, are placed in their content category and then registered in a database.

Conversely, some content filters build up a local database of sites that have been accessed by users at a particular location and store that database locally—at the customer's location or in a central database maintained by the filter company. Regardless of where the database resides, the filter is designed to first check the database to see if the site has been categorized. If it has been, the category information is returned and access is allowed or denied depending on how the local administrator has set up the filtering profile for the end user.

So far, the methodology used by these so-called content filters then is the same methodology as is used by the URL filter. The difference with a content filter is that if the URL is not contained in the database, the analysis step is conducted and the site is dynamically evaluated and categorized. Once categorized, the database is updated so the next time the same site is accessed, the analysis will not have to be repeated.

Some URL filters also incorporate an element of content analysis into the product to prevent the problem of users attempting to access a site that has not yet been classified into a content category.

For example, CyberPatrol, primarily a URL filter, uses artificial intelligence as well as keyword blocking to supplement its URL filter, known as the CyberList.

CyberPatrol's website reports: “CyberPatrol combines powerful filtering technologies. It includes the CyberList database of accurate and relevant categorized websites, as well as artificial intelligent technologies, Web Page Analysis, and CyberPatterns, to filter websites as they are visited and that are too new to be captured by the CyberList database. CyberPatrol also can filter offensive text-based words and phrases from Web-based e-mail.”

CyberPatrol 6 Fact Sheet, www.cyberpatrol.com/product/cyberpatrol_factsheet.pdf

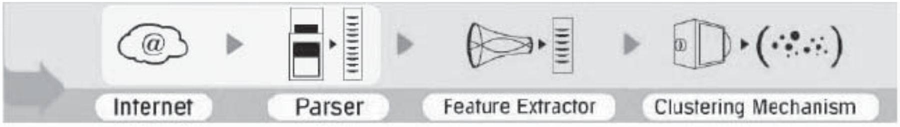

Artificial Content Recognition Engine

Artificial content recognition (ACR) technology examines each requested HTML page, and then categorizes it. The following steps describe the filtering process of ACR:

1. A Web page is requested by a user. The requested page is received by the network and sent packet-by-packet to an HTML parser.
2. The parser breaks down the HTML code of the Web page into hundreds of parameters. These parameters include (among many other) characteristics such as individual words, background color; links; number of links; banner ads; number of images and words; average number of letters in a word; color and size of font; whether a word is in a metatag, body tag, or other HTML-based tag; and the type of words.
3. The parameters make up the raw data vector (RDV), a vector that defines all information extracted from the HTML page.
4. The RDV is too large to process in real time or to be meaningful enough to provide accuracy, so the large amount of data must be processed in order to extract the relevant information. A feature extractor is implemented, an artificial intelligence algorithm capable of processing the information to create a processed data vector (PDV). Specifically, the feature extractor finds patterns in the parameters that are useful in classifying the Web page. For instance, the feature extractor might look at the color of the font as compared with the color of the background as one such distinguishable pattern. In this way, the RDV is reduced to tens of “features” from the original hundreds of parameters.
5. The processed data vector is then processed by a clustering mechanism. The clustering mechanism is also an artificial intelligence algorithm. It takes the combinations and relationships of the features and produces a unique mathematical coordinate.
6. The mathematical coordinate generated by the clustering mechanism can be grouped within a corresponding cloud of preclassified categories of Web pages. For instance, one cloud might be sex, another cloud might be hate, and yet another cloud would represent drugs. By assigning the mathematical coordinate to one of the clouds, the ACR technology identifies the type of Web page the user has requested.

Source: Allot Communications. www.allot.com/media/ExternalLink/ACR%20White%20Paper_020709.pdf.

Figure 3. Artificial Content Recognition Engine.

Pros and cons of two techniques

URL filtering and content filtering each has advantages and disadvantages. The primary advantage of URL filtering is that blocked Web pages are not allowed to even enter the network. This blocking saves bandwidth and ultimately reduces the load on the network.

With content filtering, the Web page must be retrieved for the analysis to be performed. Even those pages that will never be viewed are dragged through the network, potentially causing network congestion problems as the number of requested pages increases.

The main disadvantage of URL filters is that they rely on the size of their URL database. Sites that have not been classified may be allowed through, resulting in end users accessing sites that would normally be blocked. Given the dynamic nature of the Web, the ingenious marketing techniques employed by businesses promoting their products, and the exponential growth of the Internet, the URL databases will always remain a small fraction of the Internet sites on the Web.

Where filters are installed

Filters can be installed on each individual PC (client-based) or they can be installed on a network server (server-based) or they can stand on their own (network appliance). Some ISPs offer filtered Internet access and some filter companies can do the filtering remotely (hosted). The library's technology manager is best equipped to evaluate the pros and cons of each approach for any given library.

Client-based filters

Client-based filters are installed on each individual PC. Depending on how many PCs the library has, this job might be onerous. Maintaining the PC and filter, including making configuration changes and updating the software need to be made at each computer.

Selection tip
Don't use client-based filters on public-access computers.

Avoiding client-based filters, or any filter that requires installing software on each PC, reduces the likelihood of software conflicts with other software on that PC—especially public-access computers (PACs).

Most libraries run several layers of security on their PACs including programs such as WINSelect, Gates Security, or CleanSlate. In addition, many PACs are tied into PC reservation, session control systems, and print management programs. Avoid introducing additional software to the already complex PAC desktop environment, if possible.

Filters that have been created for home use are almost always client-based. The expectation is that the filter will be installed on the family PC. These filters tend to be easy to manage and install but are suitable only for children's computers used at home due to the focus of the content categories and sometimes simplistic blocking technology and inability to override.

If the library is planning to filter more than two or three PCs, librarians will probably be better off using a network filter (server-based or network appliance) that can be easily set up with dial-in or network access for support purposes.

Rather than leaving the entire management of the filter to the branch librarian, technology staff can install, configure, and support the filter while leaving day-to-day administration to local staff. Even a small library with four or five PACs will benefit from a server-based filter that can be centrally managed and remotely supported.

Server-based filters

Many filters designed for business or school use are designed to run on a server. Proxy servers are ideally suited for an Internet filter because they are already set up to intercept all Internet traffic that passes through the network (pass-through). Other types of servers also can be used to support filters such as Web servers, firewalls, and cache servers.

Each filter company provides a list of supported platforms for its products. Platform is an important factor to consider when selecting a filter. If the filter can be installed on an existing server, the library can save on hardware costs. Upgrading the server with more RAM or disk space may be necessary, but any upgrade would be cheaper than buying a whole new server.

Try to avoid introducing a new server platform when selecting a filter. For example, do not buy a filter that requires Microsoft Windows Internet Information Server if you are already running Apache as your Web server. If the filter cannot be installed on an existing server then select a server platform that is already supported by your staff.

Selection tip
Do not introduce a new server platform when selecting a filter.

Even more costly is the time required to train technology staff in the new software platform. Whatever hardware and software environment the library supports will likely allow many filter choices, so no reason exists to start with something new.

Selection tip
Look for a Web-based management interface unless the library has a dedicated server administrator available at the server console at all times.

Server-based filters often have a management console at the server that allows the filter administrator to set up filter profiles, manage filter settings, monitor logs, and generate reports. Sometimes these management functions, or a subset of them, can be performed from a Web-based management interface that is accessible from anywhere on the library's network. Having access to management features—without having to return to the server each time a change is needed—can be an important timesaver, especially if no dedicated filter administrator monitors filter activities from the server room.

Network appliances

Network appliances function independently of any operating system or network. These so-called black boxes are stand-alone applications that are generally managed from a Web interface. Network appliances do not require an existing server nor are they limited to the type of network they will run on. Some network appliances do not include the hardware but can be installed on any generic PC hardware platform. Network appliances can usually be added to the network at any network port or can sometimes be attached to the router or firewall or some other network device.

Being able to manage the filter from anywhere on the network, rather than being limited to a designated server console, is a big time-saver. Network appliances can almost always be managed from anywhere on the network, eliminating the need for technical staff to remain in the server room at the filter management console.

ISP-based or hosted filters

Filters also can be hosted by an Internet service provider (ISP) or the filter company. Depending on the filter used by the ISP and the degree to which local control over configuration is possible and permitted, this hosted filter might be a way to save the library money in hardware and technical staff.

ISPs use many of the same filters available to businesses so find out what filter the ISP uses and study the content categories just as you would if you were considering buying the filter yourself. In addition, evaluate the ability to configure and control unblocking and disabling the filter and whether you can control what is on the default block page. Although using a hosted filter approach might be handy, only use it if the library retains sufficient control.

Defining a good filter

When defining a good filter, the most important consideration is the suitability of the filter for your environment. A company might have a sophisticated method for finding and locating a wide range of Internet pages but if its categories don't match the library's needs, the filter isn't going to be useful.

For example, if a filter that contains URLs in its database of sex toy sites, sites with profanity or vulgar language, sex education sites, sites with nudity, safe sex sites, as well as commercial pornography sites but categorizes them all as 'sex ,' the library won't be able to minimally block pages for adult patrons. To block sexually explicit sites from children, all the content categorized as 'sex' would be blocked—including the sex education and safe sex sites. Whether the filter is a URL filter or a content filter won't matter—or if it is 99% or 80% accurate—if the categories don't work for the library.

Overblocking: Sites that are blocked because they have been placed in an inappropriate category.

Underblocking: Sites that are not blocked because they have been placed in the wrong category or have been missed by the filter.

In fact, anyone installing a filter should assume that the filter is wrong 15% of the time. Although some filters boast a higher accuracy rating, most evaluations[8] of filter accuracy have found filters to be roughly 85% accurate when considering both overblocking and underblocking mistakes.

Handling overblocking and underblocking mistakes is one of the responsibilities of using a filter. Evaluating the accuracy of a filter before buying it is key, but don't expect 100% accuracy. In fact, be careful about filters that are 100% accurate for underblocking—meaning they never miss a site—because any product that never underblocks will overblock more.

Filter accuracy is highly subjective. A site one person considers pornography may be considered nudity by someone else. Any filter that falls in the 85% accuracy range according to third party evaluations is probably adequate.

Monitoring the sites that are blocked and correcting the filter's mistakes—as you define mistakes—will cover the remaining 15%. The more important evaluation has to do with how well the content categories can be used to define usable filter profiles for the library's users.

Alternatives to commercial filters

The decision to install a commercial Internet filter is tantamount to outsourcing traditional professional responsibilities, namely selecting and categorizing content to people with no such training.

Categories into which websites are assigned do not fall into any recognized authority such as Library of Congress subject headings. And the companies doing the work do not necessarily have a commitment to commonly held library principles such as freedom of speech and the importance of free access to information for everyone.

Filters are not the solution. Filters are an imperfect approach to a complex problem. The fact that CIPA mandates the use of technology protection measures has caused many libraries to take up filtering without adequately exploring the wide range of alternatives on the market.

Not only are there more than the five or six filters that many libraries already use, but a wide range of new commercial products exist. Consider open-source products and other approaches as viable technology protection measures, too.

Restricted access

Restricted access doesn't block offensive sites; it refers to filtering by selecting websites for inclusion. Only those websites selected are available to patrons.

Most librarians acknowledge that a large amount of material available on the Web would not be chosen for their collection if that same material were available in book form. But the number of websites on the Internet and the speed with which websites are added and pages are moved or renamed makes applying traditional collection policies to the Internet impossible.

Restricted access is not a viable alternative. The only advantage to this approach is that librarians are once again in charge of collection development decisions. Only high-quality material would be part of the library's Internet collection.

But because of the amorphous and dynamic nature of the Internet, many wonderful new sites or some of the difficult-to-find sites might never make their way into the library's Internet collection.

Even more than other types of filtering, the likelihood of patrons being denied access to enormous amounts of constitutionally protected material would be high using a restricted-access approach.

PICS-rated sites

PICS, the Platform for Internet Content Selection, has developed a specification that allows Web page creators to classify their own sites based on content. The Recreational Software Advisory Council (RSAC) is the most widely used rating system available. RSAC was founded "to protect children from potentially harmful content while preserving free speech on the Internet," according to its website.

RSAC has been incorporated into ICRA, the Internet Content Rating Association. ICRA uses only a few content categories such as sexual material, violence, language, gambling, and chat. Within each broad category, levels exist.

For example, the sex category is further subdivided into passionate kissing, clothed sexual touching, nonexplicit sexual touching, and explicit sexual activity. The focus of the ICRA categories is to identify Web pages that are inappropriate for children at different ages.

PICS is not widely used although it is one way to provide filtering without overblocking and without relying on filter companies to do the job of classifying pages because the people putting up the Web page apply the appropriate rating to their own site. PICS also is easy to use and free for anyone using Internet Explorer and some other browsers. For example, Netscape has a similar program called Netwatch.

Platform for Internet Content Selection, www.w3.org/PICS

Recreational Software Advisory Council (RSAC), www.rsac.org

ICRA, www.icra.org

Netwatch, http://wp.netscape.com/comprod/products/communicator/netwatch

Content Advisor, a component of the Internet Explorer browser, uses the PICS system and allows the user to decide how restrictive the blocking will be. The most important feature of Content Advisor is that it also allows the user to decide what will happen when unrated sites are encountered.

The vast majority of sites are unrated because so few sites use any kind of PICS rating. If Content Advisor is set to allow all unrated sites, the filtering on that terminal will be minimal, but it *will* have a technology protection measure installed.

Customizing your own block list

One surefire way to accurately filter Internet content doesn't involve outsourcing the job of categorizing sites to the filter companies. Rather than buying a filter and relying on how they've decided to classify websites, libraries can develop their own list of sites to block.

Kanguard, http://skyways.lib.ks.us/KSL/libtech/kanguard

Kansas libraries are already doing just that with KanGuard[9], which began as a service of the Northeast Kansas Library System (NEKLS). Using Squidgard and a single block list designed to meet CIPA requirements, every public library in Kansas can filter CIPA-mandated content for free.

Patrons can suggest pages they'd like added to the block list and a small group of librarians determine whether the site should be added. The filter runs on Linux servers located at—and supported by—the Northeast Kansas Library system. Linux and Squidgard are free open-source products so they have no software licensing costs.

Selection tip
Consider using a customer block list for all library computers using Squidgard (free). Use a commercial product for select children's computers only.

Creating a custom block list will never be comprehensive but neither will the commercial filter vendors' lists. Lists as small as 100,000 to 300,000 sites have provided libraries with effective Internet filtering especially when combined with an effective filter monitoring program.

With even a small block list, your library can greatly reduce the likelihood of children accidentally encountering offensive sites and meet the requirements of CIPA. Although any filtering strategy undertaken by the library must be discussed with local counsel, for the purposes of CIPA compliance, the library is only required to have a filter installed and the ability to disable that filter.

When using a custom block list with a relatively small number of sites on the block list, filter administrators should monitor sites being visited by library users to see if there are sites being accessed that should be put on the library's block list. (This action can be done without seeing which specific users are visiting specific sites.)

Rather than attempting to include every offensive site on the Internet on a list, the library can narrow its focus to inappropriate sites its patrons are accessing and block only those.

Children and adults actively pursuing sexually explicit or pornographic websites will find them no matter the filtering strategy. No filter, no matter how much you pay for it or how many categories you choose to block, can prevent the determined, clever patron who wants to find sexual content on the Internet.

Filters can be effective at reducing the likelihood that patrons accidentally encounter inappropriate sites and can make finding inappropriate sites more difficult, but no filter is 100% effective.

Your library's policies determine what filtering strategy to use. Is your policy to use a broad brush and err on the side of blocking more legitimate content while reducing the likelihood of patrons accessing offensive or inappropriate material? Or is your policy to selectively block some key targets and see if that does an adequate job for your community?

If your community or library trustees require a more expansive filtering program for the children's computers, a different approach can be put in place for the children's areas while still using your custom block list for all other library computers. This combination approach (custom block list plus commercial filter on children's computers) is likely to be the most cost-effective solution for addressing CIPA and any community concerns.

Ramifications of choosing to filter

Your library's reputation will be affected by your filtering decision whether or not you decide to filter. No right answer exists that will satisfy everyone.

Many libraries, especially larger systems, have decided to forego E-rate discounts because CIPA's filtering requirement is unacceptable. These libraries may still be using filters to some extent, but not necessarily in the manner mandated by CIPA.

If the library can afford to forego E-rate discounts, they can ensure the decisions about how to handle the challenges of providing Internet access to the community are left in the hands of local decision-makers, rather than the FCC or Congress.

Unfortunately, many libraries cannot afford to give up E-rate discounts and must find a solution that satisfies their library, their community, and the FCC.

Some libraries have stated their opposition to any kind of filtering. Although good arguments exist for keeping filters out of libraries[10], libraries that choose not to use filters still face some public relations concerns.

Some community members will be distressed that their library is not adequately protecting children from the evils of the Internet. Taking the "no filters in our library" position will require the library to adequately train staff and the community to use the Internet safely[10] and will undoubtedly require patience and understanding on the part of library staff and administration who must talk with frustrated and concerned patrons.

Safe and Responsible Use of the Internet, http://csriu.org/onlinedocs

Whether libraries are implementing filtering or taking a stand against filtering, stating your position and reasons in support of your position in your Internet use policy is key (see Appendix A for examples of various IUPs).

If the library is filtering, the IUP should include the library's reasons for doing so. If the library is filtering to comply with CIPA, clearly state the goal in the policy and that CIPA is a federal mandate.

In addition, clarify the following:

- The degree to which each patron will be filtered or monitored when using any library computer
- The ways in which the library is protecting the privacy of patrons even as they filter their Internet use
- The choices patrons have to unblock pages, turn off the filter, or change the filtering level

Internet use policy as guide

The IUP serves as the guide for defining how the filter will be implemented in your library. It also should serve as a guide for selecting the right filter. The IUP should define why the library provides Internet access to their patrons and what range of computer-based activities are allowed.

ALA Intellectual Freedom Committee checklist, www.ala.org/ala/oif/iftoolkits/litoolkit/checklistcreating.htm

ALA has a developed a good checklist for creating an Internet use policy. Its checklist includes the following:

- Ensure that policies speak to access for all.
- Involve your library staff, board, and friends group in the policy writing process.
- Keep it simple. Avoid jargon. Making the policy too technical will confuse people.
- Make policies readily available and visible to the public.
- Provide an up-to-date code of conduct or etiquette guide for using the Internet at your library. Include specific suggestions for positive action. Also list prohibited behavior and the consequences of such behavior.
- Include a statement addressing patron privacy.
- Communicate clearly that users are responsible for what they access online; parents are responsible for their children's Internet use.
- Update your policy regularly; make sure it reflects the Supreme Court CIPA decision.

Consider establishing a committee to review the library's IUP and begin developing procedures to put in place when filtering is implemented. The committee should include legal counsel, the library director, the technology manager, representatives from each library department, a representative from the board, and representatives from the community. The committee should ask itself the following questions:

- Do we only want to minimally filter all PCs to comply with CIPA?
- Are we trying to reduce the likelihood that anyone in the library will encounter unwanted offers of commercial pornography, or are we trying to prevent anyone from being able to access anything gruesome, violent, or sexually explicit?
- Do we want to treat children differently? If so, what type of content do we want to prevent children from seeing? What about young adults? Is there another age group we need to filter differently?
- Do we want to have any unfiltered computers anywhere? If so, shall they be restricted to adult use only? If so, how do we verify only adults use them?
- Are any activities not allowed, such as online games, online chats, and instant messenger?
- Would we like the public access computers to filter differently at certain times of the day—for example, block games on the homework computers after school?

- Who will be allowed to request unfiltered access? Who should they ask for help?
- How do we want to handle unblocking erroneously blocked sites?
- Will there be a feedback policy for patrons who object to our blocking policy?
- Can we integrate filtering with some kind of patron authentication system and allow adults to set their own filtering level?

Knowing the degree to which a filter can be used to support library policies is difficult without a strong foundation in what filters do and how they work. Until the features of the selected filter are known, determining what procedures will be required is impossible.

The next chapter provides information about filters on the market, including details about how the various products differ.

Notes

[3] http://peacefire.org/censorware/CYBERsitter.

[4] CyberSitter Examined. www.peacefire.org/censorware/CYBERsitter/#why.

[5] WebSense Content Categories. http://infopeople.org/howto/filtering/categories/websense.html.

[6] Zittrain, Jonathan, and Edelman, Benjamin. Documentation of Internet Filtering in Saudi Arabia. http://cyber.law.harvard.edu/filtering/saudiarabia (Jan. 20, 2004).

[7] Willard, Nancy, M.S., J.D. "Filtering Software: The Religious Connection." Responsible Netizen. Center for Advanced Technology in Education, University of Oregon, College of Education. Feb. 24, 2002. http://csriu.org/onlinedocs/documents/religious1.html.

[8] The author's own work evaluating filters results in an average 85% accuracy rating across most products. See other studies done by the *Wall Street Journal* (April 5, 2002) and the Kaiser Family Health Foundation (see Derek Hansen's article "CIPA: Which Filtering Software To Use" available on WebJunction.org at http://webjunction.org/do/DisplayContent?id=2102.

[9] Kanguard: Internet Content Filter for Kansas Libraries. http://skyways.lib.ks.us/KSL/libtech/kanguard (Jan. 20, 2004).

[10] See ALA's filter and filtering pages at http://ala.org/oif/ifissues/filters.

[11] Willard, Nancy. Safe and Responsible Use of the Internet and Choosing Not To Go Down the Not-so-good Cyberstreets. Center for the Safe and Responsible Use of the Internet. http://csriu.org/onlinedocs (Jan. 21, 2004).