

Source Directed Access Control on the Internet

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

1. Abstract

This memo was developed from a deposition that I submitted as part of a challenge to the Communications Decency Act of 1996, part of the Telecommunications Reform Act of 1996. The Telecommunications Reform Act is a U.S. federal law substantially changing the regulatory structure in the United States in the telecommunications arena. The Communications Decency Act (CDA) part of this law has as its aim the desire to protect minors from some of the material carried over telecommunications networks. In particular the law requires that the sender of potentially offensive material take "effective action" to ensure that it is not presented to minors. A number of people have requested that I publish the deposition as an informational RFC since some of the information in it may be useful where descriptions of the way the Internet and its applications work could help clear up confusion in the technical feasibility of proposed content control regulations.

2. Control and oversight over the Internet

No organization or entity operates or controls the Internet. The Internet consists of tens of thousands of local networks linking millions of computers, owned by governments, public institutions, non-profit organizations, and private companies around the world. These local networks are linked together by thousands of Internet service providers which interconnect at dozens of points throughout the world. None of these entities, however, controls the Internet; each entity only controls its own computers and computer networks, and the links allowed into those computers and computer networks.

Although no organizations control the Internet, a limited number of organizations are responsible for the development of communications and operational standards and protocols used on the Internet. These standards and protocols are what allow the millions of different (and sometimes incompatible) computers worldwide to communicate with each

other. These standards and protocols are not imposed on any computer or computer network, but any computer or computer network must follow at least some of the standards and protocols to be able to communicate with other computers over the Internet.

The most significant of the organizations involved in defining these standards include the Internet Society (ISOC), the Internet Architecture Board (IAB), Internet Engineering Steering Group (IESG), and the Internet Engineering Task Force (IETF). The following summary outlines the relationship of these four organizations:

The Internet Society (ISOC) is a professional society that is concerned with the growth and evolution of the worldwide Internet, with the way in which the Internet is and can be used, and with the social, political, and technical issues which arise as a result. The ISOC Trustees are responsible for approving appointments to the IAB from among the nominees submitted by the IETF nominating committee and ratifying the IETF Standards Process.

The Internet Architecture Board (IAB) is a technical advisory group of the ISOC. It is chartered to provide oversight of the architecture of the Internet and its protocols, and to serve, in the context of the Internet standards process, as a body to which the decisions of the IESG may be appealed. The IAB is responsible for approving appointments to the IESG from among the nominees submitted by the IETF nominations committee and advising the IESG on the approval of Working Group charters.

The Internet Engineering Steering Group (IESG) is responsible for technical management of IETF activities and the Internet standards process. As a part of the ISOC, it administers the process according to the rules and procedures which have been ratified by the ISOC Trustees. The IESG is directly responsible for the actions associated with entry into and movement along the Internet "standards track," including final approval of specifications as Internet Standards.

The Internet Engineering Task Force (IETF) is a self-organized group of people who make technical and other contributions to the engineering and evolution of the Internet and its technologies. It is the principal body engaged in the development of new Internet standard specifications. The IETF is divided into eight functional areas. They are: Applications, Internet, IP: Next Generation, Network Management, Operational Requirements, Routing, Security, Transport and User Services. Each area has one or two area directors. These area directors, along with the IETF/IESG Chair, form the IESG.

In addition to these organizations, there are a variety of other formal and informal groups that develop standards and agreements about specialized or emerging areas of the Internet. For example, the World Wide Web Consortium has developed agreements and standards for the Web.

None of these organizations controls, governs, runs, or pays for the Internet. None of these organizations controls the substantive content available on the Internet. None of these organizations has the power or authority to require content providers to alter, screen, or restrict access to content on the Internet other than content that they themselves create.

Beyond the standards setting process, the only Internet functions that are centralized are the allocation of numeric addresses to networks and the registration of "domain names." Three entities around the world share responsibility for ensuring that each network and computer on the Internet has a unique 32-bit numeric "IP" address (such as 123.32.22.132), and for ensuring that all "domain names" (such as "harvard.edu") are unique. InterNIC allocates IP addresses for the Americas, and has counterparts in Europe and Asia. InterNIC allocates large blocks of IP addresses to major Internet providers, who in turn allocate smaller blocks to smaller Internet providers (who in turn allocate even smaller blocks to other providers or end users). InterNIC does not, however, reliably receive information on who receives each numeric IP address, and thus cannot provide any central database of computer addresses. In addition, a growing number of computers access the Internet indirectly through address translating devices such as application "firewalls". With these devices the IP address used by a computer on the "inside" of the firewall is translated to another IP address for transmission over the Internet. The IP address used over the Internet can be dynamically assigned from a pool of available IP addresses at the time that a communication is initiated. In this case the IP addresses used inside the firewall is not required to be globally unique and the IP addresses used over the Internet do not uniquely identify a specific computer. Neither the InterNIC nor its counterparts in Europe and Asia control the substantive content available on the Internet, nor do they have the power or authority to require content providers to alter, screen, or restrict access to content on the Internet.

3. Characteristics of Internet communications

There are a wide variety of methods of communications over the Internet, including electronic mail, mail exploders such as listserv, USENET newsgroups, Internet Relay Chat, gopher, FTP, and the World Wide Web. With each of these forms of communication, the speaker has little or no way to control or verify who receives the communication.

As detailed below, for each of these methods of communications, it is either impossible or very difficult for the speaker to restrict access to his or her communications "by requiring use of a verified credit card, debit account, adult access code, or adult personal identification number." Similarly, for each of these methods of communication, there are no feasible actions that I know of that the speaker can take that would be reasonably effective to "restrict or prevent access by minors" to the speaker's communications.

With each of these methods of communications, it is either technologically impossible or practically infeasible for the speaker to ensure that the speech is not "available" to a minor. For most of these methods--mail exploders such as listserv, USENET newsgroups, Internet Relay Chat, gopher, FTP, and the World Wide Web--there are technological obstacles to a speaker knowing about or preventing access by minors to a communication. Yet even for the basic point-to-point communication of electronic mail, there are practical and informational obstacles to a speaker ensuring that minors do not have access to a communication that might be considered "indecent" or "patently offensive" in some communities.

3.1 Point-to-Point Communications

3.1.1 Electronic Mail.

Of all of the primary methods of communication on the Internet, there is the highest likelihood that the sender of electronic mail will personally know the intended recipient (and know the intended recipient's true e-mail address), and thus the sender (i.e., the speaker or content provider) may be able to transmit potentially "indecent" or "patently offensive" content with relatively little concern that the speech might be "available" to minors.

There is significantly greater risk for the e-mail speaker who does not know the intended recipient. As a hypothetical example, if an AIDS information organization receives from an unknown individual a request for information via electronic mail, the organization has no practical or effective way to verify the identity or age of the e-mail requester.

An electronic mail address provides no authoritative information about the addressee. Addresses are often chosen by the addressees themselves, and may or may not be based on the addressees' real names. For millions of people with e-mail addresses, no additional information is available over the Internet. Where information is available (via, for example, inquiry tools such as "finger"), it is usually provided by the addressee, and thus may not be accurate (especially in a case of a minor seeking to obtain information the government has restricted to adults).

There exists no universal or even extensive "white pages" listing of e-mail addresses and corresponding names or telephone numbers. Given the rapidly expanding and global nature of the Internet, any attempt as such a listing likely will be incomplete (and likely will not contain information about the age of the e-mail addressee). Nor is there any systematic, practical, and efficient method to obtain the identity of an e-mail address holder from the organization or institution operating the addressee's computer system.

Moreover, it is relatively simple for someone to create an e-mail "alias" to send and receive mail under a different name. Thus, a given e-mail address may not even be the true e-mail address of the recipient. On some systems, for example, an individual seeking to protect his or her anonymity could easily create a temporary e-mail address for the sole purpose of requesting information from an AIDS information resource. In addition, there exist "anonymous remailers" which replace the original e-mail address on messages with a randomly chosen new one. The remailer keeps a record of the relationship between the original and the replacement name so that return mail will get forwarded to the right person. These remailers are used frequently for discussion or support groups on sensitive or controversial topics such as AIDS.

Thus, there is no reasonably effective method by which one can obtain information from existing online information sources about an e-mail address sufficient to ensure that a given address is used by an adult and not a minor.

Absent the ability to comply with the Communications Decency Act based on information from existing online information sources, an e-mail speaker's only recourse is to interrogate the intended e-mail recipient in an attempt to verify that the intended recipient is an adult. Such verification inherently and unavoidably imposes the burden of an entirely separate exchange of communications prior to sending the e-mail itself, and is likely to be unreliable if the recipient intends to deceive the speaker.

This separate preliminary communication is required because with electronic mail, there is a complete electronic and temporal "disconnect" between the sender and recipient. Electronic mail can be routed through numerous computers between the sender and the recipient, and the recipient may not "log in" to retrieve mail until days or even weeks after the sender sent the mail. Thus, at no point in time is there any direct or even indirect electronic linkage between sender and recipient that would allow the sender to interrogate the recipient prior to sending an e-mail. Thus, unavoidably, the Communications Decency Act requires that the sender incur the administrative (and in some cases financial) cost of an entirely separate exchange of communications between sender and recipient prior to the sender having sufficient information to ensure that the recipient is an adult. Even if the sender were to establish that an e-mail addressee is not a minor, the sender could not be sure that the addressee was not sharing their computer account with someone else, as is frequently done, who is a minor.

If an e-mail is part of a commercial transaction of sufficient value to justify the time and expense of obtaining payment via credit card from the e-mail addressee, an e-mail sender may be able to utilize the credit card or debit account options set out in the Communications Decency Act. At this time, however, one cannot verify a credit or debit transaction over the Internet, and thus an e-mail speaker would have to incur the expense of verifying the transaction via telephone or separate computer connection to the correct banking entity. Because of current concerns about data security on the Internet, such an e-mail credit card transaction would likely also require that the intended e-mail recipient transmit the credit card information to the e-mail sender via telephone or the postal service.

Similarly, utilizing the "adult access code" or "adult personal identification number" options set out in the statute would at this time require the creation and maintenance of a database of adult codes. While such a database would not be an insurmountable technological problem, it would require a significant amount of human clerical time to create and maintain the information. As with the credit or debit transactions, an adult code database would also likely require that information be transmitted by telephone or postal mail.

Moreover, such an adult access code would likely be very ineffective at screening access by minors. For the adult access code concept to work at all, any such code would have to be transmitted over the Internet, and thus would be vulnerable to interception and disclosure. Any sort of "information based" code--that is, a code that consists of letters and numbers transmitted in a message--could be duplicated and circulated to other users on the Internet. It is

highly likely that valid adult access codes would themselves become widely distributed on the Internet, allowing industrious minors to obtain a valid code and thus obtain access the material sought to be protected.

A somewhat more effective alternative to this type of "information based" access code would be to link such a code to the unique 32-bit numeric "IP" addresses of networks and computers on the Internet. Under this approach, "adult" information would only be transmitted to the particular computer with the "approved" IP address. For tens of millions of Internet users, however, IP addresses for a given access session are dynamically assigned at the time of the access, and those users will almost certainly utilize different IP addresses in succeeding sessions. For example, users of the major online services such as America Online (AOL) are only allocated a temporary IP address at the time they link to the service, and the AOL user will not retain that IP address in later sessions. Also, as discussed above, the use of "firewalls" can dynamically alter the apparent IP address of computers accessing the Internet. Thus, any sort of IP address-based screening system would exclude tens of millions of potential recipients, and thus would not be a viable screening option.

At bottom, short of incurring the time and expense of obtaining and charging the e-mail recipient's credit card, there are no reasonably effective methods by which an e-mail sender can verify the identity or age of an intended e-mail recipient even in a one-to-one communication to a degree of confidence sufficient to ensure compliance with the Communications Decency Act (and avoid the Act's criminal sanction).

3.2 Point-to-Multipoint Communications

The difficulties described above for point-to-point communications are magnified many times over for point-to-multipoint communications. In addition, for almost all major types of point-to-multipoint communications on the Internet, there is a technological obstacle that makes it impossible or virtually impossible for the speaker to control who receives his or her speech. For these types of communications over the Internet, reasonably effective compliance with the Communications Decency Act is impossible.

3.2.1 Mail Exploders

Essentially an extension of electronic mail allowing someone to communicate with many people by sending a single e-mail, "mail exploders" are an important means by which the Internet user can exchange ideas and information on particular topics with others

interested in the topic. "Mail exploders" is a generic term covering programs such as "listserv" and "Majordomo." These programs typically receive electronic mail messages from individual users, and automatically retransmit the message to all other users who have asked to receive postings on the particular list. In addition to listserv and Majordomo, many e-mail retrieval programs contain the option to receive messages and automatically forward the messages to other recipients on a local mailing list.

Mail exploder programs are relatively simple to establish. The leading programs such as listserv and Majordomo are available for free, and once set up can generally run unattended. There is no practical way to measure how many mailing lists have been established worldwide, but there are certainly tens of thousands of such mailing lists on a wide range of topics.

With the leading mail exploder programs, users typically can add or remove their names from the mailing list automatically, with no direct human involvement. To subscribe to a mailing list, a user transmits an e-mail to the automated list program. For example, to subscribe to the "Cyber-Rights" mailing list (relating to censorship and other legal issues on the Internet) one sends e-mail addressed to "listserv@cpsr.org" and includes as the first line of the body of the message the words "subscribe cyber-rights name" (inserting a person's name in the appropriate place). In this example, the listserv program operated on the cpsr.org computer would automatically add the new subscriber's e-mail address to the mailing list. The name inserted is under the control of the person subscribing, and thus may not be the actual name of the subscriber.

A speaker can post to a mailing list by transmitting an e-mail message to a particular address for the mailing list. For example, to post a message to the "Cyber-Rights" mailing list, one sends the message in an e-mail addressed to "cyber-rights@cpsr.org". Some mailing lists are "moderated," and messages are forwarded to a human moderator who, in turn, forwards messages that moderator approves of to the whole list. Many mailing lists, however, are unmoderated and postings directed to the appropriate mail exploder programs are automatically distributed to all users on the mailing list. Because of the time required to review proposed postings and the large number of people posting messages, most mailing lists are not moderated.

An individual speaker posting to a mail exploder mailing list cannot control who has subscribed to the particular list. In many cases, the poster cannot even find out the e-mail address of who has subscribed to the list. A speaker posting a message to a list thus has no way to screen or control who receives the message. Even if the mailing list is "moderated," an individual posting to the list still cannot control who receives the posting.

Moreover, the difficulty in knowing (and the impossibility of controlling) who will receive a posting to a mailing list is compounded by the fact that it is possible that mail exploder lists can themselves be entered as a subscriber to a mailing list. Thus, one of the "subscribers" to a mailing list may in fact be another mail exploder program that re-explodes any messages transmitted using the first mailing list. Thus, a message sent to the first mailing list may end up being distributed to many entirely separate mailing lists as well.

Based on the current operations and standards of the Internet, it would be impossible for someone posting to a listserv to screen recipients to ensure the recipients were over 17 years of age. Short of not speaking at all, I know of no actions available to a speaker today that would be reasonably effective at preventing minors from having access to messages posted to mail exploder programs. Requiring such screening for any messages that might be "indecent" or "patently offensive" to a minor would have the effect of banning such messages from this type of mailing list program.

Even if one could obtain a listing of the e-mail addresses that have subscribed to a mailing list, one would then be faced with the same obstacles described above that face a point-to-point e-mail sender. Instead of obtaining a credit card or adult access code from a single intended recipient, however, a posted to a mailing list may have to obtain such codes from a thousand potential recipients, including new mailing list subscribers who may have only subscribed moments before the poster wants to post a message. As noted above, complying with the Communications Decency Act for a single e-mail would be very difficult. Complying with the Act for a single mailing list posting with any reasonable level of effectiveness is impossible.

3.2.2 USENET Newsgroups.

One of the most popular forms of communication on the Internet is the USENET newsgroup. USENET newsgroups are similar in objective to mail exploder mailing lists--to be able to communicate easily with others who share an interest in a particular topic--but messages are conveyed across the Internet in a very different manner.

USENET newsgroups are distributed message databases that allow discussions and exchanges on particular topics. USENET newsgroups are disseminated using ad hoc, peer-to-peer connections between 200,000 or more computers (called USENET "servers") around the world. There are newsgroups on more than twenty thousand different subjects. Collectively, almost 100,000 new messages (or "articles") are posted to newsgroups each day. Some newsgroups are "moderated" but most are open access.

For unmoderated newsgroups, when an individual user with access to a USENET server posts a message to a newsgroup, the message is automatically forwarded to adjacent USENET servers that furnish access to the newsgroup, and it is then propagated to the servers adjacent to those servers, etc. The messages are temporarily stored on each receiving server, where they are available for review and response by individual users. The messages are automatically and periodically purged from each system after a configurable amount of time to make room for new messages. Responses to messages--like the original messages--are automatically distributed to all other computers receiving the newsgroup. The dissemination of messages to USENET servers around the world is an automated process that does not require direct human intervention or review.

An individual who posts a message to a newsgroup has no ability to monitor or control who reads the posted message. When an individual posts a message, she transmits it to a particular newsgroup located on her local USENET server. The local service then automatically routes the message to other servers (or in some cases to a moderator), which in turn allow the users of those servers to read the message. The poster has no control over the handling of her message by the USENET servers worldwide that receive newsgroups. Each individual server is configured by its local manager to determine which newsgroups it will accept. There is no mechanism to permit distribution based on characteristics of the individual messages within a newsgroup.

The impossibility of the speaker controlling the message distribution is made even more clear by the fact that new computers and computer networks can join the USENET news distribution system at any time. To obtain newsgroups, the operator of a new computer or computer network need only reach agreement with a neighboring computer that already receives the newsgroups. Speakers around the world do not learn that the new computer had joined the distribution system. Thus, just as a speaker cannot know or control who receives a message, the speaker does not even know how many or which computers might receive a given newsgroup.

For moderated newsgroups, all messages to the newsgroup are forwarded to an individual who can screen them for relevance to the topics under discussion. The screening process, however, does not increase the ability of the original speaker to control who receives a given message. A newsgroup moderator has as little control as the original speaker over who receives a message posted to the newsgroup.

Based on the current operations and standards of the Internet, it would be impossible for someone posting to a USENET newsgroup to screen recipients to ensure that the recipients were over 17 years of age. Short of not speaking at all, I know of no actions available to a speaker today that would be reasonably effective at preventing minors from having access to USENET newsgroup messages. Requiring such screening for any messages that might be "indecent" or "patently offensive" to a minor would have the effect of banning such messages from USENET newsgroups.

A speaker also has no means by which he or she could require listeners to provide a credit card, debit account, adult access code, or adult personal identification number. Each individual USENET server controls access to the newsgroups on that server, and a speaker has no ability to force a server operator to take any particular action. The message is out of the speaker's hands from the moment the message is posted.

Moreover, even if one hypothesized a system under which a newsgroup server would withhold access to a message until the speaker received a credit card, debit account, adult access code, or adult personal identification number from the listener, there would be no feasible way for the speaker to receive such a number. Because a listener may retrieve a message from a newsgroup days after the speaker posted the message, such a hypothetical system would require the speaker either to remain at his or her computer 24 hours a day for as many as ten days after posting the message, or to finance, develop, and maintain an automated system to receive and validate access numbers. All of this effort would be required for the speaker to post even a single potentially "patently offensive" message to a single newsgroup.

Moreover, even if such a hypothetical system did exist and a speaker were willing to remain available 24 hours a day (or operate a costly automated system) in order to receive access numbers, not all computers that receive USENET newsgroups could reasonably transmit such access numbers. Some computers that receive newsgroups do so only by a once-a-day telephone connection to another newsgroup server. Some of these computers do not have any other type of Internet connection, and indeed some computers that receive USENET newsgroups do not even utilize the TCP/IP communications protocol that is required for direct or real time communications on the

Internet. These computers would have no means by which a prospective listener's access code could be communicated back to a speaker.

It is my opinion that if this hypothetical access system ever were created, it would be so burdensome as to effectively ban from USENET newsgroups messages that might be "indecent" or "patently offensive." Moreover, the communications standards and protocols that would allow such a hypothetical access system have not as of today been developed, and no Internet standards setting body of which I am aware is currently developing such standards and protocols. Specifically, such a hypothetical access system is not part of the "next generation" Internet Protocol that I helped to develop.

3.2.3 Internet Relay Chat.

Another method of communication on the Internet is called "Internet Relay Chat" (or IRC). IRC allows for real time communication between two or more Internet users. IRC is analogous to a telephone party line, using a computer and keyboard rather than a telephone. With IRC, however, at anyone time there are thousands of different party lines available, in which collectively tens of thousands of users are engaging in discussions, debates, and conversations on a huge range of subjects. Moreover, an individual can create a new party line to discuss a different topic at any time. While many discussions on IRC are little more than social conversations between the participants, there are often conversations on important issues and topics. Although I have not personally operated an IRC server in my career, I am familiar enough with the operations of IRC servers to be able to identify the obstacles that a speaker would encounter attempting to identify other participants and to verify that those participants were not minors.

There exists a network of dozens of IRC servers across the world. To speak through IRC, a speaker connects to one of these servers and selects the topic the speaker wishes to "join." Within a particular topic (once a speaker joins a topic), all speakers on that topic can see and read everything that everyone else transmits. As a practical matter, there is no way for each person who joins a discussion to interrogate all other participants (sometimes dozens of participants) as to their identity and age. Because people join or drop out of discussions on a rolling basis, the discussion line would be overwhelmed with messages attempting to verify the identity of the participants.

Also as a practical matter, there is no way that an individual speaker or an individual IRC server operator could enforce an "adults only" rule for a selection of the discussion topics. Dozens of IRC servers are interconnected globally so that people across the world

can talk to each other. Thus, a speaker connected to an IRC server in the United States can speak directly to a listener in Asia or Europe. There is no practical way that a speaker in the United States can be reasonably certain that a given IRC discussion is in fact "adults only."

Nor can a speaker, prior to or at the time of joining an IRC discussion, ascertain with any confidence the identity of the other participants in the discussion. Individual participants in an IRC conversation are able to participate anonymously by using a pseudonym. A new speaker joining the conversation can see a list of pseudonyms of other participants, but has no possibly way of determining the real identify (or even the real e-mail address) of the individuals behind each pseudonym.

Based on the current operations and standards of the Internet, it would be impossible for someone participating in a IRC discussion to screen recipients with a level of certainty needed to ensure the recipients were over 17 years of age. Short of not speaking at all, I know of no actions available to a speaker today that would be reasonably effective at preventing minors from having access to speech in an IRC discussion. Requiring such screening of recipients by the speakers for any IRC discussions that might be "indecent" or "patently offensive" to a minor would have the effect of banning such discussions.

4.0 Information Retrieval Systems

With FTP (or File Transfer Protocol), gopher, and the World Wide Web, the Internet is a vast resource for information made available to users around the world. All three methods (FTP, gopher, and the Web) are specifically geared toward allowing thousands or millions of users worldwide to access content on the Internet, and none are specifically designed to limit access based on criteria such as the age of the Internet user. Currently much of this information is offered for free access.

4.1 Anonymous FTP

"Anonymous FTP" is a basic method by which a content provider can make content available to users on the Internet. FTP is a protocol that allows the efficient and error free transfer of files from one computer to another. To make content available via FTP, a content provider establishes an "Anonymous FTP server" capable of receiving FTP requests from remote users. This approach is called "anonymous" because when a remote user connects to an FTP server, the remote user enters the word "anonymous" in response to the server's request for a user name. By convention, the remote user is requested to enter his

or her e-mail address when prompted for a "password." The user is then given access to a restricted portion of the server disk and to the files in that area. Even though the user may have entered their e-mail address in response to the password prompt, there is no effective validation or screening is possible using the FTP server software that is currently available. Using currently available FTP software, a content provider has no way to screen access by "anonymous" users that may be minors. Even if a content provider could determine the age of a particular remote user, the currently available FTP software cannot be set to limit the user's access to non-"adult" file areas.

FTP server software can allow non-"anonymous" users to access the FTP server, and in that mode can require the users to have individual passwords that are verified against a pre-existing list of passwords. There are two major problems, however, that prevent this type of non-"anonymous" FTP access from being used to allow broad access to information over the Internet (as anonymous FTP can allow). First, with current server software each non-"anonymous" FTP user must be given an account on the server computer, creating a significant administrative burden and resource drain. If more than a limited number of users want access to the FTP system, the requirement of separate accounts would quickly overwhelm the capacity of the server to manage the accounts--the FTP server software was not designed to manage thousands or millions of different user/password combinations. Second, under existing FTP server software, each of these named users would have complete access to the server file system, not a restricted area like the anonymous FTP function supports. This would create a significant security problem. For these two reasons, as a practical matter FTP cannot be used to give broad access to content except via the anonymous FTP option (which, as noted above, does not allow for screening or blocking of minors).

As discussed below with regard to the World Wide Web, even if someone re-designed the currently available FTP server software to allow the screening of minors, the administrative burden of such screening would in many cases overwhelm the resources of the content provider.

Based on the current operations and standards of the Internet, it is not possible or practically feasible for someone operating an anonymous FTP file server to screen recipients with a level of certainty needed to ensure the recipients were over 17 years of age. Short of not operating an anonymous FTP server at all, I know of no actions available to a content provider today that would be reasonably effective at preventing minors from having access to "adult" files on the FTP server. Requiring such screening by anonymous FTP server operators to prevent minors from accessing FTP files that might be "indecent" or "patently offensive" to a minor would have the effect of banning such anonymous FTP access.

4.2 Gopher.

The gopher program is similar to FTP in that it allows for basic transfer of files from one computer to another, but it is also a precursor to the World Wide Web in that it allows a user to seamlessly jump from one gopher file server to another in order to locate the desired information. The development of gopher and the linking of gopher servers around the world dramatically improved the ability of Internet users to locate information across the Internet.

Although in many ways an improvement over FTP, gopher is simpler than FTP in that users need not enter any username or password to gain access to files stored on the gopher server. Under currently available gopher server software, a content provider has no built-in ability to screen users. Thus a content provider could not prevent minors from retrieving "adult" files.

As discussed below with regard to the World Wide Web, even if the gopher server software allowed the screening of minors, the administrative burden of such screening would in many cases overwhelm the resources of the content provider.

Based on the current operations and standards of the Internet, it is not possible for someone operating a gopher file server to screen recipients with a level of certainty needed to ensure the recipients were over 17 years of age. Short of not operating a gopher server at all, I know of no actions available to a content provider today that would be reasonably effective at preventing minors from having access to "adult" files on a gopher server. Requiring such screening of users by gopher server operators to prevent minors from accessing files that might be "indecent" or "patently offensive" to a minor would have the effect of banning gopher servers wherever there is any such material.

4.3 World Wide Web (WWW).

Fast becoming the most well known method of communicating on the Internet, the "World Wide Web" offers users the easy ability to locate and view a vast array of content on the Internet. The Web uses a "hypertext" formatting language called hypertext markup language (HTML), and Web "browsers" can display HTML documents containing text, images, and sound. Any HTML document can include links to other types of information or resources anywhere in the world, so that while viewing an HTML document that, for example, describes resources available on the Internet, an individual can "click" using a computer mouse on the description of the resource and be immediately connected to the resource itself. Such "hyperlinks" allow information to be accessed and organized in very flexible ways, and allow individuals to locate and efficiently view related information even if the information is stored on numerous computers all around the world.

Unlike with USENET newsgroups, mail exploders, FTP, and gopher, an operator of a World Wide Web server does have some ability to interrogate a user of a Web site on the server, and thus has some ability to screen out users. An HTML document can include a fill-in-the-blank "form" to request information from a visitor to a Web site, and this information can be transmitted back to the Web server. The information received can then be processed by a computer program (usually a "Common Gateway Interface," or "CGI," script), and based on the results of that computer program the Web server could grant or deny access to a particular Web page. Thus, it is possible for some (but not all, as discussed below) World Wide Web sites to be designed to "screen" visitors to ensure that they are adults.

The primary barrier to such screening is the administrative burden of creating and maintaining the screening system. For an individual Web site to create a software system capable of screening thousands of visitors a day, determining (to the extent possible) whether a visitor is an adult or a minor, and maintaining a database to allow subsequent access to the Web site would require a significant on-going effort. Moreover, as discussed above with regard to electronic mail, the task of actually establishing a Web visitor's identity or "verifying" a credit card would require a significant investment of administrative and clerical time. As there is no effective method to establish identity over the Internet, nor is there currently a method to verify credit card numbers over the Internet (and given the current cost of credit card verifications done by other means), this type of identification process is only practical for a commercial entity that is charging for access to the Web information.

Beyond the major administrative burden that would be required for a Web site host to comply with the Communications Decency Act, there are two additional problems presented by the Act. First, many Web publishers cannot utilize computer programs such as CGI scripts to process input from a Web visitor. For example, I have been informed that the major online services such as America Online and Compuserve do not allow their customers to run CGI scripts or other processes that could be a significant drain on the online services' computers as well as a potential security risk. Thus, for this category of Web publisher, the Communications Decency Act works as a ban on any arguably "indecent" or "patently offensive" speech. It is impossible for this category of Web publisher to control access to their Web sites.

Moreover, even for Web publishers who can use CGI scripts to screen access, the existence of Web page caching on the Internet can make such screening ineffective. "Caching" refers to a method to speed up access to Internet resources. Caching is often used at one or both ends of, for example, a transatlantic or transpacific cable that carries Internet communications. An example of caching might occur when a Internet user in Europe requests access to a World Wide Web page located in the United States. The request travels by transatlantic cable to the United States, and the Web page is transmitted back across the ocean to Europe (and ultimately to the user who requested access). But, the operator of the transatlantic cable will place the Web page in a storage "cache" located on the European side of the cable. Then, if a second Internet user in Europe requests the same Web page, the operator of the transatlantic cable will intercept the request and provide the page from its "cache" (thereby reducing traffic on the transatlantic cable). This type of caching typically occurs without the awareness of the requesting user. Moreover, in this scenario, the original content provider is not even aware that the second user requested the Web page--and the original content provider has no opportunity to screen the access by the second user. Nevertheless, the original content provider risks prosecution if the content is "adult" content and the second requester is a minor. The use of caching web servers is rapidly increasing within the United States (mostly to help moderate the all too rapid growth in Internet traffic), and thus can affect entirely domestic communications. For example, a growing number of universities use caching web servers to reduce the usage of the link to their Internet service provider. In light of this type of caching, efforts to screen access to Web pages can only at best be partially effective.

In light of the existence of Web page caching on the Internet, it would be extremely difficult if not impossible to for someone operating a World Wide Web server to ensure that no minors received "adult" content.

Moreover, for those Web page publishers who lack access to CGI scripts, there is no possible way for them to screen recipients to ensure that all recipients are over 17 years of age. For these content providers, short of not supporting World Wide Web access to their materials, I know of no actions available to them that would be reasonably effective at preventing minors from having access to "adult" files on a World Wide Web server. Requiring such screening by these Web publishers to prevent minors from accessing files that might be "indecent" or "patently offensive" to a minor would have the effect of banning their speech on the World Wide Web.

The Web page caching described above contributes to the difficulty of determining with specificity the number of visitors to a particular Web site. Some Web servers can count how many different Web clients, some of which could be caching Web servers, requested access to a Web site. Some Web servers can also count how many "hits"--or separate file accesses--were made on a particular Web site (a single access to a Web page that contains a images or graphic icons would likely be registered as more than one "hit"). With caching, the actual number of users that retrieved information that originated on a particular Web server is likely to be greater than the number of "hits" recorded for the server.

5.0 Client-end Blocking

As detailed above, for many important methods of communication on the Internet, the senders--the content providers--have no ability to ensure that their messages are only available to adults. It is also not possible for a Internet service provider or large institutional provider of access to the Internet (such as a university) to screen out all or even most content that could be deemed "indecent" or "patently offensive" (to the extent those terms can be understood at all). A large institution could at least theoretically screen a portion of the communications over the Internet, scanning for example for "indecent" words, but not pictures. Such a screening program capable of screening a high volume of Internet traffic at the point of its entry into the institution would require an investment of computing resources of as much as one million dollars per major Internet information conduit. In addition it would be quit difficult to configure such a system to only control the content for those users that are under-age recipients, since in many cases the information would be going to a server within the university where many users, under-age and not, would have access to it.

Based on my experience and knowledge of the Internet, I believe that the most effective way to monitor, screen, or control the full range of information transmitted over the Internet to block undesired content is at the client end--that is, by using software installed in the individual user's computer. Such software could block certain forms of incoming transmissions by using content descriptive tags in the messages, or could use content ratings developed by third parties to select what can and cannot be retrieved for display on a user's computer.

6.0 Tagging Material

I am informed that the government in this action may advocate the use of special tags or flags in electronic mail messages, USENET newsgroup postings, and World Wide Web HTML documents to indicate "adult" material. To my knowledge, no Internet access software or World Wide Web browsers are currently configurable to block material with such tags. Thus, the headers and flags the government may advocate is currently an ineffective means to ensure the blocking of access by minors to "adult" material. Even in a predictable future where there are defined standards for such tags and there are readably available browsers that are configurable to make use of those tags, a content provider--e.g., a listserv or Newsgroup poster or a Web page author--will have little power to ensure that the client software used to receive the postings was in all cases properly configured to recognize these tags and to block access to the posting when required. Thus I feel that the tagging that may be proposed by the government would in fact not be "effective" in ensuring that the poster's speech would not be "available to a person under 18 years of age," as the Communications Decency Act requires. Although I strongly support both voluntary self-rating and third-party rating (as described in the preceding paragraph), I do not feel that the use of tags of this type would satisfy the speaker's obligation to take effective actions to ensure that "patently offensive" material would not be "available" to minors. Furthermore, since it is impossible to embed such flags or headers in many of the documents currently made available by anonymous FTP, gopher and the World Wide Web without rendering the files useless (executable programs for example), any government proposal to require the use of tags to indicate "adult" material would not allow the continued use of those methods of communication for speech that might be deemed "indecent" or "patently offensive."

With the exception of electronic mail and e-mail exploders all of the methods of Internet communications discussed above require an affirmative action by the listener before the communication takes place. A listener must take specific action to receive communications from USENET newsgroups, Internet Relay Chat, gopher,

FTP, and the World Wide Web. In general this is also true for e-mail exploders except in the case where a third party subscribes the user to the exploder list. These communications over the Internet do not "invade" a person's home or appear on a person's computer screen unbidden. Instead, a person must almost always take specific affirmative steps to receive information over the Internet.

7.0 Acknowledgment

I owe a great deal of thanks to John Morris of Jenner and Block, one of the law firms involved in the CDA challenge. Without his extensive help this document would not exist, or if it did, it would be even more scattered.

8.0 Security Considerations

To be actually able to do the type of content access control that the CDA envisions would require a secure Internet infrastructure along with secure ways to determine the minor status of potential recipients around the world. Developing such a system is outside of the scope of this document.

9.0 Author's Address

Scott Bradner
Harvard University
1350 Mass Ave.
Cambridge MA 02138 USA

Phone: +1 617 495 3864
EMail: sob@harvard.edu

