Thomas Haigh

Protocols for Profit: Web and Email Technologies as Product and Infrastructure

PREPUBLICATION DRAFT

Thomas Haigh – thaigh@computer.org, www.tomandmaria.com/tom

For inclusion in "The Internet and American Business" eds. William Aspray and Paul Ceruzzi, MIT Press, 2007. This version will differ in minor respects from the published version. I have updated my submitted draft to fix many small errors and clumsy expressions marked up on the galley copied by the copyeditor. On the other hand, I have defied some elements of the MIT style guide that seem perverse and avoided introducing a few errors created by the copy editor in the editing process. And of course the pagination will be quite different. If you wish to quote or cite this paper, please use the official published version.

The Internet evolved with breathtaking speed during the 1990s, from an obscure, academic system to a mainstay of the developed world's daily routines of communication, shopping, travel, entertainment, and business. As millions of ordinary people rushed to connect their computers to the network, most were driven by a desire to use two particular kinds of program: email systems and the World Wide Web. These were the so-called "killer applications" of the Internet during its commercialization: specific application programs so compelling that people were prepared to purchase an entire computer system in order to use them.[1] Both played two roles at once: as products in their own right for a small number of companies such as

[1] The concept of the killer application was popularized in Robert X Cringely, Accidental Empires: How the Boys of Silicon Valley Make their Millions, Battle Foreign Competition, and Still Can't Get a Date (Reading, MA: Addison-Wesley, 1992).

Netscape and Microsoft, and as crucial parts of the communal software infrastructure on which countless other companies tried to build their own online business empires.

In this chapter, I explore the evolution of Internet email and the Web since the late 1980s, looking at the markets that developed for both kinds of program, efforts to remake both technologies in more commercial ways, and the ways in which the culture and technology of the pre-commercial Internet of the 1980s shaped the development of its commercialized successor. Although space does not permit me to delve deeply into the experience of Internet users during this era, I do attempt to make connections between the fundamental architecture of these Internet technologies and the opportunities available to the businesses and individuals using them.

The Internet and Its Protocols

One of the things that makes it hard to write about the history of the Internet is uncertainty over exactly what "the Internet" is and was. Internet histories generally go back to the academic ARPANET of the 1970s, even though the concept of an "Internet" dates from a later generation of networks built during the 1980s. The ARPANET first went into operation in 1969, and was funded by the U.S. Department of Defense to interconnect the computers it had purchased for researchers in different universities. No single piece of hardware, network infrastructure, or application software survived the transition from the ARPANET to the academic Internet of the early 1990s. During the 1990s the Internet changed still more rapidly, as its dominant users, applications, network operators and traffic patterns changed beyond recognition. The situation recalls the apocryphal tale of an item said to be displayed in the Tower of London with a sign "Axe, eleventh century. Head replaced in early thirteenth century, shaft replaced in fifteenth century." If every piece of something has been changed then what are we really writing about when we tell its story?

In fact the essence of the Internet lies not in hardware or software but in protocols: the agreed upon rules by which computer programs communicate with each other. When computers communicate on the Internet, it is these protocols that determine what is automatic and what is impossible; what is easy to accomplish and what requires complex or cumbersome special procedures.[2] But the protocols that define the Internet were designed by and for a group of users very different from that of the commercial Internet. The fundamental data transmission protocol suite of the Internet, TCP/IP, was developed during the late 1970s. Simple Mail Transfer Protocol (SMTP) the main protocol for Internet email transmission, is from the early 1980s. And although the World Wide Web was not created until the early 1990s, its protocols were modeled on and designed to work with the existing standards of the Internet.

Events of the 1970s and 1980s have played a profound role in shaping today's commercialized Internet. As Lawrence Lessig has argued in his book *Code*, design decisions built into computer code can play a role in shaping the development of online life just as important as the role of legal codes in shaping the development of societies off-line.[3] Decisions built into the design of protocols perform a similar role. This insight parallels ideas long familiar

---

[2] Scholars have recently paid considerable attention to the role of protocols in regulating the Internet, particularly the governance of its DNS system of domain name allocation and resolution by ICANN, the Internet Corporation for Assigned Names and Numbers. See Alexander R Galloway, Protocol: How Control Exists after Decentralization (Cambridge, MA: MIT Press, 2004), Milton Mueller, Ruling the Root: Internet Governance and the Taming of Cyberspace (Cambridge, MA: MIT Press, 2002) and Daniel Paré, Internet Governance in Transition: Who is the Master of This Domain? (Lanham, MD: Rowan and Littlefield, 2003). My concern in this chapter is not with the overtly political work of assigning domain names but with the broader influence of the Internet's pre-commercial technologies, cultures and practices on the development of its software infrastructure during the 1990s.

[3] Lawrence Lessig, Code, and Other Laws of Cyberspace (New York: Basic Books, 1999).

to scholars researching the history of technology. Following the work of Thomas P. Hughes, these scholars have come to see technologies not as individual inventions or machines but as part of broader socio-technical systems binding together elements such as standards, professional communities, institutions, patents, businesses, and users. The system outlives all its constituent parts. Successful technological systems such as the Internet develop what Hughes called "technological momentum," making them hard to displace or redesign.[4] Many historians and sociologists of technology have explored the ways in which the function and design of technologies were influenced by the choices and assumptions of their designers, which in turn reflect the values and cultures of the communities in which they worked. These scholars talk of studying the "social shaping" of technology, and of "opening the black box" of apparently inscrutable technical systems to discover the social choices hidden inside.[5] The technologies of the Internet have been clearly and powerfully shaped by the circumstances of their creation and early use, in ways that explain both their rapid spread in the early 1990s and many of the problems faced by today's network users.

As Janet Abbate has shown in her book *Inventing the Internet*, the protocols of the Internet reflect its roots in the 1970s as a closed system, used by researchers, government

---

[4] Hughes' specific ideas on the development of "large technological systems" came as a generalization of his study of the evolution of electrical power networks in Thomas Hughes, Networks of Power: Electrification in Western Society, 1880-1930 (Baltimore, MD: Johns Hopkins University Press, 1983). A variety of approaches to the topic were collected in the highly influential Wiebe Bijker, Thomas Hughes, and Trevor Pinch, The Social Construction of Technological Systems (Cambridge, MA: MIT Press, 1987).

[5] The concept of "opening the black box" was popularized by Donald MacKenzie, who explored the social shaping of missile guidance systems. Donald MacKenzie, Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance (Cambridge, MA: MIT Press, 1990)

research administrators and military bureaucrats. The Internet was created during the early 1980s by the interconnection of different TCP/IP networks. TCP/IP was designed to work with networks of all kinds, including satellite and radio systems as well as telephone lines and dedicated high speed cables. It was so simple and flexible that it could work with any kind of transmission medium as long as suitable “link layer” software had been written. TCP/IP separated transmission of data from the work of specific application programs. Likewise, computers of any kind could communicate with each other providing each was equipped with software faithfully implementing the TCP/IP protocols. This meant that someone writing a program to transmit a file or perform any other network task did not need to worry about what kind of computer was at the other end, or about the characteristics of the networks data might be sent over.

During the 1980s the most important constituent networks of the Internet included the ARPANET created for the Defense Advanced Projects Research Agency administrators and sponsored researchers, the National Science Foundation’s NSFNET, and MILNET for military use. Connections to academic networks in Europe and Asia were gradually added during the 1980s and early 1990s. TCP/IP, SMTP, and the other Internet protocols were designed to support an unprecedentedly wide range of applications, transmission mechanisms and computers. They were not, however, designed for commercial use on public networks. As a result, they acquired a characteristic set of strengths and weaknesses that together go a long way to explain not only the rapid success of the commercial Internet but also its persistent problems and weaknesses.

The special characteristics of the 1980s Internet, and their influence on its protocols and architecture, may be summarized as follows.

- **The Internet was designed for a homogeneous population of academically oriented, scientifically trained users, granted access through their employers or universities**. As a result, it relied more on social mechanisms rather than technical ones to provide security and eliminate troublemakers who engaged in antisocial behavior (public insults, resource hogging, or commercial solicitations). Breaches of decorum were punished with warnings and, eventually, by the termination of a user's access privileges by their home institution. This was impossible to do once commercial access brought a more diverse population online.
- **The Internet and ARPANET were designed as practical, working networks**. The early history of the Internet shows constant interplay between theory and practice, with protocols and designs evolving rapidly on the basis of experience. From the beginning, the ARPANET was designed to do useful work, and its users found new and unanticipated applications such as email. On the other hand, the Internet was never intended to be used commercially or set a global standard. This contrasted with long-running international efforts to set official standards for global public networks, which suffered from the need to satisfy every possible requirement and constituency.
- **The Internet was entirely non commercial, and so provided no way to charge users according to the network resources they consumed or compensate the providers of network services.** The Internet provided no way of knowing exactly who was connecting to the services offered from your computer, still less of automatically receiving money from them as a result. This was a sharp contrast with existing commercial networks, including international telephone service and commercial packet-switched networks, for which the idea of precise billing for resources used was fundamental. Of course, universities paid a certain amount each month to network providers and received a connection of a certain bandwidth in

return. But because TCP/IP was designed as a pragmatic and flexible technology for a non commercial environment, it gave networks within the Internet no way of charging users to relay packets of data, even to recoup the costs of expensive resources such as intercontinental satellite links.

- **The Internet was designed to serve a research community, rather than to perform one specific task**. From the earliest days of the ARPANET, network designers wanted to make it easy to experiment with new network applications and so aimed to make the network as flexible as possible. They initially did this by putting the network logic for each network site into a separate minicomputer known as an IMP (Interface Message Processor). With the transition to TCP/IP protocols the same goal was accomplished in software by separating the application layer protocols for things like email from the transport and network layer code of TCP/IP itself.
- **The Internet was designed to support many different machine types**. Proprietary networks such as those offered by IBM and DEC often sought compatibility by requiring all computers involved to use the same hardware and/or software. In contrast, the Internet was built around publicly available protocols and could network many different kinds of computer.
- **Any computer connected to the Internet could send and receive data of any kind**. Many commercial networks were built around the assumption that a large number of cheap terminals would be connected to a handful of powerful computers, on which all information would be centralized. The Internet was designed from the beginning as a flexible mechanism to connect computers together, and any computer connected to it could offer services to any other computer. TCP/IP worked on a peer-to-peer basis. In practical terms, that meant that a

computer could turn into a file server, an email server or (later) a Web server simply by running a new program. Users were expected to publish online information and provide services to each other, rather than to rely on a single central collection of resources.

- **The Internet integrated many different communications media.** The initial impetus behind TCP/IP came from a need to integrate data transmission over satellite and radio networks. The separation of media-specific aspects of communication from TCP/IP itself has made it possible to extend the Internet over many new kinds of link, such as dial-up telephone connections, cellular telephones, Wi-Fi networks and Ethernet local area networks which were unknown when it was originally designed. This flexibility was vital to the Internet's rapid commercialization, because once an Internet connection was established one could use exactly the same software (such as a Web browser) over a dial-up connection from home as an Ethernet connection in a university. On the other hand, because TCP/IP was designed for flexibility rather than performance the Internet does not offer a guaranteed data delivery time for critical applications, or the ability to avoid network congestion by paying a premium. This proved a huge challenge in the development of systems to provide audio and video over the Internet.

People often imagine that the success of the Internet was completely unexpected, but this is true only in a narrow sense. In fact, the breakthrough of computer networking, electronic publishing, and email into the mainstream of personal communication was confidently predicted from the late-1970s onward. The two main surprises were that this took so long to happen, and that in the end it was the Internet rather than a specially designed commercial network that accomplished the feat. Millions of home computers flooded into American homes in the early 1980s, together with a wave of interest in the idea that the newly powerful and affordable

electronic technologies based around the silicon chip would trigger a fundamental social upheaval. This was known variously as the "information society," the "microelectronic revolution" and the "home computer revolution."[6] Futurist Alvin Toffler wrote in *The Third Wave* (1980) of a new society in which electronic networks had broken apart large organizations and allowed most people to work from home.[7] Email's triumphant emergence as a replacement for most personal letters and many telephone calls took place more than a decade later than once predicted. In *The Network Nation* (1978), academics Star Roxanne Hiltz and Murray Turoff salted their scholarly examination of the online communication systems of the day with predications that by 1990 online polls would replace Congress and electronic mail would bring the postal service to its knees.[8]

For a few years in the early 1980s, telecommunications firms throughout the developed world were convinced that millions of ordinary people were about to start using their home computers or cheap terminals connected to their television screens to shop, read news, make travel bookings, search databases and trade stocks online. The same idea returned a decade later, as enthusiasm for something to be called "the information superhighway" gripped America's politicians and businesspeople, cable television and computer companies were investing

---

[6] These ideas are critiqued in Langdon Winner, "Mythinformation in the High-tech Era", in Computers in the Human Context, ed. Tom Forester (Cambridge, MA: MIT Press, 1991), 82-132. A roundup of scholarly thinking on the topic is given in Frank Webster, Theories of the Information Society (New York: Routledge, 1995).

[7] Alvin Toffler, The Third Wave (New York: William Morrow, 1980).

[8] Star Roxanne Hiltz and Murray Turoff, The Network Nation: Human Communication via Computer (Reading, MA: Addison-Wesley, 1978).

hundreds of millions of dollars to create high-bandwidth home networks accessed via computerized cable boxes plugged into television sets.[9]

Clearly, nothing like the Internet would ever have been designed or launched by a commercial telecommunications company. With no easy way to charge users for the resources they consumed, and no easy way to make money by publishing information online, the network would seem fundamentally incompatible with any plausible business plan. The Internet was decentralized and unplanned, relying on its users to fill it with content and services. It offered largely free access to a huge variety of amateur and academic content, along with a growing mass of information that companies were giving away because there was no obvious way to charge for it. Yet these very characteristics appealed to early commercial Internet users. Bizarrely, it was the Internet's fundamental disregard for the essential features of a commercial network that made it such a commercial success during this period.

EMAIL AND PERSON TO PERSON COMMUNICATION

While Internet email was practical, simple, well-proven, and universally compatible, its academic origins meant that it suffered from a number of limitations for non academic use. I suggest that the same characteristics that, in the short term, explain the initial success of Internet email for mainstream use have also, in the longer term, made it very hard to combat spam or use Internet email as a reliable mechanism for important messages. It is perhaps in the case of email that the limitations of Internet technologies for commercial use are the most apparent, and where

[9] No detailed history or analysis of the early 1990s push for home networking around cable boxes has yet been published, but an evocative sketch is given in Michael Lewis, The New New Thing (New York: W. W. Norton, 2000), 70-79, 119-20.

the technological advantages of competing systems are best developed. With hindsight, historians and economists may come to view Internet email as one of those cases, like the much-discussed Qwerty keyboard and the VHS video cassette recorder standard, in which an allegedly inferior technology gained a permanently entrenched position because of a temporary initial advantage.[10]

People have been using computers to send messages to each other even before computers were networked to each other. By the early 1960s, computers were employed in corporate communication hubs to process sales data or information requests sent via teletype, and to automatically process and retransmit telegram-style messages sent between the offices and factories of large corporations.[11] A few years later, the development of time-sharing operating systems made it possible for several users to simultaneously work interactively on the same computer. It was natural and easy to allow users to chat with each other, and to send messages to be stored in the user accounts of other system users and read when convenient.[12] Today we tend

---

[10] An accessible introduction to the economic literature on path dependence is given in Douglas Puffert, Path Dependence (EH.Net Encyclopedia, edited by Robert Whaples, 2003 [cited September 15 2006]); available from http://eh.net/encyclopedia/article/puffert.path.dependence.

[11] Westinghouse boasted one of the most advanced such networks. By 1963 its automated hub in Pittsburgh already routed more than 15,000 messages a day between 265 corporate locations. See R C Cheek, "Establishing the Telecomputer Center", Journal of Data Management 2, no. 6 (June 1964):28-33 and R. C. Cheek, "A System for General Management Control", in Control Through Information: A Report on Management Information Systems (AMA Management Bulletin 24), ed. Alex W. Rathe (New York: 1963), 5-10.

[12] The CTSS system developed at MIT from 1961 onward appears to have allowed users to send messages to each other, and is reported to have added an internal mail box system in 1965. Tom Van Vleck, The History of Electronic Mail (Multicians.org, 10 September 2004 2001 [cited October 29 2006]).

to think of email, instant messaging, and discussion forums as distinct methods of communication, but there are no obvious or inherent boundaries between them. Systems have often blended aspects of these approaches, for example mixing public discussion areas with private messaging, blending word processing with email, or offering real-time chat features as well as file exchange.

When and where networks were built to link computers together, this created an obvious opportunity for the exchange of messages between the users of the different computers on the network. Exactly when this was first done is not clear. As Abbate has shown, email unexpectedly emerged as the main application of the ARPANET in the early 1970s. Email soon accounted for most of the network traffic, becoming the communications medium of choice within ARPA itself. Although email had not been planned for when the network was being designed, its flexible nature made it easy for users to evolve their own applications. Experimental email systems were used to transfer messages between sites from 1971 onward, and in 1973 the existing standard for ARPANET file transfers was modified to support email transfer, greatly simplifying the task of transferring messages between computers of different types. Users quickly created programs to view, sort, and delete mail as well as to maintain email discussion lists. Email was effectively free to its users, because the ARPANET infrastructure was bankrolled by the Department of Defense to interconnect the academic computer scientists it was supporting. Neither universities nor individuals were billed when a message was sent.[13]

Email remained a key network application as the ARPANET was replaced by the Internet, and from the early 1980s onward, specially created protocols for Internet email ensured that people using different email programs and different types of computer would be able to

[13] Janet Abbate, Inventing the Internet (Cambridge, MA: MIT Press, 1999), ch. 3.

exchange messages freely. Mail was relayed to the recipient's machine using SMTP, designed in the early 1980s.[14] As an application protocol this ran over the Internet's native TCP/IP, meaning that all the complexity of routing messages and verifying their successful transmission was taken care of by TCP/IP. SMTP had to do little more than use the Internet's existing infrastructure to contact the recipient's mail server and then transmit the name of the sender, name of the recipient, and the text of the message itself. Email servers were connected to the Internet, and were registered with the Internet's domain name servers as the place to which all email addressed to recipients in a particular domain (e.g. ibm.com) or subdomain (e.g. wharton.upenn.edu) should be delivered. The first widely used SMTP message delivery software was Sendmail, and today SMTP is supported by virtually all email clients to send outgoing messages. While SMTP could originally transmit only plain (7-bit encoded) text, it was later supplemented with the MIME (Multipurpose Internet Mail Extensions) standard for encoding file attachments and non Western alphabets.[15]

During the 1980s and early 1990s, Internet users generally accessed email and other Internet resources by logging in with a video terminal or terminal emulator program. They would access their email by running a text based mail reader, such as the basic "mail" command built into UNIX or the more sophisticated ELM, a public domain product. These worked by directly reading the contents of the mail file into which the user's email had been deposited. These

---

[14] Jon Postel, RFC 821 - Simple Mail Transfer Protocol (Network Working Group, August 1982 [cited October 22 2006]); available from http://www.faqs.org/rfcs/rfc821.html.

[15] Glyn Moody, "E-Mail: Polishing Your Communications" The Guardian, August 18 1994, T18. A good overview of Internet email technologies in the early commercial era is given in David Wood, Programming Internet Email (Sebastapol, CA: O'Reilly, 1999).

programs ran on the computer of the ISP (Internet Service Provider) or university, meaning that email could only be sent or received when logged into the computer.

The Internet as an Email Gateway

But although the Internet provided the key technologies behind today's mass-market email systems, during the 1980s it was not open to the public or amateur computer enthusiasts. Collectively various other corporate, commercial, and academic email networks boasted an estimated six million American users by the end of the 1980s – far more than the number of people directly connected to the Internet.[16] But except in a few specialized communities such as computer science research, email had conspicuously failed to establish itself as a medium for communication between organizations, unlike the fax machine, which had spread far more rapidly during the decade to become a fixture of businesses large and small. Many problems conspired to hold back email services: they were rarely easy to use, they sometimes suffered from reliability problems, modems and computers often required an elaborate configuration process to work with them, message delivery was sometimes slow, and their complex billing systems gave users an incentive to minimize use. The biggest problem, however, was the lack of interconnection between networks. As one reporter noted, although email was "growing fast and becoming a standard business tool" within many organizations, it had "developed as a lot of small, closed systems… rather than one big network that everyone can use."[17] When the Internet first reached a reasonably broad academic audience, at the end of the 1980s, it was not through a

---

[16] John Burgess, "Electronic Mail's Delivery Problems" Washington Post, April 21 1989, F1.

[17] Ibid.

direct connection but via the interconnection of the Internet email network with other amateur and commercial networks. Internet email became a *lingua franca* between other email systems.

The amateur enthusiasts of the 1980s found refuge on local bulletin boards, which formed an archipelago of thousands of little information islands across America, dotted most closely together in major urban areas and centers of high technology. At its simplest, a bulletin board was a personal computer left running a special program and connected to a single telephone line and modem. Users could wait their turn to dial in and leave messages for each other, mimicking the asynchronous communication methods of a physical bulletin board. More ambitious boards offered multiple telephone lines for simultaneous access by several people, together with other services, such as online chat. By the mid-1980s, people had started connecting them together, most commonly through a piece of bulletin board software known as FidoNet. This worked as a poor man's version of the Internet email and Usenet discussion systems, batching together packets of mail and newsletter updates intended for the users of other FidoNet nodes, and passing them on in cheap, late night phone calls. At its peak in the mid-1990s, FidoNet had more than thirty thousand nodes worldwide (many with hundreds of users each). The FidoNet software was free, and many of the system's constituent nodes were run by volunteers (known as "sysops") who made no charge to their users. This made FidoNet a popular system for hobbyists in the United States, smaller businesses in Europe, and for institutional users in poorer regions with expensive and unreliable communication links such as Africa and the former Soviet Union.[18]

---

[18] FidoNet has so far received almost no academic or journalistic discussion. Its creator has placed some historical material online at Tom Jennings, Fido and FidoNet (n.d. [cited September 19 2006]); available from http://wps.com/FidoNet/index.html.

Another ad hoc system, BITNET, was used by academic users to pass messages between IBM mainframes. It was the original mechanism used to create "listservs," or automated email discussion lists.[19] A third system, UUCPNET, was an informal and decentralized network of university and research lab servers running free software. UUCPNET maintained its own email system, with messages passed from machine to machine via cheap, late night phone calls. It also distributed Usenet, a thriving system of thousands hierarchically arranged discussion forms running the gamut from comp.os.mac to alt.sex.swingers. Usenet discussions were synchronized between servers, with large sites such as Bell Labs, DEC, and Apple shouldering most of the long-distance communications costs.[20]

Email services were also available to individuals and small businesses, through commercial online services. By the start of the 1990s, four main firms were selling on-line services to consumers. Each offered email as one of a bundle of services, available for a monthly fee and hourly usage charges. The largest and most successful of these services, Compuserve, had been founded in 1969 as a time-sharing business. It launched a new service for microcomputer owners during the 1980s, and finished up with a solid business selling affordable email and online databases systems to smaller businesses. Genie and Prodigy were newer and smaller services, created in the 1980s for home users. The smallest of the four, America Online (AOL) had just 110,000 subscribers, far behind the more than half million each boasted by

---

[19] BITNET is discussed in David Alan Grier and Mary Campbell, "A Social History of Bitnet and Listserv, 1985-1991", IEEE Annals of the History of Computing 22, no. 2 (April-June 2000):32-41.

[20] Usenet is discussed in Howard Rheingold, Virtual Communities: Homesteading on the Electronic Frontier (Reading, MA: Addison-Wesley, 1993), ch.4.

Compuserve and Progidy.[21] More specialized online services, offering only email, were run by major American telecommunications firms such as AT&T and MCI. These mail networks imposed a monthly subscription fee, plus charges for messages sent, received, and stored. One glowing 1987 profile of MCI Mail noted that the service was “cheap” because its 100,000 subscribers would pay only a dollar to send a 7,500 character email within the United States. For an extra fee, MCI would print out the email, and then deliver it by post or courier.[22]

Large corporations built their own email networks based around servers and networks supplied by firms such as IBM, Wang, and DEC. Email was a key component of the office automation systems promoted during the early-1980s as a gateway to the paperless office of the future.[23] IBM’s mainframe-based PROFS system, for example, was widely used by large organizations such as the US federal government. DEC’s ALL-IN-1 ran on its popular VAX minicomputers. By the early 1990s, the vendor-specific corporate email systems sold by mainframe and minicomputer manufacturers were joined by a new breed of email systems able to run on standard microprocessor based servers running OS/2, UNIX or Novell Netware. Among the most successful of these was Lotus Notes, first sold in 1989, which used a novel and flexible system of replicated databases to provide email services, shared calendar functions and discussion groups, allowing corporate customers to build their own custom applications around

---

[21] Kara Swisher, aol.com: How Steve Case Beat Bill Gates, Nailed the Netheads, and Made Millions in the War for the Web (New York: Random House, 1998).

[22] T.R. Reid, "MCI Mail Lets You Set Up a Post Office Operation on Your Desk-Top" Washington Post, December 7 1987, F41.

[23] Office automation is discussed in Thomas Haigh, "Remembering the Office of the Future: The Origins of Word Processing and Office Automation", IEEE Annals of the History of Computing 28, no. 4 (October-December 2006).

these communication and textual database capabilities. Notes required users to install and configure special client software. A rival system, Microsoft Exchange, emerged as a strong competitor during the mid-1990s in conjunction with Microsoft's Outlook client software.

There was certainly no shortage of email networks and technologies. But this was actually a problem, because during the 1980s users of commercial online services or corporate email networks were generally able to send mail only to users of their own network. Someone wanting to communicate with users of Genie and CompuServe would need to purchase subscriptions to both systems, learn two completely different user interfaces, and make separate phone calls every day to each system to retrieve new mail. This dramatically limited the spread of email for ordinary home or small business users. Likewise, users of corporate email systems could reach their boss or colleagues by email, but not their customers and suppliers.

Online services and corporations thus had an incentive to allow for the exchange of emails with other systems, even as they kept their other network content proprietary. The need for this kind of email exchange between networks was recognized early, and addressed through two standards developed by CCITT, the international federation of telecommunications carriers, and adopted by the computer communication standards group OSI (Open Systems Interconnection). These were the X.400 standard for message delivery and the X.500 standard for directory access. In contrast, the Internet's mail protocol, SMTP, was little known in the business world. During the 1980s and early 1990s the OSI standards effort was supported by mandates from many national governments (including the United States) and was almost universally viewed as the future of computer networking. Plans were even made for the Internet itself to

abandon TCP/IP as its foundation, shifting instead to TP/4, a related OSI protocol.[24] Given this backing, and the clear need for email interchange, X.400 was taken very seriously. As well as the obvious base of support among national telecommunications carriers, hardware and software vendors including DEC, Novell, Hewlett-Packard, and Microsoft added X.400 capabilities to their email systems.

Yet by the end of the 1980s it was Internet email, rather than X.400, that was beginning to emerge as the leading method for email transmission between other networks. The number of computers connected directly to the Internet was still small, in part because of the restrictions placed on commercial use. But other networks were already beginning to build gateways to forward messages initiated by their users on to the Internet and to deliver messages received from the Internet. This did not require the networks to adopt either the SMTP Internet email standard or TCP/IP for their internal use, so it was quite easy to accomplish. Messages were simply reformatted and retransmitted, perhaps hours later. Email exchange took a big step forward in 1989 when MCI Mail and Compuserve began to exchange email with the Internet and

---

[24] The OSI effort, its X.25 and TP/4 protocols and their relationship to the Internet are discussed in Abbate, Inventing the Internet, 147-49. The practices and cultures of OSI and the Internet are contrasted in Andrew L. Russell, "'Rough Consensus and Running Code' and the Internet-OSI Standards War", IEEE Annals of the History of Computing 28, no. 3 (July-September 2006):48-61. As late as 1994, one Internet discussion document still insisted that "currently in the Internet, OSI protocols are being used more and more" and implied that it was no longer "predominantly a TCP/IP network." S. Hares and C. Wittbrodt, RFC 1574: Essential Tools for the OSI Internet (Network Working Group, February 1994 [cited September 19 2006]); available from http://www.rfc-archive.org/getrfc.php?rfc=1574

each other.[25] FidoNet and UUCPNET were also connected by gateways to the Internet. AOL and Genie followed suit in 1992, though these still had some limitations, such as a cap on message size at eight kilobytes for the IBM PC version of AOL.[26] Until well into the 1990s, gateways of this kind accounted for most of the email sent between users of different services.

The core technologies of the Internet were created as practical, short-term solutions to real needs, often those of the teams or organizations creating them. They did not try to design standards that would endure for decades or meet the needs of every possible user. While this led to short-sighted decisions in some areas, in many cases the simplicity and flexibility of the Internet has made it possible to evolve its technologies along with their changing uses. A comparison of SMTP, the simple and effective Internet email system, with the agreed upon world standard X.400 makes the difference clear. In the late 1980s, Internet email was capable of little more than the delivery from one server to another of textual messages (in which, using separate tools, graphics and attachments could be encoded). X.400 offered a long list of features, including security, notification when a message had been received or read, different levels of priority, protection against the faking of addresses, and the automatic conversion of messages and attachments between different formats. It was designed to support centralized directories, via the X.500 protocol, so that there was a standard way to find the email address of the person you

---

[25] Various, CompuServe <------>Internet gateway (Thread in comp.sys.ibm.pc) (Usenet, August 23 1989); available from http://groups.google.com/group/comp.sys.ibm.pc/browse_thread/thread/4d0ece2be6c1cec0/ebe6db9230b3a7d4.

[26] tom mckibben 2, GeoReps Anyone? (message in comp.os.msdos.pcgeos) (Usenet, June 8 1992); available from http://groups.google.com/group/comp.os.msdos.pcgeos/browse_thread/thread/ce74d8bf02fc70eb/de27f35a78f675a5.

wanted to contact. Internet email systems lacked all these features. Although various additional standards have been proposed to add them, they are still missing from today's typical email experience. In 2000, Microsoft Exchange switched its default protocol from X.400 to SMTP, recognizing the triumph of Internet email technologies. Despite the superior functionality promised by the X.400 standard, perhaps even because of it, X.400-based systems were large, expensive products that were harder to use and never included all the features specified by the committee.[27]

One of the clearest examples of this comes in a comparison of the email address formats used by the two standards. An X.400 email address took the form of a series of attributes, with the shortest-possible version looking something like "G=Harald; S=Alvestrand; O=sintef; OU=delab; PRMD=uninett; ADMD=uninett; C=no."[28] That was hardly an attractive thing to put on a business card or print in an advertisement. X.400 email addresses could get even more baroque. The committee stuffed every possible attribute into the address to satisfy the demands of corporate users as well as international postal and telecommunications bodies, meaning that

---

[27] Technical details on X.400 are given in Cemil Betanov, Introduction to X.400 (Boston: Artech House, 1993). Its failure to gain widespread adoption as an email gateway is chronicled in Dorian James Rutter, "From Diversity to Convergence: British Computer Networks and the Internet, 1970-1995" (Ph.D., University of Warwick, 2005), 172-200.

[28] This example is taken from the discussion of X.400 addresses in Harald T Alvestrand, X.400 addresses are ugly (June 20 1996 [cited August 29 2006]); available from http://www.alvestrand.no/x400/debate/addressing.html. A simplified addressing mechanism was adopted in a 1988 revision to the original X.400 standard, but it introduced extra complexity because the system would need to access an X.500 directory and look up the full version of the address. Internet email addresses, in contrast, simple but still included all the information necessary to deliver the message as long as DNS was working.

email addresses could include up to four different “organizational units” (companies, divisions, laboratories, etc) and even the user’s street address. Every eventuality was catered for. In contrast, an Internet email address was something in the format johndoe@xytech.edu. As we shall see later, the simplicity and lack of security in Internet email left it highly vulnerable to spam and other abuses once its technologies were transplanted into a more hostile environment.

Internet Email For the Masses

As Internet access spread during the early 1990s, direct access to Internet email became practical for millions. Most Internet Service Providers included at least one Internet email box with each personal subscription, and multiple email addresses with small business subscriptions. More mail storage space and further addresses were available for additional charges. Since the email address would include the ISP’s own domain name (for example Jane.Smith@earthlink.com), ISPs had no incentive to break out email service from their basic subscriptions and charge for it separately. Each message sent advertised the customer’s chosen ISP, while users would continue to receive email at the address only while they maintained their subscriptions. Changing an email address was a chore, and so many people continued to stick with their original ISP even after discovering the availability of cheaper, faster, or more reliable service elsewhere. (Organizations and individuals with their own Internet domains usually either set up their own mail server or contracted out this service together with their Web hosting needs).

Unlike the commercial online services and specialist email businesses of the 1980s and early 1990s, ISPs did not charge users additional fees for each message sent, or even (in most cases) for the time spent online. And an Internet email could easily be sent to anyone with an Internet email address or gateway, which by the mid-1990s included pretty much everyone with

an email account on any service. As a result, email use increased dramatically, becoming a killer application for the Internet, and indeed for the personal computer itself. Businesses brought computers, modems and service plans to communicate with their customers; parents with their college student offspring; emigrants with their friends and families across the globe; enthusiasts in every field with others sharing their interests. Between November 1992 and November 1994 the volume of Internet email almost quadrupled, to exceed a billion messages a month.[29] By the end of the decade, an estimated three billion email messages were being sent every single day, outstripping the volume of physical mail sent in the United States.[30]

Internet Email Software: The Business That Wasn't

Early ISPs followed the same patterns established by university computer centers to provide email address to their campus users. They would dial-in, and then use a text terminal window to access an email client program such as PINE or ELM running on the ISP's central computer. But as personal computers were increasingly connected directly to the Internet itself demand grew for a new kind of email client able to download messages from the central server to read, archive, organize, and compose relies directly on a personal computer.

This was achieved by coupling SMTP with a new protocol, POP (Post Office Protocol). POP was particularly useful for dial-up service users, allowing them to connect briefly and retrieve all accumulated messages without having to stay online to read them. ISPs were quick to add POP capabilities to their mail servers. Because the mail client was running on the local

---

[29] The Internet Society, Growth of the Internet -- Internet Messaging Traffic (1995 [cited 2006 August 29]); available from http://www.chebucto.ns.ca/Government/IndustryCanada/tab2.html.

[30] Joan O'C. Hamilton, "Like It Or Not, You've Got Mail", Business Week, October 4 1999.

machine it used a convenient graphical user interface, rather than the primitive text interface of earlier mail readers. Shifting work to the customers' computers also lightened the load on the ISPs servers.

This created a new market for desktop email client software. Eudora, the most widely used early POP client, was first released as a freeware package for the Apple Macintosh in 1990 by its author Steve Dorner of the University of Illinois. Its name was a humorous tribute to the writer Eudora Welty, one of whose best known stories was "Why I Live at the P.O."[31] Eudora was produced in Macintosh and Windows versions, with a freely downloadable basic version joined by a commercially distributed "Pro" one. Millions of copies of the free version were given away by universities and ISPs to their customers. Qualcomm, better known for its modems than its software expertise, sold Eudora commercially, though the program evolved slowly and gradually lost its once-dominant position. In 2006 Qualcomm gave up, dropping Eudora as a commercial product.[32]

Most users relied either on free software (including the free version of Eudora) or used one of the email clients bundled with other packages. Netscape included a clunky but functional

---

[31] For an early report on Eudora see Ron Anderson, "Qualcomm's Eudora", Network Computing, May 1 1994, 86. Dorner discusses the origins of the project in Enterprise Times, Dorner Speaks: A Conversation with the Creator of Eudora (Qualcomm, 1994 [cited August 25 2006]); available from http://www.ibiblio.org/wpercy/welty/html/dorner.html.

[32] Qualcomm, QUALCOMM Launches Project in Collaboration with Mozilla Foundation to Develop Open Source Version of Eudora Email Program (Eudora.com, October 11 2006 [cited October 20 2006]); available from http://www.eudora.com/press/2006/eudora-mozilla_final_10.11.06.html. While work on the existing Eudora program has ceased, the Eudora name is to be attached to a new, free program based on a customized version of the open source Thunderbird code managed by the Mozilla Foundation. Steve Dormer is still involved in the project.

Internet email client with its Navigator browser from version 2.0 (1995) onward. Microsoft offered an email client as a standard part of Windows 95, and bundled email and news clients with its own Internet Explorer browser from version 3.0 (1997) onward. From 1997 this was known as Outlook Express, and today it is the world's most widely used email client, though it has not received a significant update since 1999 or any new features at all since 2001.[33] In 2005 Microsoft assembled a new development team to work on the program, and has announced plans to include an improved version, renamed again to Windows Mail, in its forthcoming Windows Vista operating system.[34]

Microsoft also offered a powerful "Internet mail only" mode in its Outlook package (which despite the name was entirely unrelated to Outlook Express). Outlook was marketed as a personal information manager, a software category created during the 1980s to combine address book, calendar, and note capabilities. While many of Outlook's most powerful features worked only with Microsoft's own Exchange mail server, it was a popular package among Internet email

---

[33] At one point a Microsoft manager publicly announced that development work had ceased on Outlook Express because Microsoft preferred to steer customers toward its paid-for Outlook package. Angus Kidman, Microsoft Kills off Outlook Express (ZDNet, August 13 2003 [cited September 20 2006]); available from http://www.zdnet.com.au/0,139023166,120277192,00.htm This seemed to accurately reflect its actual policy from 2000 until 2005, though the announcement was later retracted. Angus Kidman, Outlook Express Gets Last Minute Reprieve (ZDNet Australia, August 15 2003 [cited September 20 2006]); available from http://www.zdnet.com.au/0,130061791,120277332,00.htm

[34] John Clyman, For a Better Internet Experience (PC Magazine, April 16 2006 [cited September 24 2006]); available from http://www.pcmag.com/article2/0,1759,1950279,00.asp Microsoft has also announced "Windows Live Desktop", another email client program, designed to work with Microsoft's own Windows Live webmail service (formerly Hotmail) and other email services.

users, if only because of its bundling with Microsoft's ubiquitous Office suite. Thus by the 1990s anyone using a recent version of Windows, downloading either of the two leading browsers or purchasing Microsoft Office would have acquired a perfectly usable email client. Rather than supplying Eudora or another package to their users, ISPs could simply include some instructions on how to configure Outlook Express or Netscape to access their user's email. It is little wonder that only a small minority of users chose to spend time and money on commercial competitors such as Eudora Pro.

Internet email was the first hit application of the personal computer to break the pattern established with the early firms of the 1970s and 1980s. Each new class of popular application spawned at least one major company: Microsoft and Borland from programming languages, WordPerfect and Micropro from word processors, Lotus and VisiCorp from spreadsheets, and Aldus from desktop publishing. More people spent more time using email than any of these earlier kinds of programs. Yet nobody got rich selling Internet email programs. In part this was a result of Microsoft killing the market for email client programs by giving away its own email software. But it also reflects the preexisting norms of the Internet, where software was usually written by academic or enthusiasts with no need to support themselves from royalties and people were used to downloading programs rather than going to buy them in a store.

The sale of Internet mail server software also failed to become a big business, or really any kind of business. Internet email transmission has usually accomplished either using open source products (such as sendmail and the increasingly popular qmail) or tools bundled as a standard part of all recent releases of proprietary server operating systems such as Windows Server and Novell Netware. In contrast, the market for the more capable proprietary email systems remains healthy. Even today, large organizations continue to use Microsoft Exchange

and IBM's Notes to handle internal email, relying on gateway capabilities built into these systems to translate and retransmit mail intended for external recipients. Some estimates showed that Exchange overtook Notes in 2001 to become the most widely used corporate email system. In 2006, it had an estimated 118 million users worldwide.[35] According to computer industry analyst Gartner, in 2005 Microsoft and IBM collectively controlled more than 86 percent of the market for "enterprise e-mail and calendaring" systems measured by revenue volume.[36] While many times smaller than the market for Internet email in terms of the number of users, this market remains much bigger in dollar terms, as most Internet email software and email accounts are given away free.

Webmail

With the rise of the Web an entrepreneurial opportunity did present itself, in the shape of Web mail systems. The first of these, hotmail.com, was launched in 1996 by Sabeer Bhatia, a young Indian immigrant working as an engineer in Silicon Valley, and his coworker Jack Smith. Hotmail offered users free hotmail.com email accounts, allowing them to read and compose their email using standard Web browsers. Free webmail services had two main advantages from the point of view of users: they were independent of their ISPs and so would remain valid when graduating from school, changing jobs, or moving; and they could be accessed from any Web browser. This meant that new and stored emails were accessible when traveling, and from office

---

[35] The Radicati Group, Microsoft Exchange Market Share Statistics, 2006 (March 2006 [cited October 29 2006]); available from http://www.radicati.com/brochure.asp?id=284.

[36] Tom Austin, David W Cearley, and Matthew W Cain, Microsoft E-Mail Momentum Growing at IBM's Expense (Stamford, CT: Gartner, 2006).

or public locations where other access methods might be blocked by firewalls or restrictions. Webmail also removed the need to install and configure an email client program.

Webmail was quite simple to accomplish technically: essentially the email application ran on the server like a traditional email reader such as Elm, but directed its output to a Web page rather than a terminal window. But as the first such service Hotmail received sufficient venture capital to cement its advantage by giving away eight million email accounts without having to worry about the cost. At the end of 1997 Microsoft acquired Hotmail for $400 million dollars, making Hotmail one of the emblematic success stories of the early years of the Internet boom: a fable of overnight wealth based on the confident and timely execution of a good idea.[37] Unlike most of the other Internet businesses created to give away things that usually cost money, webmail made considerable sense. Providing the service required only modest technical and network resources, but because users would visit the site frequently to check and reply to email they would spend a great deal of time exposed to advertising messages. Advertising, initially for Hotmail itself, could also be included at the bottom of each outgoing email message. Finally, as users grew more reliant on the service, some could be induced to pay monthly fees for premium services such as greater mail storage or the ability to download email to a desktop email client such as Eudora.[38]

The Hotmail idea was not hard to duplicate, and other firms such as the start-up RocketMail and the Yahoo, Excite and Lycos portals entered the free webmail business. In 1997 Yahoo acquired RocketMail, whose technology it had previously been licensing, and built up

---

[37] Po Bronson, "HotMale", Wired 6, no. 12 (December 1998).

[38] Sharael Feist, Microsoft to Charge for E-mail Retrieval (News.com, June 5 2002 [cited September 28 2006]); available from http://news.com.com/2102-1023_3-933024.html?tag=st.util.print.

email as a key feature of its popular Web portal. Yahoo's email service expanded to include calendar capabilities, discussion groups and instant messaging. Yahoo's system eventually overtook not just Hotmail (to become the most widely used free email service) but also AOL (to become the most widely used email service of any kind).[39] In 2004, Google shook up the webmail business with its announcement of Gmail.com, a new service with a powerful and uncluttered user interface, an unheard of gigabyte of online message storage (250 times the amount offered with free accounts from market-leader Yahoo), and the ability to download mail to standard clients and forward incoming mail to other accounts. Google's free email service was far superior to the premium services offered by its competitors in exchange for paid subscriptions. Initially Gmail was available by invitation only, with each user given a number of invitations to pass on to their friends and colleagues. By 2005 this policy had been relaxed. Though other services quickly moved to match Gmail's generous storage limits, Gmail has recently added calendar, voice communication and text chat capabilities.[40] Still, people are reluctant to switch email services without a very good reason, and by mid-2006 Gmail's 8.6 million active users only earned it fourth place in the US market behind Yahoo (with 78 million users), AOL and Hotmail.[41]

---

[39] Figures on usage are taken from Saul Hansell, "In the Race With Google, It's Consistency vs. the 'Wow' Factor " New York Times, July 24 2006.

[40] Jon Markoff, "Google to Offer Instant Messaging and Voice Communications on Web" New York Times, August 24 2005: 4

[41] Hansell, "In the Race With Google, It's Consistency vs. the 'Wow' Factor " These figures represent estimates of the "audience size" during June 2006 for the different services, rather than the numbers of registered users. ComScore Medix Metrix figures on global webmail usage show the services ranked in the same order, but with much higher numbers of users (more than 250 million for Yahoo) according to Ina Fried, Hotmail's New

Such was the popularity of webmail that leading mail servers such as Microsoft Exchange and Lotus Notes (or Domino as the server has been renamed) now offer it as a standard feature, meaning that ISPs and employers generally provide Web access to email accounts. It still has the disadvantage that users cannot read or compose email while off-line, but in today's world of high speed broadband connections and Wi-Fi hot spots this is less of a drawback than it once was. Similarly, new kinds of discussion and collaboration systems based on Web technologies, such as blogs, wikis, and online community sites, have largely usurped the once separate technologies used to disseminate and access Usenet newsgroups. The Web browser's role as a universal interface to online resources grows ever broader.

Spam

The other big economic opportunity opened up by email has been an unexpected one: the sending and blocking of spam (unsolicited commercial bulk email). SMTP was designed for use on a small, closed network where commercial activities were expressly forbidden, and whose user population shared an academic culture where trust and the free exchange of information were assumed. This was in sharp contrast with the ill-fated X.400 protocol, designed from the beginning for commercial use by a diverse and potentially untrustworthy user population. Internet email provided no easy way to charge users for each email they sent (which would have destroyed the economics of spam), no easy way to determine the source of a message (and hence to block or prosecute spammers), and no protection against address faking. Because Internet email systems treated composing, sending, relaying, receiving, delivering, and reading email as

---

Address (News.com, April 26 2006 [cited September 28 2006]); available from http://news.com.com/2009-1038_3-6064507.html.

six separate tasks, accomplished with different programs, most SMTP relays of the early 1990s were happy to relay messages across the network without demanding authentication from the sender. Once Internet email became the standard for hundreds of millions of users, its own strengths became almost-fatal weaknesses: the rapid, anonymous, and free transmission of an unlimited quantity of spam.

Internet lore identifies a message transmitted in 1978 as the first bulk unsolicited commercial email sent on the ARPANET. The message, sent by a DEC sales representative, invited every known network user on the West Coast of the United States to a sales presentation. This faux pas solicited rapid and widespread criticism from network users, and it was not repeated. [42] By 1994, however, the Internet was used by millions rather than thousands. On March 5, the small and rather seedy husband-and-wife law firm Canter and Siegel set aside all the rules of "netiquette" (as good online manners were then called) by posting advertisements for its immigration services in every one of the thousands of Usenet newsgroups.[43] Canter and Siegel then founded a firm called Cybersell to offer spamming services to others. Aggrieved Internet users fought back with large emails designed to overwhelm the firm's email boxes and unsolicited magazine subscriptions to fill its physical mail boxes, but it was already clear that the traditional social norms of the Internet would be insufficient to combat this new surge of unashamedly antisocial commercial activity.

The term "spam" itself, which gained currency to describe this kind of bulk solicitation, was a classic example of the geek whimsy central to the Net's pre-commercial culture. It referred

---

[42] Brad Templeton, Reaction to the DEC Spam of 1978 (n.d. [cited September 20 2006]); available from http://www.templetons.com/brad/spamreact.html.

[43] Glyn Moody, "Spam, Spam, Spam", Computer Weekly (October 03 2002):34.

to a comedy sketch from the surreal 1970s British television show *Monty Python's Flying Circus*, in which a couple entered a cheap cafe only to discover that every choice on the menu held greater or lesser amounts of Spam, a processed meat product emblematic of the poor state of British cuisine during the immediate post-war decades. No substitutions were permitted, so while Spam lovers could choose delicious items such as "spam spam spam spam spam spam spam beaked beans spam spam spam and spam," even the most spam-averse customer was forced to order "egg bacon spam and sausage" as the dish with the least Spam in it.[44] In homage to this, "spam" appears to have established itself during the 1980s as an obscure geek idiom meaning "something unpleasant forced on people indiscriminately and in large quantities" before springing into widespread usage to describe bulk email and newsgroup messages.

Spamming email accounts was a little harder than spamming newsgroups, largely because no central directories of email addresses existed. Spammers solved the problem by using software to roam the Web and newsgroups in search of email addresses, or simply by sending messages to randomly created addresses in popular domains such as AOL and Compuserve. Spammers used free email accounts, specialized bulk emailing software tools, and "zombie" computers (stealthily taken over by worm attacks) to send out billions of email messages. Because spam annoyed its recipients, the return addresses supplied were usually bogus. The advertisements promoted fraudulent or disreputable offers: invitations to embezzle funds from Nigerian bank accounts, cut-price impotence drugs, bad investment advice, and pirated software. One particularly dangerous kind of spam, so-called "phishing messages," requested that users visit fake but convincing replicas of leading Web businesses such as EBay and Citibank to enter

---

[44] Graham Chapman et al., The Complete Monty Phyton's Flying Circus: All the Words, Volume 2 (New York: Pantheon, 1989), 27-29.

their passwords. The fake website logged the account information, allowing the criminals behind the scam full control of the user's online account.[45] By the start of 2005, spam accounted for 83% of all emails sent.[46] According to one widely quoted estimate, the effort required to deal with all these messages wasted an annual average of $1,934 worth of productivity for each and every corporate employee in the United States.[47]

The rising tide of spam floated the fortunes of a whole new industry providing anti-spam products and services. The most effective of these filtered incoming email to detect spam, based on a number of characteristics, including its format, the presence of certain words such as "Viagra", and its statistical similarity to known spam messages. Spam could be blocked by software running on the personal computers of email recipients (in 2003 a simple spam filter was built into Microsoft Outlook), or detected on mail servers and deleted before ever being downloaded by users. Many anti-spam systems have been produced, both commercial and open source. Some spam blockers have even been packaged into self-contained hardware units to minimize the effort required to add them to a network. As of 2006 the anti-spam market remains fragmented, though filters appear to be working increasingly effectively. My own Gmail spam filter catches around fifty messages a day, letting only a couple through. Spams have been

---

[45] Mary Wisniewski, "Banks work on Hooking 'Phyishers': By Yearend: More Proof of Who You Are Online" Chicago Sun-Times, May 29 2006, 49.

[46] Anonymous, "Winning the War of Spam", Economist, August 20 2005.

[47] Nucleus Research, Spam: The Serial ROI Killer (Nucleus Research, 2004 [cited May 20 2006]); available from http://www.nucleusresearch.com/research/e50.pdf. While widely quoted, this number does not seem very plausible. It assumed that employees received 29 spam emails a day, which is plausible, and that each one took thirty seconds to deal with, which is not. However it did not include the costs of IT staff, software, hardware and network bandwidth devoted to transmitting, storing and clocking spam.

growing ever more surreal, in a generally unsuccessful attempt to circumvent such filters, often consisting of meaningless sequences of blank verse, interspersed with occasional Web addresses, mangled names of products such as “SOFT CIAzLIS”, or “V/AGRA,” and the opportunity to “reffnance” or to buy “s0phtw_aRe.”

Instant Messaging

The Internet is widely used for another kind of personal communication: instant messaging. Emails can be sent at any time and, like traditional mail, accumulate in the recipient’s mailbox. Instant messaging, in contrast, is more like a phone call. The people communicating are online at the same time, and send a conversational stream of short messages back and forth. Instant messaging software allows users to see which of their contacts are online and available for chat at any given time.

The basic concept was established long ago. Time-sharing systems had allowed users to “chat” with each other since the 1960s. Internet users could “finger” their friends to see whether they were currently logged in, and use the “talk” command to print short messages on their screens. A popular free protocol, Internet Relay Chat, was widely used to create Internet chat rooms from the late 1980s onward.[48] In the world of commercial online services, large part of AOL’s success has been attributed to its hugely popular public chat rooms. The term “instant messaging” was popularized by AOL, as part of its proprietary online service. Users could establish a “buddy list,” and would then see which of their acquaintances were online. Using AOL’s Instant Messenger feature they could send messages to pop up on their friends’ screens.

---

[48] IRC is specified in J. Oikarinen and D. Reed, RFC 1459: Internet Relay Chat Protocol (Network Working Group, 1993 [cited September 5 2006]); available from http://www.irchelp.org/irchelp/text/rfc1459.txt.

There was, however, no established standard for person-to-person private instant messaging over the Internet prior to its commercialization. Internet email packages such as Eudora and Outlook Express had merely to implement proven protocols for users to interface with their email accounts. Internet-based instant messaging systems, in contrst, had to establish their own protocols. They were competing not merely to offer better software but to build their own proprietary networks.

Instant messaging over the Internet was pioneered by Mirabilis, an Israeli start-up firm founded in 1996. Word of its new ICQ service (pronounced “I seek you”) spread rapidly. According to one January 1997 report, the firm gained sixty five thousands users within seven weeks of the creation of its first software before it had advertised or even put up a home page.[49] By the end of 1997 a number of firms were competing in the fast-growing market, including Excite, Yahoo, and a variety of specialist start-ups. But ICQ was by far the most popular, and ten months after its release it claimed more than three million users.[50] Fighting back, AOL released instant messenger software for users of Netscape’s Navigator browser, allowing them to send messages to each other and to customers of its online service. The companies competing in this market aimed to make money by selling advertisements to be displayed to their users, or by licensing the software to corporate users or ISPs. In June 1998 AOL announced that it was purchasing Mirabilis and merging its own Instant Messenger network with ICQ’s network of twelve million users.[51]

---

[49] Virginia Baldwin Hick, "Snowbound? Log On and Keep On Talkin'" St. Louis Post-Dispatch, January 20 1997, 1.

[50] Elizabeth Waserman, "Desktop to Desktop in a Flash" Pittsburgh Post-Gazette, October 19 1997.

[51] Jon Swartz, "AOL Buys Net Chat Company" San Francisco Chronicle, June 9 1998, E1.

Instant messaging programs gradually acquired additional features, including voice conversation over the Internet or with conventional telephones, file exchange, and integration with email. As a result the markets for instant messaging systems, Internet telephony systems (such as Skype) and webmail systems (such as Gmail) have largely converged. The instant messaging field has moved slowly toward standardization, but remains divided into several incompatible networks. Although numerous standards have been created, the most popular services prefer to keep their customers locked into their proprietary software. AOL Instant Messenger (AIM) is now the most popular, and the company has extended its reach to cellphones and other wireless devices. Apple's iChat software is able to communicate with AIM users. However, AOL took legal and technical measures to prevent other companies from producing software compatible with this service, blocking attempts by Microsoft to couple its MSN Messenger (released in 1999) with AOL's existing user network. To compete with AIM, Yahoo and Microsoft agreed in 2005 to merge their instant messaging networks. According to Internet traffic measurement firm ComScore, AIM, MSN Messenger, and Yahoo Messenger were the most widely used systems in North America during February 2006. The study showed that instant messaging was widespread in the United States with 37% of the online population sending a message that month. Interestingly, this lagged far behind the proportion in other regions, with 49% of European Internet users and 64% of those in Latin America relying on instant messaging.[52] This mirrors the exceptionally slow adoption of cellphone text messaging by North Americans.

---

[52] ComScore Networks, Europe Surpasses North America In Instant Messenger Users, comScore Study Reveals (2006).

All the major instant messaging networks remain free to users, are associated with proprietary access software, and are supported by showing advertisements. The most successful open instant messaging standard, Jabber, claims a relatively modest ten million users working with a variety of software (most notably Google's Google Talk service). Several small companies and open source teams have produced programs able to communicate with users of all the popular networks, but programs such as Trilian, though popular, account for only a small fraction of all instant messaging users.

Instant messaging appears to have found its niche among younger users people using the Internet from home. Unlike email instant messaging has been slow to emerge as a business tool. This is perhaps because email initially presented itself to business users as an extension of a known quantity: the interoffice memo. Instant messages, in contrast, are more personal, more informal, and harder to archive. As well as being a significant Internet communication technology in its own right, as an application that came of age during the commercialized era of the Internet instant messaging provides an important contrast with email. Because no dominant standard was established during the pre-commercial era, Internet instant messaging remains balkanized between closed standards established by powerful companies. This has hurt the ability of users to communicate, and slowed the mainstream adoption of instant messaging.

2: THE WORLD WIDE WEB

The ambitiously named World Wide Web was created by Tim Berners-Lee, a British computer specialist with a Ph.D. in physics working at the European particle physics lab CERN.[53] The first website went live on August 6, 1991. Berners-Lee's prototype browser and server software worked only on NeXT workstations, an obscure and commercially unsuccessful type of computer based on a heavily customized version of the Unix operating system.[54] Unlike most later browsers, it also allowed users to edit existing pages and create new ones. On the other hand, it could not display graphical elements within a Web page and used a cumbersome navigation system. Only fifty thousand NeXT computers were ever built[55], so the initial market for the Web was not large.

The Web was assembled from existing building blocks in a matter of months. Berners-Lee had neither the resources to develop Web software for more widely used systems (in fact,

---

[53] For the story of the Web's original design see Tim Berners-Lee and Mark Fischetti, Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor (San Francisco: Harper, 1999) and James Gillies and Robert Cailliau, How the Web Was Born: The Story of the World Wide Web (Oxford: Oxford University Press, 2000).

[54] In addition to the NeXT browser, a student intern working with Berners-Lee at CERN created a simple "line mode" text-only browser that could be used from anywhere on the Internet via telnet to CERN. Gillies and Cailliau, How the Web Was Born: The Story of the World Wide Web, 203-05. This allowed people to experiment with the Web, but it had many limitations including a lack of support for different fonts and the need to type numbers in order to activate hyperlinks. Its use is described in Ed Krol, The Whole Internet User's Guide and Catalog (Sebastapol, CA: O'Reilly & Associates, 1992), 227-29.

[55] Ken Siegmann, "Cannon Drops Deal to But Next Inc.'s Hardware Unit" San Francisco Chronicle, April 3 1993, D1.

CERN denied his request for assistance in creating a browser for more common versions of Unix) nor the inclination to create a company to market his invention.[56] CERN devoted a total of about twenty man-years effort to the Web during the entire course of its involvement with the project, not all of it authorized and much of it from interns.[57] This unavoidable reliance on proven, widely used, technologies was a key factor in the Web's success.

Three crucial standards defined the Web, each of which survived essentially intact from Berners-Lee's 1991 prototype into the Netscape and Mosaic Web browsers used by millions a few years later. Each built extensively on, and was only possible because of, existing Internet infrastructure.

1) HTTP, the HyperText Transfer Protocol, used by Web browsers to request pages and by Web servers to transmit them. This runs on top of TCP/IP, which handled all the hard work of sending data between the two. In this it followed the model pioneered by earlier Internet application standards, such as SMTP for mail transmission.

2) HTML, the HyperText Markup Language. After Web pages have been transmitted to the browser via HTTP, the browser decodes the HTML instructions and uses them to display the Web page on the screen. In the early days of the Web, pages were hand-coded directly in HTML using a simple text editor. HTML was an application of SGML, the Standard Generalized Markup Language, designed as a universal and extendable way of embedding information on document structure into text files. HTML included the ability to link to Web pages and other Internet resources, and "form" capabilities so that information entered into the browser could be relayed back to the server.

---

[56] Gillies and Cailliau, How the Web Was Born: The Story of the World Wide Web, 200.

[57] Ibid, 234.

3) The URL, Uniform Resource Locator (originally Universal Resource Identifier). This extended the existing DNS (Domain Name Server) system of using meaningful names (such as public.physics.upenn.edu) rather than numbers to identify computers connected to the Internet. The URL added prefixes such as nttp://, FTP://, and telnet:// to identify the particular kind of request to make to the machine in question. The prefix http:// was used to identify Web pages, in conjunction with an optional filename such as /pub/reports/newreport.html to identify the particular page required, and a TCP/IP port number such as :80 to identify how to communicate with the Web server. A full URL might therefore look like http://public.physics.upenn.edu/pub/reports/newreport.html:80.

The Web might have been a primitive hypertext system, but it was a first rate interface for the disparate resources scattered across the Internet. Berners-Lee's great contribution was to produce a simple and workable method by which it was as easy for the author of a document to link to a page on the other side of the world as to another part of the same document. Given the existing capabilities of the Internet this was technically trivial, but his invention of the URL and its use to specify links to resources on other computers made it possible for the first time to access material on the Internet without needing to know what system it was housed on.

The Web was by no means the only attempt to build a browsable interface to the ever growing mass of resources scattered across the Internet. A system named Gopher had been publicly released by the University of Minnesota in 1991. It spread much more rapidly than the Web at first, largely because Gopher software was offered for widely used computers well before fully featured Web browsers. Like the Web, Gopher displayed text documents, could index to resources held on other servers as well as indexes on the same server, and worked as a convenient means of cataloging Internet resources including telnet and FTP (File Transfer

Protocol) servers. Gopher was less flexible than the Web, because it worked on a more formal menu system, rather than allowing insertion of hypertext links within documents. But Gopher's rapid eclipse by the Web from 1993 onward has often been attributed more to attempts by the University of Minnesota to control its development and make money from it than by any technical shortcomings. CERN made it clear that anyone was free to create a Web browser or server based on the published specifications of HTTP, HTML, and the URL system. In contrast, after the University of Minnesota announced in 1993 that it planned to charge fees for its own Gopher server software many feared that it would seek to assert its intellectual property rights against rival Gopher servers.[58]

At first, the great appeal of the Web was that it provided an easy and consistent way to catalog and browse the existing resources of the Internet. Any new communication technology faces what most people would term a "chicken and egg" issue and economists would call a "lack of network externalities." Why build a website when nobody has a Web browser? Why download and configure a Web browser when there are no websites to visit? The challenge is to persuade a group of enthusiastic users to invest time and money in the new technology, so as to build up a sufficiently large and visible community to attract others. Early browsers cleverly solved this problem by making themselves into a universal interface for existing Internet resources. In 1992, very early in the Web's development, Ed Krol wrote in his seminal *Whole Internet Catalog & User's Guide* that so far "the World-Wide Web really hasn't been exploited fully yet… You can look at a lot of 'normal' resources (FTP archives, WAIS libraries, and so on), some of which have been massaged into Hypertext by a clever server…. Hypertext is used

---

[58] Philip L. Frana, "Before the Web there was Gopher", IEEE Annals of the History of Computing 26, no. 1 (January-March 2004):20-41.

primarily as a way of organizing resources that already exist…"[59] Even the line mode browser made available via telnet by CERN in 1991 integrated support for gopher, Usenet news browsing, and file downloading and browsing from FTP servers. Browsers running on personal computers could automatically open other programs as needed for kinds of resource such as telnet sessions, Postscript documents, and video files. Thus, the browser was far more useful than if its powers had been limited to viewing pages on the handful of websites then available.

In the meantime, the Web was gaining popularity following the release of browsers for other, more widely used, types of UNIX workstation. Berners-Lee's own browser vanished rapidly (though not without a trace – his "libwww" code was freely available and gave an easy starting point for later efforts). The important thing, however, was not the software but the free-to-use protocols and standards on which it was based. Because the Web relied on open, published and simple standards anyone with an Internet connection and a modicum of programming ability was free to create a Web browser or server and then hook it up to the ever growing Web. After all, the Internet had been created specifically to connect computers using different kinds of hardware and running different kinds of software. Within months, amateur teams across the world were at work on improved browsers.

The most influential of the 1992 crop of releases was the Viola browser for Unix workstations, created by a student at the University of California at Berkeley. Viola introduced several key features of later browsers: forward and back buttons, the history function, and bookmarks for favorite pages.[60] Unlike Berners-Lee's browser, Viola worked with the standard

---

[59] Krol, The Whole Internet User's Guide and Catalog, 232.

[60] Viola is discussed in Gillies and Cailliau, How the Web Was Born: The Story of the World Wide Web, 213-7.

X Windows system. Most computers with Internet connections were running Unix in the early 1990s, so it was the most important platform for both browsers and servers during the first few years of the Web.

The first browser to achieve widespread use, however, was Mosaic. Constructed by a small team at the National Science Foundation funded National Center for Supercomputer Applications (NCSA) of the University of Illinois, Mosaic was the first browser to be released for multiple platforms and the first to extend HTML to permit the display of images within a Web page. The public release of the first X-Windows version of Mosaic in early 1993 was followed by Mac and Microsoft Windows versions. In January 1994, a survey of Web users suggested that 97 percent of them used Mosaic as their primary browser, and 88 percent used Unix (Windows and Mac computers had yet to make serious inroads on the Internet).[61] By October 1994 Mosaic had an estimated two million users.[62]

Mosaic's support for Gopher servers was particularly important, because it opened the way for the rapid shift of Gopher users to the Web since they could now access the existing mass of gopher pages and the fast growing population of Web pages with a single stylish, graphical client. Users could wander from Web pages to Gopher indexes to downloading files via FTP without even being fully aware of the transition between protocols. Even the name Mosaic reflected the idea that the browser was pulling together shards of information already present on

---

[61] James Pitkow and Mimi Recker, Results From The First World-Wide Web User Survey (1994 [cited June 20 2006]); available from http://www-static.cc.gatech.edu/gvu/user_surveys/survey-01-1994/survey-paper.html.

[62] Robert H. Reid, Architects of the Web: 1,000 Days that Built the Future of Business (New York: John Wiley & Sons, 1997), 38.

the Internet, assembling them for the first time into a coherent picture. While Gopher was text based, and so could be used on simple terminals and computers without direct connections to the Internet, even this advantage was countered by the increasing popularity of the text-based Web browser Lynx, launched in 1992. By 1994, Web traffic had overtaken Gopher traffic, and over the next few years most of the leading Gopher sites shifted their content over to the Web.

The Web quickly won favor as a universal interface mechanism for online applications, such as library catalogs and institutional telephone directories. Before this, using an online database had meant either working with a plain text interface and terminal emulation software, which was ugly and hard to use, or installing a special piece of software written to work with the database in question. Special software usually required a lot of work to configure and tweak it, and would have to be updated whenever a new version of the system was released. But if the system was hooked up to a Web server then anyone with an up-to-date Web browser could work with it, regardless of what kind of computer it was running on. The complexity was all hidden away behind the Web server at the other end.

These interactive Web systems relied on dynamically generated pages. Whereas a standard, or “static” Web page is stored on the server exactly as it will be transmitted to a browser, a dynamic page is tailored to each user and filled with customized information. HTML included the capability to add “form fields” to Web pages, such as boxes to type information into or buttons to click. This information was passed back to the Web server, and used to customize the page sent back to the user. For example, if someone entered the name of an author on a library page and then clicked the “submit button” then the system might perform a database search, format the results into a Web page, and transmit this back to the searcher.

The challenge was in hooking Web servers up to databases and programming tools. The Mosaic team solved the problem by adding a simple but highly flexible feature called CGI (Common Gateway Interface) to their Web server. CGI was just a way for a Web server to take information entered into a user's browser, execute a specified program to process this data, and then receive back the results for transmission to the user. But because a CGI call could trigger any kind of software this kept the Web server itself very simple and let system builders use their preferred programming language, database system, or scripting tool to handle the task. It could be inefficient with busy systems, but made it quick and easy to hook existing applications up to the Web. CGI followed the classic Unix and Internet philosophy of building a number of flexible, specialized software tools and connecting them together as needed to solve particular problems. Here, as elsewhere, the Web succeeded by building on the mass of established technologies.

Mosaic seemed to offer the University of Illinois a huge financial opportunity. By the 1990s, major American universities had built well funded technology transfer offices, staffed with teams of lawyers to scrutinize research contracts, patent inventions made by university staff and license the rights to develop those inventions commercially. A handful of spectacular successes, most notably Stanford's patent on gene splicing and the University of Florida's rights to the Gatorade sports drink, had convinced university officials that savvy exploitation of the right invention might bring in hundreds of millions of dollars. But the successful licensing of software proved particularly difficult, because of the gulf between a lab prototype and a salable system. Computing research had given a few universities such as MIT and Stanford equity in some reasonably successful start-up firms, but it had produced no blockbusters.

Here, however, was a university with rights to the most promising new computer application since VisiCalc. While Mosaic itself remained free to use, university authorities decided to offer rights to develop its code commercially to firms interested in marketing their own browsers. The initial terms were $100,000 per license, plus a royalty of $5 for each browser sold.[63] The university licensed the Mosaic code to several firms. The most important licensee was Spyglass, Inc, a firm created to commercialize technologies developed by the NCSA. Spyglass produced its own improved Mosaic browser, and in turn licensed this code to many other companies. By 1994 a rash of commercial browsers had appeared, almost all of them based on either the original NCSA Mosaic or the enhanced Spyglass version code. One 1995 review of Web browsers included no less than twenty offerings.[64] But unfortunately for Spyglass, which had negotiated royalty agreements with its own licensees, the market for browsers would be both short-lived.[65]

Netscape and AOL Bring Browsing to the Masses

By 1995, Mosaic and its commercial descendants had themselves been marginalized with astonishing rapidity by a new browser designed from the beginning as a commercial product: Netscape Navigator. When first released in an unfinished “beta” version, in October 1994, it was

---

[63] Gary Wolfe, "The (Second Phase of the) Revolution Has Begun", Wired Magazine, October 1994.

[64] Clive Parker, "The Reckoning".Net: The Internet Magazine, September 1995, 76-81.

[65] Although Spyglass conspicuously failed to dominate the Web it still proved an exciting investment for the university. The firm remade itself as a provider of embedded Internet software for cell phones and other electronic devices, went public, and was acquired at the height of the Internet boom by cable TV box firm OpenTV, in exchange for stock briefly valued at more than two billion dollars. Claudia H Deutsch, "OpenTV, a Software Provider, To Buy Spyglass for $2.5 Billion" New York Times, March 27 2000.

already more usable, reliable, and powerful than its many competitors. Netscape was a well-funded Silicon Valley start-up firm that hired many key members of the original Mosaic team, including its cofounder Marc Andreessen, a twenty-three year old programmer who served as Chief Technology Officer. Navigator was much more efficient than Mosaic, and so gave acceptable performance over dial-up modem links and Internet connections slower than the high-speed links enjoyed by computer science labs and research centers. From the beginning, it was available in ready-to-use versions for Microsoft Windows, Mac, and several different versions of Unix. Netscape eventually claimed more than 80 percent of the browser market, crushing Mosaic and most of its licensed derivatives.[66]

The unprecedented success of Navigator made Netscape into a symbol of the so-called Internet Revolution. Andreessen himself became one of the public faces of the Web, a technology business prodigy apparently cast in the mold of the young Bill Gates.[67] By early 1995, discussion of the Web had spread far beyond campus labs and corporate research departments, into the mainstream of the computer industry, the financial press, and the international media. Overnight, the rhetoric and expectations lavished on the "information highway" during the early 1990s were shifted to Web browsers and the Internet. As *Wired Magazine* wrote in 1994, "Prodigy, AOL, and CompuServe are all suddenly obsolete - and

---

[66] Michael A Cusumano and David B Yoffie, Competing on Internet Time (New York: Free Press, 1998), 11 includes a composite chart of Internet browser market share over time.

[67] For an example of the glowing press given to Andreessen, see Elizabeth Corcoran, "Software's Surf King; Marc Andreessen, 24, Is Making Waves on the Internet's World Wide Web" Washington Post, July 23 1995, H01.

Mosaic is well on its way to becoming the world's standard interface…. The global network of hypertext is no longer just a very cool idea."[68]

The spread of the Web also provided the first really compelling reason why ordinary people might want to connect their computers directly to the Internet via TCP/IP. The leading commercial online services of 1993 already allowed users to exchange emails with the Internet, and savvy users could even use this capability to request files to be sent from Internet servers via email. For users who needed more, a few specialized services offered text-based Internet accounts. Customers used terminal emulation software to dial into to a central server, on which they could run text-based applications such as email, Internet news or programming tools. Only the server was connected to the Internet, or could execute Internet applications, but the service was affordable and everything worked fine with a slow modem and low-end personal computer. Krol's 1992 classic *The Whole Internet User Guide and Catalog* said a "dial-up Internet connection," of this kind was available for as little as twenty dollars a month.[69]

But accessing the Web required a different kind of set-up: a computer with a fast processor, ample memory, and high resolution display connected directly to the Internet. In 1993, when Mosaic first appeared, only powerful Unix workstations costing tens of thousands of dollars and hooked to high bandwidth campus networks could offer a vaguely satisfying Web experience. One 1995 report noted that "You can't pick up a newspaper or magazine without reading about the World-Wide Web…. Despite the hype, only a small percentage of the huge numbers that are touted to be 'on the Internet' are really there, except by e-mail. Right now,

---

[68] Wolfe, "The (Second Phase of the) Revolution Has Begun".

[69] Krol, The Whole Internet User's Guide and Catalog, 8&336.

relatively few have the full, direct access that allows full-screen GUI interfaces, graphics and sound--and access to the Web." [70]

This was changing rapidly, as advances in browser software coincided with the inclusion of newly capable processors (Pentium chips for Windows computers and PowerPC chips in Macintoshes) in mass market computers, improvements in operating systems, and a glut of inexpensive and relatively high-speed (up to 33.6 Kbit/s) modems. Krol dealt only briefly with the exotic and expensive idea of using the new SLIPP or PPP protocols to make a TCP/IP connection to the Internet over a telephone line, which he suggested would cost around two hundred and fifty dollars a month.[71] But by 1995 hundreds of companies, large and small, were competing to sell true Internet access over telephone lines and so turn household computers into fully-fledged Internet nodes for the duration of the connection. Suddenly it was practical, even fun, to browse the Web over an ordinary telephone line on a personal computer costing three thousand dollars.

In late summer 1995, Netscape had around ten million users and was about to ship version 2.0 of its browser. On August 9, just seventeen months after its founding, the firm made an initial public offering of stock. Its share price doubled on the first day of trading, valuing it at more than two billion dollars.[72] Netscape was hailed as the most successful start-up firm in history, and suddenly venture capitalists and stock investors were rushing to invest their money in anything Internet related. While other software companies benefited, most of the so-called

---

[70] Nancy Garman, "A New Online World", Online 19, no. 2 (March-April 1995):6-7.

[71] Krol, The Whole Internet User's Guide and Catalog, 335.

[72] Reid, Architects of the Web: 1,000 Days that Built the Future of Business, 44.

Internet stocks were companies trying to make money by providing some kind of product or service over the Internet.

Despite *Wired*'s warning that they were obsolete, the established online services quickly realized the importance of the Web. The most significant of these was now AOL. Building on its user friendly software, AOL had expanded frantically. From 1993 until the end of the decade, the company relied on a "carpet bombing" program, mailing out copies of its software and offers of a free trial by the million, attaching them to the cover of computer magazines, burying them in cereal boxes, and even handing them out on planes. By 1996 it had become the largest of the online services, with more than six million subscribers.

AOL gradually shifted itself toward the Internet, evolving piecemeal from a closed network responsible for finding its own content to a hybrid model where it eventually offered full Internet access. It dealt with the threat by buying small Internet software firms to acquire skills and code, shifting its own internal networks toward standard technologies, and gradually incorporating Internet capabilities into its existing client software. In 1994, AOL incorporated access to Usenet newsgroups into its service, unleashing what many newsgroup veterans regarded as a flood of ill-mannered plebeians into their private club. In late 1994, after Prodigy became the first of the traditional online services to offer Web access, AOL purchased an obscure Web browser company called BookLine and built its technology into a new version of the AOL client software. It became the world's biggest distributor of Internet software, bundling, modifying, and rewriting utilities into its smoothly integrated client program.[73]

---

[73] AOL distributed hundreds of millions, probably billions, of packages holding its software during the 1990s. As well as being mailed out and glued to the front of computer magazines they were, according to a history of the company, "handed out by flight attendants on American Airlines, packaged with flash-frozen Omaha Steaks"

Even after Web access capabilities were added, AOL continued to steer all but the most determined of its users to its own content rather than the broader Web. With version 3.0 of AOL, released in 1995, it became possible to use standard Internet tools and browsers with the service. AOL's own tools were sometimes limited, earning it the dismissive description of being "training wheels for the Internet." But for new users its friendly, integrated and nicely sorted interface remained a selling point. This was a real advantage, because in 1995 setting up a typical one year old computer running Windows 3.1 to use standard Internet software might require obtaining separate email, news, dialer, TCP/IP, chat and Web browser programs and configuring each of these with various obscure settings. AOL installed everything automatically, and hid the cracks so well that many users were not fully aware of the difference between AOL and the Internet. By end of the 1990s, AOL captured about half of the entire American market for home Internet users and claimed more that twenty-five million users.[74]

Remaking the Web as a Business Platform

The early Internet had been an explicitly non-commercial space. With the success of Netscape and the anointing of the Web as the ubiquitous online network of the future, businesses rushed to make money buying and selling things over the Internet. While the use of the Internet

---

and "dropped onto stadium seats at football games." Swisher, aol.com: How Steve Case Beat Bill Gates, Nailed the Netheads, and Made Millions in the War for the Web, 175. AOL's willingness to spend money to attract new subscribers threatened to ruin the company – by 1996 it was spending $270 to acquire each new customer. Amy Barrett and Paul Eng, "AOL Downloads a New Growth Plan", Business Week, October 16 1996, 85. The disks themselves are now prized by collectors looking to acquire all of the thousands of variations used.

[74] AOL, Who We Are -- History (2006 [cited September 30 2006]); available from http://www.corp.aol.com/whoweare/history.shtml.

by business is the topic of several other chapters within this volume, it is appropriate here to consider the changes made to the Internet infrastructure to support this commercial activity.

The most obvious piece of software a business needed to get online was a Web server. Traditionally, the code for server software for Internet applications such as file transfer and email exchange had been given away free of charge by its authors. Despite the commercialization of the Internet, this remained largely true. The most successful early server software was produced by the NCSA Mosaic team and distributed without charge. Plenty of companies tried to produce commercial Web server software, and both Sun and Netscape enjoyed considerable success in this market during the late 1990s. Other popular Web servers were bundled with commercial operating systems, most notably Windows. But the most widely used Web server software from March 1996 to the present has been Apache, an open source package used to host around 70% of all websites in recent years.[75] Apache is produced by the non-profit Apache Foundation, originating in 1995 as "a patchy" bundling of fixes for the NCSA server. Since then it has become a crucial part of the Web infrastructure, incorporated into commercial products from IBM, Oracle, Apple, and Novell, and used as a platform for many widely used Web applications. Because the Web relies on clearly defined standards to govern the interaction of browsers and server, few people have any idea that their Microsoft Web browser spends most of its time fetching information from open source Web servers.

Unlike commercial online services such as AOL neither the Web nor the underlying technologies of the Internet itself had been designed to support financial transactions. There was no easy way to bill users for the information they viewed, or for other goods and services ordered

---

[75] Netcraft, July 2006 Web Server Survey (2006); available from http://news.netcraft.com/archives/Web_server_survey.html.

on the Web. This had profound consequences for the development of the Web publishing and Web navigation industries. When someone ordered a book from Amazon.com (launched in 1995) or a used computer via an eBay auction (also launched in 1995) he or she generally used a credit card to pay for the purchase. This was a simple extension of practices pioneered for telephone sales, with which consumers, merchants and banks were already comfortable. But even this simple task was beyond the capabilities of Mosaic and other early Web browsers. They did not provide any secure method of transmitting credit card information – and the Internet did not help their authors by offering a standard method of encrypting or authorizing messages. This had to be fixed before the Web could be a platform for commerce, and Netscape did so when it added a then-unique feature to its first browser and server systems: the optional encryption of data entered into Web pages as they were transmitted back to the server. (Netscape displayed a lock icon on the screen to let users know that the page was secured). This advance, known as the SSL (Secure Sockets Layer) made the Web a practical platform for financial transactions and other sensitive data. SSL built on and extended existing the Internet traditions of layered protocols and open standards, to create a new protocol layer between TCP/IP and application-specific protocols such as the Web's HTTP.

To support commercial applications, Web servers needed an easy way of tracking who the user was, at least for the duration of a particular usage session. This capability was not built into the Internet itself. Its TCP/IP foundation treated each stream of data packets dispatched from one computer to another as a self-contained event, in contrast with a traditional telephone call where a connection is opened from one party to another and maintained until the call is over. Many online applications, however, naturally follow a sequence in which contact is established, a series of interactions take place, and the user finally disconnects (for example, searching and

then ordering in an online catalog, or logging into an online banking system and paying a bill). Some existing Internet application protocols, such as telnet, had been built around the concept of a session, but Berners-Lee did not build any such capability into HTTP.[76] Each request for a Web page was treated as a separate and unrelated event. Yet the Web server needed some way of tracing which user was which from one screen to another, if only to know whose bank balance or shopping cart to display. Various crude workarounds were devised for this problem by early online application developers, until Netscape introduced the idea of "HTTP cookies" with its first browser release. Cookies were essentially tracking numbers assigned by a server and then passed by back the browser with subsequent requests so that the server could more easily identify which user it was dealing with. While the Web still forced programmers to do a lot of work to create the illusion of an ongoing connection between a browser and a server, cookies removed a good deal of the messiness.

Although Netscape did much of the work to make the Web into a viable platform for online business, it generally made its enhancements public and proposed them as standards for adoption by others. Cookies, SSL, the JavaScript scripting language (discussed later), and a number of enhancements to HTML were all released first by Netscape but later codified by external groups. This ensured that Web applications would be accessible regardless of the browser being used, as long as both the creators of the pages concerned and the designers of the Web browsers adhered to the relevant standards. The most important Web standards body was the World Wide Web Consortium (W3C), headed by Tim-Berners Lee since its founding in 1994 as the guardian of HTML. In the tradition of the Internet its specifications were royalty free, with

---

[76] Jon Postel and J Reynolds, RFC 854 - Telnet Protocol Specification (Network Working Group, May 1983 [cited October 31 2006]).

the work of the consortium supported by its member companies, research grants, and host institutions. While Microsoft's interest in adherence to HTML standards has decreased dramatically since its Internet Explorer achieved market dominance, the consortium has continued its work in the new area of XML (eXtensible Markup Language) for Web-based data exchange.

Selling things over the Web demanded much more complex server facilities than simply publishing a Web site full of documents, including code to verify and bill credit cards, update inventory and maintain user accounts. As Web sites grew in scale and became more commercial, dynamic page generation was used in many different ways. Among the most obvious were up-to-date catalogs showing the current availability of goods, and the various registration and checkout pages accompanying most online retail systems. Some sites went further. Amazon.com, for example, used it to personalize recommendations so that each user would see a unique main page when connecting to the system. Newspaper sites used it to show stories of particular personal interest, and smoothly and constantly update their sites to display breaking news and link to related stories. By the end of the 1990s, a database and set of "fill in the gaps" templates sat behind almost every commercial website of any complexity.

Although they never enjoyed a majority for the market for Web servers, commercial software producers had more success in the market for other aspects of Web software infrastructure. A host of new Web-related software niches bloomed in the mid-1990s, too numerous and specialized to deal with fully here. One of the most important was the so-called middleware server and development tools designed to sit between Web servers and databases to provide a specialized environment for the easy development of Web based business applications. The flexible and easy-to-use ColdFusion, produced by a start-up firm called Allaire, enjoyed

considerable success in this market. Other firms produced tools to search corporate Web sites, tools to integrate Web servers with existing corporate applications such as Enterprise Resource Planning systems, development aids for Web programmers, bundled systems intended to simplify the process of creating a particular kind of online application, and other pieces of software infrastructure aimed at easing the burdens imposed on Web application programmers. In many of these areas, the initial advantage held by proprietary systems has gradually been eroded by open source competitors. For example, the open source package PHP is now a standard part of the Internet infrastructure and is used by an estimate twenty million Web domains.[77]

Microsoft Discovers the Web

Even as Mosaic and Netscape took the Internet by storm, and AOL opened it to millions, the top-level managers of Microsoft remained strangely oblivious. By mid-1995, Microsoft had established itself as the world's largest software company and was crushing its last serious opposition in the markets for desktop operating systems (with Windows) and office productivity applications (with its Microsoft Office suite). Throughout 1994 and early 1995, the firm was embroiled in the Herculean task of finalizing and launching its repeatedly delayed Windows 95 operating system. This distracted its leaders from noticing the Web's new prominence.

In August 1995 tens of millions of Windows 95 cartons were finally piled high in shops around the world, accompanied by the most expensive promotional campaign in the history of the computer industry. The disks inside these cartons did not include a Web browser. Windows

---

[77]PHP Group, Usage Stats for June 2006 (2006 [cited June 29 2006]); available from http://www.php.net/usage.php.

95 did include greatly improved support for networking, and recognized the growing importance of Internet technologies by providing TCP/IP support, support for dial-up connections to ISPs, and basic utilities for Internet file transfer and remote logins. This made it much easier and more reliable to configure a computer for Internet use, as ISPs no longer needed to provide or support many of the utilities demanded by earlier Windows versions.[78] The firm had even licensed the Mosaic browser code from Spyglass, after which a handful of programmers had been allowed to slap on the Microsoft logo, rename it Internet Explorer, and tinker with it to run smoothly on Windows 95.[79] But this browser was not included on the main Windows 95 disk (which somehow managed to squeeze in such essentials as multiple versions of a video of the geek rock group Wheezer singing their hit song "Buddy Holly"). Instead, Internet Explorer was relegated to the optional "Microsoft Plus!" pack, a thirty dollar impulse buy holding a lucky dip of Space Cadet Pinball, screensavers, and other trivia. The big networking push at the Windows 95 launch was for the Microsoft Network (MSN), a proprietary on-line service modeled on the pre-Internet AOL.[80] While Internet Explored languished in obscurity, MSN programs were built into Windows 95, with a prominent icon displayed on the Windows Desktop and an invitation to sign up as part of the Windows installation procedure.

---

[78] Windows 95 components replaced the separate dialer software and TCP/IP stack (usually called winsock for Windows Sockets) needed by Windows 3.1 users. These components were hard to configure without expert knowledge. Windows 95 also included telnet, FTP and Internet email clients.

[79] David Bank, Breaking Windows: How Bill Gates Fumbled the Future of Microsoft (New York: The Free Press, 2001), 59-69. The figure of four employees working on the first version of Internet Explorer comes from Kathy Rebello, "Inside Microsoft", Business Week, July 15 1996, 56.

[80] Kathy Rebello, "Microsoft's Online Timing May Be Off", Business Week, no. 3433 (July 17 1995):41.

Behind the scenes, Microsoft had begun to work with increasing urgency to incorporate the Web into its product line. The origin of this shift is conventionally attributed to May 1995 when Bill Gates sent a now-famous memo entitled "The Internet Tidal Wave" to his fellow executives. It warned that "[t]he Internet is crucial to every part of our business… the most important single development… since the IBM PC" upon which Microsoft had built its business. As a first step, Gates ordered the shift of "all Internet value added from the Plus pack into Windows 95 as soon as we possibly can."[81] The memo, and a series of other events and reports through the rest of 1995, spread the message that integration of the Internet technologies into all Microsoft products was now the most pressing task facing the firm. Gates even delayed publication of *The Road Ahead*, a ghost written book delivering the shocking news that computers networks were about to transform the world, so that references to the Internet could be inserted alongside its lengthy discussion of the "information highway." [82]

In his memo, Gates had ordered that improved Web capabilities be "the most important element" of all new releases of Microsoft application software, demanding that "every product plan go overboard on Internet features." Microsoft quickly created a download to allow Word, its word processor, to create Web pages without the need to manually code HTML tags. In 1996, it released a free Web server for use with its Windows NT server, and acquired a promising Web site creation tool called Front Page from Vermeer Technologies, a Massachusetts-based start-up

---

[81] Bill Gates, The Internet Tidal Wave (US Department of Justice, May 26 1995 [cited May 20 2006]); available from http://www.usdoj.gov/atr/cases/exhibits/20.pdf.

[82] Bill Gates, Nathan Myhrvold, and Peter Rinearson, The Road Ahead (New York: Viking, 1995). The book appeared at the very end of 1995, more than a year behind schedule.

firm.[83] The same year it launched a new tool, NetMeeting, which was soon bundled with windows to provide online audio and video conferencing features. When it shipped the next version of Microsoft Office in 1997, the company went to great lengths to promote the new-found ability of Word, Excel, and PowerPoint to output documents as Web pages. Since then, each new version of Office has added capabilities to make it easier to share documents over the Internet, edit Web documents, collaborate with remote users, and so on. Many of these features, like a number of the special features offered by later versions of Front Page, required the use of Microsoft Web browsers and servers to function properly. In this way, Microsoft tried to use the popularity of its desktop software to build a stronger position in the booming market for corporate intranets (internal networks based around Internet technologies).

Microsoft also rushed to get into the business of providing online information and entertainment. In some ways this was a natural evolution of three of its main strategic thrusts of the early 1990s. One of these was the publication of reference and educational materials on multimedia CD-ROMS. The arrival around 1993 compact disc players, audio capabilities, and high resolution graphics as standard items on high-end personal computers had opened what seemed at the time to be a huge and growing market for interactive books. Microsoft created the widely acclaimed Encarta encyclopedia, Cinemania film review archive, and dozens of other CD-ROM titles covering topics from ancient civilizations to scary wildlife. The second thrust, to overthrow AOL and the other online services with its new MSN service, brought it into the online content business, because online services licensed existing content for their subscribers (such as material from the *New York Times*) and also produced their own proprietary materials

---

[83] The founder of Vermeer tells its story in Charles H. Ferguson, High Stakes, No Prisoners: A Winner's Tale of Greed and Glory in the Internet Wars (New York: Times Business, 1999).

and services. Microsoft was also making a major effort to enter the cable television business, guided by the belief (held widely for much of the 1990s) that the forthcoming technology of digital cable television would mean that the personal computer and television industries would inevitably merge. The rise of the Web redirected all of these efforts toward the Internet, as Microsoft launched a host of news and entertainment services. These included the MSNBC cable news channel and website (launched in collaboration with NBC), the online culture and politics magazine *Slate*, the Sidewalk series of online city guides, and the travel agent Expedia. Synergies have generally proved elusive in the media business, and this was no exception. Just as Sony had discovered that there were no real advantages to be gained by owning the movie studios and record labels that produced the films and records played on its equipment, so Microsoft found few benefits from producing the Web pages viewed with its browsers. While at the time these businesses were seen as an integral part of Microsoft's Internet push, with the benefit of hindsight it became apparent that they were strategically irrelevant and all have since been disposed of.

Java and the Browser as Application Platform

The most visible aspect of Microsoft's Internet campaign was its promotion of its Internet Explorer browser as a substitute for the hugely popular Netscape Navigator. Why was Microsoft so determined to make its own browser the standard? The answer lay in the ability of the Web browser to serve as a universal access mechanism for different on-line systems. Microsoft's domination of the personal computer operating system market had become self-reinforcing, because the strength of the Windows platform ensured that application software developers would write their exciting new programs to run on Windows, which in turn would attract new users. The Web and other Internet systems worked on open standards, however. When browsing

a website or sending an email, a Windows computer had no inherent advantage over a Macintosh or Unix computer using the same website or composing the same email. They all used the same protocols. If the exciting new applications of the future were to be delivered over the Internet to a Web browser, rather than installed and run on a personal computer, then Microsoft was in danger of losing its hard-won strategic position on the high ground of the computer industry. Gates called this a "scary possibility" in his famous memo.[84] As the dominant producer of browser software, Netscape seemed liable to seize the ground lost by Microsoft.

As Netscape captured the attention of the technology world in 1995 it began to promote itself as the next Microsoft. To make its product into a ubiquitous standard, Netscape developed browsers for no less than sixteen different operating systems. Marc Andreessen is reported to have boasted in 1995 that new advances in Web technologies would make Windows 95 into nothing more than a "mundane collection of not entirely debugged device drivers" sitting between Netscape and the underlying hardware.[85] This claim might appear rather fanciful, given the rather slow, limited, and awkward nature of Web-based systems during the mid-1990s when compared with Windows applications. At the time Web browsers could do nothing more than display simple pages and accept input data. A program written for Windows could display moving images, make sounds, perform calculations, and validate or query data as the user

---

[84] Gates, The Internet Tidal Wave .

[85] Bob Metcalfe, "Without Case of Vapors, Netscape's Tools Will Give Blackbird Reason to Squawk", Infoworld, September 1995, 111. In their book on Netscape, Cusumano and Yoffie write that Andreessen claimed that "the comination of Java and a Netscape browser would effectively kill" Windows and that Andressen's famous was remake as made "time and again." Cusumano and Yoffie, Competing on Internet Time, 40. Netscape's early determination to attack Microsoft directly is discussed in Cusumano and Yoffie, Competing on Internet Time, 114-20.

entered them. It was hardly a fair fight, and in retrospect it is one Netscape would have done better to avoid picking when its own technologies were so immature.

But fight it did, and Netscape rushed to add feature after feature to Navigator to boost its credibility as an application platform. Netscape created a simple programming language, JavaScript, to make Web pages more interactive by handling tasks such as calculating a total or verifying a date without reloading the entire page. It added support for exploration of "virtual reality spaces," an idea that, for a few months in 1995, was widely discussed as the future of the Web.

The big shock came when version 2.0 of Navigator, released in September 1995, included a then-obscure programming technology from Sun Microsystems known as Java (JavaScript and Java, confusingly, are entirely different things). Over the next few years Java received a degree of avowed support, public discussion, investor hype, and general enthusiasm never granted to a programming language before or since. Java was not only a programming language but also an application platform in its own right. Java programs were written to run one specific computer. The novel thing was that this computer was never designed to be built in hardware. Instead, an additional layer of software, known as a "virtual machine," mimicked its internal functioning. Java was a kind of Esperanto: nobody was a native speaker, but it was designed as a universal second language. Java's marketing slogan was "write once, run anywhere," because any properly written Java program was supposed to run on any properly implemented Java virtual machine, and hence to be usable without modification on any computer.

The original idea behind Java was that this would avoid having to rewrite complex interactive television software for every different kind of cable box. But by the time Java was released attention had shifted to Web browsers. At this point Sun was the leading supplier of

Web server hardware. By including a Java virtual machine in every Web browser, Netscape aimed to make Java a denominator even more common than Windows. If a company needed its customers to do something they couldn't do easily in a regular Web page, such as editing formatted text or working interactively with a technical model, it would write a little Java program known as an "applet" for inclusion on the page. When the page was loaded, this code would run in any Netscape browser whether the customer used a Windows computer, a Macintosh, a Unix workstation, or a special Internet device.

Java enthusiasts hoped that within a few years it would have spread beyond these simple tasks to become the standard platform for computer programs of all kinds. Corel, the firm that acquired the faltering WordPerfect word processor, wanted to revive its fortunes by rewriting the entire package in Java to beat Microsoft into what it expected would be a huge market for Java office applications. Opponents of Microsoft, most vocally Larry Ellison of Oracle, claimed that the spread of Java would make the traditional personal computer altogether irrelevant. Rather than a complex, crash-prone personal computer running Windows, Ellison believed that the standard device for home or office computing would be a stripped-down "Network Computer" running Java software from a server. This would save on hardware costs but, more importantly, would be very much easier to manage and configure than a traditional PC.

The Browser Wars

When Windows 95 was launched, Netscape had an estimated 75% of the browser market, and its share was still rising rapidly as its early competitors faltered.[86] The first release of Internet Explorer sank without trace, weighed down by the twin burdens of running only on the

[86] Maggie Urry, "Surfers Catch the Wave of a Rising Tide" Financial Times, August 12 1995.

newly released Windows 95 (and NT) operating systems and not being very good. As the entrenched leader in the browser market, Netscape had significant advantages over Microsoft, including the goodwill and awareness of users, a steady stream of visitors to its website looking for new downloads, and the fact that many websites had been designed to take advantage of the non-standard features present in Navigator but not in official HTML standards. This meant that some websites would not appear correctly viewed in any other browser. As Michael Cusumano and David Yoffie observed in their book *Competing on Internet Time*, Netscape had a strategy of "open but not open" with respect to Internet standards, using and supporting public standards only to a point while trying to maintain a competitive edge through its de facto control of the browser market.[87]

Yet these advantages proved little match for the might, discipline, and aggression of Microsoft. Seldom in the history of the computer industry has a large, entrenched company responded so effectively to the threat posed to its established business by a radical new technology. Microsoft had three main advantages: thousands of skilled programmers and experienced managers, enormous and assured income flows from its established businesses, and control over the technical architecture of Windows itself. It exploited all three of these advantages mercilessly in its struggle against Netscape.

Like most other personal computer software companies with popular products, Netscape had hoped to make money by selling its programs. Internet users expected to be able to download software freely, and this was the only practical way for a small firm like Netscape to get its software out quickly to millions of people. So Netscape adopted a version of the classic shareware software business model: Navigator could be downloaded without payment but

---

[87] Cusumano and Yoffie, Competing on Internet Time, 133-38.

commercial users were required to register and pay for the software after a trial period. The basic price was ninety-nine dollars a copy, with discounts for corporate licensing,[88] and shrink-wrapped copies could be purchased in computer stores.[89] But Explorer could be used by anyone for any purpose without paying a penny. By giving away its own product, Microsoft cut off Netscape's main source of income and crushed the emerging market for Web browsers. Only in 1998 did Netscape finally follow suit and make its own browser officially free for all users.

Microsoft assigned large numbers of skilled programmers to a crash program designed to boost the quality of Explorer to match Netscape. The second version of Internet Explorer, released just a few months after the first, was a big improvement and ran faster than Netscape, though it still lacked some features. Abandoning its once exclusive focus on Windows 95, Microsoft produced versions of Internet Explorer for Macintosh computers, Unix systems, and computers running the older Windows 3.1 operating system, challenging Netscape on all the major personal computer operating systems of the era. By 1996 Microsoft's Internet Platform and Tools Division comprised about 2,500 people, more than twice as large as the whole of Netscape.[90] By the time the third version of Explorer was released, in mid-1996, Internet Explorer was rated by most independent reviewers as evenly matched against Netscape in technical terms.[91] Microsoft even licensed Java from Sun, matching what was perceived as a huge advantage for Netscape. (Sun later sued, after Microsoft breached its terms by promoting an

---

[88] Netscape Communciations, Netscape Communciations Offers New Network Navigator Free on the Internet (October 13 1994 [cited May 30 2006]).

[89] Richard Karpinski, "Netscape Sets Retail Rollout", Interactive Age 2, no. 16 (June 5 1995):1.

[90] Bank, Breaking Windows: How Bill Gates Fumbled the Future of Microsoft, 68.

[91] For example, Anonymous, "Best Web Browser", PC World, August 1996, 136-9.

improved but Windows-specific version of Java).[92] Navigator retained its market dominance, but its technological edge had eroded. By the time both firms released their version 4 browsers in mid-1997 it was gone entirely. Explorer was widely judged to be faster and more stable than Navigator.[93]

Netscape's other main potential source of revenue was the sale of its extensive family of Web server products. In 1996 it shifted its strategic focus from browsers to intranet products, including Web servers, email servers and collaboration support tools for internal corporate use.[94] But in the Web server and intranet markets too the firm faced competition. Microsoft also bundled a fully functional Web server into its Windows NT business-oriented operating system (later renamed Windows 2000), though as Windows had a much less dominant position in the server market than the desktop market, the effects of this on Netscape were less pronounced.

Meanwhile, estimates of the threat to Microsoft from Java proved massively overblown. Developing large applications in Java was much harder than had been expected, leading Corel to abandon its efforts without ever getting WordPerfect to work properly. Java didn't even live up to expectations as a replacement for HTML in building interactive website interfaces. Including a

---

[92] Jeff Pelline, Alex Lash, and Janet Kornblum, Sun Suit Says IE Failed Java Test (News.com, 1997 [cited September 20 2006]); available from http://news.com.com/2100-1001_3-203989.html. Microsoft eventually agreed to settle this and a range of other complaints by Sun with a the payment almost two billion dollars and an agreement to work cooperatively. Stephen Shankland, Sun Settles with Microsoft, Announces Layoffs (News.com, April 2 2004 [cited September 20 2006]); available from http://news.com.com/2100-1014_3-5183848.html.

[93] For examples of head-to-head comparisons see Mike Bielen, "Battle of the Web Browsers: Part Two", Chemical Market Reporter, October 13 1997, 16 and Yael Li-Ron, "Communicator vs. Internet Explorer. And the Winner is." PC World, 15 December 1997, 206-09.

[94] Cusumano and Yoffie, Competing on Internet Time, 28-34.

Java program as part of a Web site caused a long download for dial-up Internet users, and it would not function correctly except with a specific browser version and exactly the right Java release. Differences between different Java implementations, and the frequent release of new and incompatible versions, made the promise of universal compatibility a cruel joke. The cost of personal computer hardware dropped so rapidly that few customers ever felt tempted by the specialized "Internet appliances" sold for email and Web browsing, when for a few more dollars they could have a general purpose PC with much more flexibility. Sony, 3Com, Compaq and other major firms all suffered abject failure when they attempted to sell Internet appliances. Network computers were a little more successful in the corporate world, but remained a niche item for the kinds of jobs formerly performed on less powerful video terminals rather than a general-purpose replacement for the personal computer. This may not have been entirely inevitable. Had Microsoft not deliberately targeted Netscape's control of browser design, and taken steps to weaken Java's cross-platform compatibility by introducing its own extensions, then perhaps things would have gone differently.

The third of Microsoft's advantages, control of the desktop operating system market, was perhaps the most important. It seized upon its control of the operating system installed on desktop computers as a chance to push its own browser to the exclusion of Netscape. Its aggressive action here led to complaints from Netscape, and in turn to a major antitrust case filed against Microsoft by the United States Department of Justice in 1998. One of the major issues in this case concerned Microsoft's domination of software distribution channels. From 1996 onward, it was impossible to purchase a Windows computer without Internet Explorer preloaded onto it. At the same time, Microsoft threatened major manufacturers of personal computers with the withdrawal of their licenses to sell Windows if they did not remove Netscape from the

software bundles supplied with new computers. Microsoft also persuaded AOL, CompuServe, and Prodigy to make Internet Explorer the only browser distributed on their installation disks. So keen was Microsoft to capture the browser market that it was willing to sacrifice MSN's special place on the Windows desktop to do so. It agreed to supply software from these rival services with each copy of Windows, and even to promote them with icons placed on the desktop of each new computer.[95] People could continue to download and install Navigator themselves, but these moves ensured that Netscape was denied the two most important channels for the distribution of browser software: pre-installation on new computers, and inclusion in the promotional mailings with which AOL and its competitors saturated American households.[96]

The other main accusation against Microsoft in the antitrust case was that it had illegally bundled its Explorer browser, with which Netscape was competing, with the Windows operating system. United States law draws a sharp distinction between the competitive practices tolerated from monopolists and those acceptable for ordinary firms. Bundling two products together is normally legal, but a firm with a monopoly in one market is forbidden from using bundling to wipe out competition in another market. In this sense Microsoft's browser was not really free to its users. The thousands of employees in its Internet division still had to be paid, after all. Microsoft supported their work through its monopoly profits from the desktop operating system

---

[95] The deals were reported in Rose Aguilar, CS, Prodigy Cozy Up to MS Too (News.com, March 13 1996 [cited September 19 2006]); available from http://news.com.com/2100-1023-207369.html. The implications of the AOL deal for Netscape were explored in Cusumano and Yoffie, Competing on Internet Time, 111-8.

[96] Microsoft's behavior here was eventually judged not to breach antitrust law, on the grounds that Netscape had still been able to distribute million of copies of its browser. Ken Auletta, World War 3.0: Microsoft and its Enemies (New York: Random House, 2001), 200-02, 363.

market, effectively making the cost of Internet Explorer a tax paid on every Windows computer sold whether or not its buyer wanted Explorer. This appeared to put Netscape at a disadvantage.

Microsoft's push to tie Explorer to Windows went far beyond joint distribution. In 1996, Microsoft had been encouraging Web developers to include small Windows-only programs, dubbed "ActiveX controls" in their Web pages as a way of adding interactive features. Internal battles within Microsoft during 1997 destroyed the autonomy of the Internet team, subordinating its single-minded push to create strong Internet tools to a new strategy based on merging the browser into the heart of the Windows user interface.[97] With the launch of Windows 98, a modestly improved version of Windows 95, Microsoft made it technically impossible to remove Internet Explorer without disabling Windows itself. Its most hyped feature was Active Desktop, a buggy, cumbersome, and little-used method to display constantly updated Internet content on one's desktop. [98] Microsoft insisted that there were obvious, overwhelming technological reasons for this integration, but few outside the company found this claim credible. One technology reporter wrote that he had "been avoiding the subject of the Active Desktop because frankly, I don't really get it. I don't use it, I don't know anyone who does use it and I don't think it works very well for those who do use it."[99] Another suggested that it "sticks its tendrils into every part of your operating systems and will wreak havoc on your computer" with a "nasty habit of crashing every thirty minutes." A third wrote of the new features "don't like 'em, don't want 'em, don't need 'em and I say, to heck with 'em'."[100]

---

[97] Bank, Breaking Windows: How Bill Gates Fumbled the Future of Microsoft, 78.

[98] Ibid, 88-89.

[99] Paul Smith, "Bringing the World Wide Web Right to your Desktop" Scotsman, August 12 1998, 12.

[100] Myles White, "'Explorer' closes gap" Toronto Star, September 4 1997, J3.

Microsoft's actions were motivated by a desire to turn the disruptive power of the Web into nothing more than an extension of its proprietary world of Windows. Its insistence that Web browsers and operating systems were conceptually indistinguishable became the cornerstone of its legal defense in the case, as well as an expression of Gates's personal determination to maintain Windows as the centerpiece of Microsoft's business. If the browser was merely a legitimate and necessary feature of Windows, rather than a separate and bundled product, then nothing illegal had taken place. This was particularly important for Microsoft to prove, since earlier instances of monopoly abuse had led it to sign a 1994 consent decree promising not to tie any of its other products to the sale of Windows. The case progressed badly for Microsoft, as its executives struggled to explain away the powerful evidence mustered against them. Gates refused all compromise with the government but tarnished his reputation as a technical and business genius with testimony during which he quibbled over linguistic nuances while denying all knowledge of his own emails and memos.[101] So disastrous was his performance during the trial that in 2000 that he was quickly eased out as CEO in favor of Steve Ballmer, stepping back from business management to a new position as Chief Software Architect in which his day-to-day involvement in the business of the firm was gradually diminished in favor of charitable work.[102]

---

[101] A great deal was written at the time on the Microsoft case, as newspapers and magazines followed every twist and turn. A closely observed account is given in Auletta, World War 3.0: Microsoft and its Enemies.

[102] Bank, Breaking Windows: How Bill Gates Fumbled the Future of Microsoft, 159. Bank's account is well sourced, and agrees in substance, if not tone, with the Microsoft-authorized account of the Gates-Ballmer transition and division of labor given in Robert Slater, Microsoft Rebooted: How Bill Gates and Steve Ballmer Reinvented Their Company Portfolio, 2004), 51-71. Slater writes that a key motivation behind the transition was to remove the polarizing figure of Gates after the antitrust debacle, and that during attempts to settle the case "better a

In April 2000 Microsoft lost the case. It was ruled to have used anti-competitive means to maintain a monopoly in the market for personal computer operating systems, to have illegally tied Internet Explorer to its Windows monopoly, and to have destroyed the browser market by adopting predatory pricing to crush Netscape and stifle competition.[103] The trial judge, Thomas Penfield Jackson, wrote that "Microsoft paid vast sums of money, and renounced many millions more in lost revenue every year, in order to induce firms to take actions that would help enhance Internet Explorer's share of browser usage at Navigator's expense." Given that Microsoft had promised never to charge for its browser, Jackson continued, this could "only represent a rational investment" if its purpose was to protect the Windows monopoly from the emergence of a rival applications platform.[104] After Microsoft failed to reach an agreement with the government during mediation Jackson granted the Department of Justice's request that the firm be broken into two separate businesses: one confined to operating systems, and the other to applications.[105]

Jackson's ruling on monopolization was upheld on appeal, but the breakup was deferred for consideration of alternative remedies. Microsoft was then able to negotiate a much less dramatic settlement with the newly installed Bush administration by promising to fully document

---

clever, less emotional deal maker such as Ballmer in charge at this crucial time than someone who blindly believed his company has done no wrong." (p. 61) According to plans made public in 2006, Gates will complete his retirement from Microsoft management by 2008, handing over his remaining technical responsibilities to Ray Ozzie.

[103] Joel Brinkley, "U.S. Judge Says Microsoft Violated Antiturst Laws with Predatory Behavior" New York Times, April 4 2000, A1.

[104] Thomas Penfield Jackson, "Excerpts From the Ruling That Microsoft Violated Antitrust Law" New York Times, April 4 2000, C14.

[105] Joel Brinkley, "Microsoft Breakup is Ordered for Antitrust Law Violations" New York Times, June 8 2000, A1.

the techniques needed to write applications for Windows, make it easier for customers to remove desktop icons for Microsoft products or set competing products as defaults, and stop blackmailing computer makers into dropping products written by its competitors.[106] These measures amounted to a reprieve for the firm, after which Ballmer moved to settle the other major lawsuits against Microsoft. A private suit brought by Netscape on the basis of the antitrust findings was settled for $750 million in 2003, and a group of suits brought by Sun for almost two billion dollars in 2004.[107] A separate antitrust case launched against Microsoft by the European Union proved harder for the company to contain. Again the issue was anticompetitive bundling. After failing to comply with the terms of a 2003 order to fully document the behavior of its server products (to make it easier for competitors to achieve compatibility) and offer a version of Windows without Windows Media Player the company has been hit with a series of fines. As of October 2006 the firm had still not satisfied the European Union and had accumulated more than a billion dollars in fines while it continues to appeal.[108]

---

[106] The Bush administration's proposed settlement was given in United States of America, Revised Proposed Final Judgement (November 6 2001 [cited 2006 May 05]); available from http://www.usdoj.gov/atr/cases/f9400/9495.htm. Reaction to the settlement by those who had been following the trial was generally negative, for example Lawrence Lessig, "It's Still a Safe World for Microsoft" New York TImes, November 9 2001, A27. Nine of the eighteen states involved in the case continued to press their case separately, but with little success.

[107] Ian Fried and Jim Hu, Microsoft to Pay AOL $750 Million (News.com, May 29 2003 [cited August 28 2006]); available from http://news.com.com/2100-1032_3-1011296.html and John Markoff, "Silicon Valley Seeks Peace in Long War Against Microsoft" New York Times, April 4 2004, 1.

[108] Paul Meller and Steve Lohr, "Regulators Penalize Microsoft In Europe" New York Times, July 13 2006, C1.

But even as Microsoft's legal problems mounted during the late 1990s it was winning its commercial struggle against Netscape. By 1998 market research reports were suggesting that the usage of Internet Explorer had almost pulled level with that of Navigator.[109] Many large companies were standardizing on Explorer.[110] Meanwhile, Netscape's software developers found it increasingly difficult to keep up as they tried to match Microsoft feature for feature, support numerous versions of Navigator for different operating systems, tidy up Navigator's increasingly unwieldy code base, and supplement it with email, news, Web page editing and calendar functions. Netscape's efforts eventually foundered entirely, as programmers working on release 5.0 of Navigator labored for several years before eventually giving up. In late 1998 a diminished Netscape agreed to be acquired by AOL, consolidating two of the best-known Internet firms of the era.

Throughout the late 1990s AOL invested heavily in Internet technologies. The acquisition of Netscape, for stock worth nearly ten billion dollars when the deal was consummated, seemed at the time to make AOL into a major force in the world of Internet software. [111] A joint venture with Sun, the developer of Java and a leading supplier of Web servers, was quickly announced

---

[109] Paul Festa, Study: Netscape Share Below 50% (News.com, September 29 1998 [cited September 15 2006]); available from http://news.com.com/2102-1023_3-216043.html?tag=st.util.print.

[110] Erich Luening and Janet Kornblum, Study: Navigator Ahead in Office (October 15 1998 [cited September 16 2006]); available from http://news.com.com/2100-1023_3-216743.html.

[111] The figure of $4.2 billion was widely quoted as Netscape's price at the time of the merger. However, AOL's stock rose rapidly during the next few months, and at the point the deal closed the stock it exchanged for Netscape was worth $9.6 billion on March 17 when the deal finally closed after approval by the Department of Justice and Netscape shareholders. Shannon Henry, "AOL-Netscape Merger Official; $9.6 Billion Deal Completed After Last Hurdles Are Cleared" Washington Post, March 18 1999, E03.

with a promise that the alliance would become "the dot com software company." [112] Contemporary reports suggested that the new alliance would focus on developing the Internet as a retail environment, pushing the spread of Internet commerce to small businesses and onto wireless devices and cable television boxes.[113]

Whatever plan drove the acquisition was quickly rendered irrelevant by one of the great, all-time disasters in the history of business. Despite its push into the Internet software business, AOL ultimately decided that its future was as an entertainment company. In January 2000 it acquired Time Warner to form what the deal's architects hoped would be the perfect marriage: a virile young firm controlling tens of millions of home Internet connections spliced to a creaking but powerful media empire churning out vast quantities of "content" just waiting to be digitally delivered to eager consumers. Though AOL was a much smaller company, AOL shareholders received a majority stake in the new firm, dubbed AOL Time Warner, because theirs was perceived as a fast-growing business.

Before the new firm could even finish building its stylish New York City headquarters, AOL's advertising base collapsed and its users began to defect to high speed alternatives. Within two years of the merger all the senior AOL executives had been purged, and in 2003 the combined firm had dropped AOL. Meanwhile, the collapse of the formerly booming market for

---

[112] Kim Girard and Tim Clark, AOL, Sun Detail Software Tie-Up (News.com, March 30 1999 [cited September 15 2006]); available from http://news.com.com/2100-1023-223687.html.

[113] Louise Kehoe and Roger Taylor, "AOL to Power up E-commerce with $4.2bn buy" Financial Times 1998, 31.

Internet stocks and services crippled the prospects of AOL's alliance with Sun.[114] Starved of resources, support, and leadership Netscape's browser withered with astonishing speed. The eventual arrival of Navigator version 6.0 (in 2000) and 7.0 (in 2002), merely accelerated the collapse, as Netscape die-hards discovered their long-awaited upgrades were sluggish, prone to crashing, and labored under a clumsy user interface. Navigator's estimated market share, close to 50 percent at the time the acquisition was announced in late-1998, plummeted to well under 5 percent over the next four years.[115] Time Warner never even shifted users of its own AOL client software to the Netscape browser. In 2003, it closed its Netscape division and eliminated the few remaining employees. It had, however, finally found a use for the Netscape brand: as the name of a cut-price dial-up Internet service. The name that just a few years earlier had been attached to one of the most spectacularly successful pieces of software ever created appeared to have met a sad end.

The Web Grows Up

By the time Microsoft edged past Netscape in its share of the browser market it had a faster, more stable, and more stylish browser. Furthermore, Microsoft gave its browser away, whereas Navigator cost money. Were customers really hurt by Microsoft's triumph, whatever its

---

[114] A nice treatment of the AOL-Time Warner merger is given in Nina Munk, Fools Rush In: Steve Case, Jerry Levin, and the Unmaking of AOL Time Warner (New York: HarperCollins, 2004).

[115] Browser estimates for the late 1990s are consolidated in Maryann Jones Thompson, Behind the Numbers: Browser Market Share (CNN.com, October 8 1998 [cited September 18 2006]); available from http://www.cnn.com/TECH/computing/9810/08/browser.idg/. A survey released in late 2002 gave Netscape, and its open-source derivatives, just 3.4% of the browser market. Matthew Broersma, Tech doesn't buoy Netscape browser (News.com, August 28 2002 [cited September 19 2006]); available from http://news.com.com/2100-1023_3-955734.html.

legal transgressions? Although the question is hard to answer, a look at events since 1998 is highly suggestive. Version 5 of Internet Explorer was released in 1999, as the Netscape threat had begun to recede. Following this, work to enhance the browser almost stopped. Internet Explorer for Macintosh saw no significant improvements after 2000, and was officially discontinued three years later.[116] The Unix version of the browser was also terminated, and even the Windows version stagnated. Version 6, released along with Windows XP in 2001, added only minor improvements. In 2003 Microsoft announced that future browser versions would be made available only with the purchase of new versions of Windows (mirroring its claim during the antitrust trial that the Web browser was an operating system feature rather than a separate application).[117]

During its trial, Microsoft had argued that an enforced split between its browser and operating system businesses, as sought by the government's lawyers, would hurt its ability to delight consumers through rapid and coordinated innovation in both areas. Gates repeatedly demanded that Microsoft be given "freedom to innovate and improve our products" contending that "continuing to evolve the Internet technologies into Windows is great for consumers…"[118] The firm even set up a self-proclaimed "grassroots" organization, the Freedom to Innovate Network, complete with newsletters and press releases, to help satisfied customers lobby their

---

[116] Anonymous, "Microsoft to end Mac Browser Versions" Chicago Sun Times, June 14 June 14, 30.

[117] Paul Festa, Microsoft's Browser Play (News.com, June 4 2003 [cited September 29 2006]); available from http://news.com.com/2100-1032_3-1012943.htm.

[118] Bill Gates, Bill Gates' Antitrust Statement (ZDNet News, December 6 1998 [cited September 12 2006]); available from http://news.zdnet.com/2100-9595_22-513040.html.

elected representatives on its behalf.[119] Yet with the case settled on terms favorable to Microsoft, the promised stream of enchantments has somehow failed to find its way into the hands of users. In August 2006, Internet Explorer marked a remarkable five years without an increment of its version number from 6.0. The next month, Windows celebrated the fifth anniversary of Windows XP (also known as version 5.1) as the company's flagship personal computer operating system. Every personal computer sale added a little more money to Microsoft's famously ample stash of cash and short-term investments (which has fluctuated between thirty- and fifty-billion dollars in recent years) but little in the way of new technology was coming out of Redmond, Washington in return. Furthermore, Microsoft's rush to integrate the Web browser deep into the heart of the operating system and allow ActiveX programs to run within it, actions spurred at least in part by legal and business considerations, had the nasty side effect of opening many vulnerabilities to worms, viruses, and spyware.

The browser story has a twist. In 1998, as Netscape's browser business unraveled and its corporate strategy shifted toward the Web-portal market, its management had gambled on the novel but fashionable strategy of open source development as a way to leverage its limited resources. The full source code for its browser suite was released over the Internet, in the hope that this might lead to a flood of variant and improved versions to counter the threat from Microsoft. When the initial excitement wore off little seemed to be happening. After its acquisition by AOL, Netscape continued to support the open source project, dubbed Mozilla, but the results were disappointing. The doomed Navigator versions 6.0 and 7.0 were based on the

---

[119] The Freedom to Innovate Network's claims to be non-partisan and grass-roots are reported in Adam Cohen, ""Microsoft Enjoys Monopoly Power."", Time, November 15 1999. The group still exists as a hub for Microsoft's lobbying efforts, though it now bills itself as "Microsoft's Freedom to Innovate Network."

Mozilla code. Incremental updates improved both somewhat, but Mozilla-based browsers won support only from the most dedicated of open source enthusiasts.

Then in 2004 word suddenly began to appear in national publication such as the *Wall Street Journal* and *New York Times* of a new browser: Firefox.[120] Version 1.0 was officially released in November, but even before that the program won a large and loyal following. Firefox was based on the Mozilla code, which because of its open source license was free for anyone to modify or improve. Although the Mozilla project had produced some excellent code its own browser remained clumsy to use and was burdened with a mass of extraneous features. In 2002, a Netscape programmer and a teenager working as an intern had resolved to create a new version in which the browser was pared down to its essentials and given a stylish yet functional interface. [121] The project gained steam rapidly, recruiting experienced developers, and eventually became the flagship product of the Mozilla project. Its appearance coincided with a rash of security attacks exploiting Internet Explorer's tight integration with the Windows operating system, giving ordinary users a convincing reason to shift. Firefox boasted other simple but useful features missing from Explorer, particularly the ability to open several pages within a single browser window and shift rapidly between them. Its capabilities have been extended by dozens of add-on programs, customizing the browser to do things like block Internet advertisements.

Firefox became the first open source desktop computer application to win widespread usage by Windows and Macintosh users. Within a year of its first official release, its share of the

---

[120] Byron Acohido, "Firefox Ignites Demand for Alternative Browser" New York Times, November 10 2004, 1B and Walter S Mossberg, "Security, Cool Features of Firefox Web Browser Beat Microsoft's IE" Wall Street Journal, December 30 2004, B1.

[121] Josh McHugh, "The Firefox Explosion", Wired 13, no. 2 (February 2005).

worldwide browser market had risen above 10% in many surveys. [122] Microsoft invested an enormous amount of time and money in redesigning Windows to eliminate security holes, hoping to match Firefox's reputation as a safer browser. Firefox even rekindled browser development at Microsoft, which quickly started promoting the new features planned for version 7.0, the first significant enhancement in five years. Reversing its previous announcement, Microsoft revealed that Internet Explorer 7.0 would be available as a download for users of its existing Windows XP operating system as well as being bundled with its the forthcoming Vista version of Windows. Early reviewers found the new release welcome but unspectacular, suggesting that Microsoft had successfully duplicated many of the features pioneered by its competitors over the past five years but included few new ideas of its own.[123]

Meanwhile, a number of successful and innovative Web-based services have grown up in the past few years. By 2005, these were being collectively referred to as "Web 2.0," a term that must owe some if its success to its conspicuous vagueness. As defined by technology publishing veteran Tim O'Reilly, whose company was initially responsible for spreading the phrase, Web 2.0 involves a new wave of Web services based around user communities and collective tagging efforts (such as Wikipedia and YouTube). But the term also signifies new Web technologies to provide users with rich, interactive user interfaces comparable to those found on traditional desktop applications. This rekindled the enthusiasm for the Web browser as an application-delivery platform seen by Java enthusiasts during the mid-1990s. Like earlier bursts of Web

---

[122] Ibid.

[123] For example Paul Taylor, "The Leading Windows on the Web Receive a Fresh Polish" Financial Times, July 7 2006, 9 and Walter S. Mossberg, "Microsoft Upgrades Internet Explorer -- But Not Much is New" Wall Street Journal, October 19 2006.

creativity this is an example of what Eric Von Hippel has called “user innovation,” where technological change takes place not as the result of a strategic vision of a company building a new product but through the creative actions of their customers in reconfiguring and combining existing products.[124]

Google applications such as Gmail and Google Maps helped to demonstrate and popularize the ability of Web browsers to support attractive, powerful and responsive user interfaces. This was not the result of any single technological breakthrough, but instead rested on the creative combination of a range of existing capabilities to snappily update the data shown on webpages without reprocessing and reloading the entire page. One popular blend of technologies has been dubbed Ajax, for Asynchronous JavaScript And XML. Most of the underlying technologies (formerly called Dynamic HTML or DHTML) first appeared in browsers toward the end of the so-called Browser Wars as Microsoft and Netscape competed to add new features as rapidly as possible. Over the last few years, as high-bandwidth Web connections spread, browser bugs and incompatibilities were solved or worked around, and programmers came to grips with the nuances of these techniques, Web developers have found ways to achieve many of the objectives originally set for Java without having to deal with its weaknesses. Online applications have some real advantages in some areas, including access to large databases, the elimination of the need to configure or install software on the user’s computer, and easy data sharing between users. Google now offers a reasonably powerful online spreadsheet and word processing programs accessed through a Web browser. While unlikely to dent the sales of Microsoft’s Office suite, these applications do offer easy sharing of documents with other users, collaborative editing, secure backup, and easy accessibility of your saved documents from any

[124] Eric Von Hippel, Democratizing Innovation (Cambridge, MA: MIT Press, 2005).

computer with a Web browser. Many users are choosing online applications both for personal tasks such as tax preparation and, with the success of Salesforce.com, corporate applications such as customer relationship management

The push by Netscape, Sun, and Oracle to use Java as a direct and aggressive challenge to Windows failed miserably a decade ago. But even without major updates to Internet Explorer, the Web has gradually matured as an application platform. Few new companies have tried to enter the market for stand-alone desktop computer applications since the mid 1990s, and the markets for software such as word processors, personal finance software, and antivirus utilities are all dominated by a handful of well entrenched companies. Excitement, innovation, and investment dollars have long since shifted to Web based applications, with community networking sites such as You Tube and Facebook providing the latest round of hits.

Within the software industry, Google is widely seen as the next challenger to Microsoft's domination of the market for operating systems and office applications. Given Google's refusal to publicly outline any such strategy this may just be wishful thinking on the part of Microsoft-hating bloggers, but analysts speculate that its push into online applications and the aggressive hiring of gifted programmers portends an eventual showdown. Microsoft's leaders have also described Google as their biggest threat, perhaps recognizing the usefulness of an external enemy in motivating their own employees. Gates, his own youthful swagger long gone, said of Google, "they are more like us than anyone else we have ever competed with."[125] The Web's growing maturity as an application delivery platform has not yet hurt Microsoft's profitability or its stranglehold on its core markets, but may slowly be eroding the strategic value of that stranglehold.

---

[125] Fred Vogelstein, "Search and Destroy", Fortune, May 2 2005.

CONCLUSIONS

The Web and email technologies of the early 1990s were profoundly and directly shaped by the environment in which they had been produced. They were excellent solutions to the problem of a relatively small and homogeneous network of academic researchers. Nobody designing systems to be used by billions of people for commercial purposes would have designed anything remotely like them. Yet in a few short years, the Internet and its application protocols became the foundation for the first and only universal, commercial, and global computer network in human history. The funny thing was that their rapid success in the commercial environment had a lot to do with the very characteristics that no commercial designer, manager, or consortium would have produced: simplicity, flexibility, openness, decentralization, the omission of billing and accounting mechanisms, and so on. In particular, the Web offered itself to businesses and enthusiasts as a quick, easy, and widely accessible platform on which to establish their online publications and applications. Current battles to protect "net neutrality" (the principle that Internet Service Providers treat data from all Internet sites equally rather than favoring particular business partners) reflect the influences of principles built into TCP/IP long before the Internet was opened to business.

The Internet's strengths and weaknesses have profoundly influenced the evolution of Internet businesses, making some niches (such as commercial publishing) hard to fill and creating others (such as search engines) that did not exist on other online systems. The Internet's fundamental approach of building new systems and standards on top of old ones made it easy to extend it to new applications, including many commercial ones, but very hard to go back and change the underpinning infrastructure on which the new applications rested.

One of the most distinctive features of the pre-commercial Internet was its reliance on open standards and protocols, published and free for all to use. Its cultures, values, and practices shifted rapidly during the 1990s as the Internet commercialized. In some ways, however, they proved surprisingly resilient. Netscape remade the Web as a commercial platform, adding features such as credit card encryption and cookies, but it did so largely by building on existing standards and establishing its own innovations as new standards. Even Microsoft was initially forced to follow the same pattern to establish its own browser as a credible alternative. Since Microsoft's victory in the Browser Wars the rate of changed has slowed, but because Web browsers must continue to work with a wide variety of servers Microsoft has not been able to remake the Web as a closed system. The Web continues to be defined by open standards rather than proprietary code. Meanwhile, the focus of innovation has shifted from Web browsers themselves to Web applications, as the new approaches referred to as Web 2.0 evolved around new uses for clusters of existing technologies. The Web remains a platform for user-driven innovation, far beyond the control of any one company.

More than anything else, the Internet was a collection of protocols, and these protocols have continued to shape its evolution long after their designers lost institutional authority over the network. Over the past fifteen years, many changes have been made to the software technologies of the Internet, as the Web has been remade as a secure platform for online commerce. But the fundamental characteristics of protocols such as TCP/IP, SMTP, and HTTP continue to shape our daily lives. Whole industries have emerged to tackle problems such as email spam, yet without fundamental changes to the software architectures inherited from the pre-commercial Internet, these woes can be alleviated but not cured. In some ways today's Internet is crippled by the very design choices that provided its initial strengths. In any event, the

software technologies whose evolution I have described here formed the sometimes shaky platform on which companies of all kinds built the different types of Internet business discussed in the rest of this book.