Policy Impact Assessments: Considering the Public Interest in Internet Standards Development

John Morris and Alan Davidson[*]

August 2003

Submitted to TPRC 2003 –The 31st Research Conference on Communication, Information and Internet Policy

It is now widely understood that technical design decisions about the Internet can have lasting impacts on public policy and individual rights.[1] This raises a critical question: How can policy considerations, and the public interest generally, be best accounted for in the development of the Internet's "code?"

To answer this question, this paper first looks at the value of participation in standards bodies by policy experts and public interest advocates, and concludes that such involvement is a necessary but not sufficient response. A more systematic approach is needed to ensure that the public policy impacts are more widely considered. This paper proposes one tool to address this need – the use of ritualized "public policy impact assessments" in the Internet standards development process. The paper details the authors' proposal toward for public policy assessments within the context of one standards body – the Internet Engineering Task Force (IETF). It then examines their more general applicability, benefits, and limitations for a broader range of Internet standards efforts.

Since 2000, the Internet Standards, Technology, and Policy Project of the Center for Democracy & Technology (CDT) has explored the public interest in information and communications technology (ICT) standards. The Standards Project has participated in the work of key Internet standards bodies, engaged technologists and policy experts, and undertaken to inform and educate other policy advocates and policymakers about the ICT standards processes.[2]

[*] The authors work for the Center for Democracy and Technology, a nonprofit Internet policy and civil liberties group based in Washington, DC. John Morris is the Director of CDT's "Standards, Technology, and Policy Project," and Alan Davidson is CDT's Associate Director. This paper is based on work made possible through the generous support of the Ford Foundation. The authors wish to acknowledge**:** Jerry Berman, Matt Blaze, Scott Bradner, Carl Cargill, Lorrie Cranor, Sally Floyd, Becky Lentz, Alison Mankin, Jonathan Peizer, Allen Renear, and Stefaan Verhulst, among many others, for their input, ideas, and feedback on the work of CDT's Standards Project, much of which is found in this paper.

[1] This concept has been most recently and most accessibly explained in Lawrence Lessig, *Code and Other Laws of Cyberspace* (2000), but follows a rich literature on the impact of technology on society and policy. *See, e.g.,* Langdon Winner, *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought* (1977); Lewis Mumford, *The Myth of the Machine* (1966).

[2] The Project's experience with the pursuit of public policy concerns within standards bodies was detailed in the authors' paper submitted to the 2002 TPRC Conference, "Strangers in a Strange Land: Public

Based on the experience of the Standards Project, this paper looks at possible institutional approaches intended to ensure that public policy issues are appropriately considered in the standards setting process. By seeking to institutionalize the consideration of public policy concerns within standards development, the authors hope to ensure that the architecture of the future Internet is crafted with the public interest in mind.

Section I below briefly looks at the public interest in technical standards, and concludes that there exists a significant need to increase the consideration of public policy concerns within standards developments organizations.[3] Section II looks at the invaluable role that policy experts and advocates can play in standards development, but concludes that more systematic approaches are needed to ensure that the public interest is considered. Section III spells out a proposal for "policy impact assessments" that standards bodies themselves can use to conduct a routine assessment of the public policy impact of their design decisions. Section IV details one application of this proposal, a draft public policy assessment document that the authors have submitted to the Internet Engineering Task Force. Finally, Section V examines the benefits and limits of this approach, and suggests further areas for exploration and research.

I. The Public Interest in Technical Standards

In terms of both substance and process, the public has a strong interest in the development of the technical standards on which the Internet is based. Many standards are highly technical and arcane, and have little impact on social values or public policy. But in a small but growing percentage of cases, standards can dramatically influence policy concerns. For standards with a direct impact on issues of public concern, the development process often does not fully recognize the potential public policy impacts, much less address those impacts adequately.

A. Policy Impacts of Technical Standards

Technical standards, from building codes to "generally accepted accounting principles," can have important impacts on public policy concerns, and this is especially true in the field of information and communications technology (ICT) standards.[4] In the Internet space, standards are crucially important. Indeed, the very existence of the Internet itself originated in and is built on technical standards – the core concept of the Internet is that diverse and

Interest Advocacy and Internet Standards," at http://www.cdt.org/publications/piais.pdf or http://intel.si.umich.edu/tprc/papers/2002/97/Strangers_CDT_to_TPRC.pdf.

[3] This point is addressed more fully in the authors' 2002 TPRC paper, cited in footnote 2 above.

[4] The field of information and communications technology (ICT) broadly covers a range of computer and networking technology, with Internet technology being only one segment of ICT. This paper focuses on Internet standards development processes, but many of the points made here would apply equally to broader ICT standards efforts.

quite different computers and networks can communicate with each other through the use of well defined communications and network protocols.

Decisions about technical design standards are most commonly made in private bodies – such as the IETF and the World Wide Web Consortium (W3C) – that set technical standards for the Internet. These and other key standards bodies operate largely outside of the public eye and with little involvement of public interest groups or policymakers. Once the sole province of engineers, academics, and industry, Internet and other ICT technical decisions can increasingly have far-reaching implications on property rights, personal privacy concerns, and the public's access to information.

The broad societal embrace of Internet technologies is fueling the public's interest in the Internet's future course of development. As the Internet is used by a wider segment of society for a wider range of uses, changes to the Internet have a correspondingly wider impact. The rapid pace of change in Internet technology also makes its standards of continuing importance and broad impact.

B. The Example of IPv6 Addressing

The development by the IETF of "Internet Protocol Version 6" or "IPv6" provides one example of the ways that technical design decisions can directly impact public policy concerns. In 1998, an IETF standard describing IPv6, a new protocol for Internet addressing, set off a major controversy about user privacy and anonymity.[5] Under IPv4, the predecessor to IPv6, Internet addressing allowed a reasonable amount of privacy and anonymity, because the numeric IP address (such as 206.112.85.61) was typically not tied to any particular machine or user. With IPv6, however, the standard provided that in many cases a user's IP address would be derived from a unique number embedded in that user's Ethernet network card.[6] IPv6 would therefore enable for greater monitoring of users' online behavior since their IP address would be tied to a unique identifier. Thus, for example, a particular laptop computer would be widely identifiable and traceable when it communicated online, no matter where or how the computer was connected to the Internet.

It is not clear that the privacy implications of the new IPv6 address scheme were desired or even widely understood by its original designers; the use of a unique hardware ID in fact just seemed like a clever technical approach to generating unique IP addresses. Regardless, significant debate ensued once the concerns were raised, both in the public policy space and among technologists. The issue was ultimately resolved by the IETF with publication of an

[5] See "Transmission of IPv6 Packets Over Ethernet Networks," RFC 2464, December 1998, at http://www.ietf.org/rfc/rfc2464.txt. IPv6 was designed, among other purposes, to alleviate a growing shortage of numeric Internet addresses under the current addressing scheme, IPv4. The Internet in 2003 is in a state of transition from IPv4 to IPv6, with the great majority of computers and networks still using IPv4 but a growing number converting to IPv6 (and for the present supporting both protocols).

[6] These MAC (Medium Access Control) addresses are the 48-bit hardware addresses used to identify devices on an Ethernet network. All Ethernet-enabled devices have unique MAC addresses, so difficulties or collisions in routing can be avoided. More information about MAC addresses and how they are assigned is available at http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/mac-vendor-codes.html.

optional addressing scheme for IPv6 that added privacy-protecting alternatives to using MAC addresses.[7]

IPv6 provides just one example of the wide array of situations in which information and communications technologies (ICT) standards setting affects public policy concerns. Other recent examples[8] of policy questions raised by technical standards include (a) whether wireless location-tracking technologies will allow users to control who can track their location, (b) whether standards for electronic "e-books" will accommodate the needs of blind users, (c) whether "digital rights management" technologies to protect intellectual property will allow users to make lawful "fair use" of copyrighted content, and (d) whether third parties will be able to modify, without permission, Internet content as it is transmitted from the sender to the recipient.

C. Policy Consideration within Technical Standards Bodies

Public policy concerns do arise and are considered within standards processes, but almost always on an ad hoc basis. This ad hoc approach to public policy concerns presents at least two major problems – the lack of systematic analysis of public policy issues, and the lack of outside input into the analysis that does take place.

Though many technologists within the leading standards bodies are public-minded, few have explicit expertise in policy-making or at interpreting the public interest. Standards organizations have typically (and appropriately) emphasized technical goals over societal ones, but in the Internet's early history there was a significant overlap between the two. Openness, accessibility, anonymity, and robustness were all technical features of the network that became public values as well.

Additionally, since the Internet in its early days was small, the pressure for explicit analysis of public policy concerns was minimal – policy impacts deriving from technical choices would affect just a few people. The Internet's population and diversity of uses has grown enormously since the early days of the network, and technical design decisions how directly affect the online experiences and options of hundreds of millions of users. Although many past standards were consistent with the public's interest in a robust and flexible new mode of communication, there is little to suggest that the coincidence will continue.

The risk of divergence between standards and the public interest is significantly heightened by the commercialization of the Internet. The introduction in the early 1990s of commercial traffic to the Internet began an influx of private interests to a standards community that had been largely research-oriented. The subsequent explosion in commercial use of the Internet prefigured a significant increase in privately motivated participants in the standards process. This in turn has subtly changed Internet standards-making: While many private-sector

[7] See "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," RFC 3041, January 2001, at http://www.ietf.org/rfc/rfc3041.txt.

[8] Some other examples are detailed in the authors' paper submitted to the 2002 TPRC Conference, see footnote 2 above, and in other publications of CDT's Standards Project, see http://www.cdt.org/standards.

participants make high-quality contributions to standards, the extent to which participants can be expected to agree about the network's architecture is diminished because of diverging market interests. And because of these changes, there is a growing risk that the public interest in standards – an ethos for many of the leading Internet standard bodies – could fade into the background of discussion among private interests.

D. The Need for Early Identification of Potential Policy Impacts

In many cases, harmful impacts on policy concerns can be avoided if they are considered early in the technology design process. Raising policy issues early is essential. The standards design process often takes between 18 and 36 months, and marketplace deployment may be months later. If policy concerns are not raised until after a standard is finalized, or after products are deployed, the chance of constructive change is very low. Legislative or regulatory fiat cannot inject into a service or product technical capabilities that were not designed in the first place, and can often at best only restart a lengthy standards design process. In many cases, post-design regulation is powerless to put a harmful technological genie back in the bottle.

From a public policy perspective, the key question is how to obtain an outcome in the standards design process that appropriately balances both technical/engineering considerations and public policy concerns. The effort to obtain a desirable outcome will in many cases, as discussed in the following section, require the active participation of public policy advocates in the standards design process. But an even more threshold question is how can both the standards development and public interest communities even *identify* the design efforts that *might* impact on public policy. Even the direct recognition of a potential policy impact is alone likely to improve the handling of the policy concern.

II. Benefits of and Limits to Public Interest Engagement with Internet Technical Standards

If standards bodies impact policy, and have few mechanisms for accounting for public interests, how can those interests be included? Intervention and participation by representatives of public interests is a natural place to start. This section outlines some of the benefits and limits to these sorts of interventions as a solution.

As a threshold matter, we note that – at least in the United States – direct government intervention in standards processes remains a controversial and, at least for now, a highly disfavored approach. Government engagement in technical standards has long been a subject of controversy.[9] In particular in the world of Internet standards, government intervention has been suspect due to the complexity of the technical issues presented, the rapid pace of change, and the lack of government expertise.

[9] *See, e.g.,* Office of Technology Assessment, "Global Standard: Building Blocks for the Future" (1992), at http://www.wws.princeton.edu/cgi-bin/byteserv.prl/~ota/disk1/1992/9220/9220.PDF.

The role of non-governmental organizations, however, is more in synch with the private sector orientation of most ICT and Internet standards development. Over the past fifty years, public interest advocates and advocacy groups have played a vital role in the development of public policy. Historically, this advocacy has focused on traditional policy making venues of legislatures and regulatory agencies, in addition to advocacy in both the courts and the court of public opinion.

Technical standards setting bodies, however, are a radically different type of venue, and the traditional approaches used by public interest advocates do not easily translate to the technical fora of the standards bodies. Only recently have public interest organizations attempted to participate in Internet technical standards setting processes. These efforts suggest the potential effectiveness, the value, and the limitations of direct public interest participation in technical standards setting.

A. Can Public Interest Participation in Technical Standards Setting Be Effective?

The relatively limited experience over the past few years of public interest involvement in Internet standards setting bodies, including the work of CDT's Standards Project, indicates that direct public interest involvement in standards development efforts can be effective in identifying issues of public concern, spurring analysis of such issues, and promoting approaches and results that further the public interest.

In some cases, public interest participation has been a part of a technical design effort from its inception. In 1997 the World Wide Web Consortium (W3C) undertook to develop the Platform for Privacy Preferences (P3P) as a specification that enables web sites to express – in a machine-readable way – their practices with regard to users' personally identifiable information.[10] P3P permits users to quickly interpret privacy policies whose complexity might otherwise be disarming, and to make informed choices about disclosure. Numerous members of the public advocacy community and Internet industry participated actively in P3P's development, providing extensive input into the vocabulary P3P uses to describe all the various practices and implications for personally identifiable information. Since its release, P3P has been adopted by thousands of web sites, and is now built into the leading World Wide Web browser in use today. Public interest participation proved to be a critical element of the P3P development process.[11]

In other cases, public interest advocates have injected themselves into existing standards discussions to raise issues of public concern. In 2000 and 2001, the Internet Engineering Task Force (IETF) community struggled with the question of whether to charter a proposed working group on "Open Pluggable Edge Services" (OPES). The proposed OPES protocol

[10] More information about P3P is available at http://www.w3.org/P3P/.

[11] The public advocacy involvement in P3P has been described in detail by one of the co-chairs of the P3P development process. See Lorrie Cranor, "The Role of Privacy Advocates and Data Protection Authorities in the Design and Deployment of the Platform for Privacy Preferences," available at http://lorrie.cranor.org/pubs/p3p-cfp2002.html.

would permit operators of servers in the middle of the Internet to modify content in mid-stream from a server to a user, raising significant questions about censorship, data integrity, and user privacy. Some within IETF objected to the OPES proposals. In August 2001, as part of its Standards Project, CDT submitted extensive comments to the IETF about the public policy issues raised by OPES. In response to the concerns raised, in late 2001 the Internet Architecture Board (which provides architectural guidance to the IETF) undertook an extensive review of the OPES proposals, and in November 2001, recommended that any work on OPES include strong protections for data security and privacy.[12] The input of public interest advocates in the OPES debate helped to crystallize the issues raised by the proposal, and made clear to the IETF community that outside groups shared many of the concerns raised by some within the IETF.

In another interaction with the IETF, public policy advocates have played a major role in the development of a protocol for privacy protection in location-tracking and location-dependent services. Working within the GeoPriv Working Group, public policy advocates (including CDT) have pushed the IETF to include strong protections for privacy in any transmission that sends location information. The way that users can express their location privacy and security preferences will likely have a broad impact on user privacy and control. Although this effort has similarities to P3P described above, it is be tailored to some unique characteristics of location information, and critically, the new platform is expected to include default privacy requirements to be applied in the absence of any privacy rules created by a user.[13] Although very much a "work in progress," the GeoPriv effort shows the potential for cooperation between IETF technologists and the privacy community.

To be effective, however, public interest advocates must bring to a standards discussion a strong technical understanding of proposed standard and its context. Moreover, advocates must be prepared to commit substantial investments of time and energy necessary to follow ongoing internal discussions about a given policy proposal. To the maximum extent possible, advocates must raise public policy concerns using the procedures and terminology of the standards body.[14] When carefully done, public policy input into technical standards setting processes can make a significant contribution to the development of technology that is sensitive to public concerns.

[12] For a more detailed discussion of OPES and the issues it raises, see Standards Bulletin 1.02, August 7, 2002, available at http://www.cdt.org/standards/bulletin/1.02.shtml.

[13] The CDT Standards Project has been actively involved in GeoPriv since the working group's first meeting in August 2001. Together with other privacy advocates and technologists from private industry, CDT has drafted a variety of documents addressing important privacy priorities for the new standard. For a more detailed discussion of the GeoPriv working group, see Standards Bulletin 1.01, May 28, 2002, available at http://www.cdt.org/standards/bulletin/1.01.shtml.

[14] These and other recommendations are detailed in the authors' paper submitted to the 2002 TPRC Conference, see footnote 2 above.

B. Is Public Interest Participation Necessary?

Not only can public interest involvement make valuable contributions to technical standards design efforts, but such involvement is over the long run essential. Although technologists within standards bodies do often identify public policy concerns, public policy input is still needed, for a variety of reasons:

- Just as technical issues can be subtle, policy concerns can also be subtle, and some concerns will be overlooked entirely without directly public policy consideration of a technical proposal.
- Although some policy issues are identified by the technical participants themselves, the resolution of the concern at times requires an added level of experience with the policy concern to be able to evaluate the gravity of the policy threat and the sufficiency of proposed solutions.
- Increasingly, private and commercial agendas are being pursued within technical standards bodies (attempting, for example, to push a technology through quickly without addressing inconveniences such as privacy considerations), and public concerns will be overlooked or inadequately addressed without participants whose primary agenda is the public interest.
- More generally, technologists are often quick to acknowledge that they lack expertise in legal or policy issues, and they are thus hesitant to attempt to address a public policy concern without direct input from experts in the field.

For these and other reasons, it is vitally important that public interest advocates continue and increase their level of participation in Internet technical standards setting bodies.

C. Is Public Interest Participation Sufficient?

Although public interest participation is a vital element of the appropriate consideration of public policy concerns in technical standards bodies, it is not sufficient, for a variety of reasons. With unlimited financial and human resources, direct involvement might be able adequately to address the public interest, but given the realities of funding and resources, public interest advocates alone are not sufficient. There are a number of unrelated factors that suggest that direct public interest involvement will always be inadequate:

- On-going and active participation in a standards working group requires a very significant commitment of time. The generally accepted guideline is that meaningful participation in any active working group requires a baseline of approximately 20% of a staff person's time (1 day of work and meetings per week plus regular in-person conferences).

- Effective public advocacy within the technical standards bodies requires the right mix of technical knowledge (or ability to learn) with public policy experience, which somewhat limits the pool of possible advocates.

- The time horizons for standards development efforts is very long, and may be too long to garner the dedicated attention of many public interest organizations that are balancing scarce resources and immediate policies crises. All of the examples discussed in Section II.A above have required multi-year investments of time.

- Many standards bodies have an institutional or cultural resistance to addressing public policy issues, often based on past experience with public policy advocates who failed to tailor their message to the forum.

- The sheer size, number, and diversity of technical standards setting bodies means that public interest advocates will not be able to "cover the whole waterfront" of standards bodies – there simply are too many standards working groups and task forces for the public interest community to cover.

- Moreover, the approach of having public interest advocates monitor, in the first instance, the work of standards bodies seeking to identify public policy issues will be very hard to scale well, for the simple reason that many standards development activities do *not* have major public policy implications. The challenge will be to develop a way to identify the standards work that does impact public policy.

In light of the importance of protecting the public interest in the development of technical standards, and the lack of capacity of the public interest community to accomplish that task alone, it is clear that the technical bodies should undertake concrete internal steps to identify and begin to address public policy issues that arise within their standards development efforts.

Although historically some standards bodies have sought to avoid public policy issues and debates, public policy concerns about a standards will arise sooner or later, and the concern can be far more easily addressed if it is identified early in the design process. Although identifying and addressing public policy concerns may somewhat extend or complicate the development process, the resulting standard will likely meet with fewer post-development marketplace or regulatory obstacles. The investment of time and effort to address a policy concern early will likely pay off in terms of costs and delays avoided later.

III. Public Policy Impact Assessments: A Proposed Early Alert System

As described above:

- some, but not all, technical standards activities can have broad policy implications;
- there are major benefits to considering those impacts early in the design process; and

- participation by policy experts and interest groups is not alone sufficient, in large part because they are unlikely to have the resources to identify policy concerns in the large number of standards efforts, of which only a few may raise concerns.

Therefore, a more systemic approach to public policy issues is desirable. Such an approach is more likely to be effective at raising policy issues early in the design process, especially if it can be reasonably implemented within the existing procedures of standards bodies.

A system of "public policy impact assessments" could form the foundation of a strategy for standards bodies to identify and address public policy impacts. The core idea is fairly simple – that technical standards setting bodies should develop a procedure for a relatively brief but focused assessment of new technology proposals to identify whether public policy concerns might be affected.

Specifically, such a procedure would ideally be executed internally, usually without the direct involvement of a public policy expert or advocate. Moreover, the key purpose of the public policy assessment is to *identify* policy concerns early in the design process, not to indicate how those concerns should be addressed.

To achieve these key goals, the public policy assessment must be one that looks at technical design issues from the perspective of *the technology designer, not the public policy advocate.* In other words, the assessment must be in terms that are well understood by the community of technologists in the standards body.

For example, the public policy assessment process should *not* ask questions like "does this technology harm privacy?" Instead, the assessment process should break the technology down into components that themselves are known in some cases to harm privacy. Questions that would be more appropriate and constructive would be "does this technology expose information about an end user to a third party?" or "does this technology permit the retention of information about an end user?" Thus, one of the most challenging hurdles to overcome in the development of an effective system of public policy impact assessments is to break down abstract public policy concerns into concrete and familiar technological elements that can be evaluated.

Not all public policy issues can be easily broken into and evaluated as component parts, and the proposed system of public policy impact assessments will certainly not be perfect. But such a system of assessments can flag a wide range of public policy concerns, many of which are overlooked today.

Because of the great diversity of standards bodies (in terms of their focus, structure, and procedures), a single one-size-fits-all (or even one-size-fits-most) public policy impact assessment process is not likely to be effective, for at least two reasons. First, different standards bodies deal with different types of technologies, and thus the public policy issues most likely to arise within each standards body will be different. A look at three technical standards bodies makes this clear:

- A core focus of the Internet Engineering Task Force (IETF) is on how computers and networks communicate with other computers and networks, and in general the IETF seldom works on standards that dictate how end users will interact with their computers (or with the Internet). Thus, public policy concerns about (for example) government surveillance of communications are likely to arise (and have arisen in the past) at the IETF.

- The World Wide Web Consortium (W3C), in contrast, is heavily involved in developing how information is presented to end users on the Internet, and thus is far more likely to encounter public policy concerns about whether a given technology will (for example) be accessible to Internet users with disabilities.

- The primary focus of the Copy Protection Technical Working Group (CPTWG) is on a technical area that almost inherently affects important public policy concerns – technical methods to constrain copying of material (potentially in violation of a copyright). Thus, almost everything CPTWG does impacts on public policy, and thus there is a far more routine need to analyze the public policy impacts of any proposals at CPTWG.

Similarly, the structure and procedures of different standards bodies may suggest quite different procedural options for actually implementing a public policy assessment process. For standards bodies that are almost entirely run by voluntary participants (like the IETF), public policy impact assessments would likely be structured differently than for standards bodies with a significant full time paid staff able to work on substantive issues (like the W3C). Indeed, the W3C itself already has an entire organizational component focused on "technology and society" issues, and that component would certainly take a leadership role in the implementation of a public policy impact assessment process.

IV. Towards Public Policy Impact Assessments within the Internet Engineering Task Force

In June 2003, the authors submitted to the Internet Engineering Task Force (IETF) a first draft of a public policy impact assessment for IETF-developed standards. The "Internet-Draft," titled "Public Policy Considerations for Internet Design Decisions," appears in the Appendix of this paper. As an Internet-Draft – the primary working document of the IETF – the submission is considered a "work in progress" and has no official status within the IETF.[15]

The idea of a public policy assessment is similar to what the IETF community already does for security concerns. Security – both of communications over the Internet and of the Internet itself – is a paramount concern of the IETF, and the IETF requires that every draft

[15] As of this writing, the authors are revising the Internet-Draft based on comments from the IETF community, and will submit the revisions to the IETF for possible further consideration. The authors specifically invite any comments and suggestions on the issues raised in this paper, and on the Internet-Draft in particular. Please send comments to jmorris@cdt.org.

standards document contain a specific discussion of the security implications raised by the document.[16] The proposed public policy impact assessment is based on a very similar idea – that there are some threshold things that a technical designer should consider when creating a new technology or changing an existing one. Although the authors are not proposing that the IETF require a public policy impact assessment in every document submitted to the IETF (as is required for security), the authors do believe that a public policy assessment should be done for all standards adopted by the IETF.

A second important model for the attached Internet-Draft is a document published by the Internet Architecture Board (a committee that gives technical and architectural guidance to the IETF community) in November 2002. Entitled "General Architectural and Policy Considerations" and designated as "RFC 3426," the Architecture Board's document poses a series of questions that any Internet designer should consider in crafting a new technology.[17] The goal of RFC 3426 is to raise important architectural issues for consideration, not to suggest that the issues must be resolved in any specific way. The authors' Internet-Draft focused on public policy concerns follows the same model – it seeks to raise public policy issues without mandating how they are resolved.

A key goal of the Internet-Draft is to encourage members of the IETF community to think about the specific technical design elements that (a) have a significant potential to raise public policy concerns, and (b) are reasonably likely to arise within IETF standards. Among the design elements discussed are:

- **Bottlenecks and choke points.** Historically, the Internet does not have any single or limited number of points through which communications must flow. Any technology that creates such bottlenecks, even for narrow categories of communications, will create an opportunity for unintended third party or government censorship and control.

- **Discrimination among types of Internet traffic or classes of Internet user.** Historically, most communications on the Internet have received the same handling by the routers and networks that carry Internet traffic. Technology that allows certain traffic to receive priority over other traffic (such as some "quality of service" initiatives) might also be able to be used to discriminate against less wealthy or less powerful groups of Internet users.

- **Persistent identifiers.** Technologies under which individual users receive an identifier that remains the same over time can create concerns about users' anonymity or privacy.

[16] See “Instructions to RFC Authors,” RFC 2223, October 1997, at http://www.ietf.org/rfc/rfc2223.txt; “Guidelines for Writing RFC Text on Security Considerations,” RFC 3552, July 2003, at http://www.ietf.org/rfc/rfc3552.txt.

[17] “General Architectural and Policy Considerations,” RFC 3426, November 2002, at http://www.ietf.org/rfc/rfc3426.txt.

- **Retention of user data.** Any technology that permits the aggregation and/or retention of data about users significantly increases privacy concerns.

In all, the attached Internet-Draft raises more than 25 technical design questions, and relates those questions to seven categories of public policy concern. The first draft is certainly not yet as comprehensive as planned, but it has proved to be a useful starting point for discussions within the IETF about how to identify, consider, and address public policy concerns.

V. Assessing the Assessments, and Other Areas for Further Consideration

Although a formal public policy impact assessment process has not been implemented in the major Internet standards bodies, one can anticipate some of the likely strengths and weaknesses. In part due to the limitation of assessments, and in part due the varying character of Internet standards groups, other approaches may also prove valuable in helping Internet standards bodies to address issue of public policy concern.

A. The Value and Limits of Public Policy Assessments

The public policy impact assessment process has the potential to significantly advance and protect the public interest, and can strengthen the value of the standards development processes themselves. But the assessment process will not be perfect, and in any event it is limited in its intended reach. Assessments may be of limited value in some technology development efforts where, for example, policy implications are already widely understood.[18]

Likely benefits of the proposed assessment process include:

- assessments significantly increase the likelihood that public policy concerns will be identified easily during the technology design cycle;

- they can be effective and beneficial at an early stage of the design cycle;

- the assessment process can be implemented within the existing standards development structures and procedures, and can take place concurrently with other early steps in the design process;

[18] For example, within the OASIS standards group, the xRML effort focuses on expressing copyright and content usage limits associated with particular web content, and participants are highly aware of the policy ramifications of their work. Similarly, the Content Protections Technology Working Group (CPTWG), a technical discussion group focused on copy protection technologies for movies and video that is not properly a standards body but is the locus for many important technology recommendations, participants include technologists, lawyers, and business people who are actuely attuned to the public policy debates surrounding their work. In both cases, the groups would find limited benefit from a policy awareness effort.

- after the initial formulation of the assessment process (which will require input from both technologists and public policy experts), the basic assessment process can be executed by the technologists working in the standards process

- the system can scale well, and allow for policy impacts to be considered within a large number of standards efforts; and

- the process allows the limited resources of the public interest participants in the standards processes to focus on significant policy issues that arise in the standards bodies, rather than having to focus on "issue spotting."

At the same time, it is important to recognize that public policy impact assessments are only one tool, and they have a number of limitations:

- the individual assessments will not be done by policy experts, and their quality will likely vary;

- the quality of the assessments will also depend on the importance attached to the assessments within a standards bodies, and the resources devoted to the assessments;

- a public policy assessment will never be able to reflect all of the potential policy issues that could be impacted by a new technology;

- the assessments will not be clairvoyant, and policy concerns could still arise late in the technical design process;

- commercial, political, or other agendas could influence individual public policy assessments; and

- most critically, the proposed public policy impact assessment process will only serve to *identify* the public policy issues, and may not in and of itself point to ways to resolve the policy concern.

A critical requirement for the success of a policy impact assessment is for the standards community to perceive value in undertaking the assessments. If a standards body recognizes that its work does impact on public policy concerns, and the organization decides that it is better off identifying the concerns early in the design process, then public interest impact assessments can provide an effective and scalable way to address public policy concerns.

B. Other Possible Innovations for Standards Setting Bodies

The authors believe that public policy impact assessments would benefit many standards bodies. But such assessments are not the only steps that a standards body can take to promote the appropriate handling of public policy concerns. Given their limitations, and the variety of standards organizations, other approaches may prove beneficial.

Dedicated policy staff: By building dedicated internal policy expertise, standards efforts can improve their ability to address policy considerations. For example, the World Wide Web Consortium has expressly structured itself to recognize and respond to the reality that standards affect social and public concerns. Among its components is the "Technology & Society Domain" that specifically seeks to create standards that advance a public purpose.[19] Its mission is:

> Working at the intersection of Web technology and public policy, the Technology and Society Domain's goal is to augment existing Web infrastructure with building blocks that assist in addressing critical public policy issues affecting the Web.
>
> Technical building blocks available across the Web are a necessary, though not by themselves sufficient to ensure that the Web is able to respond to fundamental public policy challenges such as privacy, security, and intellectual property questions. Policy-aware Web technology is essential in order to help users preserve control over complex policy choices in the global, trans-jurisdictional legal environment of the Web. At the same time, technology design alone cannot and should not be offered as substitutes for basic public policy decisions that must be made in the relevant political fora around the world.

In addition to seeking to create new technologies that seek to address societal concerns, the W3C's Technology & Society Domain also works with other parts of the W3C to ensure that W3C technology is sensitive to public policy concerns.

Although the scope and size of the Technology & Society Domain – which has a staff of ten – may be beyond the capacity of many standards bodies, most standards organizations could follow the W3C's model and dedicate some resources to addressing public policy concerns. At a minimum, standards bodies could dedicate resources to implement and oversee a public policy impact assessment process.

Soliciting input: More generally, standards bodies could make efforts to solicit the participation and input of public policy advocates and organizations in the design process. Some standards and industry bodies (such as W3C) have made such efforts. Other technical bodies (such as CPTWG) are so inherently focused on public policy issues that public interest participation naturally occurs. But standards bodies that do not already have public interest involvement could take concrete steps to facilitate such involvement.

Heightening sensitivities more broadly: In the long term, the most effective approach to considering policy implications may be found in heightening awareness and sensitivity among the technologists developing new standards. Making technologists more sensitive to the social context and implications of their work has the potential to incorporate awareness about policy implications into many more technology development efforts. While an

[19] The P3P development effort discussed in Section II above was a project of the Technology & Society Domain. More information about the Domain is available at http://www.w3.org/TandS/.

ambitious idea, it is one that is increasingly called for in a society so dependant on complex technical systems and facing so many difficulties translating between technical development and policy outcomes.

VI. Conclusion

As we have outlined, in the complex and rapidly-evolving world of Internet and ICT standards technical decisions can have lasting public policy consequences, but are often made without full appreciation of those consequences. Significant social benefits can arise from consideration of those policy impacts early in the technical standards development process – and well before products are actually produced and are difficult or impossible to change.

Engagement by policy experts or public interest advocates can create important interventions and raise awareness of policy issues. While such engagement is highly useful in affecting outcomes, it is also limited by resources and does not scale well today across the large number of Internet and ICT standards efforts. Rather, more systemic approaches to raising awareness about policy are called for.

A public policy impact assessment process could be a highly useful tool for many Internet standards bodies, especially where public policy issues are not a central focus or where strong public interest involvement does not already take place. In many cases, routinely asking a set of critical policy impact questions could go a long way towards identifying and addressing potential policy consequences early in the technology development life cycle.

But as noted, good policy assessments are hard to do and face many challenges. Other methods of raising awareness – including more sweeping changes in engineering education and the sensitivities of technologists – may ultimately be needed. And assessments are only a critical first step towards identifying public policy impacts before they occur. Assessments alone will not guarantee that public concerns, once identified, will be appropriately addressed. That will likely require direct involvement by policy experts in the design process – and remains a topic in need of greater research and understanding.

Appendix

John Morris & Alan Davidson, "Public Policy Considerations for Internet Design Decisions," Internet-Draft (a work in progress) submitted to the Internet Engineering Task Force in June 2003. Available online at:

http://www.cdt.org/standards/draft-morris-policy-considerations-00.txt
(original ASCII text format)

http://www.cdt.org/standards/draft-morris-policy-considerations-00.pdf
(easier to print PDF format)

Internet Draft

J. Morris
A. Davidson
Center for Democracy & Technology

draft-morris-policy-considerations-00.txt
June 2003
Expires: December 2003

Public Policy Considerations for Internet Design Decisions

Status of this Memo

Abstract

This document suggests public policy questions that the IETF should consider and possibly address when developing new standards and protocols, and modifying or enhancing old standards and protocols. This document contains questions to be considered, as opposed to guidelines or rules that should in all cases be followed. This first draft provides a framework for identifying and discussing questions of public policy concern, and invites members of the IETF community to contribute to the questions and discussions raised here.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Table of Contents

1. Introduction and Rationale for this Document

This document suggests public policy questions that the authors believe should be considered and possibly addressed within the IETF when it is working on new or revised standards or protocols. This document offers questions to be considered, rather than guidelines to be followed. These questions are somewhat similar to the "Security Considerations" currently required in IETF documents. [RFC2316].

This document is inspired by and directly modeled on [RFC 3426], entitled "General Architectural and Policy Considerations" and published by the Internet Architecture Board in November 2002. In RFC 3426, the IAB suggests architectural questions that should be considered in design decisions, without asserting that there are clear guidelines that should be followed in all cases. This document attempts to follow in the spirit of RFC 3426 by raising questions to be considered without asserting that any particular answers must be followed.

This document is motivated by the recognition that technical design decisions made within the IETF and other standards bodies can have significant impacts on public policy concerns. One well known example of this possible impact can be found in the implementation of IPv6 on Ethernet networks. [RFC 2464], published in December 1998, specified that the IPv6 address for a computer on an Ethernet network would incorporate the unique MAC address associated with the Ethernet adapter. After the publication of RFC 2464, a significant policy concern arose because the use of the unique and unchangeable MAC address would significantly reduce a user's ability to conduct private and/or anonymous communications using IPv6. The IETF responded to those concerns by publishing in January 2001 [RFC 3041], entitled "Privacy Extensions for Stateless Address Autoconfiguration in IPv6".

The goal of this document is that potential public policy impacts of technical design decisions will be identified and considered during the initial design process. This type of policy consideration already happens in many cases within the IETF, but not in any systematic way or with any assurance that public policy concerns will be identified in most cases.

A common assertion within the IETF is that "we don't do public policy." The goal of this document is NOT to suggest that the IETF should "do" policy in the sense of intentionally conducting extensive debates on public policy issues. But, as much as the IETF appropriately tries not to "do" policy, many of its actions and decisions squarely and significantly impact on public policy concerns. This document seeks to encourage the IETF to acknowledge those times when a design decision might affect a policy concern, so that the community can make a reasoned decision on whether and how to address the concern in the particular situation.

The authors see a range of important reasons why the IETF should seek to be aware of the potential public policy impacts of its design decisions, but will only suggest one here: The chance that IETF standards will be widely deployed and then widely accepted in the market is higher if those standards minimize harmful social impacts.

To be clear, it is not the view of the authors that impact on public policy concerns must be avoided at all costs (if that were even possible), or that if a particular proposed standard adversely affects a public concern (say, privacy) that the standard should as a matter of course be rejected. Some beneficial technologies might unavoidably have secondary harmful impacts, and the benefits may outweigh the harms. More generally, some technologies (such as those that facilitate government surveillance) might intentionally compromise a public concern such as privacy. Similarly, the inherent goal of some technologies (such as those that discriminate among traffic to provide assured levels of quality of service) might simultaneously be viewed by some as beneficial and by others as harmful.

In all of these cases, there may well be good reasons to develop the technology notwithstanding the asserted harms to a policy concern. The central goal of this paper is simply to suggest that impacts on a public concern should not happen without clear recognition of the impacts, and without appropriate consideration of whether it is possible to minimize harmful impacts while still meeting the design requirements.

2. Scope of this Document

The purported scope of this document is admittedly ambitious. This document cannot possibly predict and identify all possible societal impacts of future IETF design decisions. It does try, however, to identify a broad range of possible public policy impacts that experience suggests are most likely to arise. This document is a first draft, and will require at least a few iterations before it covers a reasonably full range of potential issues. The authors invite comments, and specifically seek suggestions of specific historic examples within the IETF where a public policy concern was raised by a technology proposal. The more concrete examples this document can contain, the more useful it will be to the IETF community.

There are two broad categories of public policy impacts that this document does NOT seek to cover with any thoroughness. First, this document does not articulate the full range of concerns raised by traditional security problems in the network. The IETF is already appropriately focused on security issues, and those in the Security Area are well able to identify and articulate the types of technical design decisions that can lead to security problems. Many of the privacy concerns highlighted in this document raise related security concerns.

Second, this document does not attempt to identify the enormous range of POSITIVE societal impacts that flow from network technology. The vast majority of the work of the IETF -- from the introduction of an entirely new method of Internet use to the fine tuning of an existing routing protocol -- yields concrete and important social benefits. This document does not discuss these positive benefits, but takes as a given that technology proposals will not advance within the IETF unless at least some portion of the community views the proposals as beneficial.

This document is by no means an exhaustive list of public policy concerns that relate to the Internet. This draft has instead focused on policy issues that the authors believe are most likely to arise in the IETF context. In addition, the authors articulate a perspective that is in part a function of their country and culture. It is the authors' hope that over time this document can be expanded to include the concerns of a broader set of communities.

3. A Few Basic Definitions

This document will use a limited number of defined terms, which admittedly will not be precisely applicable in all situations:

TECHNOLOGY shall refer to a technical standard or innovation being considered within the IETF, whether it is a "new" technology or standard or a modification to an "old" technology or standard.

END USER shall refer to the user at one or the other end of a network communication, or an automated or intelligent proxy for a user located at the end of the communication. Thus, a concern over, for example, the privacy of the End User would be applicable in cases where a client-side application communicated on behalf of an End User. In some contexts, a corporation or other organized collection of human users might stand in the role of an End User. In some but not all contexts, a communication might be from one End User to another End User; in other context, a communication might be between a Service Provider (defined below) and an End User.

ACCESS PROVIDER shall refer to the entity that most directly provides network access to an End User or Service Provider. In the case of End Users on the public Internet, a Access Provider will often be an Internet Service Provider that provides dedicated or dial-up network access. In other cases a Access Provider might be a company providing access to its employees, or a university providing access to its students and faculty.

SERVICE PROVIDER shall refer to an entity (human, corporate or institutional) that provides or offers services or content to End Users over the network (regardless of whether charges are sought for such services or content). Thus, for example, a web site would be viewed as a Service Provider.

A given entity (such as a company offering content on the web) might be viewed as an Access Provider (vis-a-vis its employees), as an End User (vis-a-vis the ISP from which it obtains network access), and as a Service Provider (vis-a-vis End Users elsewhere on the Internet).

TRANSIT PROVIDERS shall refer to one or more entities that transport communications between the Access Providers at either end of a communication. Transit Providers are often thought to transport packets of communications without regard to their content (other than, of course, their destination), but increasingly some Transit Providers may handle traffic differently depending on the type of traffic.

THIRD PARTY shall refer to any individual or entity other than End Users, Access Providers, Service Providers, and Transit Providers. For a given communication, Third Parties could include, for example, governments seeking to execute lawful interceptions, hackers seeking to interfere with or intercept communications, or in some situations entities that provide, under contract, content or functionality to a Service Provider (such as, for example, an entity that serves advertisements for insertion in a web page).

In some cases the distinction between a Transit Provider and a Third Party may blur, if the Transit Provider manipulates or discriminates among traffic based on characteristics such as its content, sender, or receiver. Similarly, the line between a Service Provider and a Third Party may blur as more service functions are contracted out.

4. Questions about Technical Characteristics or Functionality

In this section we list questions to ask in designing protocols. The issues raised by the questions are discussed in more depth in Section 5 below. We are not suggesting that each of these questions requires an explicit answer -- some questions will be more relevant to one design decision than to another.

There is not a one-to-one correspondence between the questions listed in this section and the discussions in Section 5. Instead, for each group of questions listed below, there are one or more references to later substantive discussions.

Some of the questions will be easy to answer for a given technology. Others will require creative thinking to assess whether a proposed technology might be misused to achieve a result not intended by the technology proponents.

This first draft addresses the most common and well-known areas of public policy concern, focusing on areas most likely to arise in the IETF context. Subsequent drafts may include a broader range of policy concerns.

Bottlenecks, Choke-Points and Access Control:

* Would the Technology facilitate any bottlenecks or choke-points in the network through which significant amounts of particular types of traffic must flow?

* Would the Technology permit a Third Party (including a government) to exert control over End Users' use of the Internet as a whole?

* Would the Technology permit a Transit Provider or Third Party (including a government) to exert control over the use of particular content, functionality, or resources?

* Would the Technology permit an Access Provider or Service Provider to exert control over particular content, functionality, or resources, other than that known by the End User to be controlled by the Access Provider or Service Provider?

* Would the Technology permit Third Party (including a government) to require that particular content or functionality be confined (or "zoned") into, or excluded from, any particular subpart of the Internet (such as a particular Global Top Level Domain)?

> See discussions of "Content Censorship and Control," "Personal Privacy," "Discrimination Among Users and Content," "Competition and Choice," and "User Consent."

Alteration or Replacement of Content:

* Would the Technology permit a Third Party to alter any of the content of a communication without (a) the express instruction or consent of the Service Provider and the End User, or (b) the knowledge of the Service Provider or the End User?

See discussions of "Content Censorship and Control" and "User Consent."

Monitoring or Tracking of Usage:

* Would the Technology permit the monitoring or tracking by a Third Party of the use of particular content, functionality, or resources?

See discussion of "Personal Privacy."

Retention, Collection, or Exposure of Data:

* Would the Technology require or permit the retention of any information about individual packets or communications, or individual End Users, either (a) beyond the conclusion of the immediate network or communications event, or (b) for longer than a reasonably brief period of time in which a communications "session" can be concluded?

* Would the Technology permit the reading or writing of any file on an End User's computer without the explicit knowledge of the End User?

* Would the Technology permit or require that information other than location and routing information (such as, for example, personal information or search terms) be made a part of a URL or URI used for a communication?

* Would the Technology permit or require that personal or confidential information be made available to any Third Party, Transit Provider, or Access Provider?

See discussion of "Personal Privacy."

Persistent Identifiers and Anonymity:

* Would the Technology require or permit the association of a persistent identifier with a particular End User, or a computer used by one or more End Users?

* Would the Technology reduce the ability of a content provider to provide content anonymously?

* Would the Technology reduce the ability of an End User to access content or utilize functionality anonymously?

See discussion of "Personal Privacy."

Access by Third Parties:

* Would the Technology permit any Third Party to have access to packets to and from End Users without the explicit consent of the End Users?

* Would the Technology permit or require any Third Party to store any information about an End User, or an End User's communications (even with the knowledge and consent of the End User)?

See discussions of "Personal Privacy" and "User Consent."

Discrimination among Users, or among Types of Traffic:

* Would the Technology require or permit an Access Provider or Transit Provider to provide differing levels of service or functionality based on (a) the identity or characteristic of the End User, or (b) the type of traffic being handled?

* Would the Technology likely lead to a significant increase in cost for basic or widely-used categories of communications?

* Would likely implementations of a new mode of communication require such a financial or resource investment so that the mode would effectively not be available to individuals, or small or non-profit entities?

See discussion of "Discrimination Among Users and Content."

Internationalization and Accessibility

* Would the Technology function with the same level of quality, ease of use, etc., across a broad range of languages and character sets?

* Would the likely implementations of the Technology be as useful to mainstream End Users as to non-mainstream End Users (such as, for example, End Users with disabilities)?

* Would the Technology likely reduce the ability of non-mainstream End Users (such as, for example, End Users with disabilities) to utilize any common application or network functions?

See discussions of "Internationalization" and "Accessibility."

Innovation, Competition, and End User Choice and Control

* Would the Technology reduce the ability of future designers to create new and innovative uses of the Internet, or new methods to accomplish common network functions?

* Would the Technology likely reduce the number of viable competitive providers of any common application or network functions?

* Would the Technology likely reduce the ability of small or poorly-funded providers to compete in the provision of any common application or network functions?

* Would the Technology likely reduce the number or variety of methods available to the End User to accomplish common application or network functions?

* Would the Technology likely reduce the level of control the End User can exercise over common application or network functions?

See discussion of "Competition and Choice."

5. Discussion of Potential Public Policy Concerns

Below are brief discussions of common categories of public policy concern that might be raised by technologies developed by the IETF. The discussions are not intended to present comprehensive analyses of the policy concern, but are intended to assist in identifying situations in which the concern is implicated and should be considered.

A. General Comments

The fundamental design principles of the Internet, including openness, interoperability, and the end-to-end principle, have themselves been critical contributors to the value of the Internet from a public policy perspective. Thus, as a first rule of promoting healthy public policy impacts, the IETF should continue to maintain and promote the architectural goals that it has historically pursued.

Because of this congruence between architectural values and public policy values, many of the design considerations in RFC 3426, "General Architectural and Policy Considerations," directly promote an Internet that is supportive of good public policy values. As one of many examples, Section 12.1 discusses the value of user choice, and quotes [CWSB02] to say that "user empowerment is a basic building block, and should be embedded into all mechanism whenever possible." User choice is a fundamental public policy concern, discussed more below.

[CWSB02], titled "Tussle in Cyberspace: Defining Tomorrow's Internet," is itself a valuable exploration of the intersection between technology design and public policy concerns. A key premise of [CWSB02] is that "different stakeholders that are part of the Internet milieu have interests that may be adverse to each other, and these parties each vie to favor their particular interests." Many of the "tussles" that [CWSB02] analyzes are situations in which public policy considerations should be assessed in making design decisions. More broadly, [CWSB02] provides to technology designers a conceptual framework that recognizes the existence of "tussles" and seeks to accommodate them constructively within a design.

B. Content Censorship and Control

As used here, the concept of censorship can encompass both governmental and private actions.

5.1.1 Government Censorship

"Censorship" is most commonly thought of as government-imposed control or blocking of access to content. Many believe that as a matter of public policy, censorship should be minimized or avoided. For example, in May 2003 the Council of Europe stated in its "Declaration on freedom of communication on the Internet" that "Public authorities should not, through general blocking or filtering measures, deny access by the public to information and other communication on the Internet, regardless of frontiers." [COE03]. But not all censorship is viewed by all as contrary to public policy. In November 2002 in [COE02], the same Council of Europe specifically endorsed government regulation of "hate speech" on the Internet.

Some technology is intended to control access to content. The Platform for Internet Content Selection of the World Wide Web Consortium, [PICS], for example, was in part designed to facilitate the limitation of access by some users (children, for example) to certain types of content.

Harder to identify are technologies not intended for content control but which can be used to censor or restrict access to content. Any technology that creates or permits bottlenecks or choke-points in the network, through which significant traffic must pass, increases the risk of censorship. Governments seeking to censor content or restrict access to the Internet will exploit network bottlenecks (albeit often bottlenecks created by network topology not technology standards). [ZE02] documents Saudi Arabia's routing of all Internet traffic through central proxy servers, and [KB01] discusses the response of China and Cuba to the Internet, to achieve such ends.

Governments also seek to control access to content through means other than direct censorship. In the United States, for example, [CIPA] requires that libraries that receive certain government funding must use content filtering technology on Internet access they offer to patrons, and [DOTKIDS] requires the creation of a subdomain of the .US domain to be used only for children-suitable content.

5.1.2 Private Control of Content

Governments are not the only entities that attempt to restrict the content to which Internet users have access. In some cases Access Providers (commonly Internet Service Providers) seek to control the content available to their customers. Some do so with full knowledge and consent of the customers (to provide, for example, a "family friendly" online experience). Others, however, favor certain content (for example, that of contractual business partners) over competing content, and do so without the clear understanding of their customers.

Whether such private content control is contrary to public policy will turn on a host of specific considerations (including notice and alternative choice), but undeniably such content control raises policy concerns. [CMCS02] illustrates, for example, the current debate over "network neutrality" in the United States. These policy concerns are commonly phrased in terms of discrimination among content, and are discussed more fully in the next section.

C. Discrimination Among Users and Content

In a simplistic conception of the early Internet, all traffic of any kind was broken into packets and all packets were treated equally within the network. This idea has promoted a broad and strong perception of equality within the Internet -- one class of traffic will not take priority over other classes, and a lone individual's packets will be treated the same as a large corporation's packets.

Any technology that moves away from this notion of equality -- even technologies that are clearly beneficial -- raise significant public policy questions, including "who controls the preferential treatment," "who qualifies for it," "will it require additional expenditure to obtain it," and "how great a disparity will it create."

Thus, for example, quality of service and content distribution networks both raise questions about what and who will be favored, whether the rough equality of the Internet will be lost, and whether the financially strong will come to dominate the Internet and make it less useful for the less well off. [BM00], for example, explores the policy concerns raised by content distribution networks.

The concern over discrimination addresses both discrimination based on identity of user, and on type of traffic. Content distribution networks enable, for example, individual web sites able to afford the CDN services to be delivered more quickly than competing web sites that are not able to afford such services. In contrast, a core focus of quality of service efforts is on the need to provide enhanced levels of service to some types of traffic (e.g., Internet telephony).

Concern about discrimination does not suggest that technologies that handle certain categories of traffic more efficiently should never be pursued. The concern, however, may in some cases suggest that an efficiency enhancement be structured so as to be available to the broadest classes of traffic or users.

D. Competition and Choice

Critical elements of the Internet's development and success have been the ability to create new and innovative uses of the network, the relative ease in creating and offering competitive services, products, and methods, and the ability of Internet users to choose from a range of providers and methods. Anything that reduces innovation, competition, or user choice raises significant public policy concerns.

Indeed, the need for competition and user choice is perhaps greater now than in earlier days of the Internet. There is a greater divergence today in the interests and agendas of users and service providers than in the past, and that divergence makes it more important that users be able to choose among service providers (in part to seek providers that they trust the most).

[CWSB02] extensively addresses the important need for competition and user choice, and provides detailed suggestions and guidelines for Internet designer to consider.

E. User Consent

A familiar public policy concern over user consent focuses on the use of personal data (as discussed more fully below under "Privacy"). The usage here, however, has a broader meaning: the consent (or lack of consent) of a user regarding an action or function executed by or within the network.

Many actions performed using IETF protocols require the specific initiation by a user, and the user's consent can fairly be assumed. Thus, if a user transmits a request using SIP, the Session Initiation Protocol, it is safe to assume that the user consents to the normal handling and execution of the SIP request.

Other actions performed using IETF protocols are not initiated by a user, but are so inherently a part of normal network operations that consent can be assumed. For example, if in the middle of the network certain packets are slowed by congestion, it is safe to assume sufficient consent for congestion control mechanisms and rerouting of the packets.

Uncertainty about consent arises, however, in areas where IETF protocols can be viewed as deviating from some conception of "normal." A simple example relates to the evolution of caching, where as caching of various types of data became the norm, there emerged a need to be able to set flags to prevent caching, which in a sense can be thought of as a form of negative consent.

Middle boxes and other functions that deviate from the historic "norm" -- the end-to-end principle -- also can raise issues of consent. For example, section 3 of [RFC3238], titled "IAB Architectural and Policy Considerations for Open Pluggable Edge Services," explores a range of consent and data integrity issues raised by the OPES protocol proposals. As that analysis makes clear, the consent issue is not necessarily confined to the consent of the client in a client/server transaction, but may also involve the consent of the server to an action undertaken on the request of the client.

F. Internationalization

[RFC3426] calls on protocol designers to ask the key question about "Internationalization":

"Where protocols require elements in text format, have the possibly conflicting requirements of global comprehensibility and the ability to represent local text content been properly weighed against each other?"

[RFC3426] explores the significant challenges raised by the need to balances these conflicting goals, and raises the possibility that the historic preference for the use of case-independent ASCII characters in protocols may need to change to accommodate a broader set of international languages.

G. Accessibility

The concept of "accessibility" addresses the ability of persons with disabilities to use the Internet in general and the full range of applications and network functions that are commonly available to persons without disabilities.

Although focused on the World Wide Web, [W3C WAI-TA] illustrates the concern and explains that a focus on accessibility is needed "to ensure that the full range of core technologies of the Web are accessible . . . . Barriers exist when these technologies lack features needed by users with visual, hearing, physical, cognitive or neurological disabilities, or when the accessibility potential in the technology is not carried through into the Web application or Web content. For instance, in order for a multimedia presentation to be accessible to someone who is blind, the markup language for the presentation must support text equivalents for images and video; the multimedia player used must support access to the text equivalents; and the content author must make appropriate text equivalents available. These text equivalents can then be rendered as speech or braille output, enabling access to the content regardless of disability or device constraints."

Many policy concerns about accessibility relate specifically to the user interfaces used by applications, and as such these concerns generally fall outside of the province of the IETF. But in the Applications Area and to a lesser extent elsewhere within the IETF, some design decisions could ultimately constrain the accessibility of applications based on IETF protocols.

The World Wide Web Consortium's Web Accessibility Initiative [W3C WAI] reflects a very well developed and comprehensive analysis of the technical and design issues raised by accessibility concerns.

H. Personal Privacy

Individual privacy concerns are often divided into two components: First, "consumer privacy" (also termed "data protection") commonly addresses the right of individuals to control information about themselves generated or collected in the course of commercial interactions. Second, "privacy rights vis-a-vis the government" addresses individuals' protection against unreasonable government intrusions on privacy, including the interceptions of communications.

In the IETF context, a third category of privacy concern -- privacy against private interception of or attacks on data or communications -- is largely covered by the IETF's focus on security considerations. Although security considerations are crucial to privacy considerations, "consumer privacy" and "privacy vis-a-vis the government" raise significantly different issues than traditional security considerations. With security considerations, a key focus is on maintaining the privacy of information against unauthorized attack. Other forms of privacy, however, focus not on unauthorized access to information, but on the "secondary use" of information for which access was (at least temporarily) authorized. The question often is not "how can I keep you from seeing my information" but "how can I give you my information for one purpose and keep you from using it for another."

The questions raised in Section 4 above do not differentiate between the different categories of privacy, because for most purposes within the IETF, technologies that create risk for one type of privacy likely also create risk for other types of privacy. Once a potential privacy concern is identified, however, the different types of privacy concern may present different public policy considerations. Indeed, the policy considerations may well be in tension -- a technology that permits a lawful governmental interception of a communication may also increase the risk of unlawful private interception.

Privacy considerations are too numerous and multifaceted to be adequately addressed in this document. The discussion below only briefly covers the key privacy issues. A forthcoming Internet-Draft on "Privacy Considerations for Internet Protocols" will address privacy issues more thoroughly.

5.1.3 Consumer Privacy and Data Protection

Consumer privacy protection in many parts of the world is based on "fair information practices," which were authoritatively detailed in [OECD] by the Organization for Economic Co-operation and Development. Fair information practices include the following principles:

* Notice and Consent - before the collection of data, the data subject should be provided: notice of what information is being collected and for what purpose and an opportunity to choose whether to accept the data collection and use. In Europe, data collection

cannot proceed unless data subject has unambiguously given his consent (with exceptions).

* Collection Limitation - data should be collected for specified, explicit and legitimate purposes. The data collected should be adequate, relevant and not excessive in relation to the purposes for which they are collected.

* Use/Disclosure Limitation - data should be used only for the purpose for which it was collected and should not be used or disclosed in any way incompatible with those purposes.

* Retention Limitation - data should be kept in a form that permits identification of the data subject no longer than is necessary for the purposes for which the data were collected.

* Accuracy - the party collecting and storing data is obligated to ensure its accuracy and, where necessary, keep it up to date; every reasonable step must be taken to ensure that data which are inaccurate or incomplete are corrected or deleted.

* Access - a data subject should have access to data about himself, in order to verify its accuracy and to determine how it is being used.

* Security - those holding data about others must take steps to protect its confidentiality.

Many of these fair information practices are relevant to IETF protocols. Even seemingly benign data and server logs can reveal important information about individuals users. It is not sufficient to address the risk of a technical attack on a body of data, because privacy considerations must address the risk of non-technical attacks on data (through legal process, rogue employees, etc.).

5.1.4 Privacy vis-a-vis the Government

Although privacy is internationally recognized as a human right, most governments claim the authority to invade privacy through the following means, among others:

* interception of communications in real-time;
* interception of traffic data (routing information) in real-time;
* access to data stored by service providers, including traffic data being stored for billing purposes; and
* access to data stored by users.

These means of access to communications and stored data should be narrowly defined and subject to independent controls under strict standards. Real-time interception of communications should take place only with prior approval by the judicial system, issued under standards at least as strict as those for police searches of private homes.

In 1999, in the "Raven" discussions, the IETF considered whether it should take action to build wiretapping capability into the Internet. Ultimately, as detailed in [RFC2804], the community decided that an effort to build wiretapping capability into the Internet would create significant and unacceptable security risks.

6. Conclusion

This document has sought to identify a range of public policy concerns that may arise in the work of the IETF. The authors invite comments and suggestions about ways to make this document more useful and complete.

Security Considerations

This document does not propose any new protocols or changes to old protocols, and therefore does not involve any security considerations in that sense. Many of the privacy issues discussed here also raise security issues, but this document is not intended to be a comprehensive look at security issues.

References

[BM00] Berman, J. & Morris, J., "The Broadband Internet: The End of the Equal Voice?", Computers, Freedom & Privacy Conference, April 2000. URL "http://www.cdt.org/publications/broadbandinternet.pdf".

[CIPA] United States Congress, "Children's Internet Protection Act", December 2000. URL "http://www.cdt.org/legislation/106th/speech/001218cipa.pdf".

[COE02] Council of Europe, "Additional Protocol to the Convention on Cybercrime concerning the criminalisation of acts of a racist and xenophobic nature committed through computer systems," November 7, 2002. URL "http://www.coe.int/T/E/Legal_affairs/Legal_co-operation/Combating_economic_crime/Cybercrime/Racism_on_internet/PC-RX(2002)24E-1.pdf".

[COE03] Council of Europe, "Declaration on freedom of communication on the Internet," May 28, 2003. URL "http://cm.coe.int/stat/E/Public/2003/adopted_texts/declarations/dec-28052003.htm".

[CMCS02] Cooper, M., Murray, C., Chester, J., & Schwartzman, A., Letter to High-Tech Broadband Coalition, August 16, 2002. URL "http://www.mediaaccess.org/programs/broadband/chesterltr090302.pdf".

[CWSB02] Clark, D., Wroslawski, J., Sollins, K., and Braden, R., "Tussle in Cyberspace: Defining Tomorrow's Internet", SIGCOMM 2002. URL "http://www.acm.org/sigcomm/sigcomm2002/papers/tussle.html".

[DOTKIDS] United States Congress, "Dot Kids Implementation and Efficiency Act of 2002", November 2002. URL "http://www.kids.us/content_policy/kids_efficiency_act.pdf".

[KB01] Kalathil, S. & Boas, T., "The Internet and State Control in Authoritarian Regimes: China, Cuba, and the Counterrevolution", July 2001. URL "http://www.ceip.org/files/pdf/21KalathilBoas.pdf".

[OECD] Organization for Economic Co-operation and Development, "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data," 1980. URL "http://www.oecd.org/EN/document/0,,EN-document-0-nodirectorate-no-24-10255-0,00.html".

[PICS] World Wide Web Consortium, "Platform for Internet Content Selection." URL "http://www.w3.org/PICS/".

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

[RFC2316] Bellovin, S., "Report of the IAB Security Architecture Workshop", RFC 2316, April 1998.

[RFC2464] Crawford, M., "Transmission of IPv6 Packets Over Ethernet Networks", RFC 2464, December 1998.

[RFC2804] IAB & IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000.

[RFC3041] Narten, T. & Draves, R., "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[RFC3238] Floyd, S. & Daigle, L., "IAB Architectural and Policy Considerations for Open Pluggable Edge Services", RFC 3238, January 2002.

[RFC3426] Floyd, S., ed., "General Architectural and Policy Considerations," RFC 3426, November 2002.

[W3C WAI]   World Wide Web Consortium, "Web Accessibility Initiative". URL "http://www.w3.org/WAI/".

[W3C WAI-TA]   World Wide Web Consortium, "WAI Technical Activity". URL "http://www.w3.org/WAI/Technical/Activity.html".

[ZE02]   Zittrain, J. & Edelman, B., "Documentation of Internet Filtering in Saudi Arabia," September 2002. URL "http://cyber.law.harvard.edu/filtering/saudiarabia/" .

Acknowledgments

The authors would like to thank Sally Floyd and the Internet Architecture Board for devising the approach to policy considerations used in [RFC 3426], which this document tries to follow.

Authors' Addresses

John B. Morris, Jr.
Center for Democracy & Technology
1634 I Street, NW, Suite 1100
Washington, D.C. 20006
USA
Email: jmorris@cdt.org

Alan B. Davidson
Center for Democracy & Technology
1634 I Street, NW, Suite 1100
Washington, D.C. 20006
USA
Email: abd@cdt.orgs

PLEASE SEND COMMENTS AND SUGGESTIONS TO jmorris@cdt.org