CHAPTER ELEVEN

Yelling "Filter" on the Crowded Net: The Implications of User Control Technologies

Daniel J. Weitzner

I. INTRODUCTION

Alongside the debate over the V-chip and rating schemes for television smolders burning questions about the impact of blocking, filtering, and other user empowerment tools for the Internet. User empowerment technologies that give users and parents more control over information available online to their children have emerged as critical elements of protecting children from inappropriate content and as a leading alternative to government censorship of content on the Internet.[1] During the constitutional litigation that led to the demise of the Communications Decency Act, a broad cross section of the Internet and civil liberties community—from the ACLU to America Online, Microsoft to the American Library Association, and over 50,000 individual Internet users—enthusiastically supported user empowerment tools as an alternative to censorship. Indeed, in striking down the CDA, the U.S. Supreme Court found the fact that parents can shield their children from material judged inappropriate as a significant alternative to government censorship laws enacted for the purpose of protecting those same children.[2]

With the first battle over the CDA won, however, serious concerns have been raised about the impact of user empowerment tools on the free flow of information. Civil liberties advocates raise significant questions: Will blocking and filtering tools squelch the free flow of information that the Supreme Court sought to protect? Will these technologies become dominated by "mainstream" speakers and publishers, crowding out alternative, noncommercial speech? How will filters account for controversial, yet important speech such as news and public affairs? Will Internet technologies designed to enable parents to protect their children ultimately be used by repressive governments whose goal is to censor the speech of their citizens, adult and child alike?

Supporters of the CDA, on the other hand, maintain that parental empowerment alone is not sufficient to protect children, and they worry that user empowerment tools will not develop to be truly effective tools. Are the tools developing today sufficient to protect children? On what basis should parents trust commercially available filters? How should parents choose among the various offerings currently on the market? Can third-party filtering alone be effective, or is a uniform, universal self-rating approach necessary? Will Internet market forces provide all that families need to protect their children or is some government intervention necessary?

The challenge before the Internet community today is to deploy user empowerment technologies in a way that maximizes the goals of protecting children and at the same time, protects bedrock freedom of expression principles. It must also do so in a way that is sustainable over the long run in the new online environment, recognizing the unique nature and potential of the Internet—its decentralized structure, its potential for promoting access to information, and its global scope.

This is no abstract debate, however, for today the Internet market is offering users a wide variety of software tools and services that either block out content judged inappropriate for children or limit children's access to only that content which has been approved for children. Unlike the V-chip, which will only come into existence as a result of a legislative mandate, the entrepreneurial energy of the Internet is producing a growing range of software and services to give users control over content through blocking, filtering, rating, and labeling mechanisms. The ways in which these tools are deployed in the Internet environment will have a major impact on both the propensity of policymakers to impose government censorship on Internet speech and on the very character of the Net as a platform for speech. This paper begins with a review of the technology landscape and will then consider the various policy positions regarding the appropriate technological and legal outcomes.

II. INTERNET BLOCKING, FILTERING, RATING, AND POSITIVE GUIDANCE TOOLS AND SERVICES

There are a growing number of parental empowerment options available to families online. These options range from services that are part of commercial online services to stand-alone software to Web-based labeling services and filtering software. Today it is safe to say that every family using the Internet has ready access to filtering sufficient to shield themselves and their children from unwanted content. In the future, we can expect even more developments in several areas: direct access to a diversity of rating services through major Web browsers, creation of additional third-party labeling services, broader use of self-labeling options, and increased availability of positive guidance services to help families find appropriate Internet content. Tools and services that filter out certain materials when accessing the Internet fall into three distinct categories:

1. *Stand-Alone Filtering Software.* Software that runs on personal computers

together with an Internet access program and blocks access to whatever type of content the parent believes inappropriate for their child. In this category are products such as SurfWatch and Cyber Patrol, as well as Cybersitter, NetNanny, and over ten others.

2. *Commercial Online Service Blocking Features.* Several major commercial online services offer blocking features integrated into their access software. These features are easy to use and part of the regular menu of options offered to users.
3. *Web-Based PICS Filtering.* An integral part of major Web browsing software gives parents the ability to set up their Web software to block access to certain material on the World Wide Web through both self-rating systems such as RSACi and SafeSurf or independent third party systems such as Net Shepherd.

A. Stand-Alone Filtering Software

Since the introduction of the first Internet filtering software in May 1995, a wide variety of software products have been offered that gives parents (or other users) the ability to block access to various categories of objectionable content. This software is easy to use and available today to 100% of Internet-connected households, often at no charge.

A variety of stand-alone, inexpensive, and easy-to-install software blocks access to material judged inappropriate for children. Most packages give parents the option of choosing what kinds of material to block such as sexually explicit material, violence, advertising, or extremist views. Each filtering software offers different choices of content categories to be filtered. For example, one product, SurfWatch, offers users the following filtering choices:

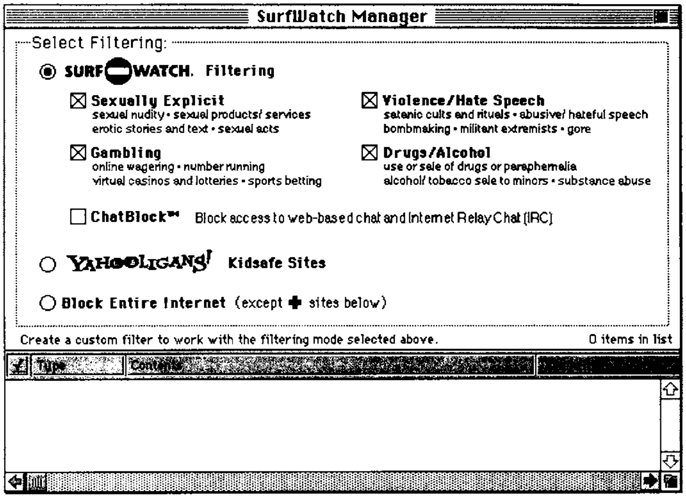

SurfWatch Manager

Select Filtering:

● SURFWATCH. Filtering

☒ **Sexually Explicit**
sexual nudity • sexual products/ services
erotic stories and text • sexual acts

☒ **Violence/Hate Speech**
satanic cults and rituals • abusive/ hateful speech
bombmaking • militant extremists • gore

☒ **Gambling**
online wagering • number running
virtual casinos and lotteries • sports betting

☒ **Drugs/Alcohol**
use or sale of drugs or paraphernalia
alcohol/ tobacco sale to minors • substance abuse

□ **ChatBlock™** Block access to web-based chat and Internet Relay Chat (IRC)

○ YAHOOLIGANS! Kidsafe Sites

○ **Block Entire Internet** (except ✚ sites below)

Create a custom filter to work with the filtering mode selected above. 0 items in list

FIG. 1. SurfWatch Setup Screen

Different software also offers additional features such as the ability to selectively unblock blocked sites, track e-mail sent and received by particular children in the household, and even monitor the amount of time spent online.

The product Cyber Patrol offers the parent this choice:

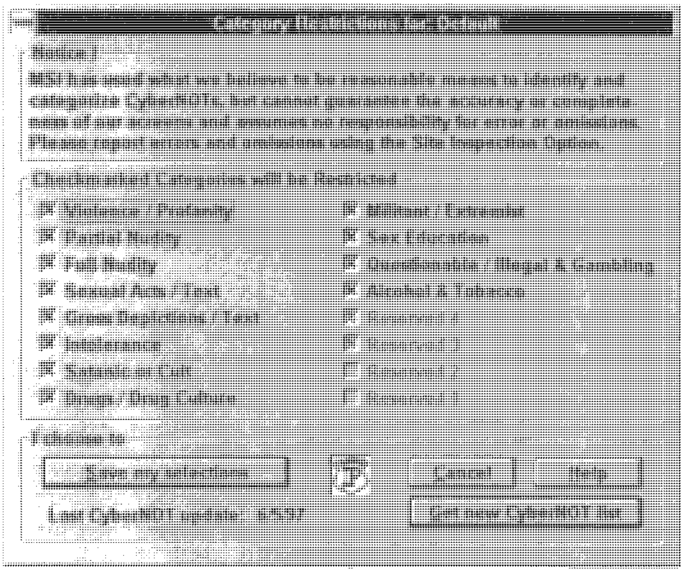

Category Restrictions for Default

Notice!

MSI has used what we believe to be reasonable means to identify and categorize CyberNOTs, but cannot guarantee the accuracy or completeness of our screens and assumes no responsibility for error or omissions. Please report errors and omissions using the Site Inspection Option.

Checkmarked Categories will be Restricted

☒ Violence / Profanity
☒ Partial Nudity
☒ Full Nudity
☒ Sexual Acts / Text
☒ Gross Depictions / Text
☒ Intolerance
☒ Satanic or Cult
☒ Drugs / Drug Culture

☒ Militant / Extremist
☒ Sex Education
☒ Questionable / Illegal & Gambling
☒ Alcohol & Tobacco
☒ Reserved 4
☒ Reserved 3
☐ Reserved 2
☐ Reserved 1

Release to:

Last CyberNOT update: 6/5/97

FIG. 2. Cyber Patrol Setup Screen

In addition to the two software products described above, over fifteen such filtering software packages exist, blocking material based on a diversity of editorial standards of the software's developers.

Many households have easy access to filtering software because it has been pre-installed on the computer they have purchased. PCs purchased through retail outlets or by mail order often come "bundled" with a variety of software products. Many consumers who have purchased a PC recently will find that it includes not only software to allow immediate Internet access, but also some form of filtering software. This is especially true of PCs sold with internal modems. Since bundled software is already loaded onto the computer's hard drive, no complicated installation is necessary and even parents who need their child's help to load software can employ blocking software as they judge necessary. Through these arrangements many millions of users around the country have ready access to filtering mechanisms.

Most filtering software vendors claim to filter based on objective criteria, but the blocking options do span the political spectrum. General interest software, such as Cyber Patrol and SurfWatch, exists along with software affiliated with conservative Christian groups such as CyberSitter, which has been endorsed and funded by Focus on the Family. Though many filtering vendors disclose their general filtering criteria, they do not reveal the actual lists of blocked sites. This lack of transparency in blocking software is a deficiency of this approach. During the relatively short lifetime of these products there have been occasions where sites are blocked inappropriately (i.e., the Center for Democracy and Technology's Web site was blocked for our discussion of bomb-making information and counterterrorism policy). For the most part, the filtering software vendors have been responsive to complaints and corrected their blocking lists based on such mistakes. However, the critical issue is that consumers be aware of such possibilities. More difficult and troublesome are situations in which filtering software publishers have blocked access to sites with information on homosexuality, feminism, or safe sex.

B. Built-In Online Service Parental Controls

Major commercial online services also offer a variety of parental controls that include site blocking, limitation on receipt of e-mail, and restriction of children's accounts to limited areas of the online service's own content. These controls are available at no cost and are easy to configure as part of establishing accounts for children.

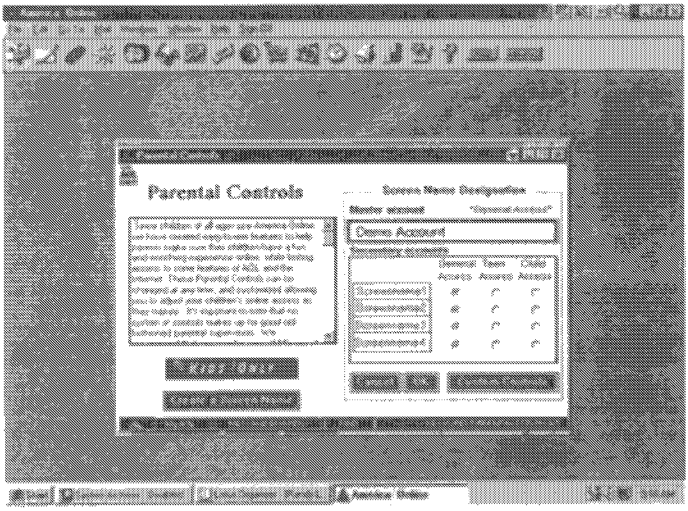

FIG. 3. AOL Parental Control Screen

Taken together with the filtering options offered by other online services, over 14,000,000 Internet households today already have easy access to filtering capability.

C. Web-Based PICS Filtering for a Diversity of Third-Party Rating

In response to a perceived need on the part of Internet parents to control children's access to inappropriate material, the Internet community undertook the development of technical standards to facilitate the growth of an unlimited variety of rating and filtering systems for the Internet. In less than two years the result is that today there are three well-established, independent rating systems accessible at no charge to all Internet families, plus a platform on which any interested party can create additional rating systems to meet the needs and values of its own community.

These three rating systems have been created using the technical tools made available by the Platform for Internet Content Selection (PICS), created through the efforts of the World Wide Web Consortium and a number of leaders in the Internet development community.[3] Since the creation of PICS and the launch of these three labeling systems, virtually all leading Internet hardware, software, and service vendors have cooperated to give families the ability to block and filter content based on PICS-formatted labels. Anyone on the Web can create third-party labels, self-label their own content, and use the labels that exist to filter Web access. Since 1996, Microsoft's Web browser, Internet Explorer, has enabled parental control through any PICS-formatted labeling service. With roughly 30% of the browser market, a substantial number of users have PICS access today. Netscape has also recently announced its commitment to implement PICS. In addition to PICS-compatible browsers, a number of stand-alone filtering products such as Cyber Patrol allow any Internet parent to filter based on PICS labels. Thus, today 100% of Internet-connected families have easy access to all PICS labeling services.

Self-Labeling. The PICS platform allows Web publishers to label their own content. Leading examples of this approach include RSACi and SafeSurf. Both RSACi and SafeSurf include standard rating vocabularies that allow Web publishers to describe the levels of sex, nudity, violence, and harsh language in a common format. To date, over 35,000 sites have rated their pages according to the RSACi labeling system and over 50,000 have rated their pages with SafeSurf. A number of major online content providers are working with RSACi to extend the reach of RSACi's ratings around the Web, including Disney, ESPN, and Playboy.

Third-Party Labeling. In addition to self-labeling, PICS also enables any individual or organization to label any content on the Web. This feature supports the creation of multiple, diverse, independent labeling on content. Recently, a Canadian group called Net Shepherd created independent labels for over 300,000 Web sites. Net Shepherd has teamed with Alta Vista to offer an Internet search service to find materials that Net Shepherd has labeled as appropriate for younger children. Using this service, children can search the Web and have returned results that include only those sites that meet rating criteria specified by parents.

III. EVALUATING THE FREEDOM OF EXPRESSION IMPLICATIONS OF VARIOUS USER EMPOWERMENT APPROACHES

Soon after the emergence of blocking and filtering technologies for the Internet, a debate ensued over the wisdom of bringing such capability to the Net. In order to assess the implications of various approaches for the Net, we can evaluate three broad categories of approaches characterizing the evolving area of user empowerment:

No Filtering—Work Toward an Unfiltered Net. This approach assumes that the evils of filtering and harmful uses of the technology outweigh any possible positive effects. In the filter-free future there may or may not be laws that censor content online.

Ubiquitous, Uniform Self-Labeling. Move toward a universal, ubiquitous content rating system: full and effective user empowerment will only come with uniform content labels that users around the world can filter based on their own values. Assuming that such labeling were to be created, it could be legally compelled, left purely voluntary for content creators, or be compelled by pressure from browsers that may block all unlabeled content and search engines that may refuse to search unrated sites.

Multiple, Third-Party Filtering. Support a diversity of third-party labeling and filtering systems as a means to provide families online with the ability to block, filter based on filters, ratings, and labels assigned by other individuals or organizations that the family chooses to trust.

This paper will judge such approaches with the following goals in mind:

Diversity. The great promise of the Internet for democracy, freedom, cultural development, and economic participation lies in its low barriers to entry for speakers, publishers, and listeners. Any solution designed to protect children must also support this fundamental value.

Trust. As a global, general purpose information and communication platform designed to serve users from different cultures and backgrounds, any approach must be an effective and trusted means of assisting parents in guiding their children's online experience.

Sustainability. Any approach must take into account the unique nature of the Internet, both as a set of technologies and a new forum for speech and market for information. While the Internet is not, and ought not be, a lawless environment, policymakers must take into account the limits of law for imposing solutions on the Internet environment.

Minimize Government Censorship of Protected Expression. Censorship of speech on the Internet is both a trespass on national free speech rights and the international right to freedom of expression and association. Solutions to the problem of children's access to inappropriate content must not result in limitations on adult's freedom of expression.

Preserving the Free Flow of Information on the Net. Finally, any approach adopted must support the free flow of information globally and enhance, or at least not harm, the democratic potential of the Internet both in the United States and around the world.

A. Option 1: The Unfiltered Net

Following the U.S. Supreme Court CDA decision, critics of Internet blocking and filtering technologies stepped up their effort to smoke out the evils of blocking and filtering technologies. Taking these concerns seriously, we must consider that the alternative to having blocking and filtering technologies is the absence of these tools. Therefore, following a look at the free expression concerns raised by these technologies, we will consider the option of working toward a wholly unfiltered Internet environment. The ACLU, one of the leading organizations to oppose the CDA in the courts, issued a strongly worded report arguing that "the dense smoke of the many rating and blocking schemes"[4] on the Net put at risk the very free speech rights won in the CDA case. Leading commentators on Internet law issues have also raised concerns about the availability of labeling, filtering, and blocking features in the Internet infrastructure. Writing in *Wired* magazine, Professor Larry Lessig declared "PICS is the devil. . . . As part of the Web's infrastructure, PICS will be an extremely versatile and robust censorship tool."[5]

While the rhetorical temperature of some filtering critics is often quite high, there are a number of substantial reasons to have concern about the impact of filtering technologies on freedom of expression. The concerns of the filtering critics include:

1. Filtering will disadvantage unpopular and controversial speech: third-party raters are more likely to filter out what the ACLU characterizes as "quirky and idiosyncratic speech" and allow only "mainstream" and commercial speech through the filters. In self-rating scenarios, the burden of rating one's own content will be high for noncommercial speakers on the Net.
2. Filtering technologies will make censorship easier for national governments: building filtering tools such as PICS into the Internet infrastructure will make it that much easier for countries who are inclined to censor the Internet.
3. Numerous problems exist with self-rating systems: aside from those already discussed, conversation and chat is hard to rate and the pressure to rate will lead speakers to homogenize their speech.
4. Third-party rating systems are not a viable alternative: the ACLU believes that the lack of development of third-party rating systems indicates that this is not a real option for the future. Without explaining why, they do not consider the over fifteen filtering software products on the market and two PICS-based third-party rating systems substantial services. Thus, the ACLU report expresses concern that third-party raters will only consider mainstream sites for rating.
5. Current rating systems fail to disclose their criteria for blocking: the ACLU and a number of other critics of filtering point out that potential users and consumers of most of the filtering products available today have a difficult time selecting among filtering options because rating and filtering criteria are not always well disclosed, nor do users have an adequate basis for deciding which filtering services to trust.

Assessment of the Unfiltered Net

Taken together, we may reasonably conclude that the filtering critics would rather see an Internet environment without filtering altogether.[6] We must now ask whether this is a desirable, viable option.

1. Effective and Trusted Solution? By definition, the unfiltered Net leaves users, including parents, with less control over the content that comes into their homes than they have today. Advocates of less filtering, or a total absence of filtering, clearly believe that families ought to rely on nontechnical means of guiding their children's Internet usage. There is general agreement that healthy, constructive use the of the Internet by children begins with parental involvement and responsibility. Yet, an unfiltered Net may leave the many parents who have chosen to use some form of filtering feeling less secure and trusting.

2. Diversity of Voices and Editorial Choices. Decreasing filtering options will reduce, not enhance, the diversity of voices online. First, both the choice to access the Net through a filter of any kind, and the act of creating a list of sites to be included or excluded in any view of the Net are each, in and of themselves, First Amendment protected activity. Second, no amount of policy debate is likely to eliminate filtering products and services from the market. Thus, those concerned about freedom of expression and diversity must work to make sure that such services serve those goals, rather than seek to eliminate them.

Creation and dissemination of blocking, filtering, and other content evaluation tools constitute important expression of editorial and critical judgments that increase the information available online. Moreover, any steps to decrease access to the tools necessary for large and small, commercial and noncommercial speakers to develop and distribute their own filtering judgments might seriously disadvantage the editorial expression of these groups, leaving only large, commercial, well-financed groups and individuals in the filtering arena.

Perhaps the technical nature of blocking, filtering, and labeling tools online lead us to forget that filtering entails important positive expression in the form of editorial judgments about other content. Even though they take a different form than in traditional media, the editorial decisions of a private entity—individual or organization—as to which content to filter in or out is a First Amendment protected activity every bit as much as is a newspaper editor's regarding which news stories to run and which letters to print or a movie reviewer's expression of opinion about a new movie.[7] The great opportunity offered by filtering in general and open specifications such as PICS is to increase the diversity of information on line. This is not an increase in primary information but in information about information, much as book reviews, movie reviews, and TV listings are not primary information yet they make important contributions to the free flow of information and ideas.

Content categorization and filtering has always been part of the Internet infrastructure and blocking software and services are unlikely to disappear. To advocate or rely on a position that hopes to eliminate blocking and filtering capabilities is to ignore key facts about the technical and market development of the Internet and online services. Filtering products and services have been in use since well before the Communications Decency Act debate. These products and services exist because at least some online users appear to want them and market forces have built upon basic open standards of the Net to provide them. To expect that these capabilities will disappear, or can be eliminated, is simply not supported by the facts of the development of the Net.

Filtering and blocking tools are based on the fundamental building blocks of the

Internet infrastructure and do not, in fact, depend on new technology specifications such as PICS. User empowerment tools, beginning with America Online's Parental Controls and products like SurfWatch and CyberPatrol, are able to block online content simply by using basic Internet addressing features—domain names, Web addresses known as URLs, numeric IP addresses that are assigned to every computer on the Internet, and newsgroup names. Specifications like PICS are important tools for making filtering easier to deploy, but blocking and filtering tools have been successfully deployed without PICS. In fact, the only foreseeable means of eliminating filtering tools would be to make them illegal. Absent such a step, which would itself raise serious freedom of expression concerns, the Internet marketplace can be expected to continue to provide user empowerment tools in response to user demand.

Shunning filtering on the Net, then, is not only an unsustainable policy option from the standpoint of the development of the Net (as described above), but it is also contrary to core First Amendment values. Blocking, filtering, and labeling tools are enabling technologies for increasing the diversity of editorial opinion on the Net. While in traditional media, the power to review a book or movie is limited to the media elite who often represent only mainstream opinion, the power to do so online can be distributed among all Internet users, thereby increasing and enriching the diversity of voices online. Since filtering is almost certain to remain part of the Internet, resisting its progress will only serve to limit its power to voices that are already powerful, a grave disservice to First Amendment values.

3. Invites Censorship: No user control invites government censorship. Decreasing the degree of user and parental control online will increase both the pressure and justification for governments to censor the Internet. For both political and legal reasons, the availability of user control is a bulwark against censorship. The United States Supreme Court decision overturning the Communications Decency Act was, without doubt, a landmark decision for freedom of expression on the Internet, inasmuch as it declared that speech online is entitled to the highest degree of First Amendment protection. The decision to overturn the statute turned on two critical factual findings: first, that the burden on speakers to limit access to minors is so great that it amounts to a total ban on speech, and second, that less restrictive alternatives such as user empowerment technologies are available to parents, so the government need not step in to protect children. From all sides of the political spectrum, policymakers agree that if parents have options that enable them to shield their own children from material they judge inappropriate based on their own family values, then the government need not step in to make the decision for the parents. All plaintiffs in the CDA challenge presented evidence and made legal arguments stressing the availability of these technologies.

Contrary to the concern that a decline in filtering will increase the call for censorship, critics of filtering worry that the availability of filtering tools will actually encourage government censorship. With such tools at hand, the reasoning goes, governments will be emboldened to act. The recent history of Internet censorship suggests, however, that the cause and effect is actually the reverse. Countries with a political tradition of censorship and repression like Singapore[8] and China began Internet censorship programs long before

PICS and other filtering options were available. Even more liberal countries, such as Australia, that are calling for the use of the PICS-based RSACi system have a long tradition of censoring film and TV. Germany, which has sought to prosecute Internet Service Providers for providing access to illegal political materials (racist and xenophobic statements), rejected proposals to employ user empowerment tools. So, countries that seek to censor the Net appear to feel able to do so without the help of blocking, filtering, and labeling tools, and are not even satisfied when such tools are offered.

Any effort that results in an actual decline in these empowerment tools would only strengthen the hand of those who would re-invoke government censorship and could even lead the courts to reevaluate their understanding of the Net. Censorship is a political decision, based on repressive impulses that extend far deeper that what is merely technically possible. The absence or presence of user empowerment tools are not likely to sway the determined censor in either direction, but they do offer many countries that option of enabling parents to protect children from what they believe to be inappropriate, without sweeping restrictions on the freedom of expression.

B. Option Two: Ubiquitous, Universal Labeling

At the other end of the policy spectrum from the unfiltered Net is the completely labeled or rated Net. Several self-labeling systems—most notably RSACi but also SafeSurf—propose to have all content on the Internet labeled by its creator according to a uniform, objective rating vocabulary. With all content thus labeled, users would be able to block access to that content that fails to meet their individual criteria. The proponents of this self-rating approach have been quite aggressive, both in the United States and around the world, in advocating their version of blocking and filtering. The Recreational Software Advisory Council has support from some of the major Internet companies in the U.S. A new International Working Group for Content Rating met with supporters in London in early October to plan a "coded means of describing Internet content, which can be used worldwide."[9] And in a recent speech, the chairman of the Australia Broadcasting Authority endorsed RSAC's self-rating approach as the best substitute for his country's traditional "subjective, state-managed classification" system, which works in traditional mass media but, according to him, will not be viable online.[10] Ubiquitous, uniform self-rating may appeal to mass media regulators, but whether it will work on the global Internet is an unanswered question. Based on the criteria used in this paper, it has several serious shortcomings.[11]

1. Limited Effectiveness and Questionable Trust: Self-rating has yet to prove effective and fails to inspire trust. The effectiveness of any objective, self-rating system turns on the feasibility of creating a single, uniform rating vocabulary that is sophisticated enough to express adequate information about Internet content to enable families around the world to make filtering decisions, while being easy enough to use and understand that Internet content creators, including individual users, can easily label their content on a real-time basis. In addition to being effective, the system must be trusted. PICS technology provides the basis to protect against fraudulent alteration of labels by unauthorized users, but will parents filtering based on labels applied by the creator of the content trust that the label is accurate?

The problems of developing an effective, objective uniform self-rating vocabulary have shown themselves most directly in the difficulty RSAC recently experienced in applying its ratings to online news sources. The RSAC system allows content providers to rate their content with numeric scores representing various levels of sex, violence, nudity, or harsh language. This classification system is in keeping with its goal of helping parents avoid online content that contains, in the parent's judgment, unacceptable levels of such material. Applying these ratings may be relatively straightforward for clearly erotic content such as pictorials from the Playboy Web site, or for obviously inoffensive material such as a geography site, however, what about the MSNBC or *Time* magazine Web site? On any given day these sites may contain reports on grisly airplane accidents or murders, discussion of sexual abuse, or reports that contain, for some reason, harsh language. Should a news story on a earthquake describing death and dismemberment receive the same rating as the Jean-Claude Van Dam fan club Web site? In a purely objective classification of the content, the answer would have to be yes, but surely many parents would want to allow access to the former but not the latter.

In reaction to the difficulty in applying RSAC's rating system to news, the organization proposed a special "news" label. This, too, produced a furor, as the debate over who is a bona fide news organization and who decides the bona fides, went public. In the end, most of the major news organizations announced that they would not self-rate their sites and RSAC put its "news" label plan on hold indefinitely.[12]

Why should parents trust the rating that content providers attach to their own material online? What incentive is there to rate honestly, rather than to rate in a way that attracts the desired surfers? RSAC's system deals with this question by providing contractual penalties for anyone who misrates using the RSAC label. Some have also proposed legislation that would provide penalties for improper or fraudulent labeling of content. It may even be that in the United States the Federal Trade Commission has the power to punish fraudulent labels under it "unfair and deceptive practices" jurisdiction. Whether enforcement is by private contract, statute, or regulation, the task of auditing labels for the thirty-five million Web pages that exist only today is certainly daunting and leaves some doubt as to whether parents can trust the labels attached by all content providers from the most scrupulous publishers to the shadiest providers of illicit material.

2. Squelching Diversity of Online Content: Uniform rating vocabulary will necessarily limit range of speech online. The limits inherent in any single ratings vocabulary go far beyond application to news and public affairs. Can a single vocabulary be developed that captures the concerns of all communities and cultures that participate in the Net, now and in the future? Keyed to Western tastes and values, the RSAC system provides tags for sex, nudity, violence, and harsh language. But what about blasphemy, racism, political extremism, excessively democratic discourse? All of these are categories of speech that some individuals or governments wish to avoid by blocking and filtering. Even if it were possible to develop a sufficiently comprehensive vocabulary, could it at the same time be made simple enough to be used by the millions of content providers (speakers) online?

Inevitably some compromises would be made to balance ease of use for publishers and those who seek to use the ratings to filter. But in these compromises will be lost the great diversity of speech that already exists on the Net today and grows daily. The result of these necessary compromises is either to limit that range of speech allowed on the Net (because unratable speech would not be permitted), or to render filtering useless for certain cultures and communities. In either case, the result is reduced diversity and homogenization of online discourse.

3. Unsustainable: Limited incentives for self-rating cast sustainability in doubt absent Internet-distorting coercion. Perhaps it is possible to design a self-rating system trusted by users and sufficiently complete to encompass all varieties of information online. Is such a system sustainable? The critical stumbling block for self-rating is that content providers, all the tens of millions of speakers who contribute content to the Internet, have no apparent reason to self-rate their sites. In the two years since RSAC was launched, roughly 43,000 sites have been rated, or 2–4% of all the million to 1,500,000 sites estimated to be on the Net today. For self-rating to work, nearly all sites must rate themselves. Proponents of self-rating have been relatively silent on how to bridge the gap from 2–4% to 100% coverage. Two options, however, have emerged: one set of infrastructure distortions, and the possibility of legal requirements to force self-rating.

First, RSAC supporters and others have suggested that critical points in the Internet infrastructure could be altered in order to force sites to rate themselves or become nearly invisible to users. Several proposals emerged that would cut off the flow of all unrated content to most users. This threat is designed to force content providers, who want their content read as widely as possible, to rate using the RSAC system. This summer, calls came for Web browsers to be configured to block access to all unrated sites. RSAC supporters even suggested that the next version of Microsoft's Internet Explorer Web browser would be shipped to customers with the default configuration set to block all sites that had not rated with RSAC. This suggestion was quickly refuted by Microsoft. At the same time, similar comments claimed that major search engines would refuse to search sites that did not have RSAC ratings. Again, the search engine operators quickly denied that such plans ever existed.

The need to distort the Internet infrastructure in order to create "incentives" for self-rating shows that this approach is entirely unsustainable and goes against the grain of the Net. Such coercive approaches would lock users into a single rating system that is quite possibly inadequate for many needs, and would force all who speak online to choose between rating every item published online or relegating their speech into certain oblivion.

4. Self-Rating Invites Censorship: Lack of incentive to rate suggests government mandates will be required. Reliance on self-rating can only be effective with legal mandates that require all content to be rated. If the above infrastructure distortions are not achieved, or not successful at forcing 100% of Internet content to be rated, the only other alternative is to make self-rating a legal requirement. In the United States, such a law would likely run afoul of First Amendment protections against "compelled speech." First Amendment doctrine holds that the government forcing someone to speak—in this case to label his or her content—is every bit as much a violation of the First Amendment as a government action that prohibits speech.[13]

U.S. legislative action alone, however, would not be sufficient. For self-rating to be an effective user empowerment strategy, nearly 100% of all content on the Internet would have to be rated, requiring an international agreement to compel self-rating according to a common ratings system. An international requirement to rate according to a single ratings code is sure to stifle the freedom of expression of users around the world, and, as was argued in point B above, unlikely to serve the needs of all the disparate cultures on the Net. Perhaps the only consolation to advocates of the free flow of information is that such a agreement is far off in the future, at best.

C. Option 3: A diversity of third-party rating services

Encouraging easy access to a multiplicity of third-party filtering systems is most likely to support the free flow of information, a diversity of ideas, and empower parents to protect their children from unwanted material online. In the last few years, over fifteen independent third-party filtering products and services have come to market. More commercial products come to market every month. The platform for third-party ratings afforded by PICS may enable the creation of an even greater diversity of third-party rating services, but this process is slow and has yet to yield encouraging results. Still, third-party services show the most promise for helping to preserve the Internet as an environment that supports the free flow of information and make it a place where parents feel empowered to protect their children as they see fit.

1. Third-Party Filtering Products Are the Most Effective Tools Available Today and Have the Potential To Build Trust Among Users. By far the most effective user empowerment tools today are the third-party blocking and filtering services for sale as stand-alone software or packaged as part of commercial online services. All of the products on the market work by creating lists of sites that are deemed inappropriate for access by children.[14] The producers of these products are continually updating these site lists so that even newly launched sites will be blocked if they meet the filtering criteria. None of the products claim to block 100% of what should be blocked, but they do claim to cover the vast majority of sites. However, because there is an active, competitive market for these tools, and because computer and consumer magazines have begun to test the relative performance of various products, market pressures are forcing competing products to do a better and better job of finding sites that should be blocked.

Not surprisingly, user empowerment technology has progressed faster than has consumer and community trust in these tools. Since trust is harder to build than software, creating trusted empowerment tools will take time. The software industry has shown that these tools are possible to develop; now trusted institutions in the community—newspapers, parent-teacher groups, librarians, religious organizations, Boy Scouts, etc.—must become involved in developing user empowerment tools that meet the needs of their constituencies and carry with them the values and trust associated with those institutions.