A Decade of Internet Evolution

by Vinton G. Cerf, Google

In 1998 the Internet had about 50 million users, supported by approximately 25 million servers (Web and e-mail hosting sites, for example, but not desktops or laptops). In that same year, the *Internet Corporation for Assigned Names and Numbers* (ICANN)[1] was created. Internet companies such as Netscape Communications, Yahoo!, eBay, and Amazon were already 3 to 4 years old and the Internet was in the middle of its so-called "dot-boom" period. Google emerged that year as a highly speculative effort to "organize the world's information and make it accessible and useful." Investment in anything related to the Internet was called "irrational exuberance" by the then head of the U.S. Federal Reserve Bank, Alan Greenspan.

By April 2000, the Internet boom ended—at least in the United States—and a notable decline in investment in Internet application providers and infrastructure ensued. Domino effects resulted for router vendors, Internet service providers, and application providers. An underlying demand for Internet services remained, however, and it continued to grow, in part because of the growth in the number of Internet users worldwide.

During this same period, access to the Internet began to shift from dial-up speeds (on the order of kilobits to tens of kilobits per second) to broadband speeds (often measured in megabits per second). New access technologies such as digital subscriber loops and dedicated fiber raised consumer expectations of Internet capacity, in turn triggering much interest in streaming applications such as voice and video. In some locales, consumers could obtain gigabit access to the Internet (for example, in Japan and Stockholm). In addition, mobile access increased rapidly as mobile technology spread throughout the world, especially in regions where wireline telephony had been slow to develop.

Today the Internet has an estimated 542 million servers and about 1.3 billion users. Of the estimated 3 billion mobile phones in use, about 15 percent are Internet-enabled, adding 450 million devices to the Internet. In addition, at least 1 billion personal computers are in use, a significant fraction of which also have access to the Internet. The diversity of devices and access speeds on the Internet combine to produce challenges and opportunities for Internet application providers around the world. Highly variable speeds, display areas, and physical modes of interaction create a rich but complex canvas on which to develop new Internet applications and adapt older ones.

Another well-documented but unexpected development during this same decade is the dramatic increase in user-produced content on the Internet. There is no question that users contributed strongly to the utility of the Internet as the World Wide Web made its debut in the early 1990s with a rapidly growing menu of Web pages.

But higher speeds have encouraged user-produced audio and video archives (*Napster* and *YouTube*), as well as sharing of all forms of digital content through peer-to-peer protocols. Voice over IP, once a novelty, is very common, together with video conferencing (*iChat* from Apple, for example).

Geographically indexed information has also emerged as a major resource for Internet users. In the scientific realm, *Google Earth* and *Google Maps* are frequently used to display scientific data, sensor measurements, and so on. Local consumer information is another common theme. When I found myself in the small town of Page, Arizona, looking for saffron to make paella while in a houseboat on Lake Powell, a Google search on my Blackberry quickly identified markets in the area. I called one of them and verified that it had saffron in stock. I followed the map on the Website and bought 0.06 ounces of Spanish saffron for about $12.99. This experience reinforced my belief that having locally useful information at your fingertips no matter where you are is a powerful ally in daily living.

New business models based on the economics of digital information are also emerging. I can recall spending $1,000 for about 10 MB of disk storage in 1979. Recently I purchased 2 TB of disk storage for about $600. If I had tried to buy 2 TB of disk storage in 1979, it would have cost $200 million, and probably would have outstripped the production capacity of the supplier. The cost of processing, storing, and transporting digital information has changed the cost basis for businesses that once required the physical delivery of objects containing information (books, newspapers, magazines, CDs, and DVDs). The Internet can deliver this kind of information in digital form economically—and often more quickly than physical delivery. Older businesses whose business models are based on the costs of physical delivery of information must adapt to these new economics or they may find themselves losing business to online competitors. (It is interesting to note, however, that the Netflix business, which delivers DVDs by postal mail, has a respectable data rate of about 145 kbps per DVD, assuming a 3-day delivery time and about 4.7 GB per DVD. The CEO of Netflix, Reed Hastings, told me nearly 2 years ago that he was then shipping about 1.9 million DVDs per day, for an aggregate data rate of about 275 Gbps!)

Even the media that have traditionally been delivered electronically such as telephony, television, and radio are being changed by digital technology and the Internet. These media can now be delivered from countless sources to equally countless destinations over the Internet. It is common to think of these media as being delivered in streaming modes (that is, packets delivered in real time), but this need not be the case for material that has been prerecorded. Users of iPods have already discovered that they can download music faster than they can listen to it.

With gigabit access to the Internet, one could download an hour's worth of conventional video in about 16 seconds. This fact certainly changes my understanding of "video on demand" from a streaming delivery to a file transfer. The latter is much easier on the Internet because one is not concerned about packet inter-arrival times (jitter), loss, or even orderly delivery because the packets can be reordered and retransmitted during the file transfer. I am told that about 10 hours of video are being uploaded to YouTube per second.

The battles over *Quality of Service* (QoS) are probably not over yet either. Services such as *Skype* and applications such as iChat from Apple demonstrate the feasibility of credible, real-time audio and video conferencing on the "best-efforts" public Internet. I have been surprised by the quality that is possible when both parties have reasonably high-capacity access to the Internet.

Technorati is said to be tracking on the order of 112 million blogs, and the *China Internet Network Information Center* (CNNIC) estimates 72 million Chinese blogs that are probably in addition to those tracked by Technorati. Adding to these are billions of Web pages and, perhaps even more significant, an unknown amount of information online in the form of large databases. The latter are not indexed in the same way that Web pages can be, but probably contain more information. Think about high-energy physics information, images from the Hubble and other telescopes, radio telescope data including the *Search for Extra-Terrestrial Intelligence* (SETI)[2], and you quickly conclude that our modern society is awash in digital information.

It seems fair to ask how long accessibility of this information is likely to continue. By this question I do not mean that it may be lost from the Internet but, rather, that we may lose the ability to interpret it. I have already encountered such problems with image files whose formats are old and whose interpretation by newer software may not be possible. Similarly, I have ASCII text files from more than 20 years ago that I can still read, but I no longer have operating software that can interpret the formatting instructions to produce a nicely formatted page. I sometimes think of this problem as the "year 3000" problem: It is the year 3000 and I have just finished a Google search and found a PowerPoint 1997 file. Assuming I am running Windows 3000, it is a fair question whether the format of this file will still be interpretable. This problem would arise even if I were using open-source software. It seems unlikely that application software will last 1000 years in the normal course of events unless we deliberately take steps to preserve our ability to interpret digital content. Absent such actions, we will find ourselves awash in a sea of rotting bits whose meaning has long since been lost.

This problem is not trivial because questions will arise about intellectual property protection of the application, and even the operating system software involved. If a company goes out of business or asserts that it will no longer support a particular version of an application or operating system, do we need new regulations that require this software to be available on the public Internet in some way?

Even if we have skirted this problem in the past by rendering information into printed form, or microfilm, the complexity of digital objects is increasing. Consider spreadsheets or other complex objects that really cannot be fully "rendered" without the assistance of application software. So it will not be adequate simply to print or render information in other long-lived media formats. We really will need to preserve our ability to read and interpret bits.

The year 2008 also marks the tenth anniversary of a project that started at the U.S. Jet Propulsion Laboratory: *The Interplanetary Internet*. This effort began as a protocol design exercise to see what would have to change to make Internet-like capability available to manned and robotic spacecraft. The idea was to develop networking technology that would provide to the space exploration field the kind of rich and interoperable networking between spacecraft of any (Earth) origin that we enjoy between devices on the Internet.

The design team quickly recognized that the standard TCP/IP protocols would not overcome some of the long delays and disruptions to be expected in deep space communication. A new set of protocols evolved that could operate above the conventional Internet or on underlying transport protocols more suited to long delays and disruption. Called "delay and disruption tolerant networking"[3, 4] or DTN, this suite of protocols is layered in the same abstract way as the Internet. The Interplanetary system could be thought of as a network of Internets, although it is not constrained to use conventional Internet protocols. The analog of IP is called the *Bundle Protocol*[5], and this protocol can run above TCP or the *User Datagram Protocol* (UDP) or the new *Licklider Transport Protocol* (for deep space application). Ironically, the DTN protocol suite has also proven to be useful for terrestrial applications in which delay and disruption are common: tactical military communication and civilian mobile communication.

After 10 years of work, the DTN system will be tested onboard the Deep Impact mission platform late in 2008 as part of a program to qualify the new technology for use in future space missions. It is hoped that this protocol suite can be standardized for use by any of the world's space agencies so that spacecraft from any country will be interoperable with spacecraft of other countries and available to support new missions if they are still operational and have completed their primary missions. Such a situation already exists on Mars, where the Rovers are using previously launched orbital satellites to relay information to Earth's Deep Space Network using store-and-forward techniques like those common to the Internet.

The Internet has gone from dial-up to deep space in just the past 10 years. One can only begin to speculate about its application and condition 10 years hence. We will all have to keep our subscriptions to *The Internet Protocol Journal* to find out!

References

[1] Cerf, V., “Looking Toward the Future,” *The Internet Protocol Journal,* Volume 10, No. 4, December 2007.

[2] `http://www.seti.org`

[3] `http://www.dtnrg.org/wiki`

[4] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, “Delay-Tolerant Networking Architecture,” RFC 4838, April 2007.

[5] Scott, K., and S. Burleigh, “Bundle Protocol Specification,” RFC 5050, November 2007.

VINTON G. CERF is vice president and chief Internet evangelist for Google. Cerf served as a senior vice president of MCI from 1994 through 2005. Widely known as one of the “Fathers of the Internet,” Cerf is the co-designer of the TCP/IP protocols and the architecture of the Internet. He received the U.S. National Medal of Technology in 1997 and the 2004 ACM Alan M. Turing award. In November 2005, he was awarded the Presidential Medal of Freedom. Cerf served as chairman of the board of the Internet Corporation for Assigned Names and Numbers (ICANN) from 2000 through 2007 and was founding president of the Internet Society. He is a Fellow of the IEEE, ACM, the American Association for the Advancement of Science, the American Academy of Arts and Sciences, the International Engineering Consortium, the Computer History Museum, and the National Academy of Engineering. He is an honorary Freeman of the City of London. Cerf holds a Bachelor of Science degree in Mathematics from Stanford University and Master of Science and Ph.D. degrees in Computer Science from UCLA. E-mail: **vint@google.com**