Enhancing Child Safety & Online Technologies:

FINAL REPORT OF THE

INTERNET SAFETY TECHNICAL TASK FORCE

To the Multi-State Working Group on Social Networking

of State Attorneys General of the United States

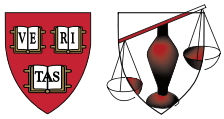

Berkman

The Berkman Center for Internet & Society

at Harvard University

DECEMBER 31, 2008

APPENDIX D:

Technology Advisory Board Report

EXECUTIVE SUMMARY

The Technology Advisory Board (TAB) solicited, evaluated, reviewed, 40 written submissions of technologies and drew conclusions from these submissions about the state of online safety technology for minors in a formal process described in detail in this document. The primary task was to assess whether and how the submitted technologies could be useful in the context of enhancing online safety for minors.

In sum, the TAB review of the submitted technologies leaves us in a state of cautious optimism, with many submissions showing promise. The children's online safety industry is evolving, and many of the technologies we reviewed were point solutions rather than broad attempts to address the children's safety online as a whole. There is, however, a great deal of innovation in this arena as well as passionate commitment to finding workable, reasonable solutions from companies both large and small. Thus, the TAB emerged from its review process encouraged by the creativity and productivity apparent in this field.

By the end of the review process, the TAB ultimately determined that no single technology reviewed could solve every aspect of online safety for minors, or even one aspect of it one hundred percent of the time. But clearly there is a role for technology in addressing this issue both now and in the future, and most likely, various technologies could be leveraged together to address the challenges in this arena.

Some critics may object to the use of technology as a solution, given the risk of failure and lack of total certainty around performance. However, the TAB believes that although it is indeed true that even the cleverest, most robust technology can be circumvented, this does not necessarily mean that technology should not be deployed at all. It simply means that – even with deployment of the best tools and technologies available to jumpstart the process of enhancing safety for minors online – there is no substitute for a parent, caregiver, or other responsible adult actively guiding and supporting a child in safe Internet usage. Likewise, education is an essential part of the puzzle. Even the best technology or technologies should be only part of a broader solution to keeping minors safer online.

As a corollary, the TAB also recommends that further evaluative work be conducted on any technology – whether or not it was among those reviewed in this process – prior to broadly recommending its use, given the potential for new risks and significant unintended consequences. The benefits of each reviewed solution need further exploration and balancing against monetary costs, possible privacy and security concerns about user information, international implications and applicability, as well as other issues. Additionally, determining which technology or set of technologies will work best for a particular child, family, school, community, or any other context in which the safety of minors on the Internet is an immediate concern will always be a highly individualized decision. It is also not a decision that can reasonably be made without a great deal of familiarity with the situation in which a technology solution would function.

Listed here, and discussed in greater detail later in this document, are the specific conclusions and recommendations generated by the TAB's review process:

- ***Technology can play a role but cannot be the sole input to improved safety for minors online.***

- ***The most effective technology solution is likely to be a combination of technologies.***

- ***Any and every technology solution has its limitations.***

- ***Youth online safety measures must be balanced against concerns for the privacy and security of user information, especially information on minors.***

- ***For maximum impact, client-side-focused technologies should be priced to enable all would-be users to purchase and deploy them.***

- ***A common standard for sharing information among safety technologies would be useful.***

- ***Developing standard metrics for youth online safety solutions would be useful.***

INTRODUCTION

The scope of the Technology Advisory Board's mandate in conducting its work for the Task Force was to review all submissions that it received detailing technology solutions for improved online safety for minors. To conduct its work, the TAB was limited to the written submission itself, written responses to several questions, and public presentations made to the Task Force. The TAB did not perform uniform, independent technical evaluations of the technologies submitted.

Based on these inputs, we discuss broad sets of technology categories that address several online safety concerns involving minors. For each category, we summarize how the technologies address one or more aspects of online safety for minors, the potential benefits of the approach, and hurdles that it must overcome to be effective.

PROCESS AND METHODOLOGY

Technology Advisory Board Members and Observers

The Technology Advisory Board comprised two teams: the TAB Members and the TAB Observers. The TAB Members team was charged with the formal review of the technology submissions from third parties. The TAB Observers team was asked to formally comment on any or all of the submissions if they so chose, but, due to potential conflicts of interest, their comments were neither part of the formal technology reviews nor part of the recommendation process to select presenters for the Berkman ISTTF Public Meeting.

The objective in building the TAB teams was to enlist people who had deep technology backgrounds, domain expertise in a field related to the Task Force's work, and a demonstrated professional interest in relevant subject areas. In addition to technology professionals, we also added representatives from other related fields to serve as Observers, so that we could draw on their areas of expertise. An additional distinction between Members and Observers is that Observers might have conflicts of interest with the review work.

Nominations for both Members and Observers came from the Task Force itself, the Task Force team at the Berkman Center, other Berkman Center affiliates, and other TAB Members and Observers. The nominations were vetted through the Berkman Task Force team, an interview and investigation of possible conflicts of interest were conducted, and then the Berkman Task Force team made the decision whether to invite the nominee to join the TAB Members or Observers team.

TAB Members *(Complete biographies are included as Exhibit 1)*
Ben Adida, Harvard Medical School, Harvard University
Scott Bradner, Harvard University
Laura DeBonis, Berkman Center, Harvard University
Hany Farid, Dartmouth
Lee Hollaar, University of Utah

Todd Inskeep, Bank of America
Brian Levine, University of Massachusetts Amherst
Adi Mcabian, Twistbox
RL Morgan, University of Washington
Lam Nguyen, Stroz Friedberg, LLC
Jeff Schiller, MIT
Danny Weitzner, MIT

TAB Observers *(Complete biographies are included as Exhibit 1)*
Rachna Dhamija, Usable Security Systems
Evie Kintzer, WGBH
Al Marcella, Webster University
John Morris, Center for Democracy and Technology
Teresa Piliouras, Polytechnic University
Greg Rattray, Delta-Risk
Jeff Schmidt, Consultant
John Shehan, National Center for Missing and Exploited Children

Soliciting, collecting, and evaluating submissions

Soliciting
The process for soliciting submissions was as follows: the TAB created a Submission Template that encompassed the various questions anticipated for any single technology. Primary areas for response included: (1) functional goals that a technology attempted to address; (2) technological detail about the technology itself; and (3) financial and other business information about the technology to inform the assessment of viability and functionality. On July 1, 2008, the Submission Template was posted to the Task Force's webpage on the Berkman website and made broadly available for download by any company, individual, or other entity that wished to submit, in writing only, a technology for consideration. (The Submission Template is included as Exhibit 2.)

The public was made aware of the Template through a Berkman Center press release and by tapping into various networks, including networks of the Berkman Center staff and affiliates, the TAB, and the members of the Task Force.

The deadline for submission was July 21, 2008, approximately three weeks after the Template was made publicly available.

Collecting
In total, the TAB received 40 written submissions from 38 companies. (An additional submission involving a registry for minors' email addresses was withdrawn from consideration by the submitting company.) Submitters were asked to include with their submission a statement indicating that they understood the Intellectual Property policy regarding submission to the Task Force. (The Intellectual Property policy is included as Exhibit 3.)

Evaluating

The TAB designed and the Berkman Task Force team approved an evaluation process that closely followed the model of other scientific reviews; in particular, that of the National Science Foundation review. Three to five TAB Members reviewed each document. Following initial discussions of the document, questions were sent to the submitting companies to clarify our understanding of their submission. All companies responded to the follow-up questions. Final review discussions considered the answers to the follow-up questions as well as all TAB Observer Comments. After final review discussions, recommendations were made to the Berkman Task Force team for companies to present at the Public Meeting of the Task Force. Many criteria were involved in determining whether a submitting company was asked to present at the Public Meeting. A recommendation to have a company present was not an endorsement of the technology. Rather, the TAB sought to have a variety of technologies, companies, and approaches discussed; to show the range of ideas extant; to inform the public; and to help foster meaningful dialogue about solutions to improving online safety for minors.

Evaluation questions were circulated to the Members of the TAB prior to their initial reading of the submissions. Members were asked to use the questions to frame their thinking in preparation for review discussions. The evaluation questions included:

- What functional requirements are met by the submission?
- What is the overall approach?
- Who is the target audience (e.g., youth under 13, teens, parents)?
- What is the target system (e.g., social networking sites, cell phones, ISPs)?
- What underlying assumptions does the proposal make? Are they reasonable?
- Does the approach require education and/or parental involvement?
- What are the strengths and weakness of the approach?
- How well does the product actually address its targeted function?
- What are unintended consequences caused by use of the product?
- Under what circumstances would the product fail? How often?
- What are the consequences of product failure?
- What other trade-offs does the product present?
- How does product work internationally?
- How does product work with different business models?

To facilitate the review process, the TAB created a list of functional goals related to online safety for minors that a technology might address. This list was included as one of the sections in the Submission Template and each company self-identified one or more of eight functional goals for the technology. For the purposes of review, the different solutions submitted were clustered according to these functional goals:

- Limit harmful contact between adults and minors
- Limit harmful contact between minors
- Limit/prevent minors from accessing inappropriate content on the Internet

- Limit/prevent minors from creating inappropriate content on the Internet
- Limit the availability of illegal content on the Internet
- Prevent minors from accessing particular sites without parental consent
- Prevent harassment, unwanted solicitation, and bullying of minors on the Internet
- Other – please specify

In addition to self-identification of functional goals, after review, the TAB also assigned one of five categories to each of the 40 technology submissions. Occasionally more than one category applied to a technology; in such situations, the primary category was the one with which the technology was associated.

The technology categories, with number of submissions received in parentheses, are:

1. Age Verification/Identity Authentication (17)
2. Filtering/Monitoring/Auditing (13)
3. Text Analysis (5)
4. Biometrics (1) (+2 with biometrics as secondary category)
5. Other (4)

A list of all submissions in alphabetical order is included as Exhibit 4. The submissions themselves as well as TAB Observer Comments are available on the Task Force's website.

ANALYSIS

Among past efforts to survey the landscape of children's online safety technologies, the 2000 COPA Commission report is one of the most relevant. For purposes of brevity we do not summarize or cite COPA or other previous reviews of technologies in our analyses below. The TAB does recognize, however, the importance of previous work in this area. Our intention with this review process is to complement previous work and not to supersede it.

Below we summarize the categories of technology solutions presented, comment occasionally on particular technologies, and discuss overall the strengths and weaknesses of each category in application to enhancing online safety for minors. In each category, some solutions help a little bit and some help more extensively. The same is true of each category of technology. We considered each proposal from the perspective of what the potential outcome would be if it were fully implemented and widely adopted. Again, no one solution can solve the entire youth online safety problem, but it was clear from the submissions that there has been excellent traction achieved.

Age Verification/Identity Authentication
Category Description

Age verification technologies seek primarily to verify the age of adults and children, while and identity technologies seek to verify individual identities. The primary goal of these technologies is to utilize age as a mechanism for limiting inappropriate contact between children and adults as well as preventing access by minors to inappropriate content. Although some technologies attempt to verify age/identity remotely, other technologies rely on a trusted third party for verification (e.g., schools, notaries, or government agencies). A submission in this category involving a registry of minors' email addresses was withdrawn from consideration by the submitting company.

We separated technology submissions in this area into four subcategories:

1. Comparison against records collected from public databases. Many records, both public and private, are available about adults, including information from credit reports, criminal history, and real estate transfers. These disparate records can be aggregated into a portfolio of data about an individual. This information can then be used, among other applications, as a basis to present challenge questions to individuals to ensure their correct identification.

2. Comparison against records collected by schools or other public entities. Records about children are difficult for third parties to collect. This subcategory of submissions commonly relies on schools or other public entities (e.g., a post office or DMV) to verify the age of a child through a designee. Permission of the parents/child is required for initial access to and use of these records.

3. Peer-based verification, which allows peers in a community to vote, recommend, or rate whether a person is in an appropriate age group based on relationships and personal knowledge established offline.

4. Biometrics. Biometric solutions involve using an individual's inherent characteristics, such as physiological traits or facial images, to verify age. These solutions are discussed in the biometrics section of this document.

Commentary

- In general, some submissions attempt to make it more difficult for minors to pretend to be adults, while others focus on making it more difficult for adults to pretend to be minors. Rarely does one technology address both problems.

- Typically, these technologies do make it more difficult for a minor to pose as an adult to whom they are not related or acquainted. Similarly, they also typically make it harder for an adult to pose as a minor who is not a family member or is otherwise unknown to them.

- Many of these technologies are designed primarily for the United States context and may not functionally optimally in international contexts.

- Peer-based methods suffer from the same basic limitation seen in many an online poll or online peer-rated merchant sites: users can vote as many times as they wish to artificially raise or lower a peer rating. Additionally, if left unchecked, users can even create multiple identities to perform the extra voting themselves. Finally, even if all identities in the system are real and unique, minors might organize against another minor in their ratings or recommendations in an online form of bullying increasingly known as cyberbullying.

- Comparison against public records is only as effective as the completeness and data quality of the public database. This approach is more suitable to verifying the age of adults as public records of minors range from quite limited to nonexistent. There are also significant privacy concerns when institutions that hold the records of minors (e.g., schools) are involved.

- The public entity–based approach, though appealing in terms of the accuracy of its data, has significant challenges from a practical perspective. Resources, incentives, legal liability and basic infrastructure are each nontrivial potential hurdles to achieving scale with this solution. For example, the coordination and participation of thousands of public entities (often resource-constrained already) would be a significant operational challenge on the aggregator side.

- More generally, in all of these approaches, the user receives digital credentials after verification that can be used across sessions without reverifying. These credentials, which are usually protected by only a user name and password, are easy to transfer from adult to child or from child to adult. Further, they can be sold, traded cooperatively, or taken under duress.

- The working assumption for technologies in this category is that age- or identity-related deception is at the center of sexual solicitation on the Internet. Some emerging research, such as that documented by the Task Force's Research Advisory Board, suggests that this may not be the central issue in online sexual solicitation. Thus, although these types of solutions do target potential risks, they may not target the most critical issues that underlie Internet-based sexual solicitation.

- Finally, there are significant potential privacy concerns and security issues given the type and amount of data aggregated and collected by the technology solutions in this category, each needing to be thoughtfully addressed and well-managed.

Conclusion

Age verification and identity authentication technologies are appealing in concept but challenged in terms of effectiveness. Any system that relies on remote verification of information has potential for inaccuracies. For example, on the user side, it is never certain that the person attempting to verify an identity is using their own actual identity or someone else's. Any system that relies on public records has a better likelihood of accurately verifying an adult than a minor due to extant records. Any system that focuses on third-party in-person verification would require significant political backing and social acceptance. Additionally, any central repository of this type of personal information would raise significant privacy concerns and security issues.

Filtering/Monitoring/Auditing
Category Description

Filtering, monitoring and auditing solutions attempt either to prevent a user from accessing inappropriate content or provide a monitoring mechanism to document this activity after it occurs. These tools are based on a set of predetermined criteria that allow dynamic monitoring of web content and on-the-fly determination of the appropriate level of access. They are usually software-based and installed on a user's computer. They can often be packaged with logging features that allow an individual to review prior Internet activity on the computer. Historically, filtering, monitoring, and auditing tools have enjoyed widespread success and have been in use by parents, schools, and other public venues in which Internet restrictions are appropriate.

Filtering, monitoring, and auditing tools are generally divided into two categories: client-side and server-side.

- Client-side software is installed locally on the user's computer and is maintained by the user. Its effectiveness is dependent on the user's installation, configuration, regular maintenance, and use of the software. Client-side filtering tools are very popular and have been deployed for over a decade. They are relatively straightforward to implement and offer parents and guardians an easy way to provide a safer Internet environment.

- In the server-side approaches reviewed by the TAB, filtering of inappropriate content is performed before the content reaches a user's computer and is bounded by the standards of the website or service platform itself. (As a note, "server-side filtering" is often used to refer to content filtering at the ISP level. The TAB received no submissions for ISP-level filtering products.) For example, a social network site can filter – or flag – user-generated content that is deemed inappropriate for some users. Thus, a website's policy, rather than individual user's preferences, dictates the level of appropriateness, with the scope limited to just that site.

Commentary

- Client-side filtering can be effective as a complementary solution to other technologies, is readily deployable by a parent or responsible adult, and is reasonably easy to use. A possible downside of client-side filtering may be that it might provide users with the illusion of total safety and problem prevention and thereby reduce critical adult vigilance and involvement. Additionally, costs may prevent families from choosing this option.

- The effectiveness of a filtering tool may vary based on its design and amount of user control. Some filters do analysis on the fly, and some filters are based on a predetermined set of criteria. For this latter group, their restrictions vary greatly based on the software manufacturer. Overly restrictive tools can filter out too much information, leaving its users frustrated and resulting in a reversion to less restrictive settings, and thereby exposure to greater risk.

- Some filtering tools address all Internet technologies, but some do not. For example, one package can restrict access to inappropriate websites but still allow unfiltered conversations to occur over instant messaging programs. Finally, although many programs offer users a varying degree of control over what they filter, frequently filtering software makes decisions that rely on its own criteria, not that of the parents, limiting parents' control over what they deem appropriate.

- Commonly, these filters can detect certain types of inappropriate content, but the focus of filtering software is more on prevention of access to pornographic content than it is to violent images and video or content involving self-harm. These tools also function more accurately with text and images than with video and audio. For continued effectiveness, it is critical that filtering tools must constantly adapt to the constant changes in Internet technologies.

- Though relatively easy to implement, filtering tools typically require a software purchase and enough technological ability to install the application. Additionally, they require the time and understanding to properly configure the software for the appropriate age level and often require regular updates via the Internet. The issue here is that responsible adults may not be computer-literate enough to be comfortable with installation, configuration, and updates, which may ultimately put minors at risk.

- Filtering software can be easily circumvented or disabled by computer-savvy users, completely eliminating their effectiveness. Frequently, parents or guardians are notified in such cases, which is beneficial. In any case, parents, guardians, and other caregivers should simply be alert to the potential for circumvention.

- Server-side filtering, though appealing for its ease of use, presents concerns about potential lack of parental control over access to content and also, at the extreme, about potential censorship.

- Auditing software typically requires regular commitment from parents or other responsible adults for effectiveness. The benefit and the challenge of auditing software is the potentially vast amount of data captured about a minor's online activity. This data, however, requires some sort of adult review, commonly available in summary fashion, for actual efficacy. There is limited impact on online safety for minors from using auditing software without the ongoing attention of a responsible adult.

- To make auditing more manageable, monitoring software often stores activity logs in a central location owned by the software provider. These records are therefore potentially at risk for compromise by hackers, have the potential to be sold to third parties seeking marketing data, and have other privacy and security issues as well.

Conclusion

Filtering, monitoring and auditing software can provide parents and other supervisory adults with a useful tool to assist in determining and limiting user access to certain types of inappropriate Internet content. Although not a total solution for minors' online safety, the effective use of these types of tools can be a key part of a holistic solution whereby parental involvement, adult supervision, and software tools work together to provide a safer Internet environment.

Text Analysis
Category Description

Text-based analysis technologies are designed to automatically detect predatory, harassing, or otherwise inappropriate conversations on the Internet. These solutions generally work by obtaining samples of the conversations to be detected, extracting a statistical signature from these conversations, and classifying them based on the measured statistic. Text analysis tools vary in their deployment schemes, ranging from local installation at Internet cafes, libraries, and other public access sites to large-scale deployments across an entire social network website. Some solutions even incorporate the automated analysis as part of a parental auditing tool, locally operating on a home computer.

Commentary

- Automated text analysis can be quite useful against inappropriate interactions including online harassment, sexual solicitation, and other types of problematic communications, as it primarily focuses on language and highlights potential problems early.

- Given the sheer volume of online interactions and communications, the development of automated techniques for analyzing text conversations seems

quite reasonable. To be effective, however, it is crucial that a statistically valid sample of representative text be collected to use as a baseline. There are two challenges to this sampling effort: millions of text-based messages are exchanged across the Internet every day, so not only does obtaining a valid "going forward" sample present a challenge, but retrospectively acquiring and tracking data to adequately identify an escalating situation would also be complicated.

- An area for further development for text analysis technologies is error rate. The current typical error rate in analyzing contextual text is problematic. Not enough research has been done yet to determine the impact of known error rates. It is likely that any large-scale implementation of text analysis technology would produce far too many false positives at this point in time, and would require additional, nonscalable manual effort to identify illicit behavior. An additional risk is that legitimate users may be denied access to Internet-based services that automatically blacklist users based on criteria. The problem also exists in the reverse. A low rate of positive identification can minimize the dangers posed on the Internet, provide a false sense of security, and actually endanger the individuals it intends to protect.

- International environments such as the Internet also present challenges to automated text analysis technology. The proposed solutions currently seemed unlikely to scale to encompass the wide variety of languages, colloquial dialects, and conversational styles present on the Internet and probably essential over time to effective text analysis. Effective systems must also evolve to take into account the various ways in which users try to circumvent the filters by altering their linguistic patterns.

- The automated text analysis technologies submitted presented some potential privacy and security concerns, particularly in the cases in which a tool proposed to track and store historical data on its servers. Internet users would be unwittingly subjected to intrusions on what may be legitimate and private conversations.

Conclusion

Text analysis technologies overall seemed to be a promising category of technology solution for improving online safety for minors, but the slate of submissions in this category were in a relatively early stage of development at this time. To accommodate for current shortcomings, certain implementations of automated text analysis could still be effective. Situations in which a parent uses the technology as a complement to other filtering, monitoring, and auditing activities may assist in the supervision of a child on the Internet. Schools and other public institutions that provide clear notice to its users, deploy the tool locally as part of an overall security program, and use consistent standards to manually review the text after identification may also find it useful. Lastly, websites that deploy the solution as part of an active monitoring and supervision program may find it assists in reducing the need for manual oversight. Although these benefits may outweigh

possible concerns, it is incumbent on an entity to thoroughly test and understand the limitations of the tool prior to its deployment and, overall, the TAB felt that text analysis tools needed to evolve a bit further prior to widespread deployment and usage.

Biometrics

Category Description

Biometric technologies attempt to identify an individual or class of individuals based upon intrinsic physical (e.g., fingerprint, iris, or DNA) or behavioral traits (e.g., walking gate or typing style). Significant research has gone into the development of biometric technologies and some have been deployed in limited commercial settings.

These tools often use a hardware-based device to accept and transmit certain biometric information through the computer. In one instance, a device attempts to determine an individual's age grouping based on a bone density analysis of that individual's hand. Another tool attempts to actually identify a specific individual through facial recognition and match the individual to a known sex offender database. Others are still more novel in their approach, attempting to identify specific individuals through the analysis of a user's typing behavior and patterns.

In each instance, information is gathered by either the hardware or software tool and submitted to determine the appropriateness of an individual using a particular service. The website or web service employing this solution incorporates the safeguard in their system and where necessary, requires the user to purchase the biometric device for their computer.

Commentary

- In limited situations, biometric techniques may provide a solution to assisting in limiting inappropriate contact between adults and minors. These solutions, however, are challenged with problems that can undermine their usefulness in addition to being expensive to deploy.

- Biometric solutions typically require supervision to be effective. A situation in which individuals are expected to self-identify through the use of a biometric device over the Internet is, at best, suboptimal. Individuals can obfuscate a facial image through the use of varied lighting, facial hair, and other indistinguishable features. Typing styles and patterns can vary drastically depending on the type of keyboard, the use of voice-recognition software, and an almost unlimited number of variables from computer to computer. Bad actors can use their own children or other individuals to submit false readings. The challenges to positive, accurate identification are numerous, especially in Internet-based deployments in which an individual is not monitored while using their biometric device.

- Accuracy rates are critical for effectiveness. The level of accuracy in the submitted tools has not yet been proven and could be problematic, resulting in

potential denial of access for legitimate users to a particular website or web service.

- The working assumption of biometric technologies is that identity deception is at the root of online safety problems. Although this may be true in some percentage of cases, the research documented by the Task Force's Research Advisory Board suggests that deception is not the central issue in online safety for minors.

- Any biometric system raises important privacy concerns and security issues, particularly if the biometric data is transmitted or stored on a central server, presenting challenges to both user and business adoption. Biometric data is, by law, considered Personally Identifying Information (PII). Servers holding large amounts of PII pose a serious security risk and would be a likely target for information theft. The retention and security of this data would present a significant business liability and might have a deterrent effect on potential users. It is possible that business risk alone would likely deter any wide scale adoption, without legislation or mandate.

Conclusion

Biometric solutions certainly have some appeal, if proven effective, and show some promise, should they continue to evolve. At present, however, there are significant challenges to widespread usage and adoption for a variety of reasons including accuracy and detection rates and a need for supervision.

Other: Individual Identification
Category Description

Submissions in the category focused on identifying or profiling individuals who have been convicted of sex offenses, for example by aggregating data from registered sex offender databases or by tracking devices and computers of registered sex offenders. These technologies then enable a website to block or otherwise prevent the individuals profiled from accessing a site or areas on a site.

Commentary

- Profiling systems are only as effective as the data they use. Not all potential problem users have been previously identified or registered in the sex offender database or other watchlists; thus, a system relying on such data will be inherently limited.

- Basing a technology solution on user-provided information is a challenge to the accuracy of any technology. It is not clear that adequate incentive exists for a user to provide accurate information in this context. Further, acquiring and using invalid personal information is a trivial exercise.

- Solutions that require a computer to be used by a single user only for effectiveness will have limited deployment options and limited effectiveness in a world where public computers with Internet access are fairly widely available. Libraries, schools, and even households can have many users that may have completely different intentions.

- Identification systems require high accuracy rates for effectiveness and adoption. Problematic accuracy rates may result in legitimate users potentially being denied access to a particular web site or service. For example, a user who shares a name or identifying information with someone in a Registered Sex Offender database might be inappropriately denied access.

- With the use of personal information essential to the functioning of many of these systems, robust data privacy and security policies and technology are critical to their success.

Conclusion
These profiling technologies represent very specific point solutions, each with its particular challenges to effectiveness but also with potentially positive benefits to usage. Should accuracy issues be addressed, these types of technologies could probably be deployed in concert with other complementary technologies to improve online safety concerns for minors.

CASE STUDY: icouldbe.org

Although icouldbe.org did not propose an explicit technology solution, but rather a general description of their enterprise, they presented a complete approach to ensuring safe interactions between teenagers and adults in their secure online community. Specifically, icouldbe.org pairs underserved teenage students with adult mentors who aid students in career development, education planning, and general mentoring. All student/mentor interactions occur online, and icouldbe.org goes to great efforts to ensure that students and mentors do not interact outside of their website or have any type of personal or physical contact. To do so, icouldbe.org has implemented a number of complementary technologies, achieving what appears to be – so far, at least – a successful and effective secure community. These technologies include text-based filtering to make sure that email addresses, personal URLs, telephone numbers, or other personal identifying information are not included in any correspondence between the mentee and mentor. Additionally, icouldbe.org does extensive verification and background searches on all mentors to allow only appropriate adults to interact with minors.

The TAB was impressed not only with the goals of icouldbe.org but also with the end-to-end solution that they have implemented. Although the scale or their community is considerably smaller than the large social network sites and the goals of their online community are fundamentally different, we believe that icouldbe.org could serve as a model for the effective implementation of complementary technologies to enhance online safety for minors.

CONCLUSIONS

At the end of the review process, the TAB was overall encouraged by the innovation and energy apparent in this emergent technology area. Although no single technology provided a total solution to the various online safety problems facing minors as identified by the Research Advisory Board, each solution had some merit and some solutions could help a great deal. Further, it is clear that technology can play a role in keeping minors safer online by limiting sexual solicitation, online harassment, and access to problematic content, but it is also clear that technology alone is not enough given the nature of the challenges at hand. We are hopeful that the submitted technologies and any others in development will continue to evolve and improve in conjunction with progress on sociological fronts to optimize the mitigation of risks to minors on the Internet. We offer some concluding thoughts and recommendations below as a result of our review process.

Technology can play a role but cannot be the sole input to improved safety for minors online. Although Internet technology presents great benefits in terms of education, access to knowledge, and commerce, it of course allows social contacts and interactions that are not as easily monitored as on a supervised playground or other public space. Fortunately, with a combination of effective child and parent education, regular parental involvement or involvement by other responsible adults, continuing and increasing corporate responsibility, and some key software tools and technologies used in complement, we can as a society work to address online safety for minors more effectively.

The most effective technology solution is likely to be a combination of technologies. To the degree that online safety for minors can be addressed by technology on a standalone basis, the most comprehensive solution will likely require a several technologies working together in concert. Many of the submitted technologies were point solutions, addressing a part but not all aspects of safety for minors online. There was no single, all-encompassing solution, but this is not surprising, as online safety for minors is a multifaceted problem. Deploying complementary technology layers or using them in an end-to-end fashion will enhance the impact of any one single technology and will serve to maximize possible effectiveness.

Any and every technology solution has its limitations. No technology should be assumed to be foolproof upon deployment. In the realm of Internet safety, this is particularly true, as bad actors are likely to be especially motivated to circumvent technologies and as the stakes are extremely high. Further, some of the technologies can be circumvented as easily as a bad actor simply obtaining previously authorized credentials from an unsuspecting child.

Youth online safety measures must be balanced against concerns for the privacy and security of user information, especially information on minors. For virtually all submissions, regardless of the functional goal or type of technology, the storage and potential exposure of personal information were a potential concern. It is critical that appropriate privacy and security measures be implemented so that this amassed user

information is secure. Further, it is also important to understand the trade-off between potentially enhanced safety and the potential cost and precedent of providing private information – particularly on minors – to a possibly vulnerable or unreliable third party.

For maximum impact, client-side-focused technologies should be priced to enable all potential users to purchase and deploy them. Price points were frequently unclear or as yet unset from many of the submitted technologies. We would strongly urge innovative thinking in how to make client-side technologies as affordable as possible. Doing so will not only encourage and enable adoption by anyone concerned by children's online safety and wishing to make technology part of their individualized solution, but will also generally encourage broad adoption, which can be critical to the effectiveness of some client-side technologies.

A common standard for sharing information among safety technologies would help. There is currently no open standard for sharing information voluntarily between users, sites, and third-party vendors interested in improving online safety for minors. It would be useful if an open data standard were defined for communication among the various classes of tools produced by different companies. This open standard should be developed with the participation of vendors, but without assuming specific server- or client-side technique and with a goal of protecting the privacy of users. To clarify, here is an example: using the standard, a server-based data-mining tool could flag conversations by sending data to the child's computer; a parental-oversight tool might then be able to process this data to alert the parents. Development of this common standard would be an excellent next step in enhancing online safety for minors.

Developing standard metrics for youth online safety solutions would be useful. Standard metrics would assist in the assessment of the relative merits and trade-offs of any potential technology solution. Development of these metrics – no doubt a challenging process – would be an excellent next step in this process of seeking to enhance safety for minors online.

Respectfully submitted to the Internet Safety Technical Task Force on behalf of the Technology Advisory Board.

Laura DeBonis, Chair, Technology Advisory Board

EXHIBITS TO APPENDIX D:

1. TAB Member and Observer Bios

2. Submission Template

3. Intellectual Property Policy

4. Alphabetical List of Technology Submissions

Exhibit 1 to Appendix D:
TAB Member and Observer Bios

EXHIBIT 1

TAB MEMBER BIOGRAPHIES

BEN ADIDA, HARVARD MEDICAL SCHOOL, HARVARD UNIVERSITY

Ben Adida is a member of the Faculty at Harvard Medical School and at the Children's Hospital Informatics Program, as well as a research fellow with the Center for Research on Computation and Society with the Harvard School of Engineering and Applied Sciences. His research is focused on security and privacy of health data, the security of web applications, and the design of secure voting systems.

Dr. Adida completed his PhD at MIT in the Cryptography and Information Security group. He is the Creative Commons representative to the W3C, working on interoperable web data as chair of the RDF-in-HTML task force. Previously, he co-founded two software startups that developed database-backed web application platforms based on free/open-source software.

SCOTT BRADNER, HARVARD UNIVERSITY

Scott Bradner has been involved in the design, operation and use of data networks at Harvard University since the early days of the ARPANET. He was involved in the design of the original Harvard data networks, the Longwood Medical Area network (LMAnet) and New England Academic and Research Network (NEARnet). He was founding chair of the technical committees of LMAnet, NEARnet and the COrporation for Research and Enterprise Network (CoREN).

Mr. Bradner served in a number of roles in the IETF. He was the co-director of the Operational Requirements Area (1993-1997), IPng Area (1993-1996), Transport Area (1997-2003) and Sub-IP Area (2001-2003). He was a member of the IESG (1993-2003) and was an elected trustee of the Internet Society (1993-1999), where he currently serves as the Secretary to the Board of Trustees. Scott is also a trustee of the American Registry of Internet Numbers (ARIN).

Mr. Bradner is the University Technology Security Officer in the Harvard University Office of the Provost. He tries to help the University community deal with technology-related privacy and security issues. He also provides technical advice and guidance on issues relating to the Harvard data networks and new technologies to Harvard's CIO. He founded the Harvard Network Device Test Lab, is a frequent speaker at technical conferences, a weekly columnist for Network World, and does a bit of independent consulting on the side.

LAURA DEBONIS, BERKMAN CENTER, HARVARD UNIVERSITY

Laura DeBonis (Berkman Affiliate for the Internet Safety Technical Task Force). Laura chairs the Technology Advisory Board, which has been asked to assess the range of technology tools that may be used to promote online safety for minors. Laura was, most recently, the Director for Library Partnerships for Book Search at Google. During her time at the company, she also worked on the launch teams for AdSense Online and Froogle and managed global operations in the early days of Book Search. Prior to Google, Laura worked at Organic Online, consulting for a variety of companies on their web strategies and design. Before attending graduate school, she spent a number of years working in documentary film, video and interactive multimedia, creating content for PBS, cable channels, and museums. Laura is a graduate of Harvard College and has an MBA from Harvard Business School.

HANY FARID, DARTMOUTH

Hany Farid received his undergraduate degree in Computer Science and Applied Mathematics from the University of Rochester in 1989. He received his Ph.D. in Computer Science from the University of Pennsylvania in 1997. Following a two year post-doctoral position in Brain and Cognitive Sciences at MIT, he joined the Dartmouth faculty in 1999. Hany is the David T. McLaughlin Distinguished Professor of Computer Science and Associate Chair of Computer Science. He is also affiliated with the Institute for Security Technology Studies at Dartmouth. Hany is the recipient of an NSF CAREER award, a Sloan Fellowship and a Guggenheim Fellowship.

From working with federal law enforcement agencies on digital forensics, to the digital reconstruction of Ancient Egyptian tombs, Hany works and plays with digital media at the crossroads of computer science, engineering, mathematics, optics, and psychology.

LEE HOLLAAR, UNIVERSITY OF UTAH

Lee A. Hollaar is a Professor in the School of Computing (formerly the Department of Computer Science) at the University of Utah in Salt Lake City. He has taught a variety of software and hardware courses, and currently teaches computer networking, operating systems, and intellectual property and computer law.

He played a major role in adding two words to the vocabulary of intellectual property law:
* "Inducement" was recognized by the Supreme Court in its unanimous Grokster opinion. The concept of liability for inducement of copyright infringement was revitalized in his paper Sony Revisited: A new look at contributory copyright infringement, and refined in his amicus brief in the case. The paper also led to the introduction of the Induce Act in the 108th Congress.
* "Foreseeability" as a limit on doctrine of equivalents in patent law is the heart of the Supreme Court's Festo. It was proposed in the amicus brief whose filing he supervised as chair of IEEE-USA's intellectual property committee.

Professor Hollaar was on sabbatical leave in Washington, DC, during the 1996-97 academic year, as a Committee Fellow in the intellectual property unit of the Committee on the Judiciary of the United States Senate, where he worked on patent reform legislation, database protection, and what eventually became the Digital Millennium Copyright Act. He has been a special master, technical expert, or consultant in a number of copyright, patent, and trade secret cases.

Professor Hollaar was one of the drafters of the Utah Digital Signature Act, which made Utah the first government in the world to recognize digital signatures as equivalent to handwritten ones. On November 19, 1997, as part of Utah's Digital Signature Day, Professor Hollaar executed the first legally-recognized digitally-signed will in the world.

He received his BS degree in electrical engineering in 1969 from the Illinois Institute of Technology, and his PhD in computer science in 1975 from the University of Illinois at Urbana-Champaign. Dr. Hollaar was on the faculty of the University of Illinois prior to joining the faculty of the University of Utah in 1980.

TODD INSKEEP, BANK OF AMERICA

Todd Inskeep has over 20 years of Information Security and Internet experience ranging from secure radio and desktop systems to Security Architecture and eCommerce Authentication strategy at Bank of America. He’s a Certified Information Systems Security Professional with a Master’s in Strategic Intelligence currently leading work on the Bank’s overall eCommerce/ATM strategy. He also teaches security & risk management part-time at the University of North Carolina at Charlotte’s NSA-Designated Center of Excellence in Information Assurance. Todd

holds a BS in Electrical Engineering from West Virginia University and a MS in Strategic Intelligence from the Joint Military Intelligence College.

BRIAN LEVINE, UNIVERSITY OF MASSACHUSETTS-AMHERST

Brian Neil Levine is an Associate Professor in the Dept. of Computer Science at the Univ. of Massachusetts Amherst, which he joined in 1999. He received MS and PhD degrees in Computer Engineering from the Univ. of California, Santa Cruz in 1996 and 1999, respectively. His research focuses on networking and security, and he has published over 60 papers on these topics. In the networking area, his research focuses on mobile systems and peer-to-peer networking. In the security area, his research is focused on privacy and forensics. His lab is currently funded by the NSF, DARPA, NSA, and ARO. He received a National Science Foundation CAREER grant in 2002 for work in peer-to-peer networking, a prestigious award for new faculty. In 2004, he was awarded a UMass Lilly Teaching Fellowship and, in 2007, his college's Outstanding Teacher Award. In 2008, he received the Excellence in Science & Technology Alumni Award from the Univ. at Albany, where he received a B.S. in 1994. Levine is currently an associate editor of the IEEE/ACM Transactions on Networking journal.

ADI MCABIAN, TWISTBOX

Adi McAbian is Managing Director of Twistbox Entertainment and currently serves on the Board of Directors of Mandalay Media (MNDL), its parent.

Since founding the company, Mr. McAbian has been responsible for facilitating strategic collaborations with over 60 mobile operators worldwide on content standards and minor protection, he has been a frequent speaker, lecturing on adult mobile content business and management issues throughout Europe and the U.S. including conferences organized by iWireless World, Mobile Entertainment Forum, and Informa.

Mr. McAbian has worked with various operators including Vodafone's Global Content Standards group on establishing best practices in minor protection for both content and contact services as well as local implementations of those standards and supporting platforms in the over a dozen local markets. Mr. Mcabian also co-authored the Content Standards Rating Matrix currently used by nearly 100 networks to rate restricted content.

Mr. McAbian is responsible for corporate strategy and carrier relationships that span the globe.

Mr. McAbian's background includes experience as a successful entrepreneur and proven executive business leader with 12+ years as Business Development and Sales Manager in the broadcast television industry. Mr. McAbian is experienced in entertainment and media rights management, licensing negotiation and production, and has previously secured deals with AOL/Time Warner, Discovery Channel, BMG, RAI, Disney, BBC and Universal among others.

Mr. McAbian currently serves on the Mobile Marketing Associations' Consumer Best Practices Committee and will chair the up coming Age Appropriate Content and Services Sub-Committee.

RL "BOB" MORGAN, UNIVERSITY OF WASHINGTON

RL 'Bob' Morgan is Senior Technology Architect for the Computing & Communications Department at the University of Washington. In this role he contributes to designing, implementing, and documenting distributed computing and security infrastructure for the UW. He is the Chair of the Middleware Architecture Council for Education (MACE), providing guidance for the Internet2 Middleware Initiative. He is a primary contributor to a number of Internet2 middleware projects, notably Shibboleth, a system for secure access to inter-institutional web

resources. He is also active in standards activities with the Internet Engineering Task Force (IETF) and the Organization for the Advancement of Structured Information Standards (OASIS), where he has helped to develop the Lightweight Directory Access Protocol (LDAP) and Security Assertion Markup Language (SAML) standards.

LAM NGUYEN, STROZ FRIEDBERG

Lam Nguyen heads Stroz Friedberg's Digital Forensics lab in Boston. With over 11 years of coding, database development and digital forensics experience for leading government and commercial entities, Mr. Nguyen is an industry leader in digital forensics for data breach, e-discovery, and cybercrime in civil and criminal litigation, as well as corporate investigations. Mr. Nguyen has investigated hundreds of criminal cases and has led forensic investigations in data breach and intrusion cases. He was the lead investigator in several searches for Personally Identifiable Information on lost and stolen computers for a large pharmaceutical company. Mr. Nguyen recently conducted a forensic examination of an employee's computer for a large investment bank. That examination led to his testimony in federal court that helped prove the employee was engaged in insider trading.

Before joining Stroz Friedberg, Mr. Nguyen was the Lead Computer Forensics Specialist for the United States Department of Justice, Child Exploitation and Obscenity Section's High Technology Investigative Unit. As the team leader, he initiated and developed online investigations of high-profile child exploitation cases; examined target computers seized in criminal investigations, and provided his expertise to federal prosecutors across the country. Mr. Nguyen is highly respected in the digital forensic community and has been qualified as an expert in federal court on a number of occasions.

Sought after for his exceptional experience and commitment, he has trained law enforcement officers and trial attorneys on computer forensic issues domestically and abroad. Mr. Nguyen was an adjunct instructor at George Mason University for several years where he developed new courses and curricula on the subject of Computer Forensics and Network Security. More recently, he has been a guest lector at Harvard Law, Harvard Extension School, and the University of Massachusetts at Amherst.

Mr. Nguyen's dedication to public service has also included coordinating and delivering technology solutions critical to the operations of the U.S. Dept. of Commerce, Bureau of the Census, U.S. Dept. of Treasury, and Internal Revenue Service. Mr. Nguyen earned his Masters of Information Technology from American Intercontinental University and his undergraduate degree in Accounting Information Systems from Virginia Tech. He is certified in EnCase.

JEFFREY SCHILLER, MIT

JEFFREY I. SCHILLER received his S.B. in Electrical Engineering (1979) from the Massachusetts Institute of Technology. As MIT Network Manager he has managed the MIT Campus Computer Network since its inception in 1984. Prior to his work in the Network Group he maintained MIT's Multics timesharing system during the time-frame of the ArpaNet TCP/IP conversion. He is an author of MIT's Kerberos Authentication system. From 1994 through 2003 Mr. Schiller was the Internet Engineering Steering Group's (IESG) Area Director for Security, responsible for overseeing security related Working Groups of the Internet Engineering Task Force (IETF). He was responsible for releasing a U.S. legal freeware version of the popular PGP encryption program.

Mr. Schiller is also responsible for the development and deployment of an X.509 based Public Key Infrastructure (PKI) at MIT. He serves as a consultant to other higher educational institution in the usage and deployment of PKI and related security technologies.

Mr. Schiller is also a founding member of the Steering Group of the New England Academic and Research Network (NEARnet). NEARnet, now part of Level3, is a major nationwide Internet Service Provider.

DANNY WEITZNER, MIT

Daniel Weitzner is Policy Director of the World Wide Web Consortium's Technology and Society activities. As such, he is responsible for development of technology standards that enable the web to address social, legal, and public policy concerns such as privacy, free speech, security, protection of minors, authentication, intellectual property and identification. Weitzner holds an appointment as Principal Research Scientist at MIT's Computer Science and Artificial Intelligence Laboratory, co-directs MIT's Decentralized Information Group with Tim Berners-Lee, and teaches Internet public policy at MIT.

As one of the leading figures in the Internet public policy community, he was the first to advocate user control technologies such as content filtering and rating to protect children and avoid government censorship of the Intenet. These arguments played a critical role in the 1997 US Supreme Court case, Reno v. ACLU, awarding the highest free speech protections to the Internet. He successfully advocated for adoption of amendments to the Electronic Communications Privacy Act creating new privacy protections for online transactional information such as Web site access logs.

Before joining the W3C, Mr. Weitzner was co-founder and Deputy Director of the Center for Democracy and Technology, a leading Internet civil liberties organization in Washington, DC. He was also Deputy Policy Director of the Electronic Frontier Foundation. He serves on the Boards of Directors of the Center for Democracy and Technology, the Software Freedom Law Center, the Web Science Research Initiative. and the Internet Education Foundation.

His publications on technical and public policy aspects of the Internet have appeared in the Yale Law Review, Science magazine, Communications of the ACM, Computerworld, Wired Magazine, and The Whole Earth Review. He is also a commentator for NPR's Marketplace Radio.

Mr. Weitzner has a degree in law from Buffalo Law School, and a B.A. in Philosophy from Swarthmore College.

TAB OBSERVER BIOGRAPHIES

RACHNA DHAMIJA, USABLE SECURITY SYSTEMS

Dhamija’s research interests span the fields of computer security, human computer interaction and information policy. She received a Ph.D. from the School of Information Management and Systems at U.C. Berkeley in 2005. Her thesis focused on the design and evaluation of usable security systems. Previously, Dhamija worked on electronic payment system privacy and security at CyberCash. Her research has been featured in the New York Times, the Wall Street Journal and the Economist.

EVIE KINTZER, WGBH

Evie Kintzer, is WGBH Educational Foundation's Director of Strategic Planning and Special Projects. For the last eight years, Evie’s work with the President and Vice Presidents has included developing the Foundation’s strategic planning agenda, assessing implications of the

competitive environment, chairing WGBH's Advanced Media Group, and advising and developing project strategy and operating plans. Evie spent 13 years in the WGBH Legal Department as Director of Business Affairs and Deputy General Counsel, handling all of the business and legal affairs issues related to documentary programs produced by American Experience, NOVA, and FRONTLINE, as well as development of the Children's Television and Interactive Departments. She holds a BA from Brandeis University and a JD from Hastings College of the Law.

AL MARCELLA, WEBSTER UNIVERSITY

Albert J. Marcella Jr., is president of Business Automation Consultants, LLC a global information technology and management-consulting firm providing information technology (IT) management consulting and IT audit and security reviews and training for an international clientele.

Dr. Marcella is an internationally recognized public speaker, researcher, workshop and seminar leader with 30 years of experience in IT audit, security and assessing internal controls, and an author of numerous articles and 28 books on various IT, audit and security related subjects.

Dr. Marcella's most recent book *Cyber Forensics: Collecting, Examining, and Preserving Electronic Evidence An Auditor's Field Manual*, second edition, focuses on issues, tools, and control techniques designed to assist audit, law enforcement, and info security professionals in the successful investigation of illegal activities perpetrated through the use of information technology.

Professor Marcella is a tenured faculty member at Webster University in Saint Louis, MO, where he is responsible for teaching information technology management courses in the University's graduate and doctoral programs.

Dr. Marcella is the Institute of Internal Auditors Leon R. Radde Educator of the Year, 2000, Award recipient. Dr. Marcella has taught IT audit seminar courses for the Institute of Internal Auditors, continues to teach for the Information Systems Audit and Control Association, and has been recognized by the IIA as a Distinguished Adjunct Faculty Member.

JOHN MORRIS, CDT

John B. Morris, Jr. is CDT's General Counsel, and the Director of its "Internet Standards, Technology and Policy Project." Prior to joining CDT in 2001, Mr. Morris was a partner in the law firm of Jenner & Block, where he litigated groundbreaking cases in Internet and First Amendment law. He was a lead counsel in the ACLU v. Reno/American Library Association v. U.S. Dep't of Justice case, in which the Supreme Court unanimously overturned the Communications Decency Act of 1996 and extended to speech on the Internet the highest level of constitutional protection. In that case, Mr. Morris was responsible for the development of the factual presentation concerning how the Internet works, a presentation that served as the foundation for the Supreme Court's landmark decision.

From May 1999 through April 2000, Mr. Morris served as director of CDT's Broadband Access Project (while on leave from his firm). The Project undertook a comprehensive assessment of the legal, policy, and factual issues surrounding the emergence of broadband Internet access technologies.

Prior to becoming a lawyer, Mr. Morris had extensive experience with computers and politics. In the mid-1970's, as a staff member on Capitol Hill, he helped to promote the use of computer software to manage and improve constituent communications. In 1981, Mr. Morris joined a D.C.-area computer company, where he was one of the lead system designers of a constituent management software system for Members of Congress. In 1985, he co-founded Intelligent

Solutions, Inc., which developed the leading constituent services product used on Capitol Hill today.

Mr. Morris received his B.A. magna cum laude with distinction from Yale University and his J.D. from Yale Law School, where he was the Managing Editor of the Yale Law Journal. Following law school, he clerked for Judge Thomas A. Clark of the Eleventh Circuit Court of Appeals, worked for three years as a staff attorney at the Southern Center for Human Rights in Atlanta, Georgia, and then joined Jenner & Block in Washington in 1990.

In addition to his work with CDT, Mr. Morris is an Adjunct Professor of Law at Cardozo Law School in New York City.

TERESA PILIOURAS, POLYTECHNIC UNIVERSITY

Teresa Piliouras is an Adjunct Professor in Computer and Information Science/Technology Management at Polytechnic University, where she has taught courses in network design, bioinformatics, network security, operations research, operations management, database design, and management of technology since 1994. The department participates in four interdisciplinary research centers and houses a number of departmental labs and research groups (http://www.poly.edu/cis/research/labs/index.php) which are funded by grants from government agencies such as the National Science Foundation, NASA, the Office of Naval Research, the Air Force, and the New York State Office of Science, Technology, and Academic Research, and private companies and foundations such as IBM, Hewlett-Packard, AT&T, the Sloan Foundation, Panasonic, Intel, and Verizon. The Information Systems and Internet Security (ISIS) Laboratory consists of heterogeneous platforms and multiple interconnected networks to facilitate experimentation in issues related to information security. ISIS was designated an NSA Center of Excellence in 2002. It is currently further being expanded with an NSF Scholarship for Service (SFS) capacity building grant and is the host laboratory for Polytechnic University's SFS program.

Dr. Piliouras is working on ways to protect children on the Internet and to promote public health. She is involved in a number of broad-based community outreach programs to bring seniors and “at-risk” youth together to address problems of health and wellness. This involves creating community wiki-webs designed to create a sense of support and community, especially among those who may have been marginalized in the past. She is founder and President of Albright Associates, a company dedicated to protecting the privacy and safety of children in digital environments. Prior to Albright Associates, she was founder of TCR, Inc., a consulting company specializing in data mining and advanced intelligent technologies. She also held executive and technical positions at Accenture, Pitney Bowes, Boehringer Ingelheim, and Pepsico. She holds a Bachelor of Science from the University of Illinois, a Masters of Business Administration from Iona College, a Ph.D. from Polytechnic University, and a Postdoctoral Fellow from the Man-Machine Institute. She has authored numerous scholarly books and articles, including "Network Design: Management and Technical Perspectives" and "CRC Press Handbook of Modern Telecommunications."

GREG RATTRAY, COL (RET), DELTA RISK

Currently, Greg Rattray is a Principal, Delta Risk Consulting, establishing risk management strategies and cyber security capacity building approaches for government and private sector clients and advising the Internet Corporation for Assigned Names and Numbers (ICANN) on approaches for enhancing global Internet security. Previously, Greg served 23 years as an U.S. Air Force officer, retiring in summer 2007. His assignments included Director for Cyber Security on the White House National Security Council staff, leading national policy development & NSC oversight for cyber security programs and oversight of Iraq telecommunication reconstruction. He commanded the Operations Group of the AF Information Warfare Center responsible for global

operations of 900 personnel/$100 million active duty and National Guard team responsible for Air Force-wide tactics, red teams, exercising, test & training. He served in a number of operational intelligence and information operations assignments from the unit to Headquarters, Air Force levels. He also served as an Assistant Professor of Political Science and Deputy Director of the USAF Institute of National Security Studies at the Air Force Academy. He is the author of numerous books and articles including Strategic Warfare in Cyberspace, a seminal work in the cyber conflict field. He received his Ph.D. from Fletcher School of Law & Diplomacy, Tufts University, his Masters in Public Policy from J. F. Kennedy School of Government, Harvard University and his B.S. from U.S. Air Force Academy. He is a full member of the Council on Foreign Relations.

JEFF SCHMIDT, CONSULTANT

Jeff Schmidt is an independent security and technology risk consultant focusing on identity-related issues. Previously, Jeff founded Secure Interiors (SI), an early provider of managed Internet security services, and Authis, a provider of innovative identity services for the financial vertical. He managed both business to successful acquisition. Jeff also assisted in the re-launch of Kleiner Perkins backed ENDFORCE (formerly SmartPipes) by managing their flagship product offering to initial revenue generation. ENDFORCE was subsequently acquired by Sophos. Jeff also served as the CIO of The Ohio State University's second largest business unit and spent time at The Microsoft Corporation where he spearheaded Microsoft's first internal malicious testing of Windows 2000.

Jeff is a founder and elected Director of the InfraGard National Members Alliance, the private sector component of the FBI's InfraGard Program (InfraGard is an FBI/private sector alliance dedicated to improving and extending information sharing between private industry and the government on matters of national security). Jeff helped the FBI create the InfraGard Program in 1998 and has received commendations from the Attorney General, the Director of the FBI, and the National Infrastructure Protection Center (NIPC - now a part of the Department of Homeland Security).

On topics of computer security, Jeff is frequently interviewed and cited by numerous national publications and news outlets. He has authored several scholarly papers and has testified before state legislative bodies and the United States Congress. Jeff is a frequent speaker at major events such as Microsoft's DevDays, ITEC, ISSA, InfraGard, and Conference Board events.

Jeff authored The Microsoft Windows 2000 Security Handbook, published by Que in four languages, and contributed to Using Windows NT 4.0, and Teach Yourself Linux in 10 Minutes, also published by Que. He received a BS CIS from The Ohio State University and an MBA Magna Cum Laude from the Fisher College of Business at The Ohio State University.

JOHN SHEHAN, NCMEC

John Shehan is the Director of Exploited Children Services (ECS) at the National Center for Missing & Exploited Children (NCMEC) in Alexandria, Virginia. He is responsible for policy decisions and the overall operations within the ECS. Mr. Shehan has been with NCMEC since February, 2000 and has participated in and presented at numerous law enforcement investigative training programs on high technology crimes, online child exploitation as well as investigative and analytical skill development. He has provided extensive technical assistance to law enforcement in the United States and abroad on cases of child sexual exploitation, especially Internet crimes against children. To raise awareness of online child sexual exploitation, he speaks regularly with media outlets such as the MSNBC, CBS World News, New York Times, CNN and others.

Mr. Shehan is an active and founding member of the Financial Coalition Against Child

Pornography. He, along with other members at NCMEC collaborated to develop CyberTipline III. This system enables participating financial institutions and law enforcement to share information with an ultimate goal of eradicating the commercial viability of child pornography. John also spearheaded and manages the NetSmartz411 program. This program educates adults on all aspects of computers, the Internet and Internet safety.

NCMEC's Exploited Children Services was established in 1996 by a mandate by the United States Congress. ECS works collaboratively with the Federal Bureau of Investigation, U.S. Postal Service, U.S. Department of Justice, and the U.S. Customs Service (now the Department of Homeland Security) in cases of child sexual exploitation. ECS serves as a resource center for the public, parents, law enforcement, and others on the issues of the sexual exploitation of children. ECS analysts process reports received on the sexual exploitation of children through the CyberTipline and disseminate the leads to federal, state, local and international law enforcement agencies for further investigation. ECS analysts provide technical assistance to federal, state, local, and international law enforcement agencies investigating child sexual exploitation cases.

Exhibit 2 to Appendix D:
Submission Template

Internet Safety Technical Task Force
Technology Submission Template

Company Name / Individual
http://www.website.com

PLEASE SUBMIT BY JULY 21, 2008

ABSTRACT
This template describes the formatting and content requirements for submissions to the Internet Safety Technical Task Force's Technical Advisory Board. (This format should be familiar to any technologist who has submitted to ACM publications.) Please follow the structure of the template below. If necessary, please repeat information to accord with the template questions and layout. *Please note: Your submission should be no longer than four pages including diagrams and bibliography.*

Keywords
Provide 1-5 keywords to describe the submitted technology. Sample keywords that might be useful in this context are: filtering, searching, identification, verification, parental controls, and forensics.

Functional Goals
Please indicate the functional goals of the submitted technology by checking the relevant box(es):
☐ Limit harmful contact between adults and minors
☐ Limit harmful contact between minors
☐ Limit/prevent minors from accessing inappropriate content on the Internet
☐ Limit/prevent minors from creating inappropriate content on the Internet
☐ Limit the availability of illegal content on the Internet
☐ Prevent minors from accessing particular sites without parental consent
☐ Prevent harassment, unwanted solicitation, and bullying of minors on the Internet
☐ Other – please specify

PROBLEM INTRODUCTION
Briefly introduce the problem being addressed, citing any relevant studies. Briefly introduce the proposed solution. If the submitted technology addresses multiple problems (e.g. has multiple goals per the subsection above), please list separately each problem-solution combination.

PROPOSED SOLUTION
Describe the technical solution being proposed. Again if the technology addresses multiple problems with each a separate solution, please address each solution separately. This solution description should include enough detail to allow an assessment of whether or not the proposed solution could solve the problem being addressed. The audience for this description will be computer scientists, security experts, and engineers. When in question, the authors should err on the side of being more technical rather than less. The submission should resemble an ACM/IEEE submission in both style and substance.

In Addition to the Above Description, Please Address Each of the Following:
- Describe the solution's technical attributes, e.g. features and functionality.
- Provide use cases.
- Specify what the technology successfully solves and what it does not. Describe how the technology's effectiveness is evaluated, measured, and tested.
- Provide a strengths-weaknesses analysis.
- Detail the implementation requirements (hardware, software, end user aptitudes).
- Describe the technical standards used in implementing the proposed technology and identify the standards bodies that are the home of existing or proposed future standards.
- Discuss the technology's reliance and use of law and policy for success.
- Discuss the viability of the technology in both the US and international context.
- Detail effectiveness to date. Please provide any information possible on "failures" of the technology.

EXPERTISE
Describe the expertise of the company/developers. If appropriate, indicate other clients and products in this space.

COMPANY OVERVIEW
Please provide a description of the company including but not limited to information about founders and key team members, sources of capital, revenue (if relevant), customer base, growth, partnerships, participation in standards bodies, etc. Information submitted in this section will vary depending on a company/organization's stage in lifecycle. Our goal is to understand the context around the technology you have submitted for review.

BUSINESS MODEL OVERVIEW
Please discuss direct and indirect costs to all potential users. Please also comment on distribution model to non-profits, start-up sites and services, and other organizations that might not be able to afford full price for this technology. Our goal is to understand financial

accessibility and cost implications for all existing and new players.

MORE INFORMATION

Feel free to provide a URL that readers can go to for more information. This may include videos, detailed specs, or anything else that might be relevant. Indicate in this document what the readers might find if they go to the URL. This is a great place for information you would like to include that does not otherwise fit the structure of this document.

CONTACT INFORMATION

The final section of this document should contain basic contact information, including a contact name, email, phone number, and address for follow up. Please send any relevant additional information about contacting the people listed here to tab@cyber.law.harvard.edu.

CERTIFICATION

At the end of your submission, you should include the following statement: "I certify that I have read and agree to the terms of the Internet Safety Technical Task Force Intellectual Property Policy." The IP Policy can be found at http://cyber.law.harvard.edu/research/isttf/ippolicy.

USE OF THIS DOCUMENT

This document should not contain information that cannot be made available to the public. (See Legal Notice below) This submission will be made available to the Technical Advisory Board, the Task Force, and the Attorneys General. Additionally, after initial review, submissions may be made public and published online for public commentary. Please note that you must be prepared, in any follow-up discussions on your submission with the Task Force, to provide sufficient, non-confidential details and explanation about how your technical solution works and upon what information it relies, in order to allow the Task Force meaningfully to evaluate your solution.

NOTE: THE SUBMISSION TEMPLATE ENDS HERE -- FORMAT INSTRUCTIONS FOLLOW BELOW. PLEASE DELETE THE FORMAT INSTRUCTIONS FROM YOUR DOCUMENT PRIOR TO SUBMISSION. THEY DO NOT COUNT AS PART OF THE FOUR PAGE SUBMISSION LIMIT.

INSTRUCTIONS

FORMAT INFORMATION

This template is modified from the template used by the Association for Computing Machinery (ACM) and, specifically, the Special Interest Group in Computer-Human Interaction (SIGCHI). By conforming to this template, we are able to provide reviewers and the public with a collection of documents that allow for easy reviewing.

All material on each page should fit within a rectangle of 18 x 23.5 cm (7" x 9.25"), centered on the page, beginning 1.9 cm (.75") from the top of the page, with a .85 cm (.33"). ***Your submission should be no longer than four pages including diagrams and bibliography.***

Normal or Body Text

Please use 10-point Times Roman font, or other Roman font with serifs, as close as possible in appearance to Times Roman in which these guidelines have been set. The goal is to have a 10-point text, as you see here. Please use sans-serif or non-proportional fonts only for special purposes, such as distinguishing source code text. The Press 10-point font available to users of Script is a good substitute for Times Roman. If Times Roman is not available, try the font named Computer Modern Roman. On a Macintosh, use the font named Times.

Title and Authors

The title (Helvetica 18-point bold), authors' names (Times Roman 12-point bold) and affiliations (Times Roman 12-point) run across the full width of the page – one column 17.8 cm (7") wide.

Abstract and Keywords

Every submission should begin with an abstract of about 100 words, followed by a set of keywords. The abstract and keywords should be placed in the left column of the first page under the left half of the title. The abstract should be a concise statement of the problem and approach of the work described.

Subsequent Pages

For pages other than the first page, start at the top of the page, and continue in double-column format. Right margins should be justified, not ragged. The two columns on the last page should be of equal length.

References and Citations

Use the standard Communications of the ACM format for references – that is, a numbered list at the end of the article, ordered alphabetically by first author, and referenced by numbers in brackets [1]. See the examples of citations at the end of this document. Within this template file, use the style named references for the text of your citation. References should be published materials accessible to the public. Internal technical reports may be cited only if they are easily accessible (i.e. you can give the address to obtain the report within your citation) and may be obtained by any reader. Proprietary information may not be cited. Private communications should be acknowledged, not referenced (e.g., "[Robertson, personal communication]").

Page Numbering, Headers and Footers

Do not include headers, footers or page numbers in your submission.

SECTIONS

The heading of a section should be in Helvetica 9-point bold in all-capitals. Sections should be unnumbered.

Subsections

The heading of subsections should be in Helvetica 9-point bold with only the initial letters capitalized. (Note: For sub-sections and subsubsections, a word like the or a is not capitalized unless it is the first word of the header.

Subsubsections

The heading for subsubsections should be in Helvetica 9-point italic with initial letters capitalized.

FIGURES

Figures should be inserted at the appropriate point in your text. Figures may extend over the two columns up to 17.8 cm (7") if necessary. Each figure should have a figure caption in Times Roman.

LANGUAGE, STYLE AND CONTENT

Please write for a well-informed, technical audience, but try to make your submission as clear as possible:

- Briefly define or explain all technical terms.
- Explain all acronyms the first time they are used in your text.
- Explain "insider" comments. Ensure that your whole audience understands any reference whose meaning you
- do not describe (e.g., do not assume that everyone has used a Macintosh or a particular application).
- Use unambiguous forms for culturally localized concepts, such as times, dates, currencies and numbers (e.g., "1-5- 97" or "5/1/97" may mean 5 January or 1 May , and "seven o'clock" may mean 7:00 am or 19:00).

REFERENCES

1. Anderson, R.E. Social impacts of computing: Codes of professional ethics. Social Science Computing Review 10, 2 (Winter 1992), 453-469.
2. CHI Conference Publications Format. Available at http://www.acm.org/sigchi/chipubform/.
3. Conger., S., and Loch, K.D. (eds.). Ethics and computer use. Commun. ACM 38, 12 (entire issue).
4. Mackay, W.E. Ethics, lies and videotape, in Proceedings of CHI '95 (Denver CO, May 1995), ACM Press, 138-145.
5. Schwartz, M., and Task Force on Bias-Free Language. Guidelines for Bias-Free Writing. Indiana University Press, Bloomington IN, 1995.

The columns on the last page should be of equal length.

PLEASE SUBMIT YOUR FINAL DOCUMENT AS A PDF

LEGAL NOTICE

Exhibit 3 to Appendix D:
Intellectual Property Policy

Intellectual Property Policy
for the Internet Safety Technical Task Force

This IP policy is intended to state the manner in which intellectual property presented or submitted to the Task Force will be handled and to clarify that no confidentiality obligations will be imposed on Task Force members.

No Confidentiality of Contributions

No contribution or presentation by any Task Force member or non-member contributor to the Task Force regarding any research, technology or service (hereinafter "Submission") will be treated as confidential. Task Force members and the Technical Advisory Board, including its members and observers, shall have no duty to maintain the confidentiality of, and shall not execute or be subject to any confidentiality agreement for, such Submissions. Contributors should not present, and the Task Force will not accept, any information in a Submission that is confidential, proprietary or otherwise not for public dissemination. Contributors should submit only information that they are willing to have made public. Contributors must be prepared, in any follow-up discussions with the Task Force or the Technical Advisory Board to their initial Submission, to provide sufficient, non-confidential details and explanation about how their proposed technology or service works and upon what information it relies to allow the Task Force meaningfully to evaluate their Submission; otherwise the Task Force may not be able to continue to assess that Submission and include it in any reports.

Copyrighted Materials

Task Force members and non-member contributors will retain copyright in their Submissions to the Task Force. By providing your Submission to the Task Force, you are granting the Berkman Center and the Task Force a non-exclusive, royalty-free, perpetual, irrevocable and worldwide license to use your Submission for the sole purposes of carrying out the Task Force's work and developing the Task Force's reports, including, without limitation, the license rights to store, copy, distribute, transmit, publicly display, publicly perform, reproduce, edit, translate and reformat your Submission, and/or to incorporate it into a collective work. The Berkman Center and the Task Force shall have no obligation to publish, disseminate, incorporate in Task Force reports, or make any other use of any Submission.

Task Force members and non-member contributors understand that they may currently or in the future be developing internally information eligible for copyright, or receiving such information from other parties, that may be similar to the materials furnished in Submissions. Participation in this Task Force shall not in any way limit, restrict or preclude any Task Force member from pursuing any of its present or future copyright activities or interests or from entering into any copyright agreement or business transaction with any person.

Patents

Task Force members and non-member contributors will retain all pre-existing patent rights in their Submissions to the Task Force. No license, express or implied, of any patent owned by the contributors disclosed during this Submissions process is granted. Task Force members and non-member contributors understand that they may currently or in the future be developing patentable information internally, or receiving patentable information from other parties, that may be similar to the patents disclosed during this process. Participation in this Task Force shall not in any way limit, restrict or preclude any Task Force member from pursuing any of its present or future patent activities or interests or from entering into any patent agreement or business transaction with any person.

Trade Secrets

Because Task Force members and the Technical Advisory Board, including its members and observers, will be under no obligation to maintain the confidentiality of Submissions, any material that a contributor considers to be a trade secret or otherwise confidential or proprietary should not be submitted to the Task Force or the Technical Advisory Board.

Intellectual Property Created by the Task Force

All intellectual property in any Task Force report, except that in Submissions by Task Force members contained in such reports, shall be owned by the Berkman Center. The Berkman Center will grant to each Task Force member an appropriate, non-exclusive, royalty-free, perpetual, irrevocable and worldwide license to store, copy, distribute, transmit, publicly display, publicly perform, reproduce, edit, translate, and/or reformat the contents of any Task Force report for the purposes of facilitating or carrying out that member's participation in the Task Force and activities related to the work of the Task Force.

Exhibit 4 to Appendix D:
Alphabetical List of Technology Submissions

EXHIBIT 4

SUBMISSION LIST

1. ALIAS
2. Appen Speech Langauge Technology: Data Stream Profiling
3. Appen Speech Language Technology: Text Attribution Tool
4. Aristotle
5. AssertID
6. Been Verified
7. Chatsafe - Carmichael Group
8. Chatsafe-Crystal Reference Systems
9. CheckMyAge
10. Choicepoint
11. Covenant Eyes: Accountability
12. CovenantEyes: Accountability and Fllter
13. CredInt
14. DeepNine
15. eGuardian
16. EthoSafe
17. Gemalto
18. GenMobi Technologies
19. Icouldbe.org
20. IDology
21. Infoglide
22. InternetSafety.com
23. Keibi
24. Kidsnet
25. McGruffSafeGuard
26. Microsoft
27. Net Nanny / Content Watch
28. NetlDme
29. Portcard
30. Privo-Parity: Privacy Vaults Online
31. Privo-Parity:KidCards
32. PureSight
33. RedStarhs
34. RelyID
35. Saferspace
36. Sentinel: ADAPT
37. Sentinel: SAFE
38. Spectorsoft
39. Symantec
40. Verificage