The Broadband Internet: The End of the Equal Voice?

John B. Morris, Jr. [*] **and Jerry Berman** [**]
Broadband Access Project

The Center for Democracy & Technology [***]

As the fast-moving and hard-fought "open access to cable" debate continues – in both the United States and Canada – and perhaps moves toward resolution, it is vital to recognize that there are significant "openness" and free speech issues concerning broadband Internet access that have little or nothing to do with the cable debate. This essay looks at one such issue – an issue that is only now beginning to take shape. As described below, the emerging content distribution model on the Internet could diminish or eliminate the rough "equality of voice" between small and large speakers that is a key characteristic of the narrowband Internet. Unless those involved in creating and shaping the Internet – from network engineers to corporate leaders to public policy advocates – take steps to address this issue, we risk seeing changes in the Internet that could threaten the legal conclusion that speech on the Internet deserves the highest level of protection that the United States Constitution can afford.

When the United States District Court for the Eastern District of Pennsylvania undertook in 1996 the first comprehensive assessment of the narrowband Internet by an American court, it found what it termed "a unique and wholly new medium of worldwide human communication."[1] One key characteristic of the Internet that led the court to its conclusion was the rough "equality of voice" that exists in the narrowband Internet between small speakers and large corporate or government-controlled speakers. As Judge Stewart Dalzell of the Eastern District phrased it, "the Internet provides significant access to all who wish to speak in the medium, and even creates a relative parity among speakers."[2]

The broadband Internet, as it is now evolving, may undermine this equality of voice between small and large speakers on the Internet, and that may in turn chip away at the foundation of the sweeping First Amendment protections that speech on the Internet has been afforded by courts in the United States. The World Wide Web of the future may be one in which only large and wealthy speakers can afford to offer broadband, bandwidth-intensive, speech, while smaller speakers and publishers are relegated to offering more static and passive speech.[3]

This essay looks at the narrowband Internet and its legal context, reviews the development of the distributed broadband content delivery model, and assesses its potential impact on the ability of small speakers to speak and be heard. The essay raises questions that both policy advocates and network

engineers must address, and urges the development of a consultative process to ensure that the unique characteristics of the Internet are protected.

The Narrowband Internet and the *Reno* Court's Conclusions

In the mid-1990's, the Internet moved beyond its academic and governmental origins, and became a popular and commercial medium, with the vast majority of individual Internet users accessing the network over "narrowband" dial-up connections. The top speed of users' "last mile" connection inched up from 9.6 to 14.4 to 28.8 to 56 kbps, but the relative slowness of these connections imposed practical limits on how bandwidth-intensive any given site on the World Wide Web could be. Although sites could offer graphics-intensive Web pages, the length of time it would take to receive and view all of the graphics would often deter listeners.

The relative narrowness of the last-mile connection to Internet listeners in turn led to a "relative parity" among Internet speakers – there simply was no great advantage that money could buy. Smaller, start-up Web sites could offer content just as flashy and current as the largest corporate speaker. An individual critical of a corporation, for example, could post a web page with just as much impact as that posted by the corporation itself. Although a corporate Web site may well have more server capacity and greater bandwidth to the Internet backbone (enabling the site to respond to more simultaneous visitors), the two web sites could nevertheless speak with the same basic quality and impact.[4] This rough "equality of voice" made the Internet unique among means of mass communications – for the first time, individual and small speakers and publishers could speak to vast numbers of listeners, and could do so with content able to compete with the largest speakers. Moreover, this rough equality of voice could be achieved for a very low amount of money – individuals could post personal web pages for little or no cost, and a web site with a unique domain name could be hosted for a very low investment.

The opportunity of small and underfunded speakers and publishers to reach a wide audience on the Internet – and to do so with a rough equality of voice – made a significant impression on the three judge U.S. court that evaluated the Internet in the 1996 challenge to the Communications Decency Act.[5] In its Findings of Fact, the court concluded:

> 75. The Internet is not exclusively, or even primarily, a means of commercial communication. . . . For the economic and technical reasons set forth in the following paragraphs, the Internet is an especially attractive means for not-for-profit entities or public interest groups to reach their desired audiences. . . .
>
> 76. Such diversity of content on the Internet is possible because the Internet provides an easy and inexpensive way for a speaker to reach a large audience, potentially of millions. The start-up and operating costs entailed by communication on the Internet are significantly lower than those associated with use of other forms of mass communication, such as television, radio, newspapers, and magazines. This enables operation of their own Web sites not only by large companies . . . but also by small, not-for-profit groups . . . .
>
> . . . .
>
> 79. Because of the different forms of Internet communication, a user of the Internet may speak or listen interchangeably, blurring the distinction between "speakers" and "listeners" on the Internet. . . .

> 80. It follows that unlike traditional media, the barriers to entry as a speaker on the Internet do not differ significantly from the barriers to entry as a listener. Once one has entered cyberspace, one may engage in the dialogue that occurs there. In the argot of the medium, the receiver can and does become the content provider, and vice-versa.
>
> 81. The Internet is therefore a unique and wholly new medium of worldwide human communication.[6]

One of the judges on the District Court panel, Judge Dalzell, further explored the significance of the ability of small speakers to speak on the Internet, and concluded that the "Internet is a far more speech-enhancing medium than print, the village green, or the mails."[7] Judge Dalzell summarized the most critical factual findings of the three judge court:

> Four related characteristics of Internet communication have a transcendent importance to our shared holding that the CDA is unconstitutional on its face. We explain these characteristics in our Findings of fact above, and I only rehearse them briefly here. First, the Internet presents very low barriers to entry. Second, these barriers to entry are identical for both speakers and listeners. Third, as a result of these low barriers, astoundingly diverse content is available on the Internet. Fourth, the Internet provides significant access to all who wish to speak in the medium, and even creates a relative parity among speakers.[8]

In considering the Communications Decency Act that was before the court, Judge Dalzell sought to avoid "an Internet that mirrors broadcasting and print, where economic power has become relatively coterminous with influence."[9]

A critical issue facing both the three judge District Court panel and the United States Supreme Court on appeal was the level of First Amendment protection that should be afforded to the Internet. The debate centered on whether the Internet should receive the very high level of constitutional protection given to the print medium, or whether it should be subjected to a lower level of protection, as is the broadcast medium. Based on his analysis of the Internet, and in particular its speech-enhancing characteristics, Judge Dalzell concluded that under applicable First Amendment jurisprudence, the Internet deserved a level of constitutional protection *even higher* than that afforded to print; he concluded that protected speech on the Internet simply could not be regulated by Congress.[10]

Consistent with the U.S. Supreme Court's preference to decide cases as narrowly as possible, the Court did not reach the ultimate issue raised by Judge Dalzell – whether the Internet deserved protection even higher than that afforded to print. Writing for the Supreme Court, Justice Stevens acknowledged Judge Dalzell's conclusions, and expressly indicated that the high Court was not reaching the question.[11] Thus, the Supreme Court left open for another case the question of the full scope of the First Amendment protection that should be afforded to the Internet.

The Emerging Broadband Internet

Broadband technology is fundamentally changing the distribution of content over the Internet. In the narrowband Internet, there was little incentive to optimize – in terms of network infrastructure – the distribution of content to the end user, because the end user's "last mile" connection was so "narrow" that content could be served by a single Web server[12] to end users around North America[13] about as fast as the end users could receive it.[14] In the emerging broadband world, however, a single server in Reston or San Jose or Peoria can no longer efficiently and effectively serve high-bandwidth content to users all around the Internet.

There are a number of interrelated reasons for the inadequacy of a single server for broadband content rich in multimedia and graphics.[15] First, by definition a single server is distant from many of the end users of the content, and the latency or lag time that is inherent in delivering content a great distance over the Internet reduces the quality of delivery of broadband content to the users. Second, attempting to serve – from a single server – high-bandwidth content simultaneously to many users all around the country would require a greater investment in server capacity than even some large corporations would choose to make. Finally, the cost of transmitting high-bandwidth content over the Internet backbone to users located across the Internet can be substantial.

Thus, for a host of reasons, broadband content is leading to the development and refinement of a new model of content distribution – the distributed content model. Under the distributed content model, a skeleton HTML page is commonly served from the same "single" server discussed above (located in Reston, San Jose, Peoria, or wherever), but the broadband content that is embedded with the HTML page (graphic images, video and audio clips, dynamic content, etc.) is served from a distributed network of servers located all around the United States and across the world, so that the high-bandwidth content is served to a particular user from a server located as close as possible to the user. Thus, for a hypothetical company located in Peoria, the company might create and serve a Web site from its home city, but the company would contract with a distributed content network to serve the company's video and audio content from servers located close the end users.

Two graphics from the Web site of one of the leading content distribution companies – Akamai Technologies, Inc. – illustrate the "before" and "after" pictures of the Internet. In Figure 1[16] (showing traffic flow without the benefit of Akamai's "FreeFlow" distributed content service), Web content is served using the traditional narrowband "single server" approach:

Figure 1: "Internet Content Delivery Without FreeFlow"

In contrast, in Figure 2 (showing an implementation of Akamai's distributed content services), the high-bandwidth content is served to the end user from a local Akamai server instead of from the original company's server located near the company's home office:

Figure 2: "Internet Content Delivery With FreeFlow"

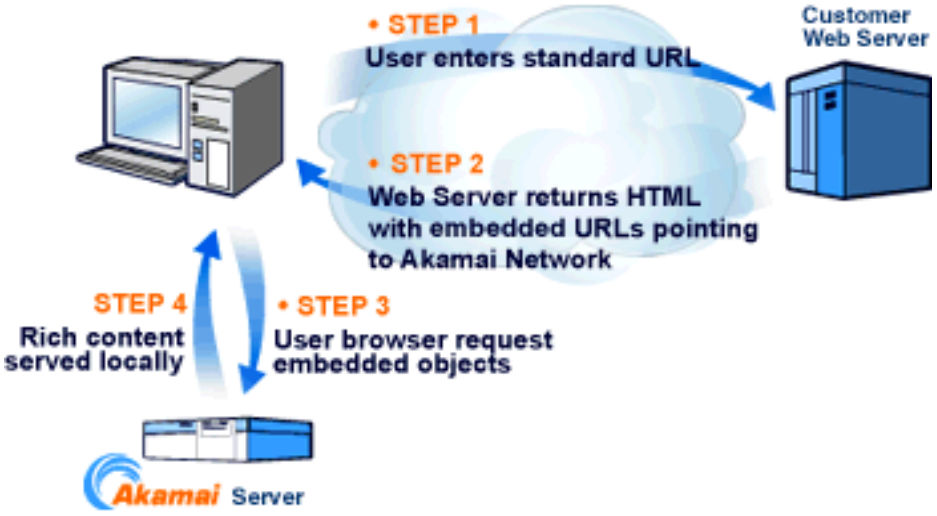

Under this distributed content model, the high-bandwidth video and audio content is served quickly and reliably to the end user. The ultimate "Web viewing" experience of the end user is likely to be significantly better using the distributed content model than with the more traditional single server model. Akamai is far from the only company offering distributed content services. Numerous other large and small companies are offering such services, including for example the Intel Corporation and INTERVU, Inc. (in which the Microsoft Corporation has made significant investment).

Moreover, the largest "last mile" broadband provider in the United States – the At Home Corporation providing Excite@Home services to over one million cable subscribers in North America – utilizes the distributed content model in its nationwide network of cable modem systems. According to its Web site, "Excite@Home uses a hierarchical, distributed network architecture with proprietary caching and replication technology to ensure that the information a user wants is always 'as close as possible' within the network."[17]

Increasingly, content on the World Wide Web will be served from a combination of a "single" server directly controlled by the Web site owner with a network of distributed content servers controlled by Akamai, Excite@Home, or another distributed content company. This new distribution model will likely mean a smoother and more efficient experience for the Web surfers visiting those Web sites. But, as discussed below, it could leave out, and thus disadvantage, the speech of smaller speakers and publishers.

A possible – but by no means inevitable – result of the emergence of broadband and distributed content networks is the loss of the rough "equality of voice" between large and small speakers. The distributed content model is still in its infancy, and numerous questions about whether and how different (and often competing) distributed content networks will interconnect have yet to be answered. It is quite easy to envision, however, a world in which it is relatively costly to have one's high-bandwidth multimedia-rich broadband content efficiently and smoothly distributed to Internet users, and only the better-funded speakers will be able to afford to have their broadband content "distributed." In stark contrast to today's narrowband Internet (where small speakers can publish low-bandwidth content at little or even no cost), the broadband Internet may require a significant investment to publish and distribute – efficiently and effectively – high-bandwidth content.[18]

As the Internet becomes an increasingly important vehicle for companies, political figures, governments, activists, and individuals to speak to large groups of people, a disparity of quality of voice could significantly skew debates, and could undermine the Internet's contribution to open and democratic discourse. It is quite possible, for example, that an oil company would transmit multimedia and other broadband content showing that it protects the environment, but environmental activists would not be able to afford to respond with video evidence of harm done by the company. Similarly, mainstream political parties and candidates will certainly be able to use the Internet to its fullest potential to produce and distributed high quality political advertisements, but underdog and third party candidates may not be able to respond with the same quality of presentations. A President will be able to deliver a complex, multimedia presentation to the nation, but such delivery may be out of reach of presidential critics. As the Internet as a medium becomes more central, and as social and political advocacy utilizes more multimedia and broadband content, the ability of the small or underfunded speaker to afford to speak and be heard will be vital to ensuring a full and robust debate.

The "relative parity" between large and small speakers in the narrowband Internet does make the Internet a unique medium. Although the narrowband Internet will continue in the future to exist, the primary focus of Internet users will shift to broadband content and applications. As the Internet grows and becomes more popularly available, it is certainly possible that small speakers will get lost amid a sea of large and corporate voices trying to reach (and sell to) the millions of new Internet users that will "get online" in the coming years. If small and underfunded speakers cannot offer and deliver reliable and efficient broadband content at a reasonable cost, then the "speech-enhancing" qualities of the Internet may wither.

If that happens, then the Internet may lose some of the characteristics that led the District Court and the Supreme Court in the *Reno* case to afford the Internet such high constitutional protection. Other key factors – such as the Internet's general lack of scarcity – will weigh in favor of maintaining a high level of First Amendment protection. But the possibility that Judge Dalzell offered (but the Supreme Court never reached) – that the Internet deserves *even higher* constitutional protection than is afforded to print – may be lost.

Questions for the Future

This possible reduction of the ability of small speakers and publishers to be speak and be heard is not, and does not need to be, inevitable. Moreover, a reduction of speech does not serve anyone's interest (except perhaps those who might want to squelch small speakers). If small speakers cannot speak and be heard, then everyone loses, in at least three ways:

- the value of the Internet is diminished;
- the diversity of available content is reduced; and
- the risk that a government will decide to step in to enforce openness and access significantly increases.

To avoid this, however, we must strive to inject into the network architecture the ability of the small speaker to deliver – effectively and efficiently – broadband content in competition with the mass of such content offered by large speakers. More generally, we must create a mechanism for this type of issue to be considered and resolved with the public interest represented and protected. To avoid governmental imposition and management of such a mechanism, the Internet industry and community in general must develop such a consultative process.

Critically, these issues must be addressed now. Key decisions about the structure of the broadband Internet are being made now, and those decisions being carried out in the design of the network

architecture. Once those design decisions are initially made, and facilities and equipment are deployed, then retrofitting the network with an architecture that enhances speech could be very difficult. If we fail to inject public interest considerations into the design decisions at this stage, the Internet as we now know it may be lost.

In particular, the distributed content model as it is currently evolving raises questions of whether and how distributed content servers in a given local area will interconnect. In practical terms, this issue poses two interrelated questions (or, perhaps more accurately, one question viewed from two different perspectives):

> 1. **Will a speaker (a content provider) have to contract with more than one distributed content service in order to reach – effectively and efficiently – all Internet users within a given geographic area?** In other words, will a speaker be able to reach all Internet users – especially all users who have broadband "last mile" service – just by signing up with (hypothetically, for example) Akamai? Or, alternatively, will the speaker be required to sign up with Akamai, and Intel, and Excite@Home, and . . . ? Having to contract with more than one distributed content networks will likely lead to higher costs, and those higher costs will exclude some, if not many, small speakers and publishers.
>
> 2. **Will a "last mile" broadband subscriber (say a DSL subscriber with a local phone company or a cable modem subscriber with a cable company) be able to get fast and efficient access to all broadband content that resides on content servers in the subscriber's local area, or will the subscriber only have access to the distributed content affiliated with the subscriber's "last mile" broadband provider?** In other words, will an Excite@Home subscriber have fast access to the broadband content served up by a distributed content provider affiliated with a local DSL provider, and vice versa? Limiting an end user to fast access only to a subset of broadband content would be a significant step backwards from openness of the narrowband Internet.

In trying to anticipate the answers to these questions, there are at least two plausible business models that one can envision: (A) "last mile" broadband providers could strive to connect their users to as much broadband content as possible; alternatively, (B) "last mile" broadband providers could attempt to use particular broadband content as a competitive weapon against competing broadband providers (such that, hypothetically, one might be able to get superfast access to a broadband ESPN sports site over a DSL connection, while cable modem subscribers might instead have access to a broadband Sports Illustrated site). The first model would likely lead to some type of local interconnection (possibly "local peering") between broadband content servers, and that might in turn maximize the possibility that small speakers could speak and be heard (because they would only have to get on one distributed content server in an area). The second model would move away from the Internet's traditional assumption that everyone has access to most content, and might make it harder for smaller speakers to speak and be heard.

Alternatively, one could envision the ISP community, non-profit organizations, or even local governments creating "non-profit" broadband content servers and offering the content to all "last mile" providers in a local area. If such "non-profit" content servers develop across the country and then are themselves networked or interconnected, a small speaker might be able to speak with a rough "equality of voice" with the large, well funded speakers.

There are certainly no clear answers, and indeed the questions themselves are far from clear. What is clear, however, is that care and attention must be paid – by network architects, the public interest community, and others – to the question of how small speakers and publishers can continue to reach the entire Internet. As the distributed content model is refined and many as-yet-unanswered questions about the broadband Internet are addressed, it is in everyone's long term interest to preserve and carry forward the unique and speech-enhancing characteristics of the narrowband Internet.

The Internet industry (including content providers, access providers, and equipment manufacturers) must work with consumer groups and public interest advocates to ensure that First Amendment values are enhanced, not reduced, by the development and refinement of the broadband Internet. Simply put, we must find a way to ensure that these issues are raised and given appropriate weight wherever and whenever decisions about network architecture are made.

* * * * *

The Center for Democracy and Technology and The Broadband Access Project

The Center for Democracy and Technology (CDT) is dedicated to ensuring that democratic values and constitutional liberties are a central feature of the new digital age. With its unique mix of expertise – in law, technology and public policy – CDT works for practical, real-world solutions that enhance free expression, privacy, open access and democracy in the rapidly evolving global communications technologies. CDT endeavors to build consensus among all parties interested in the future of the Internet, finding common ground among activists, nonprofit groups, Internet businesses and government policymakers.

Following the passage of the Communications Decency Act in 1996, CDT helped to organize the Citizens Internet Empowerment Coalition, including leading members of the Internet industry, to challenge the constitutionality of the Act in *American Library Association/ACLU v. Reno*. This coalition wired the courthouse in Philadelphia, and the coalition's counsel argued the case in the U.S. Supreme Court.

In undertaking its Broadband Access Project, CDT seeks to ensure that the characteristics of the narrowband Internet that were so critical in *Reno,* and the resulting legal principles, continue to thrive as the Internet moves into the broadband world. The Project is looking at all forms of broadband access that are emerging as ways to reach the Internet, including cable modems, digital subscriber lines, satellites, and terrestrial wireless services. Working closely with a broad cross-section of the Internet, computer and communications industries, as well as with consumer groups and other interested parties, CDT is developing a comprehensive and balanced assessment of where the technology is today, where it can be tomorrow, and what impact (if any) the new technology will have on speech and access to content on the Internet. Among other continuing activities, the Project expects to release a report on policy and factual issues relating to openness and access in the broadband Internet in late winter, 2000.

---

[*] John Morris is a Partner with the Washington, D.C. office of the law firm Jenner & Block. In 1996, he was one of the lead trial counsel representing the industry and association plaintiffs in *American Library Association v. Reno/ACLU v. Reno*, the constitutional challenge to the Communications Decency Act. During 1999-2000, he took a leave of absence from his firm to direct the Broadband Access Project for the Center for Democracy & Technology.

[**] Jerry Berman is the Executive Director of the Center for Democracy & Technology. In 1996, he was a principle organizer and leader of the Citizens Internet Empowerment Coalition, the industry coalition that was a plaintiff in *American Library Assocation v. Reno/ACLU v. Reno*.

[***] The Center for Democracy & Technology and its Broadband Access Project are described in greater detail at the conclusion of this essay.

Footnotes

[1] *American Civil Liberties Union v. Reno*, 929 F. Supp. 824, 844 (E.D. Pa. 1996) (hereafter "*Reno* District Court Opinion") (available at http://www.ciec.org/victory.shtml).

[2] *Reno* District Court Opinion, 929 F. Supp. at 877 (Dalzell concurring).

[3] E-mail and newsgroups utilize "store and forward" protocols that do not require real time connections between speaker and listener, and thus those means of communication are less sensitive to differences in the bandwidth available to a speaker or a listener on the Internet.

[4] Without question, large and wealthy speakers on the World Wide Web have always been able to reach *more* listeners simultaneously than small speakers can commonly reach. A large corporation will often invest in a high-powered Web server and a broad pipe to the Internet backbone, and thereby gain an advantage over less well-funded speakers. The advantage, however, is not overwhelming, and a small speaker would nevertheless be able to deliver the same basic type of content, albeit a second or two more slowly than the large speaker.

[5] In 1996, the U.S. Congress passed the Communications Decency Act of 1996, which purported to regulate lawful, but "indecent," speech on the Internet, but did so in a manner that was simultaneously ineffective and very burdensome on speakers. Two separate lawsuits challenged the Act as unconstitutional. The two suits – one led by the American Civil Liberties Union and one led by the American Library Association and American Online, Inc. – were consolidated for trial and appeal. A three judge District Court struck the law down as unconstitutional. In *Reno v. ACLU*, the Supreme Court agreed, and upheld the lower court's conclusions that the Internet deserves a very high level of First Amendment protection.

[6] *Reno* District Court Opinion, 929 F. Supp. at 843-44.

[7] *Id.* at 882 (Dalzell concurring).

[8] *Id.* at 877 (Dalzell concurring).

[9] *Id.* at 878-79 (Dalzell concurring).

[10] *Id.* at 877 (Dalzell concurring).

[11] *Reno v. American Civil Liberties Union*, 521 U.S. 844, 863 n.30 (1997) (available at http://www.ciec.org/SC_appeal/decision.shtml).

[12] References to a "single" Web server are intended to encompass multiple interconnected or coordinated Web servers located at the same place. The key concept is that in the narrowband Internet, a Web site is typically served entirely from a single location.

[13] There is, unavoidably, a North American focus to this analysis of the narrowband Internet. For narrowband Web sites based in the U.S. and primarily aimed at North American listeners, there was little incentive to serve Web content from multiple locations around the continent. For U.S.

Web sites that sought to develop a strong overseas audience, there was incentive to "mirror" sites on a European and/or Asian server, and such mirroring occurred long before the distributed content model emerged to speed the delivery of broadband content.

Caching has also been widely implemented as a way to speed delivery of content to users and reduce the need for ISPs to pay backbone charges to repeatedly retrieve popular content from across the Internet. Significantly, however, caching typically speeds the most *popular* content, regardless of whether it is high-bandwidth or not, and caching will speed content from small speakers just as it does for larger speakers – if the content is popular.

[14] Over time in the narrowband world (especially since the rise of streaming audio and video content), the assumption that a single server was sufficient has been increasingly questioned. It is the emergence and widespread deployment of broadband "last mile" services, however, that has made the inadequacy of the single server model clear.

[15] For a detailed discussion of the economic and technical factors that make serving broadband content from a single server location, see Kim Maxwell, Residential Broadband at 106-21 (John Wiley & Sons 1999).

[16] Both Figures 1 and 2 are drawn from http://www.akamai.com/service/howitworks.html (viewed Feb. 3, 00), and are Copyright © 1999-2000 Akamai Technologies Inc.

[17] http://www.home.net/about/network.html (viewed Feb. 3, 00).

[18] None of this is to criticize the distributed content model. Distributing content to local servers is certainly one of the most efficient and effective approaches to delivering high-bandwidth content to users located across the Internet. This essay does not argue against distributing content, but instead advocates trying to shape the distributed content model to maximize the ability of small and underfunded speakers to be able to distribute broadband content along side larger and wealthier speakers.